Minding the Gap: Optimizing Airport Schedule Displacement and Acceptability

Konstantinos G. Zografos[1*], Konstantinos N. Androutsopoulos[2], Michael A. Madas[3]

[1] Lancaster University Management School, Department of Management Science,
Centre for Transport and Logistics (CENTRAL)
Lancaster, LA1 4YX, United Kingdom
k.zografos@lancaster.ac.uk

[2] Athens University of Economics and Business, Department of Management Science and Technology
47A, Evelpidon Str. & 33, Lefkados, 11362, Athens, Greece
kandro@aueb.gr

[3] University of Macedonia, Department of Applied Informatics
Information Systems and e-Business Laboratory (ISeB)
156, Egnatia Str., 54636, Thessaloniki, Greece
mmadas@uom.gr

*Corresponding author

Abstract

Serious congestion problems at slot-controlled airports worldwide call for some action. Slot scheduling related research has mainly focused on scheduling models allocating airport capacity by optimising scheduling efficiency. However, existing literature does not capture the effect of slot allocation decisions on the acceptability of slot schedules. The objective of this paper is to investigate the trade-off between scheduling efficiency and the airlines' dis-utility of slot schedules expressed by various metrics of schedule displacement. We develop and solve two bi-objective scheduling models considering different combinations of total and maximum acceptable slot displacement objectives. The proposed models are applied to real-world scheduling data. Substantial improvements in schedule acceptability metrics are achieved without sacrificing a lot in terms of scheduling efficiency. Sacrifices in scheduling efficiency increase the capability of the airport coordinator to allocate slots that are eventually acceptable and hence more intensively used.

Keywords: *airport scheduling, slot allocation, demand management, bi-objective optimisation*

1. Introduction

Congested airports constitute a serious bottleneck of the overall efficiency of the air transport system with congestion implications moving from ground to air and vice versa (SESAR JU 2015). The sustainability of the air transport system closely depends on the evolving relationship between demand and supply of air transport services. Long-term forecasts (Eurocontrol 2013), adjusted for economic downturn effects, anticipate lower air traffic growth rates, which are however accompanied by reduced airport capacity expansion plans due to weaker economic outlook. According to the same forecast (Eurocontrol 2013), more than 30 airports are expected to be operating at 80% of capacity or more for 3 or more hours per day by 2035. Intensive use of saturated airport capacity and increasing imbalances between demand and capacity will adversely impact predictability and punctuality of the air transport system (SESAR JU 2015). Current evidence and forecasting trends discussed above necessitate some form of intervention aiming to better control the so called demand-to-capacity ratio.

Supply-side interventions aiming to build new capacity are capital intensive solutions, require significant implementation time, and are often subject to heated political debates. The need for an immediate relief to seriously congested airports calls for short to medium-term, demand-side solutions that are based on the optimum allocation and use of available airport capacity (Zografos et al. 2017). The dominant demand-side mechanism currently applied at about 170 of the busiest and most congested airports worldwide (outside the United States)[1] is driven by a set of rules, priorities, and voluntary guidelines set out by IATA (IATA 2014); later adapted and further complemented by relevant EU regulation (European Commission 1993; 2004; 2009).

The IATA-driven mechanism prescribes a strategic, pro-active (few months before operations) schedule coordination process that aims to build a viable flight schedule by controlling (actually limiting) the maximum number of scheduled operations during a unit of time at each airport. A necessary prerequisite of the schedule coordination process is to optimally set its declared capacity, i.e., the capacity that can be made available for allocation and use to airport users (Ball et al. 2007a; Odoni and Morisset 2010; Jacquillat and Odoni 2015). The declared capacity of a schedule coordinated airport is assigned by a coordinator to airport users through the allocation of slots, i.e., a time interval during which a flight can use the airport infrastructure for landing or take-off. Declared capacity is rationed according to a set of criteria and rules assigning different priorities to: i) requests with historical usage rights ("grandfathered slots" - GFR), ii) requests with new entrant status, and iii) all remaining requests. Slots are mainly allocated in series[2] (at least five slots for the same time and day of the week regularly) for the entire scheduling season (i.e., winter/summer). In order to cope with strong slot complementarity at the airport network level, the initial slot allocation outcome at each airport is further streamlined by airlines in biannual worldwide scheduling conferences.

