Multi-Objective Air Traffic Flow Management Through Lexicographic Optimisation

Ramon Dalmau, Gilles Gawinowski & Jacek Kopec
European Green Sky Directorate
EUROCONTROL
Brétigny-Sur-Orge, France

Abstract—**Air transportation is challenged by rising air traffic demand and limited capacity in airspace and airports. This demand-capacity balance problem is inherently multi-objective, with sometimes conflicting goals. These objectives include minimising air traffic flow management delays, mitigating the operational impacts of reactionary delays, reducing environmental impact, etc. Additionally, specific regions of the network at particular times, referred to as "spots" in this paper, must be considered when they exhibit specific characteristics. This paper introduces a comprehensive mixed-integer programming model for the demand-capacity balance problem, addressing its multi-objective nature through a lexicographic approach. In this approach, the multiple objectives are ranked by importance, and the problem is solved sequentially to minimise each objective without exceeding the optimal value of the previous objectives. The versatility and effectiveness of the lexicographic approach in air traffic flow management are demonstrated using historical air traffic data over France and Spain.**

Keywords—**Air traffic flow management; lexicographic optimisation; spot control**

I. Introduction

The skies are busier than ever, with the demand for air traffic soaring to unprecedented levels. According to the International Air Transport Association (IATA), the number of global air passengers is expected to reach 8.2 billion by 2037 [1]. Amid this growing demand, air traffic flow management (ATFM) plays a crucial role in ensuring that the increasing number of flights are accommodated safely by a limited capacity in airspace and, particularly, airports. Addressing this demand-capacity problem is no small feat, as it involves a range of often conflicting objectives, from minimising ATFM delays and mitigating the cascading effects of reactionary rotational delays to reducing environmental impact, among others [2].

Traditional models for optimising ATFM delays and/or other ATFM flow measures, like rerouting or level capping, often struggle to manage multiple objectives concurrently due to their reliance on combining these objectives into a single function using a weighted sum [3]. These models face challenges in assigning appropriate weights to different objectives, especially when the objectives are measured in diverse units (e.g., minutes, kilograms). Furthermore, it becomes unclear whether the optimal solution genuinely reflects the decision-makers' preferences or is distorted by the chosen weights.

This paper addresses the multi-objective nature of the problem with a comprehensive mixed-integer programming (MIP) model for ATFM that leverages lexicographic optimisation [4]. Lexicographic optimisation involves ranking objectives by their importance. The optimisation process begins with the highest-priority objective, fully optimising it before addressing the next objective, and so forth. This sequential approach ensures that the most critical goals are not compromised by the need to balance them against less critical ones. For instance, if minimising ATFM delay is deemed the highest priority, the model will first optimise for this goal. Once the best possible solution for minimising ATFM delay is found, the model then optimises the second priority objective, such as fuel consumption, without altering the optimal ATFM delay.

Furthermore, the proposed model also considers specific regions of the network during certain periods, referred to as "spots", which require special attention due to their unique characteristics. For instance, in a *protection* spot, it is crucial to minimise additional traffic to prevent overloads, especially during times of uncertain air traffic demand. The effective management of protection spots helps to prevent localised congestion that can ripple out to affect the broader network. The various types of spots considered by the proposed model will be discussed in greater detail in the subsequent sections.

Another innovative aspect of the proposed model is its integration of frequently overlooked objectives and constraints to enhance practical implementation in operational settings. For example, the model encompasses the capability to minimise the difference of ATFM measures across successive solutions, monitor occupancy counts alongside entry counts, and promote traffic smoothness to mitigate peak congestion. Additionally, it takes into account data currently or potentially accessible to the European Network Manager (NM), which could enhance its relevance and applicability beyond the research context.

II. Background

This section provides information on European ATFM and the lexicographic approach to optimising multiple objectives.

A. *European Air Traffic Flow Management*

The role of ATFM involves assigning optimal ground delays and, in some cases, flight plan alternatives to each flight. These flight alternatives can include adjustments in vertical (i.e., flight level) and/or lateral (i.e., sequence of waypoints) dimensions. The primary goal of the ATFM measures is to minimise ground delay and fuel consumption while ensuring that airspace and airport capacities are not exceeded.

In Europe, the most common ATFM measure consists of limiting the rate at which aircraft enter a congested traffic volume (TV)[1] during a given period of time, referred to as TV-period hereafter, i.e., to activate a regulation. The flights subject to one ore more regulations are issued with a ground (ATFM) delay that is assigned on a *first-come, first-served* basis by the Computer-Assisted Slot Allocation (CASA) system [5], so that the maximum entry rate in the regulated TVs is not exceeded during the regulated period of time. Currently, local flow management positions (FMPs) are responsible for detecting potential overloads and determining regulations in coordination with the NM. In addition to regulations, minor ground delays, flight level capping, and rerouting of a small number of flights via short-term ATFM measures (STAMs) are among the solutions used in Europe to mitigate overloads.

In the research community, the demand-capacity balance challenge is commonly modelled as a MIP problem, where a potential solution involves the allocation of ground delays and/or the selection of alternative flight plans [6], [7].

B. *Lexicographic optimisation*

The lexicographic optimisation Algorithm 1 begins by initialising the feasible set to include all potential solutions. For each objective, ranked by importance, the algorithm performs the following steps: (1) solve the current optimisation problem to find the optimal value of the objective, considering only the solutions from the current feasible set; and (2) update the feasible set to include only those solutions that achieve this optimal value. This updated feasible set is used in subsequent iterations for optimising the next objective function. The process continues until all objective functions have been optimised in sequence. The final feasible set contains solutions that are optimal with respect to all objective functions, in the specified order of priority. Any solution from this final set can be chosen as the lexicographical optimal solution, ensuring that it satisfies the hierarchical optimisation criteria.

