Transportation Research Part C 124 (2021) 102914

Contents lists available at ScienceDirect

Transportation Research Part C

journal homepage: www.elsevier.com/locate/trc

Modelling and solving the airport slot-scheduling problem with multi-objective, multi-level considerations

Fotios A. Katsigiannis, Konstantinos G. Zografos *, Jamie Fairbrother

Centre for Transport and Logistics (CENTRAL), Lancaster University Management School, Department of Management Science, Lancaster LA1 4YX, United Kingdom

ARTICLE INFO

Keywords:
Airport slot allocation
Multi-objective programming
Airport demand management
Airport scheduling
Multi-level decision-making

ABSTRACT

In overly congested airports requests for landing and take-off slots are allocated according to the IATA World Scheduling Guidelines (WSG). A central concept of these guidelines is the prioritization of the satisfaction of the requested slots according to a hierarchy that recognizes historic usage rights of slots. A number of criteria have been proposed in the literature to optimize airport slot allocation decisions. Multi-objective programming models have been proposed to investigate the trade-offs of the slot allocation objectives for the same level of the slot hierarchy. However, the literature currently lacks models that can study in a systematic way the trade-offs among the scheduling objectives across all levels of the hierarchy and the airport schedule as a whole. To close the existing literature gap, we are proposing a new tri-objective slot allocation model (TOSAM) that considers total schedule displacement, maximum schedule displacement and demand-based fairness, and we introduce a multi-level, multi-objective algorithm to solve it. We are using real world slot request and airport capacity data to demonstrate the applicability of the proposed approach. Our computational results suggest that the systematic consideration of the interactions among the objectives of the different levels of the slot hierarchy, results to improved schedule-wide slot scheduling performance. In particular, we found that small sacrifices made for the attainment of the scheduling objectives of the upper echelons of the slot hierarchy, result in significant improvements of the schedule-wide objectives.

1. Introduction

The most recent air transport demand forecasts for 2040 (EUROCONTROL, 2018a) predict an average annual passenger and aircraft movement growth of 2.5% and 1.9% respectively. It is also predicted that the growth rate of demand for transport services will outpace the growth rate of the airport infrastructure development and of the available airport capacity (EUROCONTROL, 2018b). As a result, the current demand-capacity imbalance, the associated economic, societal, and environmental costs are expected to be further exacerbated. It is worth noting that, in Europe in 2013 the demand-capacity mismatch was responsible for €6.72 billion in passenger costs, €4.5 billion in airline costs, and significant CO2 emissions (IATA, 2013). Supply side interventions can be used to restore this demand-capacity imbalance. However, the size of the airport capacity expansion investments, the reaction of the communities located in the vicinity of proposed airport sites, the scarcity of space, and the long horizon associated with the implementation of capacity

* Corresponding author.
E-mail addresses: f.katsigiannis@lancaster.ac.uk (F.A. Katsigiannis), k.zografos@lancaster.ac.uk (K.G. Zografos), j.fairbrother@lancaster.ac.uk (J. Fairbrother).

https://doi.org/10.1016/j.trc.2020.102914
Received 12 March 2020; Received in revised form 4 September 2020; Accepted 30 November 2020
0968-090X/

expansion projects suggest that we should turn our attention to demand side interventions for short term solutions to the demand-capacity imbalance problem (Zografos et al., 2017b). Airport slot allocation is acknowledged by researchers and practitioners as an effective tool for scheduling slots and managing demand at congested airports. Outside United States (U.S.)[1], the IATA Worldwide Scheduling Guidelines (WSG) (IATA, 2019) are used to schedule slots at congested airports.

A key policy concept of the IATA WSG is the consideration of a well-defined hierarchy for satisfying airport slot requests. This hierarchical structure is based on the recognition of historical airport usage rights for different airlines. Specifically, this hierarchy classifies the slot requests to Historic, Changes to Historic, New Entrants, and Other. For a detailed description of the guidelines defining the different classes and the associated rights of slot requests the reader is referred to the latest version of the IATA WSG (IATA, 2019). The optimum allocation of slots at congested airports has also attracted the attention of the research community and a variety of optimization models have been proposed incorporating the IATA hierarchical structure for the satisfaction of slots (for more details on the most related literature we refer the readers to section 2). In addition to considering the IATA WSG hierarchical allocation requirements, existing models have used objectives and constraints that reflect other aspects of the IATA WSG such as the need to allocate the slots for an entire scheduling season[2] considering *series of slot requests*[3].

Existing studies have introduced single and multi-objective formulations to optimize slot allocation decisions. Single objective formulations (Zografos et al., 2012) aimed to minimise schedule displacement, i.e. the total absolute difference between the requested and allocated times across all slot requests. Multi-objective formulations considered the schedule displacement objective in combination with one or more objectives, i.e. maximum displacement, the number of displaced/violated/rejected requests and fairness metrics. The decision-making environment of the slot allocation problem is characterized by multiple levels corresponding to the different echelons of the hierarchy that dictates the priority for the satisfaction of different types of slot requests, and multiple objectives reflecting the views of the different stakeholder groups involved in and affected by the slot scheduling process. The current state of the art (a thorough literature review is presented in the following section) includes models that have analysed trade-offs between the objectives of the same level of the hierarchy, as well as models that can examine trade-offs among the slot scheduling objectives when all slots are satisfied simultaneously without considering different priorities for their satisfaction.

Therefore, the literature currently lacks solution approaches which can provide a systematic way for examining the trade-offs among multiple objectives across all levels of the slot allocation hierarchy and the airport schedule as a whole. The examination of these trade-offs is of importance in making slot scheduling decisions, since it can produce additional solutions that may facilitate a more cooperative decision making, and eventually result in schedules that are more acceptable by all stakeholders. In addition, the analysis of the relationship between the schedules of the individual priorities and the schedule-wide[4] solution, may indicate the slot priorities and the extent that they must compromise their slot scheduling objectives in order to improve the efficiency of the airport schedule as a whole. This analysis is particularly useful for airport slot coordinators since it suggests a range of alternative slot allocations that can improve the overall efficiency of the airport slot schedule. The slot coordinators in cooperation with other stakeholders, i.e. airlines, airport operators, can choose the schedules that minimise the problems arising from the conflicting requirements of airlines' slot requests. In a nutshell, the generation of multiple schedule-wide alternatives supports slot coordinators on two of their main roles (IATA, 2019): (a) the offering of advice to airlines and relevant authorities regarding all areas which are likely to better manage demand and congestion; and (b) the resolution of problems arising by the differing needs of airlines requesting slots at the airport.

This gap is addressed in this study, which builds upon work presented previously in the extended abstract of (Katsigiannis and Zografos, 2019) by: *(i)* proposing and implementing a solution methodology that introduces multi-level considerations and produces a wide spectrum of non-dominated slot schedules while considering IATA's slot request prioritisation; *(ii)* formulating and solving a tri-objective mixed integer program which considers slot scheduling efficiency and demand-based fairness while simultaneously bounding worst-case displacement across all requests; *(iii)* demonstrating the implications of the proposed framework for policy and practice.

The remainder of this paper is organised as follows. Section 2 provides an overview of previous related work and further elaborates on the arguments justifying the need for the proposed model and solution approach. Section 3 presents the formulation of the proposed model. Section 4 includes the proposed solution algorithm as well as the rationale behind its architecture and design. Section 5 describes our experimental setup and presents the application and the implications of the proposed model and algorithm for scheduling slots at a real-world slot coordinated airport. Finally, Section 6 concludes this work and provides recommendations for future research.

2. Previous related work

Airport slot allocation methods fall into one of the following three categories: (i) administrative; (ii) market-based; and (iii) hybrid,

[1] Airports within the U.S. do not manage demand for airport capacity (with the exception of airports that follow the High-Density Rule (HDR)). Instead, airlines schedule their flights based on a mechanism which considers expected delays.

[2] The airport slot scheduling calendar is based on two scheduling seasons. The Summer season typically starts on the last Sunday of March and lasts till the last Saturday of October. Accordingly, the winter season starts on the last Sunday of October and finishes on the last Saturday of March.

[3] "*A series of slots is at least 5 slots requested allocated for the same or approximately same time on the same day-of-the-week, distributed regularly in the same season*" (IATA, 2019).

[4] The term schedule-wide is herein used to signify the performance of the airport slot schedule across all the requests that it contains and differentiate it from the performance of the schedules of each priority level.

i.e. approaches that combine administrative and market-based methods. For more information on the description and classification of available airport slot allocation and scheduling methods we refer the reader to (Zografos et al., 2017b). The research reported in this paper deals with the scheduling of slots at a single airport using administrative approaches deriving from the IATA WSG. Therefore, our literature review is focussed on this part of the literature. However, where appropriate, we are also reviewing models that have extended the single airport slot scheduling problem to network-wide slot allocation (Corolli et al., 2014; Pellegrini et al., 2017).

Zografos et al. (2012) first formulated the airport slot scheduling problem as a single objective integer optimization model. The proposed model considered the IATA requirement for allocating series of slots for the entire scheduling season. The objective of this model was to minimize total schedule displacement, subject to aircraft turn-around, and airport capacity constraints. Zografos and Jiang (2016, 2019) proposed two bi-objective formulations of the slot scheduling problem by extending the model presented in Zografos et al. (2012) through the consideration of fairness, and total weighted schedule displacement as an additional slot scheduling objective. The proposed fairness objective was based on a fairness index which postulates that schedule displacement should be allocated to each airline proportionally to the number of requested slots. Under a perfectly fair schedule the fairness index of all airlines should be equal to one. This value represents the absolute schedule fairness. However, due to operational and capacity constraints absolute schedule fairness cannot be achieved. Therefore, an alternative measure of fairness was used to express the fairness objective (Zografos and Jiang, 2016, 2019; Jiang and Zografos, 2017). The fairness objective in the study of (Zografos and Jiang, 2016, 2019) seeks to minimise the maximum difference between the fairness index of each airline from the average fairness index of all airlines. (Jiang and Zografos, 2017) proposed a fairness metric that minimises the maximum distance from the absolute fairness. The weighted schedule displacement objective introduced in Zografos and Jiang (2016) weights the displacement according to the size of the aircraft operating the slot as well as the flight length. The incorporation of such weights in the displacement objective favours (within the same level of hierarchy) slot requests that correspond to flights with longer distance and more passengers. Zografos and Jiang (2019) found within the same level of the hierarchy of slot requests, trade-offs between schedule efficiency (displacement) and schedule fairness. Furthermore, they found that the price of fairness, i.e. the total extra displacement that should be encountered by all slot requests (of the same slot priority) in order to achieve a given level of fairness, differs among the different levels of the hierarchy of slots. It is worth noting that fairness has been incorporated in network wide slot allocation models (Pellegrini et al., 2017).

Zografos et al. (2017a) formulated two additional bi-objective models. The first model minimizes total and maximum schedule displacement, while the second minimizes total displacement and the number of violated slot assignments. In the latter model, violated slot assignments are defined as the slot assignments that exceed a threshold of acceptable displacement. The maximum displacement objective minimises the worst case of displacement across all slots requested and can be perceived as a ‘guaranteed measure of the quality of the schedule ’ that can be provided by a given airport (Zografos et al., 2017a) . Zografos et al. (2017a) and Zografos and Jiang (2019) solved the bi-objective slot allocation problem considering historic, new entrant and other requests and investigated trade-offs between the objectives of the proposed bi-objective formulations under the following two regimes: (i) a regime that considers historical slot usage rights (hierarchical slot allocation); and (ii) a regime that considers all slot requests simultaneously (non-hierarchical slot allocation). It was found that the simultaneous slot scheduling outperforms the sequential approach leading to lower total schedule displacement. However, the resulting performance improvement is on the expense of the historic slot requests (Zografos and Jiang, 2019).

Fairbrother et al. (2019) and Fairbrother and Zografos (2018a), proposed a two-stage slot allocation mechanism which incorporates fairness and airline preferences. In the first stage a fair reference schedule is constructed using a new fairness measure called demand-based fairness. This builds on the work of Zografos and Jiang (2016, 2019), and Jiang and Zografos (2017) by defining a fairness metric which considers slot requests made during peak demand periods. The amount of displacement allocated to an airline in the first stage is referred to as the displacement budget. In the second stage, airlines must specify how this displacement should be distributed along each of their slot requests according to their operational needs and priorities. This mechanism then adjusts the reference schedule to meet airlines’ needs. A key aspect of the proposed mechanism is that it does not require airlines to disclose information regarding their internal valuation of slots, and that it produces airline Pareto Optimal (non-dominated) schedules.

Ribeiro et al. (2018) introduced a formulation which models the IATA hierarchical allocation in a more detailed manner. In addition to historic, new entrant, and other slot requests, the proposed model explicitly considers the “changes to historics” as a distinct slot priority, providing a more accurate representation of the hierarchical slot allocation. They formulated the model using a weighted quadr-objective function minimising the number of rejected slots, maximum displacement, schedule displacement, and the number of displaced slot requests. Pre-emptive objective function weights were used resulting in a lexicographic optimisation model. The resulting model was tested on two airports. For the tested problem instances the resulting models were reduced to tri-objective formulations since the capacity profile of the test airports did not lead to rejected slots. The reported computational results study the bi-objective trade-offs between total displacement and maximum displacement; and total displacement versus the number of displaced slot requests.

The slot scheduling problem has been also addressed in the US slot scheduling context (Jacquillat and Odoni, 2015; Jacquillat and Vaze, 2018; Pyrgiotis and Odoni, 2015). There are several commonalities between the IATA-based and the U.S-based models, i.e. the constraints and the objective functions. However, a distinguishing feature of the literature concerning the U.S, decision context (except for the priority-based structure), is that they schedule slots for a single day of operations, while IATA-based models consider the entire scheduling season. Pyrgiotis and Odoni (2015) proposed a model for scheduling slots at network level. This model minimises a scalarised objective function by jointly considering maximum and total displacement. Jacquillat and Odoni (2015) proposed an integrated approach which optimises simultaneously the airport slot schedule and utilisation of airport capacity. For the optimisation of the airport slot scheduling interventions they optimised lexicographically the maximum and total displacement objectives. The capacity utilisation part controls runway configuration so as to alter the arrival and departure service rates and eventually minimise tactical congestion costs. More recently, Jacquillat and Vaze (2018) introduced a multi-level model which studies the trade-off between efficiency and inter-airline equity, i.e. a fair balancing of displacement among airlines under on-time performance considerations. They minimise lexicographically maximum and total displacement metrics. Interestingly, on the second level of their solution approach they minimise lexicographically the maximum and total schedule displacement of each airline in a descending order, i.e. the airline with the largest maximum or total displacement are optimised first. Under this solution approach they are able to introduce fairness considerations applying the min–max fairness scheme proposed in Kalai and Smorodinsky (1975).

Ribeiro et al. (2019) proposed a weighted total displacement function which considers the total displacement objective for the historic, changes to historic, new entrants and other requests. However, this approach does not provide a systematic way for studying the trade-offs among the multiple levels of the slot request hierarchy regarding multiple objectives.

The literature review revealed that the state-of-the-art lacks solution approaches that can systematically investigate the multi-objective trade-offs among the slot request priorities and the schedule-wide performance of the airport. Existing research studies examine the effect of compromising the objectives of the different slot scheduling priorities either implicitly, by solving for all requests simultaneously (Zografos et al., 2017a; Ribeiro et al., 2018; Fairbrother et al., 2019; Zografos and Jiang, 2019) or by solely studying compromises using a single objective (Ribeiro et al., 2019).