One issue that arises frequently at schedule coordinated airports is that the temporal distribution of demand expressed in slots requested by the airport users does not necessarily coincide with the temporal distribution of the available (declared) capacity. Therefore, although the daily available capacity may be sufficient to cover the daily demand, the airport users may not be able to obtain their preferred slots since demand may exceed supply during certain sub-intervals (e.g., hour, 15-min. interval). In this context, the slot scheduling problem seeks to satisfy slot requests by allocating available capacity. This means that a slot request may not be entirely satisfied in terms of the time that will be eventually scheduled. As a result, the flight will experience a displacement (scheduled at an earlier or later slot than the one originally requested). This displacement, also known as "schedule delay", is expressed by the absolute value of the distance between the requested and allocated slot times (Koesters 2007; Zografos et al. 2012). The slot scheduling problem may provide a "solution" to the optimum utilisation of available capacity by displacing flights from time intervals where demand exceeds capacity to time intervals where capacity exceeds demand. However, it may produce solutions that are not acceptable or even practical at all. This is because the displacement of a flight to an undesirable slot may have a detrimental effect on

[1] The situation is quite different in the United States; following the phase-out of the High Density Rule (HDR) between 2000 and 2006, only the three New York area airports and Washington DC's Reagan airport currently operate under certain scheduling limits ("slot caps").

[2] Non regularly planned operations, called *ad hoc* services, can be allocated, but series of slots have allocation priority, while *ad hoc* operations are not eligible for historical usage rights (GFR) (IATA 2014).

the feasibility of the entire flight schedule of the airline's network or the commercial viability of the flight. As a result, certain slots may not be attractive enough to be actually operated by the assigned airport users, a fact that may lead to waste of a really scarce resource. Under these circumstances, the capacity shortage problem is further sharpened by low utilisation levels and severe misuse (e.g., late return of unwanted slots, "no shows") of that scarce capacity (Madas and Zografos 2010). Studies suggest that even at congested airports, over 10% of the allocated slots go unused (Steer Davies Gleave 2011) with significant economic consequences reaching €20 million per season at large congested European airports (ACI Europe 2009).

Schedule optimisation signifies a challenging stream of research dealing effectively with the complexity and size of the resulting airport scheduling problem with a promising potential to deliver quick capacity utilisation improvements (Zografos et al. 2012). The strategic airport slot scheduling problem has been recently addressed in the literature (Zografos at al. 2017). Existing models have basically considered the following objectives: i) total schedule displacement (Zografos et al. 2012; Castelli et al. 2012), ii) total and maximum schedule displacement (Pyrgiotis and Odoni 2015; Jacquillat and Odoni 2015), iii) schedule displacement and expected operational (queuing) delays (Corolli et al. 2014), and iv) fairness and total schedule displacement (Zografos and Jiang 2016). However, the issue of acceptable schedule displacement and alternative ways of accommodating airlines' preferences and tolerance levels against schedule displacements have not been sufficiently addressed in existing literature. Most importantly, existing research and practice ignores the effect of slot allocation decisions on the potential acceptability and utilisation of scarce airport resources or the expected number of unused slots as a result of "unacceptable" schedule displacement.

This paper aims to investigate alternative ways of better accommodating airlines' preferences and tolerance levels in compliance with the existing IATA-based strategic scheduling process. The objective of the paper is to develop and solve two bi-objective single-airport slot scheduling models for studying the trade-off between total and acceptable schedule displacement. The proposed models consider combinations of the following objectives: i) total schedule displacement, ii) maximum schedule displacement, and iii) maximum displacement thresholds above which allocated slots "unacceptably" violate the corresponding slot requests ("violated slot assignments"). The latter is in line with the IATA-based allocation scheme that urges airlines to indicate a "Timing Flexibility" range within which slots are considered acceptable. The underlying conjecture is that allocated slots, being closer to airlines' preferences and tolerance levels, will be more intensively operated. This eventually delivers substantial benefits not only for airlines provided with preferable slots to build their flight schedules, but also for the airport community as a whole through increased utilisation of scarce airport resources and flight service frequencies, as well as reasonable access to a broad number of airport destinations. The proposed models are applied to real-world scheduling data.