Algorithm 1 Lexicographic optimisation

Require: $F = \{f_1, f_2, \dots, f_k\}$: Ranked objectives
X: Feasible set of solutions
Ensure: Optimal solution x^*
1: Initialise feasible set $S_0 = X$
2: **for** $i = 1$ to k **do**
3: Solve $\min_{x \in S_{i-1}} f_i(x)$ to get optimal value f_i^*
4: Update feasible set $S_i = \{x \in S_{i-1} \mid f_i(x) \leq f_i^*\}$
5: **end for**
6: $x^* =$ any element of S_k
7: **return** x^*

It should be noted that, in line (4) of Algorithm 1, some absolute and/or relative tolerance ϵ could be allowed to degrade the optimal objective and increase the flexibility for the following objectives, e.g., $S_i = \{x \in S_{i-1} \mid f_i(x) \leq f_i^* + \epsilon_i\}$.

[1]A traffic volume is related to a single geographical *entity* (either an airport, a set of airports, an airspace sector or a point), and may consider all traffic passing through that entity or only specific flows.

III. Literature review

Approximately 30 years ago, [5] established a comprehensive MIP model for optimal ground delays to enhance the heuristic (first-planned, first-served) algorithm used by CASA, which is still in use today. Their results indicated that optimisation methods could reduce ATFM delays by approximately one-third. Building on this, [6] also incorporated the assignment of optimal routes in addition to ground delays, demonstrating that for small instances, this problem could be solved relatively quickly using standard optimisation solvers.

Remarkable advances in optimisation solvers have enabled the solution of large instances in reasonable amounts of time, and over the last decade, thanks to a very active research community, a new wave of model variations with new features has emerged. For instance, [8] proposed a model that also minimises delay propagation to subsequent flights, simultaneously increasing flight adherence to departure slots at coordinated airports. In parallel, [9] proposed a model that can be solved in short computation times for large instances during the strategic phase; and a variant that also considers aircraft rotations through the turn-around time constraints was proposed by [10].

Few years later, [11] proposed a collaborative ATFM framework aiming to improve the cost-efficiency for airspace users, which also considers alternative trajectories with rerouting and/or level capping in addition to ground delay. In a more recent study, [12] examined the flexibility of each flight in terms of feasible delay windows within a solution. A convenient way of presenting the results to FMPs was proposed in a complementary publication [13]. Last but not least, in [14] the authors compared the performance of the MIP-based approach to current mechanisms for ATFM delay allocation.

Given the vast number of models developed for the demand-capacity balance problem and the extensive history of research in this area, this section could easily span several pages. However, three common areas for improvement persist across most models: (1) addressing the multi-objective nature of the problem; (2) accommodating operational constraints, such as occupancy (peak and sustained) counts and traffic smoothness; and (3) maximising the stability of consecutive solutions when the problem is solved on a regular basis (e.g., every 5 min).

In addressing the first area for improvement, it is crucial to acknowledge that the use of lexicographic optimisation in air transportation applications is not a novel concept. [15] employed a lexicographic approach to optimise departure procedures, considering multiple objectives such as noise and fuel consumption. Similarly, [16] utilised this method to enhance airport scheduling interventions, incorporating inter-airline equity objectives. Additionally, [17] applied the lexicographic approach to reschedule arrivals and departures, focusing on efficiency and equity. The trade-off between efficiency and equity in flight schedules and its relation to lexicographic optimisation is thoroughly discussed in the series of papers by [18], [19]. These examples demonstrate the versatility and effectiveness of lexicographic optimisation in addressing the numerous objectives inherent in air transportation problems.

IV. Spots

This section describes the various spots considered by the MIP model that will be presented in the next section. These spots provide a structured framework for managing air traffic.

A. Hot spot

A hot spot refers to a TV-period during which the current demand exceeds its declared capacity. To address this imbalance, the solver should mitigate the congestion by assigning ATFM delays and/or rerouting flights outside the TV-period.

B. Cold spot

A cold spot is just the opposite: a TV-period characterised by spare capacity. The primary objective in a cold spot is to maximise the number of flights passing through it while adhering to capacity constraints. Essentially, cold spots allow for the absorption of traffic from hot spots.

C. Protection spot

A protection spot denotes a TV-period where the demand is nearing or exceeding the available capacity. The primary objective in a protection spot is to minimise the number of additional flights passing through it to prevent overloads, particularly during periods of high demand uncertainty.

D. Opti spot

An opti spot is a TV-period where the time over of various flights has been optimised by a local solver (e.g., arrival manager). The system in charge of optimising ATFM measures on a network level should only make changes to these locally optimised target time overs (TTOs) if absolutely necessary to maintain operational integrity and minimise disruptions.

V. The model

This section present the mathematical model, including the decision variables, the constraints and the various objectives.

A. Variables

At its core, binary variables $x_{fad} \in \mathbb{B}$ determine the assignment of flights $f \in \mathcal{F}$ to specific alternatives $a \in \mathcal{A}$ and delays $d \in \mathbb{Z}^{\geq}$, thereby defining the ATFM measures of the solution. For instance, $x_{AB3} = 1$ indicates that flight A is assigned alternative B with a delay of 3 minutes. Accordingly, the model assumes that each flight f has a set of alternatives $\mathcal{A}_f$ and that each alternative $a \in \mathcal{A}_f$ can be assigned, at most, $d_a^{\max}$ minutes of delay. The set of alternatives could be obtained, for instance, from the routing assistance request service of the NM B2B service, from the airspace users themselves, or be the result of a rerouting group measure of the form "avoid airspace A" that generates alternative flight plans avoiding specific airspace. Regarding the maximum delay per alternative, it could be based on the current status of the flight (e.g., flights close to departure may have a lower value).

It should be noted that other variables are required to define some objectives and constraints. These variables are not introduced in this section as they can be treated as auxiliary. They will be presented as the corresponding constraints and objectives are listed in the subsequent sections, respectively.