3. Model formulation

In this section, we formulate a Tri-Objective Slot Allocation Model (TOSAM) which jointly considers total displacement, maximum displacement and demand-based fairness. Total displacement represents a commonly accepted measure of scheduling efficiency (Zografos et al., 2012), while maximum displacement acts as a metric of the quality of service of the airport (Jacquillat and Odoni, 2015; Jacquillat and Vaze, 2018; Ribeiro et al., 2018; Zografos et al., 2017a) minimising the worst-case displacement across all slot requests. Demand-based fairness is also considered, requiring that the displacement experienced by each airline should be proportional to the peak requests that it submits (Fairbrother et al., 2019; Fairbrother and Zografos, 2018a). Fairness is an important aspect of the airport slot allocation process, since IATA WSG require the fair and non-discriminatory treatment of the submitted requests (IATA, 2019). Interestingly, IATA and Airports Council International (ACI) have jointly agreed on the development of a slot allocation procedure that ensures the 'transparency, certainty, consistency, fairness and non-discrimination' of the airport slot allocation process (ICAO, 2016). In what follows we present the mathematical structure and notation of the TOSAM.

Input data sets

A : set of airlines requesting slots denoted by a
M : set of request series denoted by m with cardinality $|M|$
M_a : set of request series of airline a
$M^{Arr(Dep)}$: $M^{Arr} \cup M^{Dep} = M^{Total}$, set of arrival (departure) series
$P \subseteq M^{Arr} \times M^{Dep}$: set of paired requests (m_{Arr}, m_{Dep}) indexed by p
D : set of days in scheduling season denoted by d
D_m : $\{1, \cdots, |D_m|\}$, set of days that slot request m is to operate
$C = \{15, 30, 60\}$: set of capacity constraint durations indexed by c
$T = \{1, \cdots, |T|\}$: set of time intervals per day indexed by t, t'
K : $\{Arr, Dep, Total\}$ set of movement types denoted by$k$
M^k : set of request series of movement type k

Parameters

t_m : requested time for slot series m

$$A_t^d (D_t^d) = \begin{cases} 1, & \textit{if time } t \textit{ of day } d \textit{ is an arrival (departure) peak period} \\ 0, & \textit{otherwise} \end{cases}$$

$T_{max,p}, T_{min,p}$: maximum and minimum turnaround times of paired request p
$u^k_{d,t,c}$: capacity for movement k for period $[t, t+c)$ on day d based on time scale c

$$\eta_{d,m} = \begin{cases} 1, & \text{if request series } m \text{ is requested on day } d \\ 0, & \text{otherwise} \end{cases}$$

$\rho_a = \frac{\sum_{m\in M_a^{Arr}}\sum_{d\in D_m} A^d_{tm} + \sum_{m\in M_a^{Dep}}\sum_{d\in D} D^d_{tm}}{\max\left\{\sum_{m\in M^{Arr}}\sum_{d\in D_m} A^d_{tm} + \sum_{m\in M^{Dep}}\sum_{d\in D} D^d_{tm}, 1\right\}}$: the number of peak requests of airline a in relation to the total number of peak requests

Decision variables and expressions

$$x_{t,m} = \begin{cases} 1, & \text{if request } m \text{ is allocated to time } t \\ 0, & \text{otherwise} \end{cases}$$

$\mu_a = \frac{\sum_{m\in M_a}\sum_{t\in T}|t-t_m|x_{t,m}}{\sum_{m\in M}\sum_{t\in T}|t-t_m|x_{t,m}}\rho_a^{-1}$: the demand-based fairness index of airline a (Fairbrother et al., 2019; Fairbrother and Zografos, 2018a) expressing the share of total displacement of each airline should be proportional to the share of peak requests that it submits (ρ_a)

Base model

The objective function of TOSAM is:

$$\text{minimise } Z = \left\{ \overbrace{\sum_{m\in M}\sum_{t\in T}|D_m||t-t_m|x_{t,m}}^{Z_1}, \overbrace{\max_{\forall m\in M}|t-t_m|}^{Z_2}, \overbrace{\max_{\forall a\in A}|\mu_a - 1|}^{Z_3} \right\} \tag{3.1}$$

Subject to:

$$\sum_{t\in T} x_{m,t} = 1, \forall\ m\in M \tag{3.2}$$

$$\sum_{m\in M^k}\sum_{t^{`}\in[t,t+c-1]} \eta_{d,m}x_{t^{`},m} \le u^k_{d,t,c}, \forall\ k\in K,\ d\in D,\ c\in C,\ t\in[1,|T|-c+1] \tag{3.3}$$

$$T_{min,p} \le \sum_{t\in T}\left(x_{t,m_{Dep}}t\right) - \sum_{t\in T}\left(x_{t,m_{Arr}}t\right) \le T_{max,p}, \forall\ p\in P \tag{3.4}$$

Expression (3.1) minimizes the total displacement[5] (Z_1), maximum displacement (Z_2) and maximum deviation from the absolute value of demand-based fairness (Z_3). When the fairness index (μ_a) is less than one, then airline a is experiencing less displacement in relation to the peak requests that it has submitted(ρ_a). On the other hand, for values of μ_a greater than one, the displacement that the airline will experience is greater than the proportion of its slot requests at peak times. Therefore, objective function Z_3 is minimised, since we would like μ_a to take values close to one (value of absolute fairness). Constraints (3.2) ensure that each of the requests will be allocated to a time slot (Fairbrother and Zografos, 2018a; Fairbrother et al., 2019). Moreover, constraints (3.3) are rolling capacity constraints for each type of movement-constraint (arrival, departures, or total movements) meaning that the total number of movements scheduled within various time intervals (e.g.15 or 60 min), must not exceed the capacity of the airport for this movement and interval. Constraints (3.4) are turnaround time constraints which require that the time difference between two paired requests, should not be less than the lower turnaround time ($T_{min,p}$) or larger than the maximum turnaround time ($T_{max,p}$).

4. Solution methodology

In this section we present the multi-level, multi-objective airport slot allocation solution framework. Before diving deeper into the proposed solution methodology, herein we provide some key concepts of multi-objective programming underpinning the understanding of the algorithms and concepts presented in the remainder of the document.

4.1. Prerequisites: Concepts and definitions

In multi-objective programming the notion of optimality is expressed through the concept of *non-dominated solutions* (differently stated as Pareto optimal solutions).

[5] In (3.1), $|t-t_m|$ measures the absolute deviation from the requested time (t_m). Therefore, if request m is scheduled at time t', then $x_{t^{`},m} = 1$ and the displacement of m will be equal to $|t\text{'} - t_m|$.

Definition 1:. ***Non-dominated solution/schedule***: *A solution x' belonging to the feasible space of the problem* (X) *is said to be non-dominated or Pareto optimal if* $\nexists\ x \in X/x' : Z(x) \leq Z(x')$, *and* $Z_i(x) \leq Z_i(x')$ *for at least one objective* (i) *in the objective set.*

In addition, there are algorithms that do not necessarily result in non-dominated solutions having useful properties for practice (please note that we use the terms solution/point interchangeably). These are the weakly non-dominated and dominated points which are defined as follows.

Definition 2:. ***Weakly non-dominated solution/schedule***: *A solution x' belonging to the feasible space of the problem* (X) *is said to be weakly non-dominated or weakly Pareto optimal if* $\nexists\ x \in X/x' : Z(x) < Z(x')$.

As per Definition 2 a weakly non-dominated point exists if there is no other point that improves all objective functions simultaneously. If there is such a point x, then x' is said to be *dominated.*

4.2. Multi-level, multi-objective solution framework

To provide additional insights on the slot allocation problem, we will situate it in a multi-objective multi-level solution framework.

In a multi-level decision-making process, decisions of different entities belonging to different levels are applied sequentially with each of them optimising their respective objective functions independently. Decision entities whose decisions follow each other are referred to as *leaders* and *followers*. This architecture implies that while the leader is free to determine their own decision, the followers' decisions are constrained by the leaders' decisions. Even though leaders' decisions greatly affect followers' decision space, the decisions of the followers' do not explicitly affect the feasible decision space of the leaders. In a multi-objective multi-level problem, each decision entity is attempting to find an efficient solution with respect to multiple objectives. We observe that the IATA-based airport slot scheduling problem shares similar characteristics with the multi-level multi-objective decision-making problem defined above, since:

- The decision process is composed of interacting decision-making levels with a hierarchical structure, i.e. slot priorities;
- The decisions of the lower levels, i.e. the slots available to accommodate the levels' requests, are defined after and only after the decisions of the leading levels have been made;

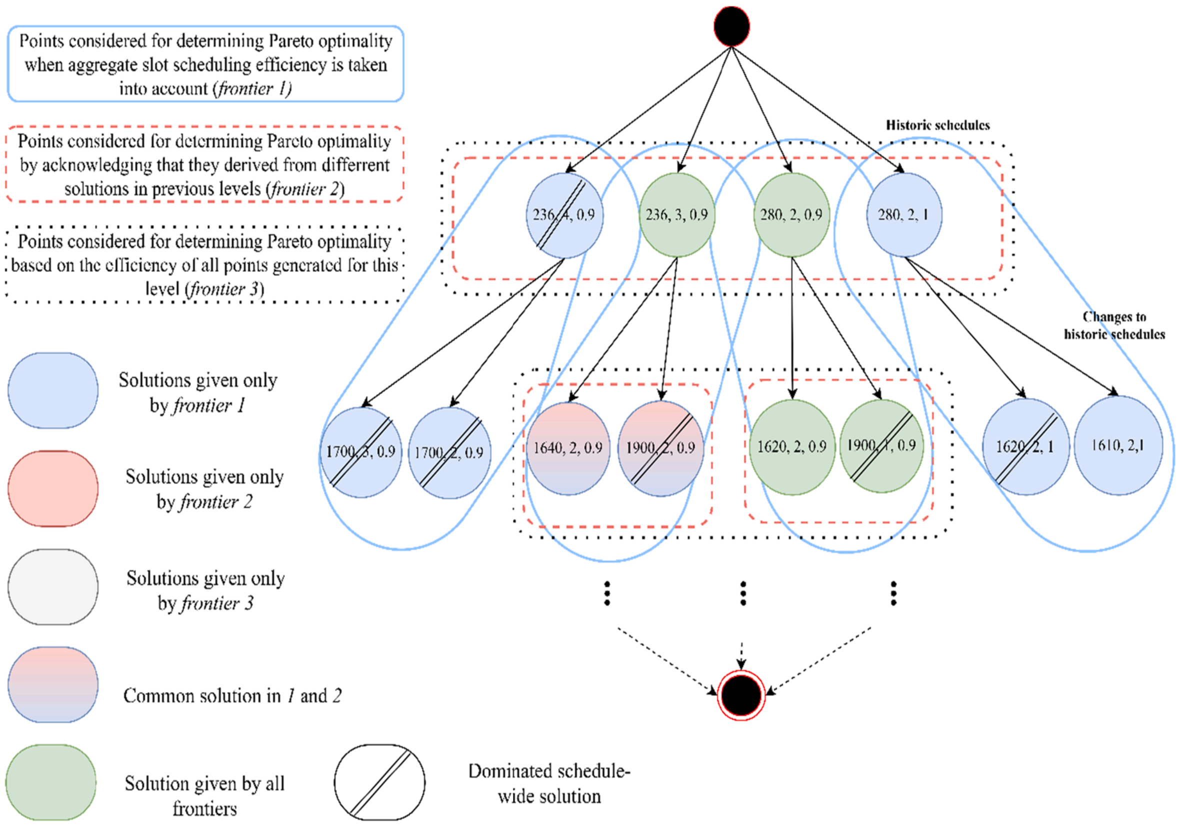

Fig. 1. An example illustrating the advantages of the proposed multilevel solution approach.

- For each level, the objectives are optimised independently without considering lower levels' actions, hence lower levels are constrained by the decisions of the leaders; and
- The quality of the schedule depends on multiple criteria, i.e. schedule displacement, maximum displacement etc.

The solution methodology that we propose differentiates from existing airport slot scheduling solution algorithms, since instead of selecting a single reference solution for each level (Fairbrother et al., 2019; Zografos and Jiang, 2019; Fairbrother and Zografos, 2018a), it reports a set of solutions with respect to multiple objectives and solves the corresponding optimisation problems. By acknowledging that the decisions made for the lower priority levels of the slot hierarchy depend on the decisions made at the higher levels, we report additional schedule-wide, non-dominated solutions than those explored by solely optimising the objective functions of each level under the hierarchical and lexicographic approaches.

To better illustrate how multilevel considerations may yield additional non-dominated slot scheduling solutions, we provide an illustrative toy-example (Fig. 1) considering the first two levels (historic and changes to historic requests) of the slot allocation process. The example illustrates that schedule-wide efficient solutions may occur from branches which in previous levels did not result in non-dominated schedules (outer right node of historic's level). In particular, in the first level (historic) of the example, only the middle two schedules are non-dominated (green shaded). However, the first and last points of the historic level are weakly dominated with respect to the maximum displacement and fairness objectives respectively. As a result, by considering only non-dominated schedules at the first level, we report only two non-dominated schedule-wide solutions. By discarding the weakly-dominated solutions at the upper level (historic), non-dominated solutions at the lower level are discarded as well. Hence, it may be worth examining weakly dominated solutions, if and only if they lead to non-dominated schedule-wide results, since they may result in additional schedule-wide solutions that are preferred by the stakeholders. When considering the weakly dominated solutions of the historic level, an additional schedule-wide non-dominated point is reported (outer right node of the changes to historic). Therefore, Fig. 1 suggests that by comparing solutions at the schedule-wide level, we obtain non-dominated schedule-wide solutions that would otherwise be eliminated (outer right node of changes to historics' level). Therefore, solution methods that consider points which are not necessarily non-dominated, can produce additional schedule-wide non-dominated solutions and result in frontiers of higher cardinality. This observation is formalised in the following proposition and remark.

Proposition 1.. *Non-dominated schedule-wide solutions may be generated from solutions which are not necessarily non-dominated at the preceding hierarchical levels.*

Proof.. *Assume that non-dominated schedule-wide solutions occur only from previous levels' solutions which are non-dominated. In addition, assume two airport slot schedules consisting of four decision-levels concerning historic (H), changes to historic (CH), new entrants' (NE) and others' (O) requests, denoted by s_1 and s_2. Now assume that s_1 is weakly non-dominated at the NE level by s_2 such that $Z_1^{NE}(s_2) < Z_1^{NE}(s_1)$ and $Z_i^{NE}(s_2) = Z_i^{NE}(s_1), i = 2, 3$ and share the same parent solutions at the H and CH levels. Having solved for the O level based on s_1 and s_2 we observe that $Z_1^O(s_1) < Z_2^O(s_2)$ and $Z_i^O(s_1) = Z_i^O(s_2)$, $i = 2, 3$. This means that s_2 is weakly dominated by s_1 if $Z_1^{NE}(s_1) - Z_1^{NE}(s_2) < Z_2^O(s_2) - Z_1^O(s_1)$. However, if this logical condition is true, it contradicts the initial hypothesis that non-dominated schedule-wide solutions occur only from solutions which are non-dominated at all decision-levels. Therefore, non-dominated schedule-wide solutions may occur from points that are weakly non-dominated at a priority level.*

Remark.. *Based on Proposition 1, solutions which are weakly non-dominated at a slot request priority are acceptable, if and only if they result in larger benefits in the following priority levels.*

The proposed multi-level framework is further detailed in section 4.4 where we present our solution algorithm.