The remainder of this paper is structured into five sections. Section 2 discusses previous relevant research in the area of strategic airport slot scheduling. Sections 3 and 4 present the proposed slot scheduling models and the solution approach, respectively. Section 5 presents the tested scenarios and the underlying data, as well as the model results under various scenarios. Finally, Section 6 provides the concluding remarks of the paper and discusses some emerging modelling needs and future research directions.

2. Previous Related Work

The focus of this paper is on slot scheduling at a single airport. In this section, we review mainly existing single-airport slot scheduling models. However, for completeness we are also making reference, without elaborating on the details of model formulation and/or solution, to relevant slot scheduling models at network level. An extensive and comprehensive review and assessment of existing slot scheduling models both at single airport and network level can be found in (Zografos et al. 2017).

Zografos et al. (2012) proposed an integer linear, deterministic, single-airport scheduling model allocating series[3] of slot requests to coordination time intervals during the entire scheduling season[4] by minimising total schedule displacement. The proposed model assumes rolling declared capacity constraints in addition to turnaround time and slot/flight assignment constraints. Series of slots are allocated hierarchically for each priority class (e.g., grandfathered slots, new entrants). By comparing the results of their model against existing allocation practice at three Greek airports, the authors demonstrated that there is a large room for improvement (between 14% and 95%) of the existing allocation outcome (Zografos et al. 2012).

Zografos and Jiang (2016) introduced two new models incorporating fairness and accessibility objectives in addition to the total displacement objective (Zografos et al. 2012). In these formulations, a fairness index is introduced. The fairness index requires that the total displacement that an airport user will incur for all its slot requests should be proportional to the number of slots it has requested. The fairness objective is then expressed by the minimisation of the variance of the fairness index. Thus, the minimisation of the fairness index ensures that no airport user will suffer disproportionally high displacement as compared to the rest of the airport users. The accessibility objective is used to ensure that the displacement that will be imposed to small aircraft connecting the congested (hub) airport with remote regional airports will not exceed a threshold value. This objective ensures that the connectivity of small regional airports with a hub airport will be accommodated within an acceptable level of service. The first of the two models considers the total displacement objective introduced by Zografos et al. (2012) along with the fairness and accessibility objectives. The second model introduces the idea of the weighted displacement. In this formulation, the displacement of each slot is weighted by the aircraft size and the length of the flight. Then, the total weighted displacement is minimised. The motivation behind the weighting of the displacement is to reduce the displacement of flights carrying a lot of passengers at a long distance.

Slot allocation models for the airport network problem have been also introduced in the literature (Castelli et al. 2011; Castelli et al. 2012; Corolli et al. 2014). These models take into consideration the strong complementarity of slots between origin and destination airports, and extend the single-airport allocation framework to a network of interconnected airports. Existing scheduling models for the airport network mainly adopt total schedule displacement metrics (e.g., "shift cost", "schedule delay") (Castelli et al. 2011; Castelli et al. 2012; Corolli et al. 2014). More recent work has considered the sum of total schedule displacement and expected operational (queuing) delays,

[3] Slots are allocated as series of at least five slots requested by an airline at an airport for the same time on the same day of the week regularly in the same period (European Commission 1993).
[4] Typically April to October for the summer and November to March for the winter scheduling season.

taking into account airport capacity, flight connectivity, and turnaround time constraints (Corolli et al. 2014).