B. Constraints

The most straightforward constraint is that each flight must be assigned one and only one alternative and delay, i.e.:

$$\sum_{a \in \mathcal{A}_f} \sum_{d=0}^{d_a^{\max}} x_{fad} = 1 \quad \forall f \in \mathcal{F}. \tag{1}$$

As mentioned before, some of these alternatives may stem from rerouting group measures implemented by Network Management (NM). In such cases, all alternatives generated by the same rerouting group must be either selected or rejected entirely (i.e., they are a pack). This requirement can be operationally enforced by introducing an auxiliary binary variable $z_m \in \mathbb{B}$ for each rerouting group measure $m \in \mathcal{M}$, which explicitly indicates whether the rerouting group measure m is active in the solution. To formalize this condition, we define the following constraint:

$$\sum_{d=0}^{d_a^{\max}} x_{fad} = z_m \quad \forall (f,a) \in \mathcal{R}_m, \forall m \in \mathcal{M}, \tag{2}$$

where the set $\mathcal{R}_m$ includes all combinations of flight-alternative pairs generated by the rerouting group measure m. This constraint ensures that if any alternative within a rerouting group is selected, z_m will be set to 1, enforcing the selection of all alternatives associated with that group. Conversely, if $z_m = 0$, all alternatives in the group are rejected.

As the model represents a demand-capacity balance problem, capacity constraints form a critical component. These constraints are typically formulated in terms of entry counts:

$$\sum_{(f,a,d) \in \mathcal{X}_{vse}^{\text{entry}}} x_{fad} \leq c_{vse}^{\text{entry}} \quad \forall (v,s,e) \in \mathcal{T}^{\text{entry}}, \tag{3}$$

where $\mathcal{T}^{\text{entry}}$ is the set of TV-periods where entry capacity is monitored, c_{vse}^{entry} is the entry capacity of TV-period (v,s,e), being v the TV identifier, s the start time (inclusive) of the period, and e the end time (exclusive), and $\mathcal{X}_{vse}^{\text{entry}}$ refers to the set of flight-alternative-delay combinations (i.e., variables) that would enter the TV-period if selected. Mathematically:

$$\mathcal{X}_{vse}^{\text{entry}} = \big\{(f,a,d) \mid s \leq \tau_{av}^{\text{entry}} + d < e \\ \forall d = 0, \ldots, d_a^{\max}, \forall a \in \mathcal{A}_f, \forall f \in \mathcal{F}\big\}.$$

where τ_{av}^{entry} is the estimated entry time of alternative a at TV v, without delay. It should be noted that capacity constraints can be also formulated in terms of occupancy counts, which are commonly monitored in relation to peak occupancy:

$$\sum_{(f,a,d) \in \mathcal{X}_{vse}^{\text{occ}}} x_{fad} \leq c_{vse}^{\text{peak}} \quad \forall (v,s,e) \in \mathcal{T}^{\text{occ}}, \tag{4}$$

and sustained occupancy:

$$\sum_{(f,a,d) \in \mathcal{X}_{vse}^{\text{occ}}} x_{fad} \leq c_{vse}^{\text{sust}} + y_{vse}^{\text{sust}} M \quad \forall (v,s,e) \in \mathcal{T}^{\text{occ}}, \tag{5}$$

where $\mathcal{T}^{\text{occ}}$ is the set of TV-periods where occupancy capacity is monitored. c_{vse}^{peak} and c_{vse}^{sust} are the peak and sustained occupancy capacities of TV-period (v, s, e), respectively, and

$$\mathcal{X}_{vse}^{\text{occ}} = \{(f, a, d) \mid \tau_{av}^{\text{exit}} + d > s \wedge \tau_{av}^{\text{entry}} + d < e \quad \forall d = 0, \ldots, d_a^{\max}, \forall a \in \mathcal{A}_f, \forall f \in \mathcal{F}\}.$$

In Eq. (5), the variable $y_{vse}^{\text{sust}} \in \mathbb{B}$ acts as an auxiliary variable, indicating whether the sustained capacity can be exceeded during the TV-period (v, s, e). The constant M represents a large number. To ensure effective monitoring of sustained occupancy, it is necessary to introduce a new constraint. This constraint guarantees that the capacity is not exceeded more than $n_{vse}^{\max}$ times across the following windows, including the period (v, s, e):

$$\sum_{(v', s', e') \in \mathcal{W}_{vse}^{\text{sust}}} y_{v's'e'}^{\text{sust}} \leq n_{vse}^{\max} \quad \forall (v, s, e) \in \mathcal{T}^{\text{occ}}. \tag{6}$$

where $\mathcal{W}_{vse}^{\text{sust}} = \{(v, s + kw, e + kw) \mid k = 0, 1, \ldots, K_{vse}\}$, with w being the duration of each window and K_{vse} denoting the number of windows monitored for sustained occupancy.

C. *Objectives*

The model incorporates a variety of objectives. During the lexicographic optimisation process, these objectives are prioritised and optimised sequentially according to their importance. The ranking of these objectives will be discussed later in the paper. Below is a list of the objectives included in the model.

An objective consists of minimising the total **ATFM delay**:

$$\sum_{f \in \mathcal{F}} \sum_{a \in \mathcal{A}_f} \sum_{d=0}^{d_a^{\max}} x_{fad} d. \tag{7}$$