4.3. Single objective transformation

Before proceeding to the multi-level algorithm for solving the hierarchical slot allocation problem with multi-objective considerations, there is need to provide linear transformations for the non-linear objective functions of TOSAM (Z_2 and Z_3), so that this problem can be solved by standard integer linear programming methods. Although the maximum displacement (Z_2) can be linearized through an additional auxiliary decision variable and constraints (Ribeiro et al., 2018; Zografos et al., 2017a), the fairness objective (Z_3), can only be linearized as a constraint (Fairbrother et al., 2019; Zografos and Jiang, 2019; Fairbrother and Zografos, 2018a, 2018b). Fortunately, efficient multi-objective optimisation algorithms such as the ε-constraint (Haimes, 1971) work by optimising one objective while constraining the others. Therefore, we transform TOSAM (rTOSAM) with the use of the following expressions:

$$minimise \quad Z_1 = \sum_{m \in M} \sum_{t \in T} |D_m||t - t_m| x_{t,m} \tag{4.1}$$

Subject to (3.2)-(3.4) and:

$$|t - t_m| x_{t,m} \le \varepsilon_{Z_2}, \forall t \in T, m \in M \tag{4.2}$$

$$\left(1 - \varepsilon_{Z_3}\right) \rho_a \sum_{m \in M} \sum_{t \in T} |t - t_m| x_{t,m} \le \sum_{m \in M_a} \sum_{t \in T} |t - t_m| x_{t,m} \le \left(1 + \varepsilon_{Z_3}\right) \ \rho_a \sum_{m \in M} \sum_{t \in T} |t - t_m| x_{t,m}, \forall \ a \in A \tag{4.3}$$

Expression (4.1) minimizes the total displacement objective. Constraints (4.2) and (4.3), aid in the linearization of objectives Z_2 and

Z_3 by taking as input upper bounds to their values $(\varepsilon_{Z_2}, \varepsilon_{Z_3})$.

4.4. Multi-level solution approach

Having provided linear expressions for our non-linear objectives (Z_2 and Z_3) we now propose an algorithm (Algorithm 2) solving the rTOSAM according to the hierarchical slot scheduling regime. In doing so, we denote $\Gamma : \{H, CH, NE, Oth\}$ to be the set of priorities where H, CH, NE and Othrepresent the standard slot coordination priorities, i.e. historic, changes to historic, new entrant and other slot requests respectively. Denoting by M^{γ} the set of requests of priority level $\gamma \in \Gamma$ the whole set of slot requests can be expressed as $M = \bigcup_{\gamma \in \Gamma} M^{\gamma}$.

In multi-level programming, finding efficient schedule-wide solutions, requires the consideration of all the feasible combinations of the decision variables that yield minimum values for the objectives at each level (Zhang et al., 2010). However, the exploration of all such combinations of feasible solutions, even for single objective problems would be intractable. Indeed, there are no general exact algorithms able to aid in the efficient solution of multi-level, multi-objective problems (Cappanera and Scaparra, 2010). As a result, previous research attempts solving bi-level, multi-objective problems have employed heuristic algorithms so as to provide sets of non-dominated solutions in more tractable solution times (Deb and Sinha, 2010).

Our approach differentiates from the existing heuristic approaches and is described as follows (Algorithm 2). For a fixed fairness level, our method calculates a set of schedules with schedule and maximum displacement objectives. For each of the generated schedules the remaining capacity of the airport is calculated (and u is updated), and weakly-dominated/non-dominated schedules are generated for the next priority (using Algorithm 1). Once all priority levels are processed, a schedule-wide solution constructed from the solutions of each priority level is added to list. This process is repeated for a given set of fairness levels (ε_{Z_3}) that are used to express the linear constraints for demand-based fairness (Fairbrother et al., 2019; Zografos and Jiang, 2016, 2019; Fairbrother and Zografos, 2018a). As a final step, the generated schedule-wide solutions are filtered out using the dominance criterion so that we end up with a list of schedules which are non-dominated with respect to the schedule-wide objectives.

Algorithm 1. *Sub-process for generating solutions at level γ based on fairness threshold $\varepsilon_{Z_3} = \delta$*

	input:M^{γ}, u, δ *# slot requests of level γ, capacity, fairness value*$(\varepsilon_{Z_3} = \delta)$
	output: Y^{γ}_{δ} *# list of solutions in level γ based on fairness value*δ
1	$Y^{\gamma}_{\delta} \leftarrow []$; *#initialise the set of solutions of level*γ
2	$LB^{\gamma}_{Z_2} \leftarrow rTOSAM(Z_2, cap = u, reqs = M_{\gamma}, Z_1 = None, Z_2 = None, Z_3 = \delta)Z_2$; *#lower bound of*$Z_2$
3	**if** $\gamma = NE$ **then**:
4	$UB^{\gamma}_{Z_2} \leftarrow max(1hour, LB^{\gamma}_{Z_2})$; *#for NE the upper bound of Z_2 is set equal to 1 h or equal to the lower bound*
5	**else**:
6	$UB^{\gamma}_{Z_2} \leftarrow rTOSAM(Z_1, cap = u, reqs = M_{\gamma}, Z_1 = None, Z_2 = None, Z_3 = \delta)Z_2$; *# upper bound of*$Z_2$
7	**for** $i \in \left[LB^{\gamma}_{Z_2}, UB^{\gamma}_{z_2}\right]$ **do**: *# iterate between the lower and the upper bound of*Z_2
8	$s \leftarrow rTOSAM(Z_1, cap = u, reqs = M_{\gamma}, Z_1 = None, Z_2 = i, Z_3 = \delta)$; *# optimise Z_1 subject to the current value of Z_2*
9	**if** s is feasible **then**:
10	**if** $i \neq s.Z_2$ **then**: *#i not binding the value of*Z_2
11	**Break**; *# stop loop*
12	**else**:
13	append s to Y^{γ}_{δ};
14	**return** Y^{γ}_{δ};
Notes:	# stands for comments on the algorithmic process; in **bold** are common algorithmic functions; solution of the transformation of TOSAM minimising objective Z_1 with rolling capacity u, priority requests γ, and ε constraints for the three objectives$\varepsilon_{Z_i}, i = 1, 2, 3$ respectively $\{rTOSAM(Z_1, cap = u, reqs = M_{\gamma}, Z_1 = \varepsilon_{Z_1}, Z_2 = \varepsilon_{Z_2}, Z_3 = \varepsilon_{Z_3})\}$; The dot (.) after lines 2 and 6 denotes the utilisation of the value of objective Z_2.

The process of finding schedules at a given priority level which are efficient with respect to schedule displacement and maximum displacement is detailed in Algorithm 1. Algorithm 1 extends the ε-constraint approach (Haimes, 1971) and builds on other solution approaches which include fairness metrics (Fairbrother et al., 2019; Zografos and Jiang, 2019; Fairbrother and Zografos, 2018a, 2018b). However, in contrast to previous research attempts, 'new entrants' maximum displacement is set to be between the lower bound obtained in line 2 of the algorithm and one hour. This limit to the maximum displacement of this slot priority is according to IATA's WSG, since new entrants' requests who are accommodated within one hour from their initially requested time, should either accept or lose the new entrant status (IATA, 2019).

Therefore, by limiting the displacement of new entrants to one hour (lines 3,4 of Algorithm 1) we serve two purposes. First, we address the aforementioned policy rule and second, the number of solutions to be considered is limited, thus reducing the complexity of the overall solution process. For a priority level γ and fairness value δ, Algorithm 1 is initiated by calculating the lower and upper bounds of maximum displacement (Z_2) under the given fairness level (lines 2, 6). Then, for each value of maximum displacement defined by these bounds, total displacement is optimised (line 8), thus providing a weakly dominated schedule. Building on

Proposition 1 and in accordance to the IATA's hierarchical allocation of slots, we have designed Algorithm 1 so as to introduce reasonable sacrifices (expressed in the form of weakly dominated schedules) at the upper levels of the hierarchy. This property of Algorithm 1 is formalised in the following proposition (Proposition 2).

```
input: M, u, Δ # slot requests, capacity, list of fairness values
output: Y # list of efficient schedule-wide solutions
Y←[]; # initialise the empty list of solutions
for δ ∈ Δ do:
    Y_δ^H←Algorithm 1(M^H, u, δ); # solve for historic requests
    for s_H ∈ Y_δ^H do:
        u^CH←copy of u; # re-initialise capacity
        update u^CH based on solution s_H;
        Y_δ^CH←Algorithm 1(M^CH, u^CH, δ); # solve for changes to historic requests
        for s_CH ∈ Y_δ^CH do:
            u^NE←copy of u;
            update u^NE based on solutions s_H, s_CH;
            Y_δ^NE←Algorithm 1(M^NE, u^NE, δ); # solve for new entrant requests
            for s_NE ∈ Y_δ^NE do:
                u^Oth←copy of u;
                update u^Oth based on solutions s_H, s_CH, s_NE;
                Y_δ^Oth←Algorithm 1(M^Oth, u^Oth, δ); # solve for other requests
                for s_Oth ∈ Y_δ^Oth do:
                    s_SW ←combine s_H, s_CH, s_NE, s_Oth into a schedule-wide solution;
                    append s_SW to Y;
filter out schedule-wide dominated solutions from Y;
return Y;
```

Notes: # stands for comments on the algorithmic process; in **bold** are common algorithmic functions; SW: schedule wide.

Algorithm 2. *A multi-level algorithm for the rTOSAM*

Proposition 2.. *If s (see line 8 of Algorithm 1) is feasible, then it is a weakly non-dominated solution.*

Proof.. *Suppose s is feasible but it is not a weakly non-dominated point. Then there must exist a solution* $s' \in Y_\delta^\gamma$ *having* $Z_i(s') < Z_i(s), i = 1, 2, 3$. *Therefore, based on the bounds set by the ε-constraints for the maximum displacement and demand-based fairness objectives (see line 8 of Algorithm 1),* $Z_i(s') < Z_i(s) \le \varepsilon_{z_i}, i = 2, 3$ *and* $Z_1(s') < Z_1(s)$, *which however contradicts the optimality of s.*

Additionally, when the maximum displacement of an airport slot schedule is known, we reduce the number of decision variables ($x_{t,m}$) needed, since some of them, will not be used (see Proposition 3 below). Therefore, we embed this property in line 8 of Algorithm 1, where instead of generating $x_{t,m} \forall t \in T, m \in M$ we only generate $x_{t,m} \forall t \in (t_m - i, t_m + i), m \in M$.

Proposition 3.. *Knowing the maximum displacement* (ε_{Z_3}) *that a schedule will receive, we may reduce the number of decision variables by* $|M||T| - 2|M|\varepsilon_{Z_3}$.

Proof.. *The maximum displacement of the problem solved in line 8 of Algorithm 1 is known, since* $\varepsilon_{Z_2} = i \in \left[LB_{Z_2}^\gamma, UB_{Z_2}^\gamma\right]$. *Therefore, in any feasible solution that has* $Z_2 \le \varepsilon_{Z_2}$, *it is true that* $\forall m \in M \nexists t > t_m + \varepsilon_{Z_2}, t < t_m - \varepsilon_{Z_2}$ *such as* $x_{t,m} > 0$. *Hence,* $\forall m \in M$ *only* $x_{t,m} \forall t \in (t_m - i, t_m + i)$ *are required. As a result, only* $2|M|\varepsilon_{Z_3}$ *variables are generated, reducing the number of decision variables by* $|M||T| - 2|M|\varepsilon_{Z_3}$.

The overall algorithm for generating a set of schedule-wide non-dominated solutions is detailed in Algorithm 2. The algorithm iterates over a list of fairness values (Δ) and generates schedules for each priority level of the scheduling hierarchy of IATA. Note that by using the same fairness values across all priority levels, we reduce the complexity of the search for efficient solutions and schedule all requests with the same fairness considerations no matter what their priority level is. Since all requests must be treated in a fair and transparent manner (IATA, 2019), that is to ensure that within each schedule all requests will be treated under a harmonised fairness threshold no matter what their priority level is.

At this point, it is important to note that each feasible solution provided by each iteration of the proposed algorithm is solved to optimality based on the current maximum displacement and fairness thresholds (optimality gap smaller that 1e-4).

5. Application and results

The application of Algorithm 2 provides valuable insights regarding: (i) the impact of schedule-wide performance considerations on the quality and spectrum of airport slot scheduling alternatives; (ii) the schedule-wide trade-offs between total schedule displacement, maximum displacement and demand-based fairness at a schedule coordinated airport; and (iii) the trade-offs between the objectives of the slot priorities (historic, changes to historic, new entrant, other) and their impact on the objective values of the reported schedule-wide slot scheduling solutions.

It is also important to underline that the goal of our analysis is to demonstrate the potential benefit of compromising the slot scheduling objectives of the upper-level slot request priorities in favour of the schedule-wide solution. Therefore, a better understanding of the interactions between the requests belonging to different levels of the slot hierarchies and the way that they influence the overall slot schedule of the airport is also analysed. The work presented in this paper, serves as a discussion tool that can provide multiple non-dominated solutions from which the stakeholders of the airport system can later choose the one that better expresses their interests. Therefore, the detailed solution approach can facilitate and inform discussions among the interested stakeholders and therefore reduce their conflicts when negotiating (e.g. during the bi-annual slot coordination conference).

In what follows this section describes the experimental setup and the data used for the allocation of slots at a real-world coordinated airport (Section 5.1); presents our experiments regarding points (i)-(iii) mentioned above (Section 5.2); and discusses the potential implications of our findings for decision and policy-making (Section 5.3). Finally, in section 5.4 we discuss criteria that can be used to filter out solutions based on the needs of the participating stakeholders.

5.1. Data and experimental setup

The data used concern the summer scheduling season of 2009 (from the 29th of March to the 26th of October) at a medium-sized regional European airport. The declared capacity of the airport is expressed using 15-minute and 1-hour intervals and can be seen in Table 1. Please note that in the following experiments the results concerning total and maximum displacement are expressed in 15-minute intervals.

In the airport that is studied, the declared capacity imposes that within an hour timeframe (e.g. 10:45 – 11:45) no more than ten movements can be scheduled in total, while for the 15-minute time scale, no more than three. The declared capacity also puts an upper bound on the number of arrivals and departures that can be effectuated per hour (four and six accordingly). Please note that between Friday and Sunday, the capacity of the airport increases allowing either 5 arrivals, 6 departures or 11 movements in total.

The second dataset that is essential for the application of our model is the initial request set. Table 2 provides a snapshot of the request set that exemplifies the structure of slot requests.

The distribution of requests within the request set per priority type is given in Table 3, where we observe that the new entrants' requests are the fewest among all request types and they include the smallest number of individual slots per request.

Furthermore, the absence of year-round requests[6] (action codes I, Y, V) indicates that the airport under analysis is schedule coordinated only during the summer period. This is also explained by the fact that the airport is in a remote touristic area, which is characterised by summer seasonality. To conclude, even though the airport that we analyse is small, it is coordinated, and the available request set concerns over 1,300,000 passengers and more than 7,600 flight pairs which are distributed over 449 series of paired requests or more than 15,000 individual slot requests.

We test our model and algorithm using data from the airport case study described above. For all experiments presented in the following sections, we use Gurobi 8.1.0 (Gurobi Optimization, 2018) as our integer programming solver. The model and the proposed solution algorithm are implemented in Python 3.7 programming language (Rossum, 1995) using the Anaconda distribution. The reported computational experiments were conducted on a computer having a 1.9-GHz Intel® i7-8650U central processing unit and 31.8 GB of RAM, running on Windows 10 pro edition. For each iteration we solve rTOSAM looking for exact solutions (optimality gap less than 1e-4).