Based on the scheduling practice applied at the four U.S. High Density Rule (HDR) airports, Pyrgiotis and Odoni (2015) proposed a schedule optimisation model formulated as an integer program. They aim to lexicographically minimise the sum of the maximum and total (aggregate) displacement subject to scheduling limits, as well as flight time, aircraft and passenger connectivity constraints at the airport network. At a second stage, the scheduling model is combined with a network queuing model to estimate both the local and nationwide effects of introducing alternative scheduling limits at a major airport. The authors demonstrated the model at one of the busiest days in 2007 at Newark Liberty International (EWR) airport. It is interestingly shown that demand can be accommodated even for the busiest day with only small schedule shifts that do not exceed 30 minutes for any single flight. Furthermore, the smoothed demand at EWR results in considerable delay savings in the order 20% at both airport and network level. Jacquillat and Odoni (2015) proposed a single-airport model formulation based on the strategic scheduling model of Pyrgiotis and Odoni (2015). The model considers also a two-stage lexicographic objective; first minimising the maximum displacement and then, among all feasible flight schedules under this objective, select the one that minimises the total (aggregate) displacement. A notable difference in this model formulation, as compared to Pyrgiotis and Odoni (2015), lies in the form of scheduling parameters that are defined similarly to the formulation of Bertsimas and Stock Patterson (1998) for the ATFM problem. The model application at John F. Kennedy International (JFK) airport confirms that a moderate rescheduling of about 75-90% of the flights at no more than 15-30 minutes for any flight can mitigate substantially congestion at busy U.S. airports. Both of the aforementioned models reflect the U.S. scheduling practice, but do not inherently capture several practicalities or properties of the IATA-driven allocation process (e.g., allocating slots for the entire scheduling season, rolling capacity constraints, series of slots).

It is interesting to observe that although there have been recent research efforts exploring the strategic slot scheduling problem at network level, there is still ample room for further research and improvements in single-airport slot scheduling (Zografos et al. 2017). Given the size and complexity of slot allocation at network level, it would be plausible to experiment with alternative single-airport scheduling models in order to come up with more comprehensive, efficient, and robust modelling formulations. These can thereafter act as the main building block for network-wide modelling extensions. Moreover, it is worth investigating advanced models capable of improving the way that airlines' preferences and their potential acceptability are incorporated and modelled in the slot scheduling process.

Starting from these observations, in this paper we develop and solve two bi-objective formulations extending the single-objective scheduling model initially proposed by Zografos et al. (2012). The underlying research motivation is that existing literature does not capture the effect of slot allocation decisions on the acceptability of alternative slot schedules, and consequently the utilisation of scarce airport resources. Our main research contribution lies on the analysis of the trade-off between scheduling efficiency (expressed in the form of total schedule displacement) and the airlines' utility (or dis-utility) of alternative slot schedules expressed by various measures and levels of tolerance against schedule displacement. Total schedule displacement is modelled as linear cost function of the difference between requested and allocated slot times (Koesters 2007; Zografos et al. 2012). It constitutes the common optimality criterion for both models and it is used as a reasonable proxy of the dis-utility that airlines encounter due to schedule displacement. In addition to total schedule displacement as a central tendency measure, the first model considers the maximum displacement objective as a measure of the worst-case service level provided to airlines. The maximum displacement metric can be implicitly viewed as a "guaranteed service level" on the grounds that airlines can be reasonably ensured in advance of the scheduling season that slots will be either granted as requested or rescheduled in a time range that will not definitely exceed this worst-case service level. In the second model, we consider thresholds for unacceptable displacement of requested slots, above which slots are practically rejected, hence remaining under- or even unutilised. This additional objective, termed as "violated slot assignment", resembles the "Timing Flexibility" option envisaged by the IATA-based allocation scheme and accounts for slots that are allocated beyond the airlines' tolerance limits.

Both models proposed in this paper are solved simultaneously with respect to both objectives. Furthermore, we investigate potential trade-offs among the objectives and determine the set of non-dominated solutions. This is achieved by transforming each model to a set of single-objective problems through the ε-Constraint Method. The proposed models address the operational constraints arising from the existing slot allocation framework (e.g., rolling capacity constraints), as well as the airlines' and airports' operational requirements (e.g., turnaround time, ground handling operational requirements, flight connectivity). Finally, both models fully conform to the fundamental regulatory properties of the slot scheduling problem under consideration: i) explicit consideration of multiple priority classes (with explicit priority to historic slot holdings or "grandfather rights"), ii) allocation of series of slots (rather than individual slots), and iii) allocation of slots within a scheduling season (rather than a single day).