The result of minimising Eq. (7) only accounts for *primary* delays caused directly by ATFM measures. However, flights assigned an ATFM delay may arrive so late at their destination that the scheduled turn-around time is insufficient, leading to *reactionary* rotational delays. These reactionary delays, although not explicitly considered by Eq. (7), inevitably occur, resulting in additional costs for airspace users and distorting the plan expected by the model. This distortion can also cause further overloads in other parts of the network. To address this limitation, the model allows to minimise the expected **reactionary rotational delay** by defining the following objective:

$$\sum_{f \in \mathcal{F}} k_f, \tag{8}$$

alongside the following set of turnaround constraints:

$$\tilde{\tau}_f^{\text{arr}} + \delta_{ff'}^{\min} - \tilde{\tau}_f^{\text{dep}} \leq 0 \quad \forall (f, f') \in \mathcal{L},$$

where $\mathcal{L}$ represents the set of inbound-outbound flights, f' is the flight following f in a rotation, $\delta_{ff'}^{\min} \in \mathbb{Z}^+$ is the minimum turnaround time between these two flights, and the delayed arrival and departure times are, respectively:

$$\tilde{\tau}_f^{\text{arr}}(\mathcal{X}) = \sum_{a \in \mathcal{A}_f} \sum_{d=0}^{d_a^{\max}} (\tau_a^{\text{arr}} + d)\, x_{fad} + k_f \quad \forall f \in \mathcal{F},$$

$$\tilde{\tau}_f^{\text{dep}}(\mathcal{X}) = \tau_f^{\text{dep}} + \sum_{a \in \mathcal{A}_f} \sum_{d=0}^{d_a^{\max}} x_{fad} d + k_f \quad \forall f \in \mathcal{F},$$

It is important to note that the knock-on delays incorporated into the model are not punitive measures enforced on flights. Airspace users often have strategies such as tail swapping to mitigate these delays. Therefore, imposing a delay that could be effectively avoided by the user would be unfair.

The purpose of the rotational reactionary delay objective is to minimise the expected number of occurrences and support airspace users in their operations. In other words, only d will result in ATFM measures, but not k_f. The parameters $\delta_{ff'}^{\min}$ could be defined by the NM as a function of the airspace user, airport and/or aircraft type, based on historical data of actual turnaround times. For instance, one could take the 5$^{\text{th}}$ percentile of the conditioned turnaround time distribution.

A further option is to reduce the number of flights assigned an ATFM delay that exceeds a predetermined threshold $d_a^{\text{impact}} \in \mathbb{Z}^{\geq}$. Some flights may tolerate delays below a certain threshold without causing significant issues, while delays beyond that threshold could lead to **operational disruptions**. In such cases, the model could prioritise allocating flights within their "operationally acceptable" delay range:

$$\sum_{f \in \mathcal{F}} \sum_{a \in \mathcal{A}_f} \sum_{d=d_a^{\text{impact}}}^{d_a^{\max}} x_{fad}, \tag{9}$$

This objective can also help minimise delays in the initial rotations of the day by assigning a very low value to d_a^{impact}. In 2024, EUROCONTROL emphasised the importance of prioritising the first rotation, noting that each minute of delay in the first rotation multiplies fourfold by the last rotation, exacerbating knock-on delays and risking night curfew violations. Indeed, prioritising the first rotation is one of the five key priorities EUROCONTROL urges the aviation sector to focus on while preparing for the busy summer traffic season[2].

As the solver has the capability to adjust both ATFM delay and flight plan alternatives for a flight, optimising solely the previous objectives might yield solutions that are highly efficient in terms of delay-related costs but potentially detrimental to the environment. Aiming to promote environmentally-friendly ATFM measures, an additional objective aimed at minimising **fuel consumption** has been introduced:

$$\sum_{f \in \mathcal{F}} \sum_{a \in \mathcal{A}_f} \sum_{d=0}^{d_a^{\max}} x_{fad} \phi_a, \tag{10}$$

where $\phi_a \in \mathbb{Z}^+$ is the fuel consumption of alternative a.

[2] https://www.eurocontrol.int/news/all-together-now-2024

The position of the environmentally friendly ATFM objective within the lexicographic order holds significant importance and warrants early consideration in the decision-making process. Establishing its rank early on ensures that environmental concerns are given appropriate priority and are not overshadowed by other objectives. Actually, this transparent multi-objective optimisation approach facilitates the integration of eco-friendly practices into ATFM strategies.

An novel feature introduced in Section IV is the concept of spots. The hot spots are treated by the models as regular capacity constraints. However, the remaining spots in the model are defined as the objectives discussed below.

When optimising for **cold spots**, the objective is to promote traffic through under-loaded TV-periods, i.e., to maximise:

$$\sum_{(v,s,e)\in\mathcal{S}^{\text{cold}}} \sum_{(f,a,d)\in\mathcal{X}^{\text{nodelay}}_{vse}} x_{fad}, \tag{11}$$

where $\mathcal{S}^{\text{cold}}$ is the set of cold spots in the network and

$$\mathcal{X}^{\text{nodelay}}_{vse} = \{(f,a,d) \in \mathcal{X}^{\text{entry}}_{vse} \,|d \le d^{\text{cur}}_f\}$$

is the set of flight-alternative-delay combinations that would enter the cold spot (v,s,e) if selected, while not exceeding their current ATFM delay value d^{cur}_f.

Protection spots pertain to TV-periods that are not yet overloaded but where additional traffic is undesirable due to the likelihood of potential overloads.This goal can be achieved by enforcing the following set of constraint:

$$\sum_{(f,a,d)\in\mathcal{X}^{\text{entry}}_{vse}\setminus\mathcal{X}^{\text{cur}}_{vse}} x_{fad} - \|\mathcal{X}^{\text{cur}}_{vse}\| \le \Delta_{vse} \quad \forall (v,s,e) \in \mathcal{S}^{\text{prot}},$$

where $\mathcal{S}^{\text{prot}}$ is the set of protection spots in the network,

$$\mathcal{X}^{\text{cur}}_{vse} = \{(f,a,d) \in \mathcal{X}^{\text{entry}}_{vse} \,|a = a^{\text{cur}}_f \wedge d = d^{\text{cur}}_f\}$$

is the flight-alternative-delay combinations currently entering the TV-period (v,s,e), a^{cur}_f is the current alternative of flight f, and $\Delta_{vse} \in \mathbb{Z}^{\ge}$ represents the extra traffic in the protection spot. Accordingly, the following expression describes the total **extra traffic in protection spots**:

$$\sum_{(v,s,e)\in\mathcal{S}^{\text{prot}}} \Delta_{vse}. \tag{12}$$

Finally, entry times in an opti spot should only be changed when absolutely necessary. Accordingly, the solver's objective is to minimise **deviations from the (locally optimal) TTOs in the set of opti spots** within the network, $\mathcal{S}^{\text{opti}}$:

$$\sum_{(v,s,e)\in\mathcal{S}^{\text{opti}}} \sum_{f\in\mathcal{F}^{\text{opti}}_{vse}} \left|\tau^{\text{opti}}_{fv} - \tilde{\tau}^{\text{entry}}_{fv}\right|, \tag{13}$$

where, $\tilde{\tau}^{\text{entry}}_{fv}$ represents the entry time of flight f at TV v, considering the selected alternative ATFM delay[3], i.e.:

[3]Equation (13) is non-linear. However, it can be straightforwardly linearized by introducing auxiliary variables and constraints to maintain computational tractability.

$$\tilde{\tau}^{\text{entry}}_{fv}(\mathcal{X}) = \sum_{a\in\mathcal{A}_f} \sum_{d=0}^{d^{\max}_a} \left(\tau^{\text{entry}}_{av} + d\right) x_{fad}$$
$$\forall f \in \mathcal{F}^{\text{opti}}_{vse},\ \forall (v,s,e) \in \mathcal{S}^{\text{opti}}.$$

and $\mathcal{F}^{\text{opti}}_{vse}$ is the set of flights which TTO has been optimised for the opti spot (v,s,e). Please note that the delayed entry time does not account for the knock-on delay due to the inability to guarantee its occurrence, as discussed previously.

An often overlooked aspect in the literature is the stability of consecutive solutions. The ATFM model is not intended to be solved once and then halt. Realistically, one would expect the solver to generate a solution, then shortly afterwards, with new inputs such as updated flight plans, flight cancellations, or flights becoming airborne and therefore exempted from ATFM measures, generate a new solution. In consecutive runs, one might anticipate similar solutions if the traffic situation hasn't changed significantly, although this remains a speculation.

The model presented herein addresses this issue by allowing for the minimisation of changes (or, equivalently, maximisation in similarity) with a previous solution, explicitly considering the volatility of proposed ATFM measures. This is achieved through three distinct objectives:

$$\sum_{f\in\mathcal{F}} \sum_{a\in\mathcal{A}_f} \sum_{d=0}^{d^{\max}_a} x_{fad} \left|d - d^{\text{cur}}_f\right|, \tag{14a}$$

$$\sum_{f\in\mathcal{F}} \sum_{a\in\mathcal{A}_f} \sum_{d=d^{\text{cur}}_f+1}^{d^{\max}_a} x_{fad} \left(d - d^{\text{cur}}_f\right), \text{ and} \tag{14b}$$

$$\sum_{f\in\mathcal{F}} \sum_{a\in\mathcal{A}_f\setminus\{a^{\text{cur}}_f\}} \sum_{d=0}^{d^{\max}_a} x_{fad}. \tag{14c}$$

The three objectives related to stability and presented in Eq. (14) can be used to minimise the **ATFM delay change** (independently of the sign), **ATFM delay increase**, and **flight plan change** with respect to the current situation, respectively.

Additionally, while MIP models may deliver impressive results in terms of delay, they may not fully account for operational practicality from the perspective of FMPs. Specifically, if traffic spread is not explicitly considered, some optimal solutions might have the form where flights are scheduled at the very first minute of each monitored capacity window. For example, if the capacity for a 20-minute period is 20 flights, the solver might allocate all 20 flights to the first minute, leaving the remaining 19 minutes empty.

The model aims to maximise the traffic spread within specific TV-periods, denoted as $\mathcal{T}^{\text{spread}}$. Let us define $e^{\max}_{vse} \in \mathbb{Z}_{\ge 0}$ and $e^{\min}_{vse} \in \mathbb{Z}_{\ge 0}$ as the maximum and minimum number of entries in a TV-period (v,s,e), respectively. Additionally, let $\mathcal{W}^{\text{spread}}_{vse} = \{(v,s,s{+}r), (v,s{+}r,s{+}2r), \ldots, (v,e{-}r,e)\}$ represent the set of windows within which traffic is to be evenly distributed throughout that TV-period, with a resolution of r minutes. The **traffic spread** can be enhanced by minimising:

$$\sum_{(v,s,e)\in\mathcal{T}^{\text{spread}}} e^{\max}_{vse} - e^{\min}_{vse}, \tag{15}$$

subject to the following constraints for each $(v, s, e) \in \mathcal{T}^{\text{spread}}$:

$$e^{\max}_{vse} \geq \sum_{(f,a,d)\in\mathcal{X}^{\text{entry}}_{v's'e'}} x_{fad} \geq e^{\min}_{vse} \quad \forall (v', s', e') \in \mathcal{W}^{\text{spread}}_{vse} \tag{16}$$

Table I categorises the various objectives by topic and includes links to their mathematical expressions.

TABLE I. TOPICS, OBJECTIVES AND EXPRESSIONS

Topic	Objective	Expression
Delay	ATFM delay Rotational reactionary delay Operational impact of delay	Eq.(7) Eq.(8) Eq.(9)
Environmental impact	Fuel consumption	Eq.(10)
Spot control	Traffic through colds pots Extra traffic in protection spots TTO change in opti spots	Eq.(11) Eq.(12) Eq.(13)
Stability	ATFM delay change ATFM delay increase Flight plan alternatives change	Eq.(14a) Eq.(14b) Eq.(14c)
Safety	Traffic spread	Eq.(15)

VI. ILLUSTRATIVE EXAMPLE

This section presents an illustrative example to demonstrate the application and effectiveness of the lexicographic approach in addressing the demand-capacity balance problem.

A. Scenario

For this analysis, July 19th, 2024, was selected as a notably busy day along the South-West axis. All traffic demand and capacity data used in this example were sourced from the NM.