We consider incremental changes (increment of 0.1) for discretised fairness levels (δ) ranging between the absolute value of fairness ($\delta = 0$) and the maximum fairness value (1.7 in this case) that generates non-dominated solutions. Regarding the priority levels, we study a set Γ including the standard slot scheduling priorities of IATA, i.e. $\{H, CH, NE, Oth\}$ (IATA, 2019). In addition, to reduce the size of the problem, we restrict the schedule-wide maximum displacement to 3.5 h (by modifying line 6 of Algorithm 1) since we observe that the trade-offs between schedule-wide maximum and total displacement become rather weak for all fairness levels when schedule-wide maximum displacement is greater than 2.5 h. Finally, regarding the turnaround time parameters used in constraints (3.4), i.e. $T_{max,p}$ and $T_{min,p}$, we consider them to be equal to the initially requested arrival and departure times ($t_{m_{Dep}} - t_{m_{Arr}}$). Please note that in graphs presented in the subsequent sections, the values of the maximum and total displacement objectives (Z_1, Z_2) are expressed in 15-minute intervals.

[6] With the term year-round operations, IATA (2019) refers to requests for operations to extend an existing operation into a year-round basis operation, i.e. for two scheduling seasons.

Table 1
Airport's declared capacity.

Runway movement type	15 min (1 interval)	60 min (4 intervals)
Arrivals	–	4 (5)
Departures	–	6
Total	3	10 (11)

Notes: Capacity during Friday and weekends in parentheses.

Table 2
Structure of the request set.

ID	1	2	3	4	5	6	7	8	9	10	11	12	13	14
Label	*AC*	*ANU*	*DC*	*DNU*	*HF*	*MHF*	*HT*	*MHT*	*M*	*T*	*W*	*H*	*F*	*S*
1	A1	1111	A1	1112	16	JUN	13	OCT	0	2	0	4	0	0
…	…	…	…	…	…	…	…	…	…	…	…	…	…	…
|*M*|	AN	9998	AN	9999	30	MAY	03	OCT	0	0	0	0	0	6
ID	15	16	17	18	19	20	21	22	23	24	25	26	27	28
Label	*U*	*SEN*	*TYP*	*AFR*	*BFR*	*AH*	*AM*	*DH*	*DM*	*V*	*ADE*	*BDE*	*FY*	*Q*
1	0	231	737	PRG	PRG	12	55	13	55	0	PRG	PRG	JJ	3
…	…	…	…	…	…	…	…	…	…	…	…	…	…	…
|*M*|	0	167	321	BLL	BLL	08	40	9	40	0	BLL	BLL	CC	1

Notes: Arrival / Departure Company (AC/DC), Arrival/ Departure Number (ANU/DNU), first/ last day of operations (HF/HT), first/ last month of operations (MHF/MHT), Monday (M), Tuesday (T), Wednesday (W), Thursday (H), Friday (F), Saturday (S), Sunday (U), Seats Expected (SEN), type of aircraft (TYP), airport of origin (AFR), last stopover airport (BFR), Arrival/Departure Hour (AH/DH), Arrival/ Departure Minute (AM/DM), Overnight indicator (1 if the aircraft will depart the next day, or 0 if it departs the same day) (V), next stopover airport (ADE), destination airport (BDE), Service codes for the arrival and departure flights (FY where J/F: schedule passenger/ cargo flight, C/H: chartered passenger/cargo flight, P: positional, X: technical, D: general or private, N: Business aviation/ air taxi), frequency indicator (Q).

Table 3
Requests per priority and action code.

		Request series		Individual requests	
Priority	Action Code	#	%	#	%
Historic	F	126	28.1%	4304	28.0%
Other	N	222	49.5%	7412	48.1%
Changes to historic	R	55	12.2%	2264	14.7%
	L	22	4.9%	748	4.9%
New entrant	B	24	5.3%	660	4.3%
All	Total	449	100%	15,388	100%

Notes: Changes to historic requests that are willing to accept slot times between the historic or the requested time (R), changes to historic requests that will only accept the historic slot if the requested time is not available (L) (IATA, 2019), percentage (%), number (#).

5.2. Results

The purpose of this section is to discuss the results of the multi-level, multi-objective decision framework proposed in this study. The section comprises three sub-sections. Section 5.2.1 compares the proposed solution methodology with an alternative multi-level approach and an existing algorithm used to obtain non-dominated sets of airport slot schedules (points). 5.2.2 discusses the schedule-wide trade-offs among the three considered objectives, while section 5.2.3 presents and analyses the trade-offs among the objectives of the different priority levels and the schedule-wide objectives. The decision and policymaking implications obtained by the computational analyses of this section are discussed in 5.3.

5.2.1. Comparison of alternative solution approaches

To demonstrate the potential improvement in the range of schedules found by considering solutions which are not necessarily non-dominated with respect to the objectives of the hierarchy, we compare the performance of Algorithm 2 with two alternative solution approaches.

(a) A frontier (identified as "*frontier 2*") obtained by moving to lower priority levels only if the solutions at a given priority are non-dominated in comparison to other points which share the same solutions in each of the leading priority levels (triangular cyan points). This approach is an alternative to Algorithm 2 that allows less interactions between the objectives of the slot priorities and the schedule-wide solution. The approach for generating *frontier 2* is detailed in Algorithm 3 (see Appendix)

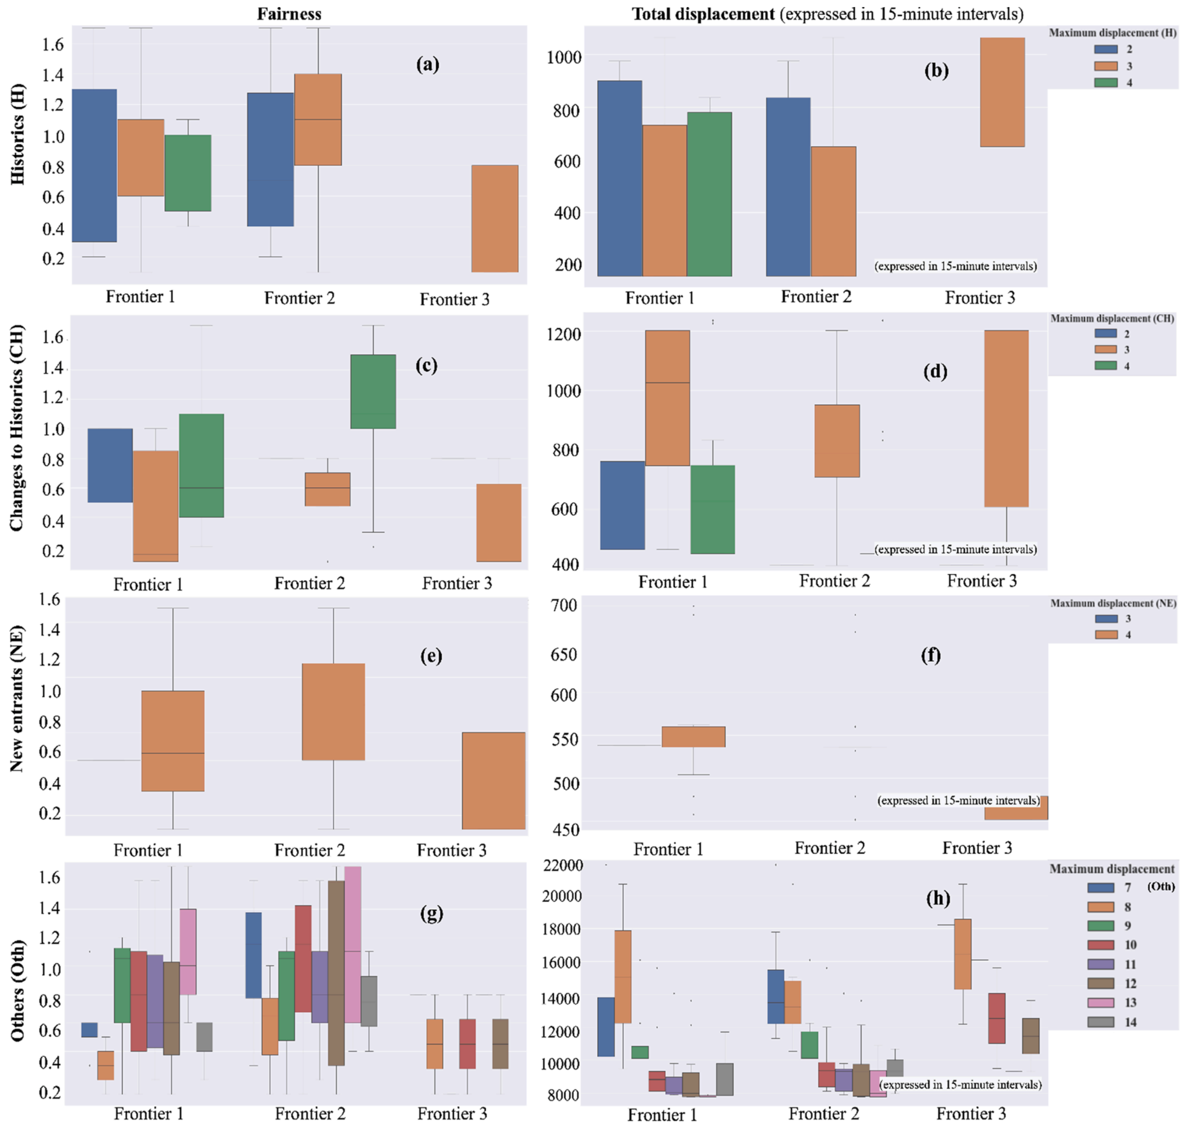

Fig. 2. Boxplots summarising the solutions reported by the three frontiers.

(b) A frontier (identified as "*frontier 3*") obtained by moving to lower slot priority levels only if all leading levels result in non-dominated solutions (black crossed points). This approach is detailed in Algorithm 4 (see Appendix) and is similar to the considerations of multi-objective papers existing in the literature (Zografos et al., 2012, 2017a; Ribeiro et al., 2018; Zografos and Jiang, 2019; Ribeiro et al., 2019; Fairbrother et al., 2019).

In Fig. 2 we provide boxplots summarising the solutions reported by Algorithms 2–4. We first observe that the cardinality (number of points) of *frontiers 1* and *2* is significantly larger than that of *frontier 3*. Especially, *frontier 3* is dominated by *frontiers 1* and *2* since they report solutions which would have never been reported by *frontier 3.* In particular, *frontier 3* reported non-dominated solutions that solely correspond to two demand-based fairness values (equal to 0.1 and 0.8). Fig. 2 demonstrates the increased cardinality

Table 4
Comparative table of the three frontiers.

Frontier	Time to obtain (mins)	Cardinality	Hyper-volume
frontier1	4870.1	57	139,186.5
frontier 2	1723.9	59	126,949.7
frontier 3	471.3	9	98,450.6

Notes: For the calculation of the hyper-volume, we use the common nadir point of all frontiers.

resulting from the consideration of the multi-level interactions among the objectives of the slot scheduling priorities. However, in order to establish the dominance of the proposed multi-level solution approach (Algorithm 2), in a subsequent part of this section we provide quantitative evidence using a suitable aggregate metric, i.e. the hyper-volume indicator. For the historics' level (sub-plots (a) and (b) of Fig. 2) we observe that *frontier 1* dominates its counterparts in terms of cardinality and solution quality. Besides, in sub-lot (a) we observe that *frontiers 1* and *2* outperform approaches that consider only non-dominated solutions at each level *(frontier 3)*. Also, we observe that *frontier 1* not only matches the quality of *frontiers 2* and *3* but in addition provides additional solutions (hued by green colour). The total displacement (sub-lot (b)) of *frontiers 1* and *2* is significantly lower (more than 60%) than pre-existing approaches (*frontier 3*) that reported similar maximum displacement (reduced from 600 X 15-minute intervals to 200). Similar findings are reported for the changes to historics' (*CH*) level (sub-plots (c) and (d) of Fig. 2). At the new entrants' level (*NE*), *frontier 1* reports the same values of maximum displacement with frontiers 2 and 3 but requires less deviation from the absolute value of fairness (the median value of demand-based fairness for maximum displacement equal to 4 is by 33% smaller in comparison to frontier 2 and by 25% smaller in comparison to frontier 3). Furthermore, the sacrifices made at the upper priority levels (H, CH) allowed the generation of solutions with improved maximum displacement (33% reduction). At the other's level, the improved trade-offs reported by *frontier 1* (for both demand-based fairness and maximum displacement) are visually evident (see sub-plots (g) and (h) of Fig. 2), since the boxplots representing *frontier 1* match the quality of the ones of *frontiers 2* and *3* and in several cases lie below the median or the 1st quartile of their counterparts; hence, reporting a frontier of improved quality.

Additionally, at the schedule-wide level, the sacrifices that the upper slot priorities (historics', changes to historics' levels) accept in their objective functions yield improved overall airport slot scheduling efficiency, resulting in significantly improved total displacement and demand-based fairness values for all values of maximum displacement. By comparing *frontiers* 1–3 we observe that demand-based fairness is significantly reduced, i.e. having lower values for all considered values of maximum displacement. For instance, when maximum displacement is equal to 7 (unit of measurement for displacement metrics is 15 min), *frontier 1* results in an average fairness value of 0.62 in contrast to *frontier 2* which reports a significantly increased value equal to 1.07 (72.5% increase). In contrast, *frontier 3* contained only one solution with maximum displacement equal to 7. The same impact is also observed for Z_1 since for most reported values of maximum displacement, *frontier 1* yields lower average total displacement values than *frontiers 2* and *3.* Using as an example schedules with maximum displacement equal to 7, we observe that *frontier 1* results in 4.45% improvement (15271.8) in terms of Z_1 in comparison to *frontier 2* (15951.9).

In Table 4 we consider three measures which are suitable for comparing three-dimensional efficient frontiers. These are the computational times required, the cardinality and the hyper-volume covered by each front. Hyper-volume can be used to assess efficient frontiers having more than two dimensions (Zitzler et al., 2003), since it is able to capture in a single scalar the distance of the solutions from the reference point and their spread across the objective space. In this occasion that the complete set of schedule-wide non-dominated points is not known a-priori, we must compare the sets of efficient solutions based on a pre-defined reference point. The reference point that is usually chosen is either the nadir point of the investigated frontiers, or a point with slightly worse coordinates for all objectives (Cao et al., 2015). In this approach we calculate the hyper-volume of each frontier based on their distance from their common nadir point using the LebMeasure algorithm discussed in While et al. (2006).

Comparisons based on metrics used for assessing the quality of the three frontiers, illustrate that *frontier 1* is better with respect to the quality and diversity of its schedule-wide solutions (Table 4). By observing the reported computational times for each approach, it is obvious that *frontier 1* is far more expensive than its counterparts demanding 3.38 days and the generation of 1052 schedule-wide solutions before terminating. However, it should be noted that the generation of each schedule-wide solution considered for *frontier 1* (number of schedule-wide solutions contained in Y before filtering), required on average less than 5 min to solve. For the generation of *frontiers 2* and *3,* the associated solution algorithms (Algorithm 3 and Algorithm 4) examined 373 and 12 schedule-wide solutions respectively. The length of the decision horizon available to the coordinators (22 days) and the fact that the resulting frontier demonstrated improved hyper-volume in comparison to *frontiers 2* (9.6%) and *3* (41.4%), renders the additional computational costs viable for the considered problem instance. We observe that there is an issue with the scalability of Algorithm 2, since for larger instances the increased number of schedule-wide solutions and the subsequent computational times would be impractical. This observation supports the need for suggesting suitable heuristic and schedule/solution elicitation solution algorithms.