3. Proposed Slot Allocation Models

The main setting of the strategic airport slot allocation problem involves a set of slot requests (M), a scheduling time period defined by the set D of calendar days, and the set of coordination time intervals per day $T = \{0, 1, \ldots, n-1\}$. In this paper, the length of each coordination time interval is assumed to be equal to five minutes. Each slot request (m) is associated with a movement (arrival or departure), the calendar days (D_m) that the movement must be scheduled, and the corresponding requested time interval (τ_m). $M_d \subseteq M$ is the set of slot requests involving departures and $M_a \subseteq M$ the set of slot requests involving arrivals. For convenience, any slot request involving an arrival (departure) will be referred to simply as arrival (departure). It should be highlighted that the requested time interval (τ_m) of a slot request applies to all calendar days in D_m. Moreover, the time interval that will be actually allocated to a slot request will also apply to all calendar days in D_m. This is a standard constraint in the slot allocation process which implies that a movement is assigned to a series of slots (referring to the same time interval) rather than a single slot.

$M_G \subseteq M$ is the set of slot requests that are subject to Grandfather Rights. Any slot request $m \in M_G$ (with Grandfather Rights), must be scheduled in one of the pre-specified alternative coordination intervals $t \in T_m$ set by the relevant airline. Although in the current slot allocation regime each slot request involves a single time interval, we relax this hard constraint in our paper by allowing the airline to provide a range of time intervals. Hence, we assume that T_m includes a range of consecutive time intervals reflecting the flexibility offered by an airline with view to facilitating the scheduling process. Thus, a set of time intervals T_m is specified for every slot request m ($T_m = T$ if $m \notin M_G$). However, it is worth noting that the current scheduling practice may be also reflected on the proposed formulation by forcing T_m to include a single time interval.

The basic slot allocation problem involves the optimum assignment of slot requests to the coordination time intervals. Two categories of constraints are relevant to the emerging slot scheduling problem: i) the airport capacity constraints and ii) the precedence constraints (Zografos et al. 2012). A typical capacity constraint $c \in C = \{0,1,2\}$ in slot scheduling is expressed by the maximum number of movements ($u_c^{\delta s}$) in $M_c \subseteq M$ that can be scheduled within any time period $[s, s + t_c)$ (denoted by T_c^s) of length $t_c > 0$ for $s \in [0, n - t_c + 1)$. Three subsets of movements where the above constraint applies include $M_0 = M_a$ (in which case $c = 0$), $M_1 = M_d$ ($c = 1$) and the entire set of movements $M_2 = M$ ($c = 2$). This type of constraint imposes a maximum number of movements in M_c within any possible time period of length t_c within the time horizon T. The emerging constraints are called rolling capacity constraints, since the time period over which the capacity constraint is checked rolls throughout the entire time horizon.

The precedence constraints of the airport slot scheduling problem arise from pairs of arrival and departure slot requests (m_1, m_2) where $m_1 \in M_a$ and $m_2 \in M_d$ for which either the corresponding flights are connected or they are operated by the same aircraft on the same day. In both cases, a minimum time deviation (time lag) is imposed between the time intervals allocated to m_1 and m_2. Let P be the set of linked movement pairs (m_1, m_2), $m_1 \in M_a, m_2 \in M_d$. The slots assigned to movements m_1 and m_2 should have a minimum time difference of $t_{m_1 m_2} \geq 0$.

A schedule performance criterion used for the basic slot scheduling problem relates to the minimisation of the total earliness-tardiness metric expressing the total schedule displacement. The emerging basic slot scheduling problem belongs to the category of resource-constrained scheduling problems with partially renewable constraints, which has been proved to be NP-hard (Bottcher et al. 1999). At the single-airport level, arrival and departure slot requests constitute the jobs to be processed by a single constrained resource type, the airport, over a given planning horizon, subject to airport capacity (resource) constraints. Given that the airport capacity constraints must be satisfied for a series of calendar days within the scheduling season, the utilisation of the airport at any given calendar day is considered as a different resource. Moreover, due to rolling capacity constraints (the capacity is defined for a time period, e.g., 1 hour, rather than a fixed time unit), each resource is treated as partially renewable (Bottcher et al. 1999).

Based on the setting defined above, two slot allocation models are proposed and discussed in what follows. Each of the models emerges by adding different optimisation criteria (in addition to total schedule displacement) to the basic slot scheduling problem described above.