The traffic demand sample includes the 8179 flights crossing French and/or Spanish airspace, regardless of whether they departed from or arrived at airports outside these countries, with EOBT between midnight and 6 PM UTC. This region and time frame were chosen due to the high volume of air traffic and the presence of multiple critical traffic spots, making it an ideal test scenario for the model. For each flight crossing any hot spot, a set of synthetic flight plan alternatives was generated, consisting of all unique flight plans flown for the same city-pair and aircraft type over the past seven days. On average, each flight had 1.4 flight plan alternatives.

To optimise the reactionary rotational delay, the minimum turnaround time for each leg of an aircraft rotation, except for the first, was set to the lesser of either 45 min or the scheduled turnaround time. While this uniform assumption simplifies the complexities of real-world operations, the primary goal in this section is to demonstrate the principles of the lexicographic approach within a reasonably realistic scenario.

Regarding capacity, 93 active and monitored TVs in French and Spanish airspace on that day were considered to generate the capacity constraints. Capacity monitoring was based on entry counts within 60-min windows, with slices every 20 min.

To illustrate the concept of spots, 77 cold spots were identified in TV-periods where demand was below 20% of declared capacity, indicating significant under-loading. Additionally, 29 protection spots were defined in TV-periods where demand exceeded 75% of declared capacity (but remained below it). Lastly, an opti spot was created for the TV capturing 73 arrivals at Barcelona Airport between 6 and 9 AM UTC, using scheduled arrival times as hypothetical locally optimal TTOs.

B. Model

The MIP model for this scenario includes 547,797 variables (543,996 binary and 3,801 integer) and 12,985 constraints. Solving the optimisation problem for a single objective takes between 30 sec and 3 min, depending on the objective and its priority in the lexicographic order. This suggests that finding a complete solution with three objectives may require between 1.5 and 9 min. Naturally, the execution time is influenced by factors such as the problem's size and complexity, the solver used, and the specifications of the machine on which it is run. In this case, the SCIP solver version 8.0.4 was used, and the problem was solved on a server equipped with 12×Intel® Xeon® W-2235 CPUs @ 3.80GHz, 256GB RAM.

C. Results

The model was solved for various rankings of objectives, including all possible orders of ATFM delay (D), rotational reactionary delay (K), and fuel consumption (F). Additionally, the optimisation spot-related objectives was explored in a context where ATFM delay and fuel consumption were prioritised.

Figure 1 illustrates the values of various key performance indicators (KPIs) when optimising for objectives D, K, and F under different priority rankings. Each cell in the figure corresponds to a unique ordering of these priorities. For instance, in the top-left cell, D is given the highest priority, followed by K, and then F. The red line represents the KPI values when optimising exclusively for D. The grey line shows the outcomes when optimising for D first, followed by K, with the constraint that the optimal D value is maintained. The blue line incorporates all three objectives, optimising them in the specified order of priority. It is important to note that the grey line always dominates the red line, and the blue line always dominates the grey line. In this context, "dominates" means that the optimised objectives are equal to or better.

Several conclusions can be drawn from this figure. First, the order in which objectives are prioritised significantly influences their resulting values. For example, compare the top-left and top-centre cells. In the top-left cell, where D is the highest priority, its optimal value is 2.47K. However, in the top-centre cell, where D is the second priority, its optimal value increases to 2.52K. In the most extreme case, when F and K are prioritised over D, treating D as a less critical objective, its value can soar to 10.84K – more than four times higher.

Second, even when ϵ is 0, there is often potential to optimise other objectives while maintaining the optimal value

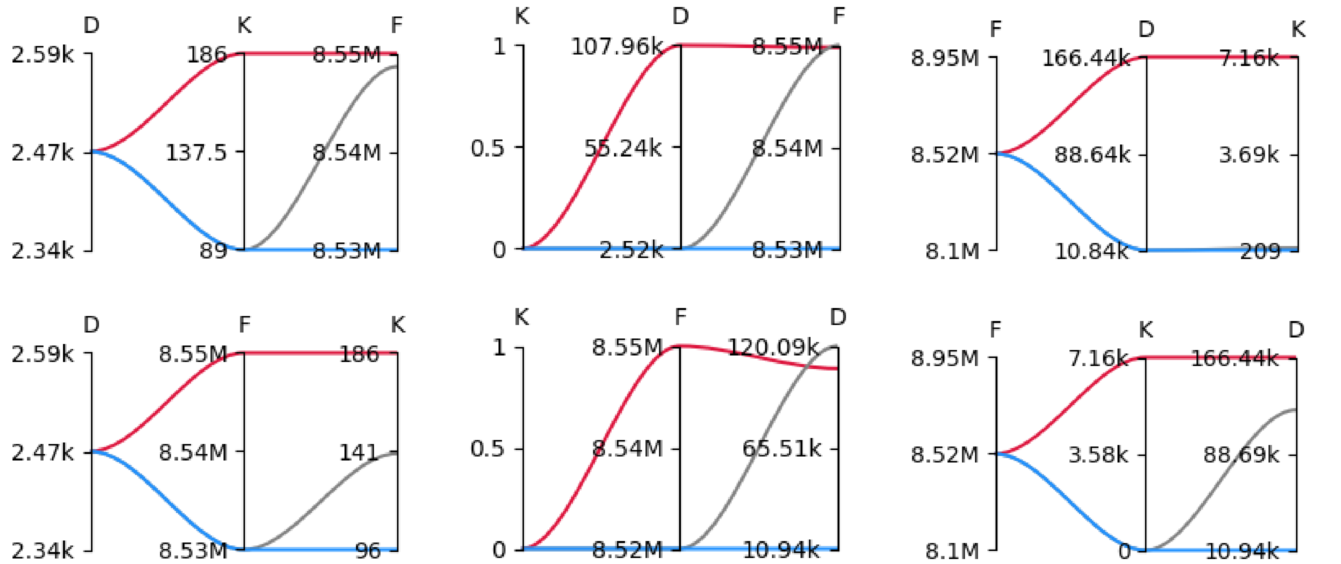

Figure 1. KPIs when ϵ is 0 and spots are not considered. D: ATFM delay; K: Reactionary rotational delay; F: Fuel consumption. Within each cell, the lexicographic order progresses from left to right. Red: optimised for the first objective; grey: for the first and second objectives; blue: for all three objectives.

of a primary objective. For example, in the top-right cell, the red line represents the scenario where F is optimised exclusively, resulting in D and K values of 166.44K and 7.16K, respectively. However, in the blue line, which optimises all three objectives in sequence, F retains its optimal value, but D and K are reduced to 10.84K and 209, respectively.