Despite its computational complexity, Algorithm 2 is essential to obtain the frontier of the highest quality (reporting the best hyper-volume) and therefore can be used as a benchmark for assessing the performance of heuristic algorithms. In addition, the proposed approach requires no prior information about the preferences of the interested stakeholders and generates a spectrum of efficient alternative solutions which explore the trade-offs of the objectives, the dynamics and the structure of the problem, thus guiding stakeholders to more informed decisions. Given the insights provided by the generated set of efficient solutions, the decision makers can express afterwards their preferences (*a-posteriori* articulation of preferences) so as to select the airport slot schedule that better suits their needs. The increased number of non-dominated solutions provides a wide spectrum of options to all interested parties and facilitates the choice of the most preferable solution. Therefore, the fact that the proposed framework assumes no preferences with respect to the values of the objectives and produces multiple solutions, is in accordance with the structure of the current decision-making process and can support the activities during the initial slot allocation. We further discuss potential criteria for selecting slot schedules from a set of alternatives in section 5.4. Alternatively, the informed views of the stakeholders can be articulated and used in a bi-level hybrid algorithm that will report reduced computational times and improved scalability for larger airport instances. The output of the proposed algorithm can be used to guide the opinions of the stakeholders which will be then used as bounds or goals in algorithms which are designed to make use of *a-priori* information producing a single solution in viable computational times, i.e. goal programming approaches. For more information on the synergies between a-posteriori and a-priori multi-objective solution algorithms we refer the reader to the review papers of Marler and Arora (2004) and Ehrgott (2005). It is worth noting that in the absence of stakeholders' preferences Algorithms 3 and 4 and the subsequent frontiers produced by their application (*frontiers 2* and *3*) are valuable for larger/harder cases of airport instances and provide a 'starting point' supplying limited yet informative insights.

This proposed multi-level, multi-objective solution approach and the analyses enabled by the information that it provides, allow the systematic study of the trade-offs among the schedule-wide airport slot scheduling objectives. In addition, by proposing multiple non-dominated schedules, it provides alternatives that can study the trade-offs between the objectives of each priority and the schedule-wide performance across all priority levels. This decision-support capability may enable a more collaborative slot allocation process which can cater the differing requirements of the participating actors, i.e. airlines, coordinators and airport/local authorities. For instance, in the case that a certain airline (or multiple airlines) is/are not satisfied with its/their allocated slots, the coordinator may use an alternative solution (given as a non-dominated point) that is more acceptable to them. In addition, the proposed methodology can be used by all interested parties to study the impact of alternative allocations (with potential sacrifices at a certain slot scheduling priority) on schedule-wide results and allocations of lower priority levels. Besides, the architecture of the proposed algorithm and the explicit recognition of schedule-wide performance allow stakeholders to cooperate and achieve the improved utilisation of airport capacity. This is enabled through the seamless observation of the materialisation of their sacrifices which can inform decisions and lead to a mutually beneficial consensus. Further insights on the implications of the proposed methodology for policy and practice are provided in section 5.3.

In this section we have proposed, implemented and applied three alternative multi-objective solution algorithms for producing airport slot schedules. Albeit, from the analyses presented in the previous sections, we have demonstrated that *frontier 1* (generated by the multi-level considerations proposed in this work) dominates its counterparts. Hence, in the following sections, frontier 1 is discussed because of its diversity in terms of solutions, its improved hyper-volume (Table 4) and and its ability to provide decision support regarding the inter-priority objective trade-offs. *frontier 1* is obtained by considering the interactions among the objectives of the different slot request priority levels (using Algorithm 2) and provides insights on trade-offs that are not studied previously by the literature, which are further detailed in the sections that follow.

5.2.2. Trade-offs between the schedule-wide objectives

We now examine in more detail the trade-offs between the objectives of the schedule-wide solutions, i.e. schedule displacement (Z_1), maximum displacement (Z_2) and demand-based fairness (Z_3). These trade-offs are illustrated in two scatter plot matrices hued by the maximum displacement (Fig. 3) and demand-based fairness (Fig. 4) of each point.

The intensity of the trade-offs between the objectives fluctuates according to their values. For each fairness level, the trade-off between maximum (Z_2) and total displacement (Z_1) is rather weak for maximum displacement values above 9 (Z_1 is improved by 3.1% for each unit of Z_2 increased). However, when maximum displacement moves from 7 to 8 (from 8 to 9) we observe on average 17.1% (17.7%) improvement of total displacement for all fairness levels (see sub-plots (b) and (d) of Fig. 3). Another observation is that the trade-off between maximum and total displacement is reinforced when demand-based fairness is taking smaller values, therefore being more binding. By observing sub-plot (b) of Fig. 4 we see that as demand-based fairness improves, the trade-off between total and maximum displacement intensifies. Namely, when demand-based fairness is equal to 0.1, the average benefit to total displacement

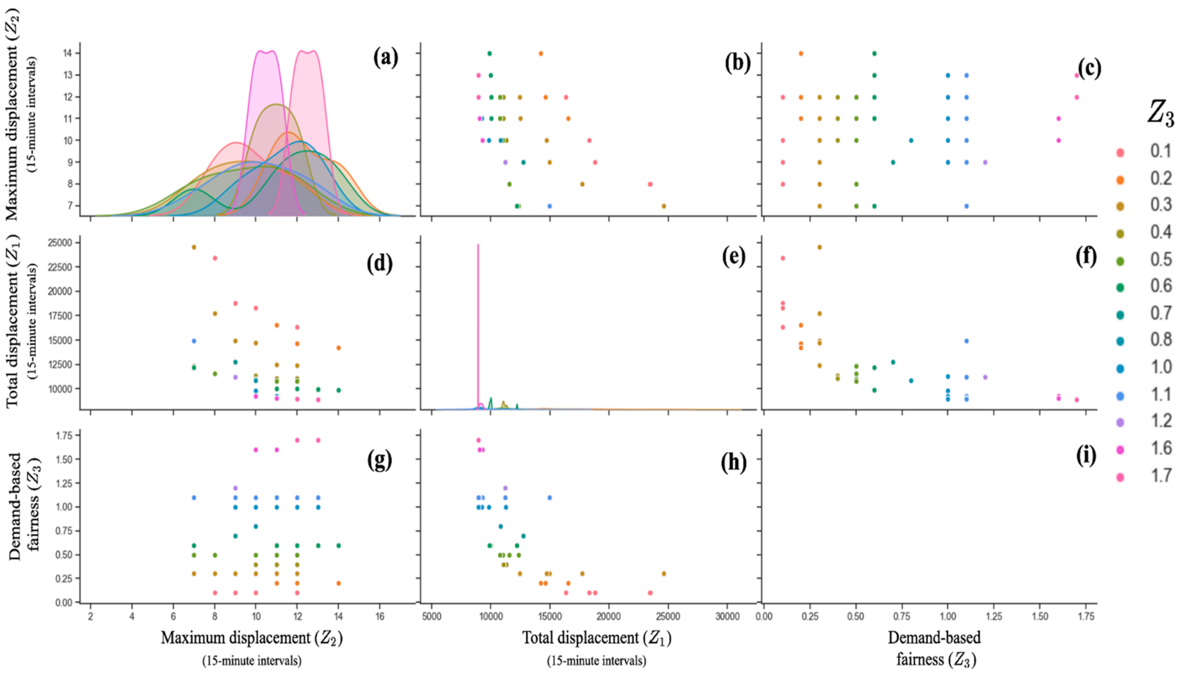

Fig. 3. Schedule-wide trade-offs hued by maximum displacement.

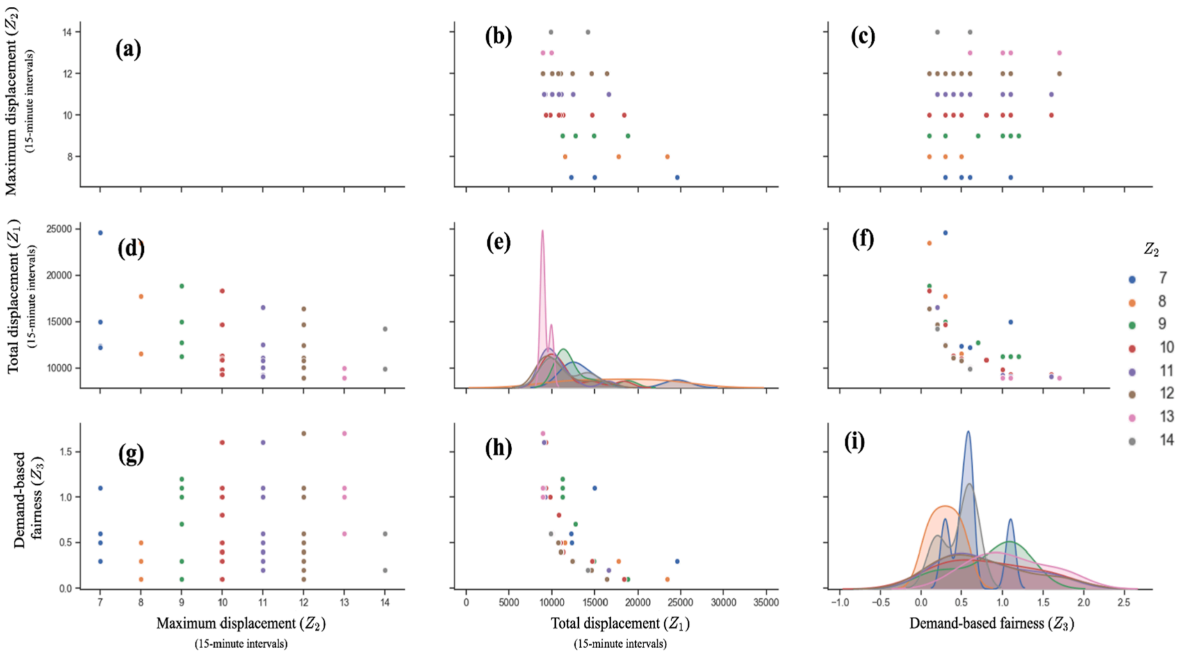

Fig. 4. Schedule-wide trade-offs hued by fairness.

from the marginal increases of maximum displacement is 11% (2360 units) while when demand-based fairness is 1.0 this benefit drops to just 5%. Such observations regarding the schedule-wide relationship between total and maximum displacement extend the findings of Zografos et al. (2017a) and Ribeiro et al. (2018) which reported the existence of strong trade-offs among the objectives of the slot scheduling priorities without considering the schedule-wide effects. This is an advantage of the proposed framework over existing studies, since instead of studying the trade-offs among the objectives of each priority level in isolation, it considers only the solutions (and the respective trade-offs) that result in improved schedule-wide performance.

In addition, we observe that the trade-off between demand-based fairness and total displacement is rather strong for low (better) fairness values (below 0.4) having an average benefit of almost 15.2% in terms of total displacement when fairness is marginally deteriorated by 0.1 no matter what the level of maximum displacement is. In contrast, for higher fairness values there are imperceptible improvements regarding total displacement. For example, when we move from 1.6 to 1.7 the average benefit for all maximum displacement values in terms of total displacement is just above 2.5%. This relationship is illustrated on sub-plot (f) in Fig. 4. The trade-offs between the three objectives, suggest that when demand-based fairness is allowed to take larger values, maximum displacement can receive reduced values at a small expense (5%) with respect to total displacement. This observation implies that in occasions where tight fairness considerations are not meaningful or result in significant total-displacement expenses, solutions with worst (higher) demand-based fairness values can be selected and to a certain extent satisfy both total and maximum displacement objectives.

5.2.3. Trade-off relationships between the objectives of the priorities and the schedule-wide efficiency

Another significant set of observations stems from the examination of the trade-offs between the objective values obtained for each of the levels (Table 5). This table also acts as a heat map that illustrates the relative percentage gap of each of the reported non-dominated solutions. Warmer colours represent bigger relative percentage gaps and therefore more intensive sacrifices. In addition, for discussion purposes we include a column named 'AD' that sums the deviations of the aggregate objectives. To provide more insights on the performance of the produced schedules, we include columns 'DS' and 'Z_1/DS' which report the number of displaced slot requests and the average displacement per displaced slot request accordingly. The relative percentage gap of each cell $i$belonging to column j of the table is calculated as (with $f(i,j)$ being the value of solution i with respect to metric j):

$$100 \times \left[f(i,j) - \min_{\forall i \in j}\{f(i,j)\}\right]\left[\max_{\forall i \in j}\{f(i,j)\} - \min_{\forall i \in j}\{f(i,j)\}\right]^{-1} \tag{5.1}$$

Table 5 shows that several schedule-wide non-dominated solutions of the reported efficient frontier, are given by solutions which are dominated in some levels (e.g. schedules 2 are 3 are dominated at the changes to historics' level by schedule 5, schedule 4 is dominated at the historics' level by schedule 6). It is also interesting that the schedule with the least displacement per displaced slot (schedule 35) was obtained by accepting a sacrifice of 2 units (30 min) of maximum displacement at the *CH*level. Similarly, schedules with the minimum number of displaced slots (see schedule 42) were given by accepting a deterioration of 30 min at the maximum displacement of *CH*.

Fig. 5 sheds light on the relationships between the slot scheduling objectives of the slot priorities and their influence on various

Table 5
Heat map with the trade-offs among the objectives of the levels and the aggregate slot schedule.