3.1 Slot Allocation Model for Minimising Maximum and Total Schedule Displacement Objectives (SAM-I)

The maximum schedule displacement expresses the worst case displacement realized over all the scheduled slot requests. Incorporating this metric in the slot allocation problem gives rise to a new bi-objective slot scheduling model which aims to explore the trade-off between the maximum schedule displacement and the total schedule displacement for a given airport capacity level. The proposed model enhances the basic slot allocation model developed by Zografos et al. (2012) and it is formulated as a bi-objective integer programming model, similar to the model proposed by Jaquillat and Odoni (2015). The model examined here differs from the related model in Jaquillat and Odoni (2015) in that the former model includes rolling capacity constraints and series of slot requests which are not covered in the latter.

Variables $x_m^t \in \{0,1\}$ for $m \in M$, $t \in T$ take value 1 if slot request m is scheduled on time interval t , and 0 otherwise. Variable $z \geq 0$ expresses the maximum displacement across all slot requests. Assigning slot request m to time interval t is associated to a cost f_m^t defined by deviation $|t-\tau_m|$. Let a_m^δ be one if request m operates on day $\delta \in D$ (i.e., $\delta \in D_m$), zero otherwise; this allows handling both individual movements (slots), as well as series of movements defined for more than one day. Movement m consumes b_m^c units of the capacity of constraint $c \in C$; this consumption equals one if the movement is an arrival and the constraint applies to arrivals only or total movements, and zero otherwise (and similarly for departures).The proposed model (denoted SAM I) is defined as follows:

$$Minimise \qquad Z_1 = \sum_{m\in M}\sum_{t\in T} f_m^t x_m^t \tag{1}$$

$$Z_2 = z \tag{2}$$

subject to:

$$\sum_{t\in T_m} x_m^t = 1, m \in M \tag{3}$$

$$\sum_{m\in M} a_m^\delta b_{mc} \left(\sum_{t\in T_c^s} x_m^t\right) \leqslant u_c^{\delta s}, c \in C, \delta \in D, s \in T_c^s \tag{4}$$

$$\sum_{t\in[0,k)} x_{m_2}^t + \sum_{t\in\left[k-t_{m_1m_2},n\right)} x_{m_1}^t \leq 1, (m_1, m_2) \in P, k \in \left[t_{m_1m_2}, n\right) \tag{5}$$

$$\sum_{t\in T} x_m^t t \leqslant \tau_m + z, \qquad m \in M \tag{6}$$

$$\sum_{t\in T} x_m^t t \geqslant \tau_m - z, \qquad m \in M \tag{7}$$

Objective functions (1) and (2) express the minimization of the total schedule displacement (total 'cost' of schedule delay) and the minimization of the maximum schedule displacement. Objective function (1) is similar with the one used by Zografos et al. (2012). Constraint (3) implies that every slot request must be assigned to exactly one slot. Constraint (4) expresses the declared capacity constraints of the airport. Constraint (5) implies that a departure movement must be scheduled at least ($t_{m_1m_2}$) units ahead of the corresponding arrival movement. It is worth noting that constraints (4) and (5) are similar with the corresponding capacity and precedence constraints used in the model proposed in (Zografos et al. 2012). Constraints (6) and (7) facilitate the definition of variable z . The left part of both inequalities represents the time interval assigned to any slot request m . Thus, constraints (6) and (7) imply that the absolute value of the scheduled time of any slot request (i.e., $\left|\sum_{t\in T} x_m^t t - \tau_m\right|$) should not exceed z (or equivalently, z represents the maximum schedule displacement across all slot requests). The Grandfather rights can be included in the model described above by narrowing T_m (used in constraint 3) to a single or a few requested time intervals for every $m \in M_C$.

Minimising the worst schedule displacement across all movements implies that schedule deviation will be spread more evenly among the slot requests. This intuitively implies that decreasing the proposed measure creates a tendency towards a more uniform distribution of schedule deviation among movements, and thus a possible increase in the total schedule displacement delay (due to limited capacity). Example 1 that follows illustrates the trade-off between the maximum and the total schedule displacement. Five slot requests (m_1-m_5) must be scheduled from time interval 1 to time interval 20. The time interval requested by m_1-m_3 is 7, while for requests m_4 and m_5, the requested time interval is 10. The rolling capacity constraint is 2 movements per three time intervals, (i.e., it is forbidden to schedule more than two movements within any time period of length 3). No turnaround time constraint is considered. Figure 1 illustrates the solution (solution 1) of the slot allocation problem that optimises the total schedule displacement. The time horizon of the problem is represented by the first numbered sequence of cells, which is replicated as a sequence of empty cells for every request separately. The grey cell represents the requested time interval for the corresponding slot request, while the allocated slot to each request is represented with an "x" in the corresponding cell.