Figure 2 presents the results under an ϵ of 5%, meaning that the optimal values of previously optimised objectives in the lexicographic optimisation process can be exceeded by up to 5%. For instance, if the optimal value of D were 100, a solution with a D value of 105 would be considered feasible.

This figure demonstrates that, in this situation, a line does not necessarily dominate the previous one. For example, in the bottom-left cell, the optimal value of F during the second optimisation (the grey line) is 8.53M, but it increases above 8.54M during the third optimisation (the blue line), where K can reach its theoretical optimal value of 0 (as opposed to 96 in the analogous plot with $\epsilon = 0$). It is worth noting that increasing ϵ does not always mean that the solver will fully utilise the available extra margin for the objective, nor does it guarantee that the objective values for all subsequent objectives will improve compared to a smaller ϵ.

Similar discussions apply to Figure 3, which shows the values of various KPIs when optimising for spot-related objectives with $\epsilon = 0$ in a context where D and F are the priority.

Figure 3 illustrates that when D and F are prioritised over minimising the number of flights crossing cold spots (C), the optimal value for C is 379. Additionally, when D and F are taken into account, the extra traffic in protection spots (P) can be reduced to 40, and the change in TTOs for opti spots (O) can be minimised to 1.09K. It is crucial to note that prioritising spot-related objectives earlier in the ranking, rather than last, can result in significantly different outcomes.

Finally, it is important to note that the potential to optimise remaining objectives is not always significant. In some instances, the priorities set by earlier objectives can so tightly restrict the feasible space that further optimisation of subsequent objectives becomes challenging.

VII. Conclusions

In this paper, we address the complex issue of balancing rising air traffic demand with the limited capacity of airspace and airports through a novel mixed-integer programming model. Our approach leverages a lexicographic method to manage the inherent multi-objective nature of the problem, allowing us to systematically prioritise and address conflicting goals such as minimising air traffic flow management delays, mitigating reactionary delays, and reducing environmental impacts.

The application of our model to historical air traffic data from France and Spain underscores its practical relevance and versatility in real-world scenarios. Furthermore, the case studies highlight the model's ability to address specific regional challenges – referred to as "spots" – in the network, which can vary significantly in their characteristics and demands.

Future research could expand on this work by integrating additional objectives and validating the model in various operational contexts. A particular focus will be on developing objectives that address currently overlooked aspects, such as the fairness or equity of solutions. For example, an objective to minimise the number of flight reversals (i.e., violations of the first-planned first-served principle of CASA) could be introduced. Additionally, the authors plan to investigate the lexicographic max-min approach to ensure that the most equitable solutions are achieved considering all airspace users.

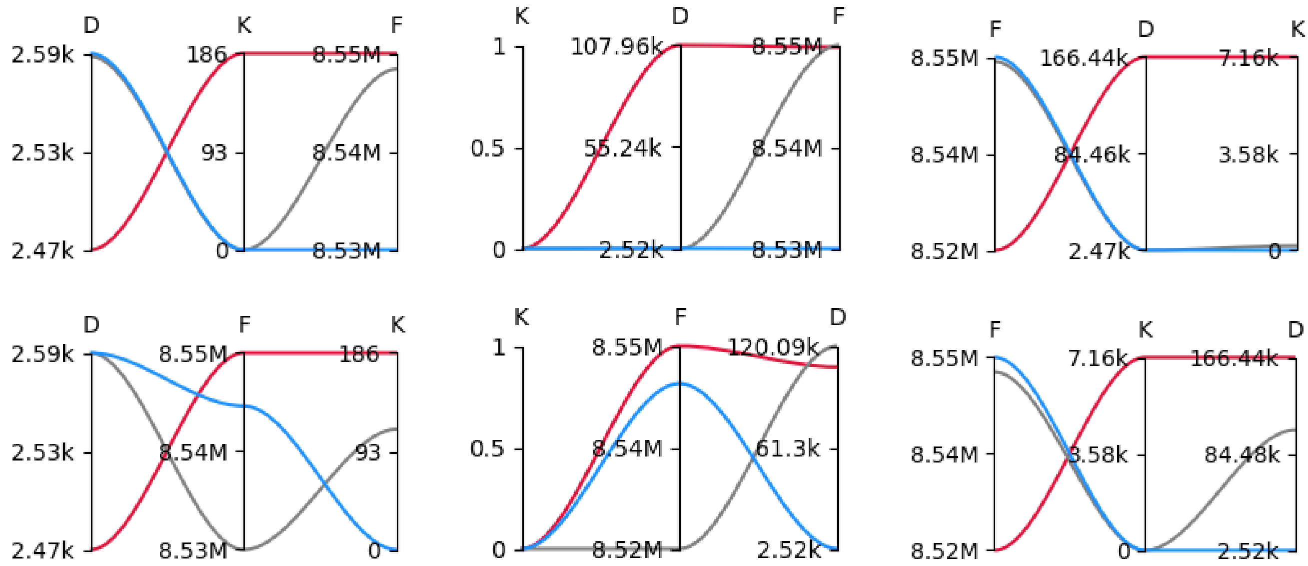

Figure 2. KPIs when ϵ is 5% and spots are not considered.