ID	H		CH		NE		Oth		$Schedule\text{-}wide$					
	Z_1	Z_2	Z_1	Z_2	Z_1	Z_2	Z_1	Z_2	Z_1	Z_2	Z_3	AD	DS	Z_1/DS
1	68.7	100.0	39.6	0.0	43.0	100.0	17.6	0.0	21.9	0.0	25.0	46.9	18.0	59.6
2	63.4	50.0	38.0	100.0	33.1	0.0	17.2	0.0	21.1	0.0	31.3	52.3	19.4	52.9
3	63.4	50.0	38.0	100.0	33.1	100.0	17.2	0.0	21.1	0.0	31.3	52.3	19.1	53.9
4	168.7	100.0	39.6	0.0	43.0	100.0	11.9	14.3	16.9	14.3	25.0	56.2	20.5	39.9
5	74.9	100.0	22.5	100.0	41.3	100.0	10.9	42.9	15.4	42.9	18.8	77.0	20.5	36.6
6	81.9	0.0	100.0	100.0	42.1	100.0	31.7	28.6	38.5	28.6	12.5	79.5	28.1	63.5
7	68.7	100.0	39.6	0.0	43.0	100.0	7.9	42.9	13.2	42.9	25.0	81.1	16.9	42.5
8	81.9	0.0	100.0	100.0	42.1	100.0	51.5	14.3	56.3	14.3	12.5	83.1	64.3	18.0
9	74.9	100.0	22.5	100.0	41.3	100.0	9.0	57.1	13.7	57.1	18.8	89.6	19.7	35.1
10	58.7	50.0	98.7	100.0	0.0	100.0	18.4	28.6	24.4	28.6	37.5	90.5	22.8	50.1
11	100.0	50.0	95.7	50.0	8.7	100.0	58.9	28.6	63.2	28.6	0.0	91.8	81.2	4.9
12	81.9	0.0	100.0	100.0	42.1	100.0	30.0	42.9	36.9	42.9	12.5	92.2	43.5	24.3
13	81.9	0.0	100.0	100.0	42.1	100.0	14.2	57.1	22.7	57.1	12.5	92.4	34.6	17.6
14	68.7	100.0	39.6	0.0	43.0	100.0	6.4	57.1	11.9	57.1	25.0	94.0	16.3	40.9
15	63.4	50.0	38.0	100.0	33.1	0.0	1.7	57.1	7.1	57.1	31.3	95.5	18.3	23.7
16	63.4	50.0	38.0	100.0	33.1	100.0	1.7	57.1	7.1	57.1	31.3	95.5	17.4	26.0
17	54.4	100.0	51.0	50.0	19.0	100.0	7.4	42.9	12.2	42.9	43.8	98.8	14.6	47.2
18	0.0	100.0	1.8	0.0	32.2	100.0	16.4	28.6	14.9	28.6	56.3	99.7	4.2	100.0
19	0.1	0.0	0.0	100.0	32.2	100.0	42.8	0.0	38.6	0.0	62.5	101.1	52.8	11.0
20	100.0	50.0	95.7	50.0	8.7	100.0	55.4	42.9	60.1	42.9	0.0	103.0	67.7	18.1
21	74.9	100.0	22.5	100.0	41.3	100.0	9.0	71.4	13.7	71.4	18.8	103.9	18.0	40.0
22	0.0	100.0	0.8	100.0	32.2	100.0	6.4	42.9	5.8	42.9	56.3	104.9	0.8	89.0
23	0.1	0.0	0.0	100.0	32.2	100.0	16.4	28.6	14.8	28.6	62.5	105.8	5.3	93.6
24	0.0	50.0	0.0	100.0	32.2	100.0	16.4	28.6	14.8	28.6	62.5	105.8	4.2	99.6
25	81.9	0.0	100.0	100.0	42.1	100.0	13.9	71.4	22.4	71.4	12.5	106.3	30.6	25.4
26	100.0	50.0	95.7	50.0	8.7	100.0	91.6	14.3	92.7	14.3	0.0	107.0	100.0	14.7
27	0.0	100.0	0.0	100.0	32.2	100.0	2.5	42.9	2.2	42.9	62.5	107.6	2.2	69.5
28	68.7	100.0	39.6	0.0	43.0	100.0	6.4	71.4	11.8	71.4	25.0	108.3	14.9	45.4
29	63.4	50.0	38.0	100.0	33.1	0.0	1.6	71.4	7.1	71.4	31.3	109.7	17.4	25.9
30	63.4	50.0	38.0	100.0	33.1	100.0	1.6	71.4	7.1	71.4	31.3	109.7	19.4	20.5
31	0.1	0.0	0.0	100.0	32.2	100.0	16.4	28.6	14.7	28.6	68.8	112.1	4.2	99.6
32	0.0	50.0	0.0	100.0	32.2	100.0	16.4	28.6	14.7	28.6	68.8	112.1	4.8	96.5
33	90.2	0.0	48.7	100.0	95.9	100.0	44.6	57.1	48.8	57.1	6.3	112.2	55.3	21.6
34	81.9	0.0	100.0	100.0	42.1	100.0	100.0	0.0	100.0	0.0	12.5	112.5	99.4	22.2
35	90.2	0.0	48.7	100.0	95.9	100.0	30.9	71.4	36.5	71.4	6.3	114.2	58.4	0.0
36	0.0	100.0	0.8	100.0	32.2	100.0	2.3	57.1	2.1	57.1	56.3	115.5	0.0	81.0
37	100.0	50.0	95.7	50.0	8.7	100.0	41.4	71.4	47.5	71.4	0.0	119.0	58.4	15.1
38	0.0	100.0	0.0	100.0	32.2	100.0	0.9	57.1	0.8	57.1	62.5	120.5	2.5	63.5
39	63.4	50.0	38.0	100.0	33.1	0.0	1.1	85.7	6.6	85.7	31.3	123.6	19.1	20.2
40	63.4	50.0	38.0	100.0	33.1	100.0	1.1	85.7	6.6	85.7	31.3	123.6	18.3	22.5
41	0.0	100.0	1.8	0.0	32.2	100.0	0.1	71.4	0.2	71.4	56.3	127.9	0.8	69.7
42	0.0	100.0	1.8	50.0	32.2	100.0	0.1	71.4	0.2	71.4	56.3	127.9	0.0	74.2
43	0.0	100.0	0.0	100.0	32.2	100.0	0.1	71.4	0.1	71.4	62.5	134.0	0.6	70.8
44	63.4	50.0	38.0	100.0	33.1	0.0	0.6	100.0	6.2	100.0	31.3	137.4	17.4	23.8
45	63.4	50.0	38.0	100.0	33.1	100.0	0.6	100.0	6.2	100.0	31.3	137.4	17.4	23.8
46	0.1	0.0	0.0	100.0	32.2	100.0	2.5	42.9	2.2	42.9	93.8	138.8	2.2	69.4
47	0.0	50.0	0.0	100.0	32.2	100.0	2.5	42.9	2.2	42.9	93.8	138.8	2.5	67.9
48	90.2	0.0	49.7	50.0	100.0	100.0	27.9	100.0	33.9	100.0	6.3	140.2	53.9	2.6
49	0.0	100.0	1.8	0.0	32.2	100.0	0.0	85.7	0.1	85.7	56.3	142.1	2.5	61.1
50	0.0	100.0	1.8	50.0	32.2	100.0	0.0	85.7	0.1	85.7	56.3	142.1	3.4	57.1
51	0.0	100.0	0.0	100.0	32.2	100.0	0.0	85.7	0.0	85.7	62.5	148.2	0.6	70.7
52	0.1	0.0	0.0	100.0	32.2	100.0	0.9	57.1	0.8	57.1	93.8	151.7	2.5	63.4
53	0.0	50.0	0.0	100.0	32.2	100.0	0.9	57.1	0.8	57.1	93.8	151.7	2.8	62.0
54	0.1	0.0	0.0	100.0	32.2	100.0	0.1	71.4	0.1	71.4	100.0	171.5	1.4	66.4
55	0.0	50.0	0.0	100.0	32.2	100.0	0.1	71.4	0.1	71.4	100.0	171.5	1.7	64.9
56	0.1	0.0	0.0	100.0	32.2	100.0	0.0	85.7	0.0	85.7	100.0	185.7	1.1	67.6
57	0.0	50.0	0.0	100.0	32.2	100.0	0.0	85.7	0.0	85.7	100.0	185.7	1.1	67.6

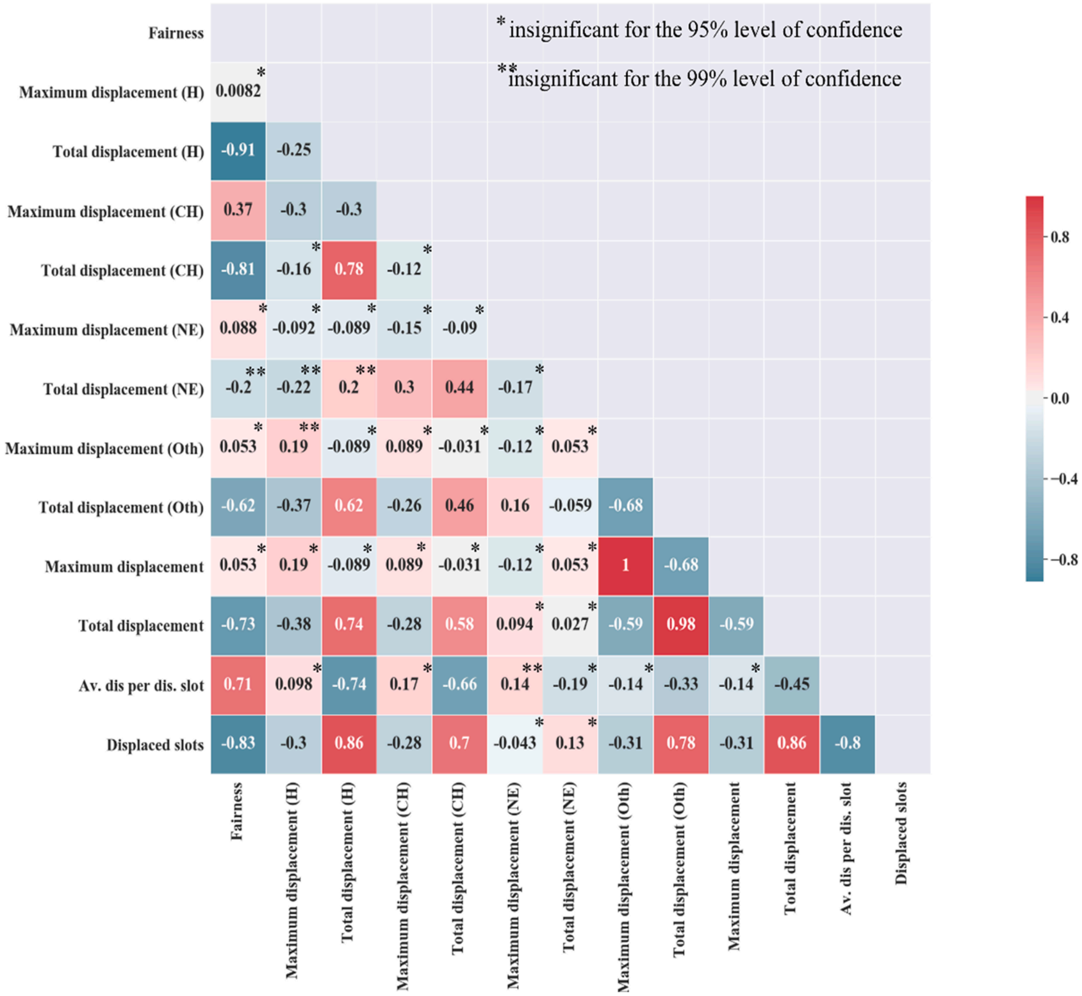

Fig. 5. Correlation heatmap between the measures of the reported efficient frontier.

metrics that are representative of the quality of the slot schedules produced for the studied airport. From this figure, slot coordinators and other interested stakeholders understand which slot priorities and objectives have to be compromised so as to improve any metric of interest. For example, in the examined airport, slot coordinators may increase the total displacement of *H* and *CH* in order to improve the average displacement per displaced slot or decrease it in order to get fewer displaced slots. Furthermore, we observe that the minimisation of the total displacement (Z_1) for the *H* level helps to a certain extent the minimisation of this objective for the *CH* and *Oth* levels (since they are positively associated), while it does not impact the *NE* level. A strong positive correlation exists between the number of displaced slot requests and total displacement (Z_1) of *H*. Similar but weaker correlations are reported for the total displacement of *CH*. The influence of *NE*'s objective functions to the other levels is minor and statistically insignificant due to the small number of requests falling into this slot priority (Table 3). Yet, a positive correlation with the objectives of *CH* is reported. The implications of our findings for practitioners and policymakers are further discussed in the following section.

5.3. Policy and decision-making implications

The increased cardinality and quality of the reported frontier that resulted from the multi-level considerations allow the extraction of useful insights that can support policy and practice. In the following sections, we provide both aggregate and detailed analyses and their implications.

5.3.1. Potential impact of findings on practice and policymaking

The information (that would otherwise be difficult to process by decision-makers) is summarised using suitable visualisations and aggregation metrics (Fig. 5) and is further discussed below. The analyses included below can be used by decision-makers to draft policy rules that optimise schedule-wide efficiency without compromising the interests of airlines possessing requests with historical rights or airlines that have received the 'new entrant' status. Furthermore, the intensity of the reported trade-offs that is discussed below may provide the based for the development of policy rules that apply to different types of airports and improve efficiency and capacity utilisation.

Implications of demand-based fairness Fairness in airport slot scheduling has been a central issue in airport slot allocation regulations (see IATA WSG) and deliberations among theinterested groups of stakeholders (ACI Europe, 2017). Our original findings associated with this metric are grouped and reported below.

(a). Inter-priority effects: From the summarised findings of Fig. 5, it appears that the improvement of fairness is associated with increased total displacement for all priority levels. However, the intensity of these associations is not homogeneous. For instance, we note that the effect of fairness is more significant at the upper priority levels of the slot scheduling hierarchy (*H*, *CH*), where we demonstrate significant and strong negative correlations (-0.91 and −0.81 respectively). On the contrary, the lower levels of the hierarchy (Oth, NE) report smoother trade-offs for the two objectives. Notably, the correlation between the two metrics for the new entrants' level is weak and shows that the considered fairness metric may result in more favourable outcomes for *NE* than the other priority levels. Concerning the maximum displacement objective, it appears that the two metrics do not demonstrate significant associations regardless of the priority. The only correlation that is statistically significant concerns the maximum displacement for the *CH* level. This implies that improved (or even optimal) values of fairness (lower deviation from the absolute value of fairness) can be achieved without requiring increases in terms of maximum displacement.

(b). Schedule-wide effects: At a schedule-wide level, an insight provided by this study is that demand-based fairness is not associated with maximum displacement. Hence, there may exist schedules that minimise both metrics simultaneously. From a technical perspective, the formulation of demand-based fairness affects the number of displaced requests significantly, demonstrating a strong and important negative correlation. However, the increased number of displaced requests for each improved unit of fairness coupled with our findings relating to maximum displacement suggest that the average displacement per displaced request is improved when fairness is reduced/optimised.

Implications of maximum displacement Herein we present and discuss insights relating to the maximum displacement objective.

(a). Inter-priority effects: We highlight some interesting findings that are associated with the maximum displacement of different priority levels that have not been reported before. At the historics' (*H*) level, we observe that there are significant negative correlations with the maximum and total displacement of all priority levels (and the schedule wide objectives except fairness). Hence, the increase of maximum displacement at the historics' level may have an important effect on the allocations all lower request priorities. Interestingly, it appears that there is a weak positive relationship among *H*'*s*, *Oth*'*s* and the schedule-wide maximum displacement. Similar findings are reported for the *CH*'*s* maximum displacement. However, there is a weak but significant positive relationship between *CH*'*s* maximum displacement and *NE*'*s* total displacement. *NE*'*s* maximum displacement does not have a significant or important relationship with other metrics. That could be justified by the limited number of new entrant requests (a common characteristic of most airports). At the *Oth* level maximum displacement appears to have insignificant relationships with the metrics of the upper request priorities (H, CH, NE). This suggests that regardless of the slot allocations at the upper priority levels, the already scarce capacity of the airport that has been allocated to the upper priority levels generates large displacements for *Oth* requests. By considering that an airline may submit requests that fall into different request priorities, this observation may motivate airlines with requests of multiple priorities, to reconsider their request portfolio belonging to the upper levels and receive improved allocations at the lower levels.

(b). Schedule-wide effects: In contrast, to the maximum displacement at each priority, the relationship of maximum displacement with other objectives/metrics, i.e. schedule-wide total displacement, demand-based fairness, priority-based total and maximum displacement, at the schedule-wide level is almost always insignificant or weak. The only exemption is the negative correlation with total displacement which highlights the existence of trade-offs between the two objectives at the schedule-wide level. A similar observation is also reported for *Oth*'s total displacement (which has the largest share of the schedule-wide total displacement).

Implications of total displacement Total displacement is a commonly used objective in airport slot allocation. It provides an aggregation of the overall displacement of the allocations provided by the coordinator. As with the previous objectives, we obtain some interesting findings that are new to the airport slot allocation literature.

(a). Inter-priority influence: An important finding that is enabled by the multi-level architecture of the current study, is that the total displacement of each level is positively associated with the total displacement of the other levels. At the *H* level there is a strong and significant positive correlation with the total displacement of the *CH* level. This implies that by minimising the total displacement of *H*requests, the total displacement of changes to historic requests is also improved. Similar but weaker relationships are reported between the total displacement of *H* and the total displacement of the *Oth*'s level. Interestingly, there is no strong or important relationship with the total displacement of *NE* requests. On the contrary, the total displacement of *CH*requests has a stronger relationship with the total displacement of*NE*. This set of findings may support the ongoing deliberations that seek to improve the efficiency of the administrative rules of IATA WSG. For instance, our findings suggest potential benefits in aggregating the *CH*and *H* priorities in a single level or allocating the *NE* requests before the *CH* requests.