Figure 2 presents an alternative solution (solution 2) of the problem described in example 1, which optimises the maximum schedule displacement. Comparing solutions 1 and 2, it can be verified that shifting request m_3 (in solution 1) to the right by one time unit requires shifting requests m_1 (or m_2) and m_5 (or m_4) to the right by one time unit in order to retain feasibility. Thus, solution 2 represents one of the optimal solutions of the problem of optimising maximum schedule displacement in the expense of larger total schedule displacement as compared to solution 1. This example indicates that total schedule displacement and maximum schedule displacement may be viewed as conflicting scheduling objectives.

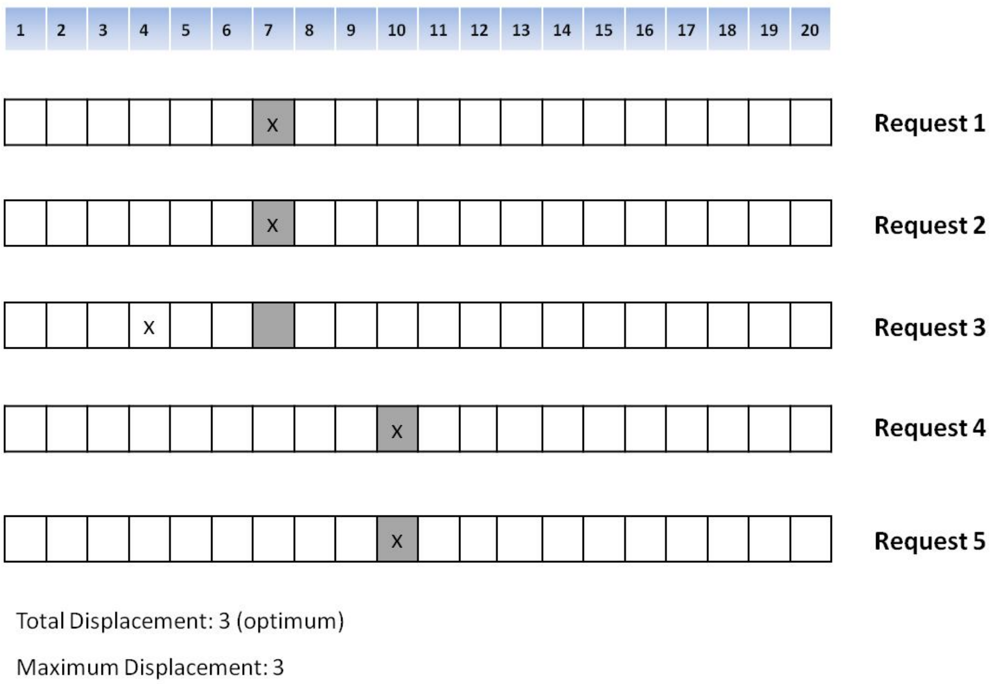

Figure 1: Optimal Solution of the Problem in Example 1 (optimising total schedule displacement).