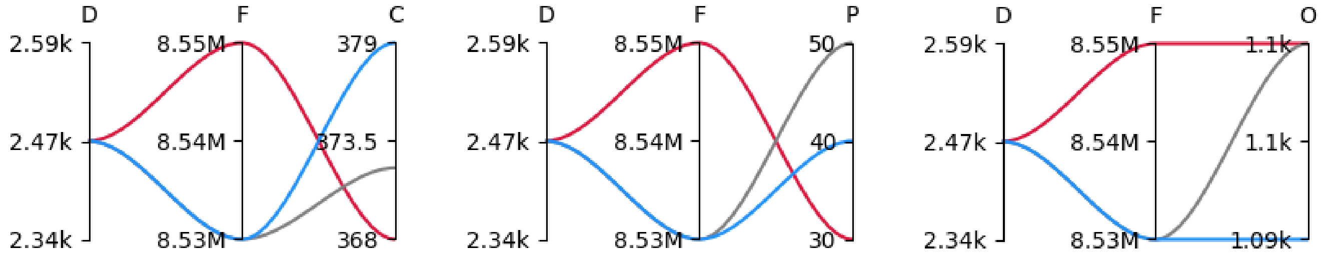

Figure 3. KPIs when ϵ is 0 and spots are considered. C: Cold spots; P: Protection spots; O: Opti spots. Within each cell, the lexicographic order progresses from left to right. Red: optimised for the first objective; grey: for the first and second objectives; blue: for all three objectives.

Lastly, the set of synthetic flight plan alternatives used in the experiment is somewhat limited, averaging only 1.4 alternatives per flight. Exploring or inferring how a more extensive set of alternatives might impact the solution could provide valuable insights.

Acknowledgement

This project has received funding from the SESAR 3 Joint Undertaking (JU) under grant agreement No 101114675. The JU receives support from the European Union's Horizon Europe research and innovation programme and the SESAR 3 JU members other than the Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or SESAR 3 Joint Undertaking. Neither the European Union nor the granting authority can be held responsible for them.

References

[1] International Air Transport Association (IATA), "IATA forecast predicts 8.2 billion air travelers in 2037." https://www.iata.org/en/pressroom/2018-releases/2018-10-24-02/, 2018. Accessed: 2024-09-30.

[2] EUROCONTROL, "Challenges of growth 2018," 2018.

[3] S. Xu *et al.*, "A fast-time demand-capacity balancing optimizer for collaborative air traffic flow management," *Journal of Air Transport Management*, vol. 90, p. 101968, 2021.

[4] D. P. Bertsekas, *Dynamic Programming and Optimal Control*. Athena Scientific, 1995.

[5] P. Vranas, "Optimal slot allocation for european air traffic flow management," *Air Traffic Control Quarterly*, vol. 4, no. 4, pp. 249–280, 1996.

[6] D. Bertsimas and S. Stock Patterson, "The air traffic flow management problem: An integer optimization approach," *Operations Research*, vol. 46, no. 3, pp. 406–422, 1998.

[7] S. Xu *et al.*, "A synchronized demand-capacity balancing approach within a collaborative atfm framework," *Transportation Research Part C: Emerging Technologies*, vol. 120, p. 102837, 2020.

[8] N. Ivanov, F. Netjasov, R. Jovanović, S. Starita, and A. Strauss, "Air traffic flow management slot allocation to minimize propagated delay and improve airport slot adherence," *Transportation Research Part A: Policy and Practice*, vol. 95, pp. 183–197, 2017.

[9] T. Bolić, L. Castelli, L. Corolli, and D. Rigonat, "Reducing atfm delays through strategic flight planning," *Transportation Research Part E: Logistics and Transportation Review*, vol. 98, pp. 42–59, 2017.

[10] P. Pellegrini, T. Bolić, L. Castelli, and R. Pesenti, "Sosta: An effective model for the simultaneous optimisation of airport slot allocation," *Transportation Research Part E: Logistics and Transportation Review*, vol. 99, pp. 34–53, 2017.

[11] Y. Xu, R. Dalmau, M. Melgosa, A. Montlaur, and X. Prats, "A framework for collaborative air traffic flow management minimizing costs for airspace users: Enabling trajectory options and flexible pre-tactical delay management," *Transportation Research Part B: Methodological*,

vol. 134, pp. 229–255, 2020.

[12] T. Bolić, L. Castelli, L. Corolli, and G. Scaini, "Flexibility in strategic flight planning," *Transportation Research Part E: Logistics and Transportation Review*, vol. 154, p. 102450, 2021.

[13] T. Bolić, L. Castelli, G. Scaini, G. Frau, and S. Guidi, "Flight flexibility in strategic traffic planning: visualisation and mitigation use case," *CEAS Aeronautical Journal*, vol. 12, no. 4, pp. 847–862, 2021.

[14] R. Dalmau, G. Gawinowski, and C. Anoraud, "Comparison of various temporal air traffic flow management models in critical scenarios," *Journal of Air Transport Management*, vol. 105, p. 102284, 2022.

[15] X. Prats, V. Puig, J. Quevedo, and F. Nejjari, "Lexicographic optimisation for optimal departure aircraft trajectories," *Aerospace Science and Technology*, vol. 14, no. 1, pp. 26–37, 2010.

[16] A. Jacquillat and V. Vaze, "Interairline equity in airport scheduling interventions," *Transportation Science*, vol. 52, no. 4, pp. 941–964, 2018.

[17] H. Erkan, N. K. Erkip, and Özge Şafak, "Collaborative decision making for air traffic management: A generic mathematical program for the rescheduling problem," *Computers & Industrial Engineering*, vol. 137, p. 106016, 2019.

[18] D. Bertsimas, V. F. Farias, and N. Trichakis, "The price of fairness," *Operations Research*, vol. 59, no. 1, pp. 17–31, 2011.

[19] D. Bertsimas, V. F. Farias, and N. Trichakis, "On the efficiency-fairness trade-off," *Management Science*, vol. 58, no. 12, pp. 2234–2250, 2012.