(b). Schedule-wide influence: The relationship of each priority's total displacement on schedule-wide quality metrics is also insightful for policy and practice. For instance, we observe that the schedule-wide total displacement is more correlated with the total displacement of the historic requests than with any other priority level. Another interesting result is that there is a strong negative correlation between the average displacement per displaced request and the total displacement of the *H* and *CH* level, while the lower levels present weaker correlations. This suggests that the more we try to reduce the total displacement of the upper levels, the more displacement per displaced request we will receive across all priority levels. Therefore, the efforts of slot coordinators that wish to optimise the average displacement should focus on the upper priority levels (where there may be reduced flexibility) and not at the *Oth* and *NE* levels. Such original inherent trade-offs may have important influence on the outcome of the initial slot allocation and should be considered by practitioners and policymakers. Regarding the number of displaced requests, similar findings are reported. An interesting finding is that the total number of displaced requests is highly correlated with the total displacement of requests for historic operations (more than the total displacement for *Oth* operations). This contradicts the initial expectations, i.e. the displaced requests should be highly correlated with the metrics of *Oth*, and suggests that the reduction of displaced requests can be achieved by placing special emphasis on the minimisation of *H's* total displacement.

The implications presented above are extracted based on the summarisation of our findings in Fig. 5. By considering the proposed multi-level solution approach, policy and decision makers can investigate the inherent trade-offs among the objectives of the slot scheduling priorities and unveil the elements that influence the schedule-wide performance of a slot scheduling solution. The improved slot scheduling performance and the multiple schedule-wide non-dominated solutions of our approach (resulting from the consideration of dominated or weakly-dominated solutions at the upper priority levels) enable the extraction of aggregate metrics that can help decision/policy-makers to understand the impact of multiple alternative slot schedules on the objectives of the different stakeholders, i.e. airlines, airport operators etc, and thus facilitate a more collaborative decision-making process. In addition, the interested parties may seamlessly explore the benefit of trading-off the allocations of different slot scheduling levels with respect to different slot scheduling metrics; and consider non-dominated schedules that improve the allocations of their own slot request portfolio or the allocations at the schedule-wide level. More detailed examples and insights that study the trade-offs among the objectives of the slot request priorities are presented in the following section.

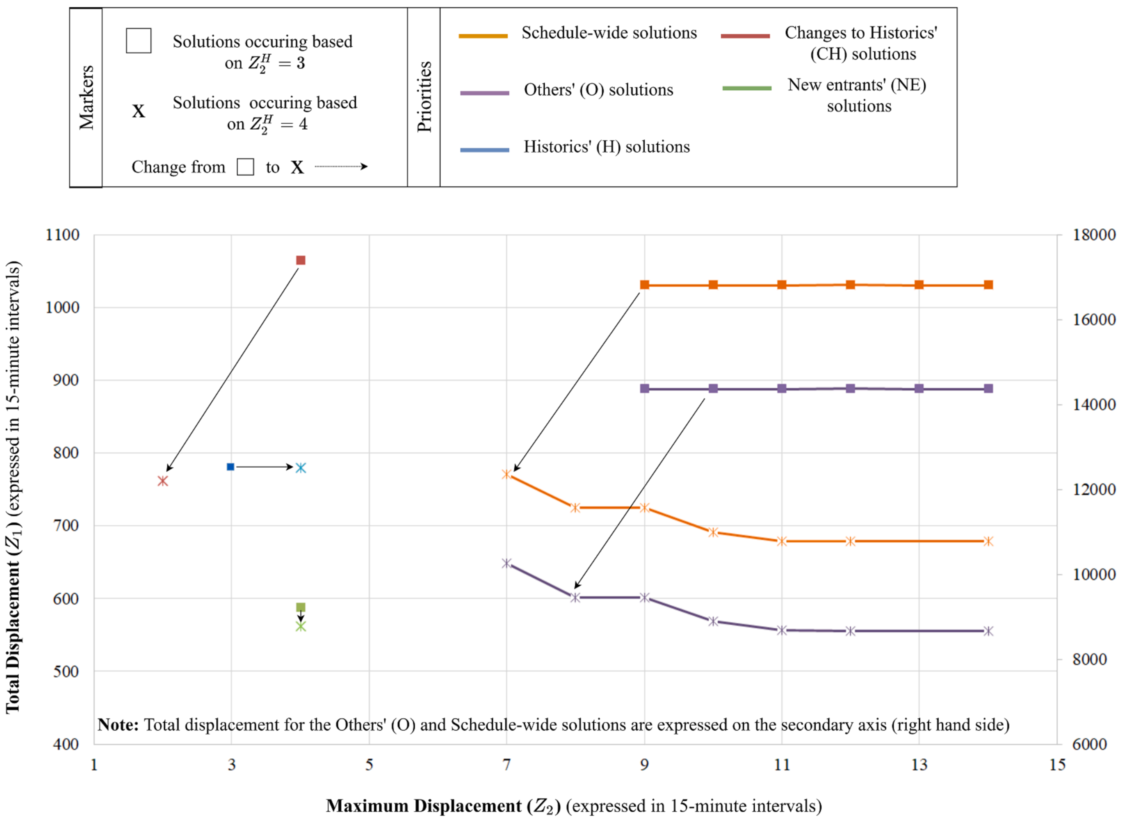

Fig. 6. Slot scheduling solutions when $Z_3 = 0.5$.

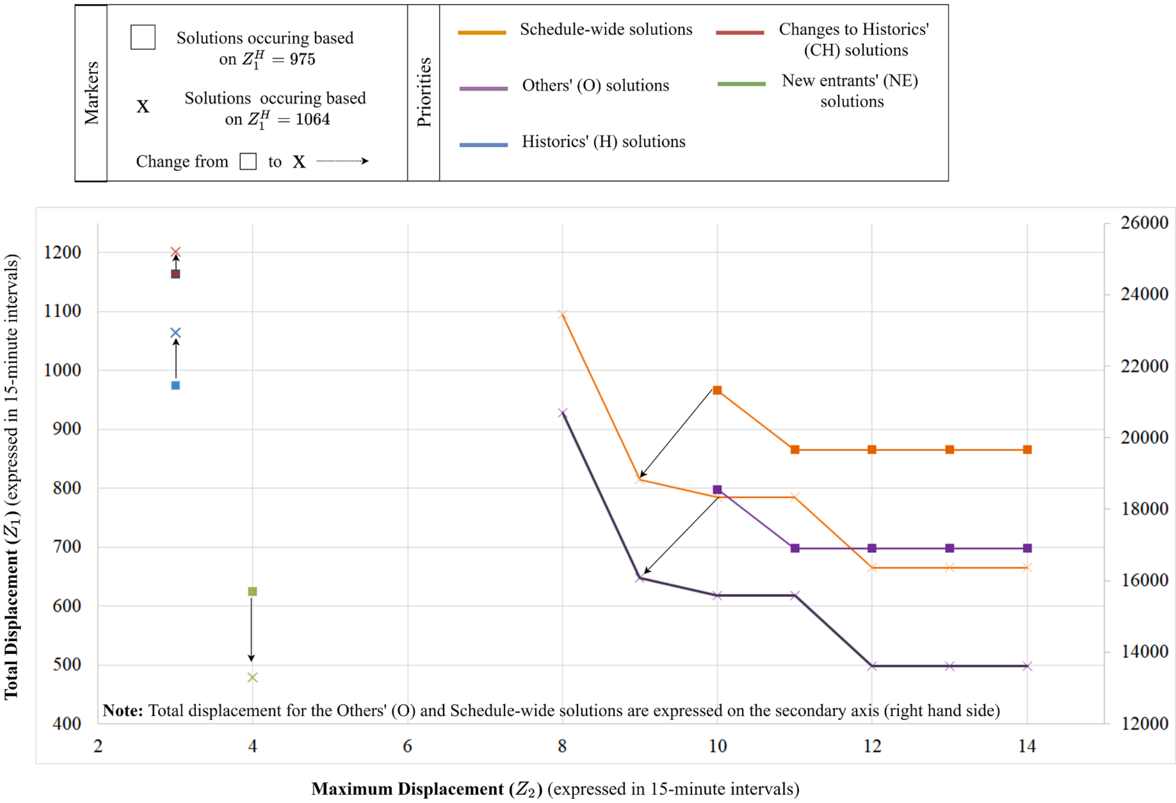

Fig. 7. Slot scheduling solutions when $Z_3 = 0.1$.

5.3.2. In-depth insights and examples

In depth insights on the influence of each level's compromises can be extracted by analysing the output of Algorithm 2 in the form of what-if scenarios. In the case that the stakeholders wish to focus on individual schedules detailed analyses can be conducted to illustrate the added value of the multi-level considerations. Herein we provide two examples of such what-if analyses which illustrate how compromises with respect to the maximum (Z_2) and the total displacement (Z_1) of the Historics' (H) or the Changes to Historics (CH) can influence policy making.

In Fig. 6 we illustrate that if the Historic's requests (H) tolerate an increase of 1 unit (15-minutes) of maximum displacement (Z_2^H) *ceteris paribus* then all lower slot priorities (CH, NE, O) will receive significant improvements regarding their objectives, thus leading to improved schedule-wide performance. In particular, by accepting this minimal deterioration at the H level's maximum displacement objective, requests at the CH level benefit by a reduction of 2 units of maximum displacement (30-minutes) and 303 (from 1065 to 762) units of total displacement (28.45%). Additionally, new entrants' requests receive an improvement of 26 units in terms of total displacement (4.4%). Interestingly, we observe significant improvements for the others' level (O) for both objectives. Total displacement is reduced by 38.4% on average (from 14,864 to 9157) and the minimum maximum displacement for O is decreased by 2 $\times$ 15-minute units from 9 to 7 (22.2% reduction). Therefore, for a given fairness level, by accepting a minimal increase of maximum displacement at the H level (1 unit), the total displacement of the schedule-wide solution is decreased by 34% on average (from 17,174 to 11,261 units) and the maximum displacement across all slot priorities is improved by two units (from 9 to 7).

In Fig. 7 we illustrate how small compromises (less than 10%) in the total displacement objective for the upper levels of the slot scheduling hierarchy (H, CH) may result in significant improvements with respect to the objectives of the lower levels of the hierarchy and the schedule-wide performance. The example illustrated in Fig. 7 shows that under a given fairness threshold ($Z_3 = 0.1$), increases of 9.1% and 3.2% at the total displacement of the H level (from 975 to 1164 units) and the CH level (from 1164 to 1202 units) respectively, result in significant improvements for the new entrants' and others' requests. Specifically, the total displacement at the NE level is decreased from 624 to 479 (23% improvement) and for the O level we observe an average improvement across all maximum displacement levels of 9.8% (from 17230.6 to 15551.1). Additionally, by introducing such increases/deteriorations at the total displacement of the H and CH levels, we report points which reduce the minimum of the maximum displacement for others' (O) requests by 22.2% (a reduction of 2 $\times$ 15-minute units). Therefore, the schedule wide performance is improved by 8.5% (from 19993.6 to 18296) and 22.2% (from 9 to 7) with respect to the total and maximum displacement objectives accordingly.

The examples discussed above, suggest that small compromises at the Historics and/or Changes to Historics' objectives can result in significant improvements at the objectives of the subsequent request priorities and therefore improve the schedule-wide performance of the examined airport. Analyses similar to the ones presented in Fig. 6 and Fig. 7 may contribute to the airport slot allocation policy debate by supporting amendments to the existing regulations aiming to improve the allocations of lower hierarchies and the performance of the overall slot schedule. Our results suggest that there is no need to resort to dramatic deviations from the existing slot prioritisation regime in order to improve new entrants' and others' schedules. Instead, it suffices to allow *H* or *CH* requests to receive slightly increased maximum or total displacements under the condition that they improve following level's allocations. The proposed algorithm is guaranteed to produce solutions (based on Proposition 1) which accept such sacrifices only if they result in schedule-wide non-dominated solutions. From the airlines' perspective, airlines that submit requests which fall into multiple slot priorities may explore how the compromises at the objectives of their *H* or *CH* requests affect the performance of their *NE* and *O* requests and their overall request portfolio.

5.4. Solution space reduction

The aim behind the use of the multiple objectives and the multi-level considerations is to provide information that will improve slot allocation decisions. The trade-offs between the objectives of the different slot priority levels (H, CH, NE, O) and the aggregate airport slot schedule justify the choice of a multi-objective approach with multi-level considerations. However, as our results suggest, multi-objective programming may result in many efficient solutions (57 in our case) which are difficult to process when a final slot scheduling decision is due to be made. It is therefore essential to address this decision-making need by proposing criteria and measures that will enable the elicitation of fewer alternatives which are better suited to the needs of the problem's context (Dal Sasso et al., 2019). The following criteria can be used to filter out solutions based on the needs of the interested stakeholders:

- Limiting the values of objectives at a specific level to be below a certain value or within a certain range (e.g. $UB^{Oth}_{Z_2} \leq 10$);
- Limiting the values of objectives for the aggregate airport slot schedule (e.g.$Z_1 \leq 14000$);
- Limiting the added deviation for all schedule-wide objectives to be below a threshold (e.g. below 70%);
- Limiting the allowable relative percentage gap for each level and objective to be within certain limits; and
- Limiting the amount of the allowable relative percentage gap for quality metrics such as the number of displaced slots or the average displacement per displaced slot.

The joint or independent application of the above criteria leads to a reduction of non-dominated solutions for consideration, thus facilitating the choice of the slot scheduling solution to be implemented. For example, by limiting the added deviation for all objectives below 70%, the number of efficient solutions to be considered, drops to just 4 (see Table 5). Interestingly, all solutions having an added deviation below 70% would not be reported without the multi-level considerations since they occurred from weakly non-dominated schedules at the upper levels of the slot scheduling hierarchy. Since the schedules reported with multi-level considerations are of higher quality with respect to the schedule-wide objectives than other solution approaches, the filtered solutions will always have improved schedule-wide objectives no matter what the criteria of the stakeholders are.

Having prior knowledge regarding the preferences of the stakeholders on the above criteria, we may incorporate them to the multi-level solution algorithm so as to constraint the feasible space of the problem to obtain a set of solutions of smaller cardinality that abide by their interests. Such approaches lead to reduced computational times since they require fewer iterations so as to search the problem's feasible space.

6. Concluding remarks

A multi-objective model for the allocation of airport slots according to IATA's WSG is modelled and solved using a multi-level solution algorithm. The model considers three slot scheduling objectives previously proposed in the literature, i.e. total displacement, maximum displacement and demand-based fairness. The proposed modelling and solution approach provide a framework that can be used to systematically investigate the trade-offs between the objectives of the different slot scheduling hierarchies and their influence on the airport slot schedule as a whole. The proposed framework results to improved schedule-wide performance and sets of non-dominated points of increased cardinality.

Through the analyses conducted on a schedule-coordinated airport, we demonstrate that small tolerances regarding the different objectives of the higher levels of the slot hierarchy may produce more efficient airport slot schedules. Multiple figures and suitable aggregation metrics distil the generated information and summarise our findings, thus easing the utilisation of our approach for policy and decision-making. For instance, when Historics' requests accept a 15-minute increase at their maximum displacement, the total displacement of Changes to Historics, New Entrants and Others is improved by 28.45%, 4.4% and 38.4% respectively. Meanwhile, the maximum displacement for both Changes to Historics and Others is reduced by 30 min (50% and 22.2% reduction). Furthermore, the schedule-wide total and maximum displacement are reduced by 34% and 22% respectively. Another interesting finding is that when Historics and changes to Historics requests tolerate small increases (approximately 9%) with respect to their total displacement, New Entrants' requests receive reductions of 23% for this objective while Ohers' requests benefit from significant improvements for both total and maximum displacement (approximately 10% and 22% respectively).

The application of the proposed model and solution algorithm to a case involving the allocation of 15,388 slots resulted to significant computational times. Although such solution times are tolerable for the tested instance it is expected that larger problems will not be solvable within practical time limits. Future research should focus on developing specialised heuristic algorithms for solving the multi-level, multi-objective airport slot allocation problem. The proposed modelling and solution framework can be used to incorporate other combinations of slot scheduling objectives (e.g. total displacement, maximum displacement, number of displaced slot requests) without loss of its generality, future research may employ this methodology to study the trade-offs of different triplets of objectives.

Acknowledgments

We are grateful for the comments of the associate editor and the reviewers on the initial version of the manuscript. The work reported in this paper has been supported by the UK's Engineering and Physical Sciences Research Council (EPSRC) through the Programme Grant EP/M020258/1 "Mathematical models and algorithms for allocating scarce airport resources (OR-MASTER)". Due to confidentiality concerns, supporting data cannot be made openly available. Additional details relating to the data are available from the Lancaster University data archive at https://nam03.safelinks.protection.outlook.com/?url=https%3A%2F%2Fdoi.org%2F10.17635%2Flancaster%2Fresearchdata%2F407&data=04%7C01%7CC.Corrections%40elsevier.com%7C756caf4ad37a4396a21c08d8a281af50%7C9274ee3f94254109a27f9fb15c10675d%7C0%7C0%7C637438025623222113%7CUnknown%7CTWFpbGZsb3d8eyJWIjoiMC4wLjAwMDAiLCJQIjoiV2luMzIiLCJBTiI6Ik1haWwiLCJXVCI6Mn0%3D%7C2000&sdata=q3ErV7uc%2BvRmIOPNRpZ4lesnPn5Acu7%2FI2oRxV12OrE%3D&reserved=0. The opinions expressed in this article reflect the authors' views.

Appendix

Herein we describe two additional multi-level solution algorithms which are used to generate *frontiers* 2 and 3 in section 5.2.1. Both algorithms (Algorithm 3, Algorithm 4) are based on Algorithm 2 but generate schedule-wide slot schedules/solutions based on alternative dominance considerations. The core process for generating non-dominated schedules within each level is again Algorithm 1.

Algorithm 3. *A variant of Algorithm 2 considering leading priority levels*

input: M, u, Δ # *slot requests, capacity, list of fairness values*
output: Y # *list of efficient schedule-wide solutions*
1 $Y \leftarrow []$; # *initialise the empty list of solutions*
2 **for** $\delta \in \Delta$ **do**:
3 $Y_\delta^H \leftarrow$ Algorithm 1(M^H, u, δ);
4 filter out dominated solutions in Y_δ^H;
5 **for** $s_H \in Y_\delta^H$ **do**:
6 $u^{CH} \leftarrow$ copyof$u$;# *re-initialise capacity*
7 update u^{CH} based on solution s_H;
8 $Y_\delta^{CH} \leftarrow$ Algorithm 1(M^{CH}, u^{CH}, δ);
9 filter out dominated solutions in Y_δ^{CH};
10 **for** $s_{CH} \in Y_\delta^{CH}$ **do**:
11 $u^{NE} \leftarrow$ copyof$u$;
12 update u^{NE}based on solutions s_H, s_{CH};
13 $Y_\delta^{NE} \leftarrow$ Algorithm 1(M^{NE}, u^{NE}, δ);
14 filter out dominated solutions in Y_δ^{NE};
15 **for** $s_{NE} \in Y_\delta^{NE}$ **do**:
16 $u^{Oth} \leftarrow$ copyof$u$;
17 update u^{Oth}based on solutions s_H, s_{CH}, s_{NE};
18 $Y_\delta^{Oth} \leftarrow$ Algorithm 1(M^{Oth}, u^{Oth}, δ);
19 filter out dominated solutions inY_δ^{Oth};
20 **for** $s_{Oth} \in Y_\delta^{Oth}$ **do**:
21 append s_{Oth} to Y^{Oth};
22 $s_{SW} \leftarrow$combine $s_H, s_{CH}, s_{NE}, s_{Oth}$ into a schedule-wide solution;
23 append s_{SW} to Y;
24 filter out schedule-wide dominated solutions from Y;
25 **return** Y;

Notes: # stands for comments on the algorithmic process; in **bold** are common algorithmic functions; SW: schedule-wide.

The first variant (Algorithm 3) proceeds to the generation of schedules for following slot priorities, if and only if a solution concerning level γ (s_γ) is not dominated by other solutions generated during the current iteration (lines 4, 9, 14 and 19 of Algorithm 3). Hence, this approach filters out dominated solutions only if they occurred from the same solutions in leading levels. As a final step, Algorithm 3 filters and reports only the schedule-wide non-dominated schedules.

In contrast, Algorithm 4 does not consider the leading levels' schedules that led to the generation of each solution s_γ. In particular, Algorithm 4 generates all solutions at level γ and then filters out dominated solutions. Then, for each of the remaining non-dominated schedules, it proceeds to the solution of the next level. This process is repeated until the 'others' level, where the solutions that led to each non-dominated point of this priority level are combined to create a schedule-wide solution. Please note that Y^H, Y^{CH}, Y^{NE}, Y^{Oth} are auxiliary lists which in addition to storing the solutions of each level γ, also keep track of the leading levels' solutions that led to each solution. For instance, for each s_{NE}, list Y^{NE} contains the schedules of H and CH levels that led to s_{NE}. Similar to Algorithm 2, Algorithm 4 terminates after filtering out dominated schedule-wide solutions.

Algorithm 4. *A variant of Algorithm 2 without considering leading levels' solutions*

input:M, u, Δ *# slot requests, capacity, list of fairness values*
output: Y *# list of efficient schedule-wide solutions*
$Y \leftarrow []$; *# initialise the empty list of solutions*
$Y^H, Y^{CH}, Y^{NE}, Y^{Oth} \leftarrow []$; *# initialisation of auxiliary lists of solutions for each level*
for $\delta \in \Delta$ **do**:
    $Y^H_\delta \leftarrow$ Algorithm 1(M^H, u, δ)
    append Y^H_δ to Y^H;
filter out dominated solutions in Y^H;
for $s_H \in Y^H$ **do**:
    $u^{CH} \leftarrow$ copyof$u$;*# re-initialise capacity*
    update u^{CH} based on solution s_H;
    $Y^{CH}_\delta \leftarrow$ Algorithm 1($M^{CH}, u^{CH}, \delta = s_H.\delta$); *# δ=$s_H.\delta$ sets fairness equal to the H level's fairness*
    append Y^{CH}_δ to Y^{CH};
filter out dominated solutions in Y^{CH};
for $s_{CH} \in Y^{CH}$ **do**:
    $u^{NE} \leftarrow$ copyof$u$;*# re-initialise capacity*
    update u^{NE} based on solution obtained at level H and currents_{CH};
    $Y^{NE}_\delta \leftarrow$ Algorithm 1($M^{NE}, u^{NE}, \delta = s_{CH}.\delta$); *# δ=$s_{CH}.\delta$ sets fairness equal to the CH level's fairness*
    append Y^{NE}_δ to Y^{NE};
filter out dominated solutions in Y^{NE};
for $s_{NE} \in Y^{NE}$ **do**:
    $u^{Oth} \leftarrow$ copyof$u$;*# re-initialise capacity*
    update u^{Oth} based on solutions obtained for levels H and CH and currents_{NE};
    $Y^{Oth}_\delta \leftarrow$ Algorithm 1($M^{NE}, u^{NE}, \delta = s_{NE}.\delta$); *# δ=$s_{CH}.\delta$ sets fairness equal to the NE level's fairness*
    append Y^{Oth}_δ to Y^{Oth};
filter out dominated solutions in Y^{Oth};
for $s_{Oth} \in Y^{Oth}$ **do**:
    $s_{SW} \leftarrow$combine $s_{Oth}$with the solutions of previous levels that led to s_{Oth} into a schedule-wide solution;
    append s_{sw} to Y;
filter out schedule-wide dominated solutions from Y;
return Y;

Notes: # stands for comments on the algorithmic process; in **bold** are common algorithmic functions; SW: schedule-wide.

References

ACI Europe, 2017. ACI Europe - Position on allocation of slots.

Cao, Y., Smucker, B.J., Robinson, T.J., 2015. On using the hypervolume indicator to compare Pareto fronts: Applications to multi-criteria optimal experimental design. Journal of Statistical Planning and Inference 160, 60–74. https://doi.org/10.1016/j.jspi.2014.12.004.

Cappanera, P., Scaparra, M.P., 2010. Optimal Allocation of Protective Resources in Shortest-Path Networks. Transportation Science 45, 64–80. https://doi.org/10.1287/trsc.1100.0340.

Corolli, L., Lulli, G., Ntaimo, L., 2014. The time slot allocation problem under uncertain capacity. Transportation Research Part C: Emerging Technologies 46, 16–29. https://doi.org/10.1016/j.trc.2014.05.004.

Dal Sasso, V., Djeumou Fomeni, F., Lulli, G., Zografos, K.G., 2019. Planning efficient 4D trajectories in Air Traffic Flow Management. Eur. J. Oper. Res. https://doi.org/10.1016/j.ejor.2019.01.039.

Deb, K., Sinha, A., 2010. An Efficient and Accurate Solution Methodology for Bilevel Multi-objective Programming Problems Using a Hybrid Evolutionary-local-search Algorithm. Evol. Comput. 18, 403–449. https://doi.org/10.1162/EVCO_a_00015.

Ehrgott, M., 2005. Multicriteria Optimization, 2nd ed. Springer-Verlag, Berlin Heidelberg.

EUROCONTROL, 2018a. EUROPEAN AVIATION IN 2040 - Annex 1: Flight forecast to 2040 [WWW Document]. URL https://www.eurocontrol.int/sites/default/files/content/documents/official-documents/reports/challenges-of-growth-annex-1-25092018.pdf (accessed 5.20.19).

EUROCONTROL, 2018b. European Aviation in 2040 Challenges of Growth Report [WWW Document]. URL https://www.eurocontrol.int/sites/default/files/content/documents/official-documents/reports/challenges-of-growth-2018.pdf (accessed 5.20.19).

Fairbrother, J., Zografos, K.G., 2018a. On the Development of a Fair and Efficient Slot Scheduling Mechanism at Congested Airports. Presented at the 97th Transportation Research Board (TRB) Annual Meeting, Washington, D.C, United States.

Fairbrother, J., Zografos, K.G., 2018b. Introducing flexibility and demand-based fairness in slot scheduling decisions. Presented at the Odysseus conference, Sardinia, Italy.

Fairbrother, J., Zografos, K.G., Glazebrook, K., 2019. A slot scheduling mechanism at congested airports which incorporates efficiency, fairness and airline preferences. Transport. Sci. https://doi.org/10.1287/trsc.2019.0926.

Gurobi Optimization, L., 2018. Gurobi Optimizer Reference Manual.

Haimes, Y.Y., 1971. Modeling and Control of the Pollution of Water Resources Systems Via Multilevel Approach1. JAWRA Journal of the American Water Resources Association 7, 93–101. https://doi.org/10.1111/j.1752-1688.1971.tb01681.x.

IATA, 2019. Worldwide Slot Guidelines - 9th Edition [WWW Document]. URL https://www.iata.org/contentassets/4ede2aabfcc14a55919e468054d714fe/wsg-edition-9-english-version.pdf (accessed 2.16.20).

IATA, 2019. Worldwide Slot Guidelines - 9th Edition [WWW Document]. URL https://www.iata.org/contentassets/4ede2aabfcc14a55919e468054d714fe/wsg-edition-9-english-version.pdf (accessed 2.16.20).

ICAO, 2016. Draft text for the report on agenda item 39, in: ICAO Assembly - 39th Session. Presented at the Economic commission EC/41.

Jacquillat, A., Odoni, A.R., 2015. An Integrated Scheduling and Operations Approach to Airport Congestion Mitigation. Oper. Res. 63, 1390–1410. https://doi.org/10.1287/opre.2015.1428.

Jacquillat, A., Vaze, V., 2018. Interairline Equity in Airport Scheduling Interventions. Transportation Science. https://doi.org/10.1287/trsc.2017.0817.

Jiang, Y., Zografos, K.G., 2017. Modelling fairness in slot scheduling decisions at capacity-constrained airports. Presented at the 96th Transportation Research Board (TRB) Annual Meeting, Washington, DC.

Kalai, E., Smorodinsky, M., 1975. Other Solutions to Nash's Bargaining Problem. Econometrica 43, 513–518. https://doi.org/10.2307/1914280.

Katsigiannis, F.A., Zografos, K.G., 2019. Optimizing multi-level, multi-objective airport slot-scheduling decisions. Presented at the tenth Triennial Symposium on Transportation Analysis. TRISTAN X, Hamilton, Queensland, Australia.

Marler, R.T., Arora, J.S., 2004. Survey of multi-objective optimization methods for engineering. Struct Multidisc Optim 26, 369–395. https://doi.org/10.1007/s00158-003-0368-6.

Pellegrini, P., Bolić, T., Castelli, L., Pesenti, R., 2017. SOSTA: An effective model for the Simultaneous Optimisation of airport SloT Allocation. Transportation Research Part E: Logistics and Transportation Review 99, 34–53. https://doi.org/10.1016/j.tre.2016.12.006.

Pyrgiotis, N., Odoni, A., 2015. On the Impact of Scheduling Limits: A Case Study at Newark Liberty International Airport. Transportation Science 50, 150–165. https://doi.org/10.1287/trsc.2014.0564.

Ribeiro, N.A., Jacquillat, A., Antunes, A.P., Odoni, A., 2019. Improving slot allocation at Level 3 airports. Transportation Research Part A: Policy and Practice 127, 32–54. https://doi.org/10.1016/j.tra.2019.06.014.

Ribeiro, N.A., Jacquillat, A., Antunes, A.P., Odoni, A.R., Pita, J.P., 2018. An optimization approach for airport slot allocation under IATA guidelines. Transportation Research Part B: Methodological 112, 132–156. https://doi.org/10.1016/j.trb.2018.04.005.

Rossum, G. van, 1995. Python tutorial (No. CS-R9526). Centrum voor Wiskunde en Informatica (CWI), Amsterdam.

While, L., Hingston, P., Barone, L., Huband, S., 2006. A faster algorithm for calculating hypervolume. IEEE Trans. Evol. Comput. 10, 29–38. https://doi.org/10.1109/TEVC.2005.851275.

Zhang, G., Lu, J., Montero, J., Zeng, Y., 2010. Model, solution concept, and Kth-best algorithm for linear trilevel programming. Inf. Sci. 180, 481–492. https://doi.org/10.1016/j.ins.2009.10.013.

Zitzler, E., Thiele, L., Laumanns, M., Fonseca, C.M., da Fonseca, V.G., 2003. Performance assessment of multiobjective optimizers: an analysis and review. IEEE Trans. Evol. Comput. 7, 117–132. https://doi.org/10.1109/TEVC.2003.810758.

Zografos, K.G., Androutsopoulos, K.N., Madas, M.A., 2017a. Minding the gap: Optimizing airport schedule displacement and acceptability. Transportation Research Part A: Policy and Practice. 10.1016/j.tra.2017.09.025.

Zografos, K.G., Jiang, Y., 2019. A Bi-objective Efficiency-Fairness Model for Scheduling Slots at Congested Airports. Transportation Research Part C: Emerging Technologies 102, 336–350. https://doi.org/10.1016/j.trc.2019.01.023.

Zografos, K.G., Jiang, Y., 2016. Modelling and solving the airport slot scheduling problem with efficiency, fairness, and accessibility considerations. Presented at the ninth Triennial Symposium on Transportation Analysis. TRISTAN IX, Aruba.

Zografos, K.G., Madas, M.A., Androutsopoulos, K.N., 2017. Increasing airport capacity utilisation through optimum slot scheduling: review of current developments and identification of future needs. J Sched 20, 3–24. https://doi.org/10.1007/s10951-016-0496-7.

Zografos, K.G., Salouras, Y., Madas, M.A., 2012. Dealing with the efficient allocation of scarce resources at congested airports. Transportation Research Part C: Emerging Technologies 21, 244–256. https://doi.org/10.1016/j.trc.2011.10.008.