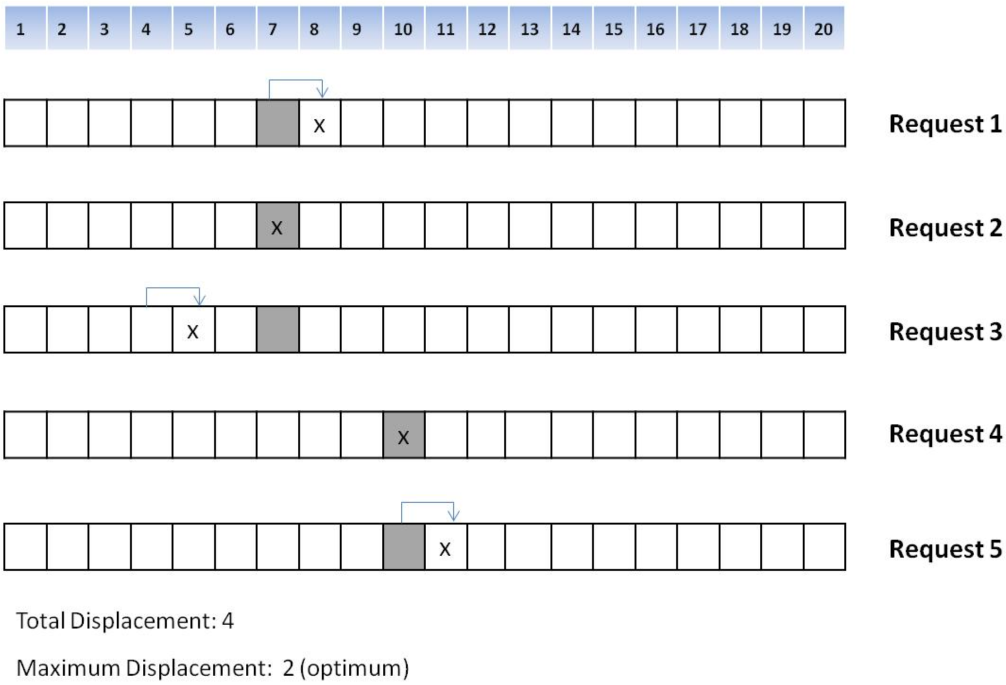

Figure 2: Optimal Solution of the Problem in Example 1 (optimising max schedule displacement).

3.2 Slot Allocation Model for Minimising Violated Slot Assignments and Total Schedule Displacement (SAM-II)

A practical implication of the current slot allocation process is that it may schedule a slot request far away from its initially requested time interval. This, in turn, may decrease substantially the actual utility realized by the relevant airline in using the allocated slot, hence resulting in poor (if any at all) slot utilization rates. This reaction of the airline is not taken into account by existing models leading to underutilisation of the airport capacity. An airline's tolerance in accepting a slot offered by the airport coordination authority can be expressed by a maximum acceptable deviation (q_m) from the requested time interval (τ_m). If the slot allocated to a given request ($m \in M$) is not aligned with this tolerance limit (i.e., the slot assigned for this movement lies before time $\tau_m - q_m$ or after time $\tau_m + q_m$), then the corresponding slot assignment is called "violated slot assignment". The number of violated slot assignments constitutes an aggregate measure of the dissatisfaction of airlines for the allocation of slots at a schedule coordinated airport. It also provides an estimate of the portion of slots that will not be used or alternatively the portion of demand that will not be serviced by the airport. Determining and using a schedule that minimises the number of violated slot assignments decreases the expected number of cancellations (i.e., cases where the slot allocated to a request is not used by the airline). A new bi-objective slot allocation model is proposed that aims to determine alternative solutions that trade-off violated slot assignments with total schedule displacement. The proposed model uses decision variables $x_m^t \in \{0,1\}$, $m \in M, t \in T$ and $y_m \in \{0,1\}$ $m \in M$. Variables x_m^t obtain value 1 if the requested movement m is assigned at time interval t , and 0 otherwise. Variables y_m obtain value 1 if the slot request m is violated (i.e., it is allocated a time interval outside the tolerance limits), and 0 otherwise. Moreover, the model involves the following additional variables: i) binary variables $w_m \in \{0,1\}$ $m \in M$ which take value 0, if the slot allocated to request m lies before (below) the upper limit ($\tau_m + q_m$) and otherwise and ii) binary variables $v_m \in \{0,1\}$ $m \in M$

which take value 0 if the slot allocated to request m is after (above) the lower limit ($\tau_m - q_m$) and 1 otherwise. Thus, if one of w_m or v_m take value 1 then the relevant request m is violated.

The model (denoted by SAM-II) is defined by (8)-(15) and (3)-(5):

$$\textit{Minimise} \qquad Z_1 = \sum_{m\in M}\sum_{t\in T} f_m^t x_m^t \tag{8}$$

$$Z_2 = \sum_{m\in M} y_m \tag{9}$$

subject to constraints (3)-(5) and :

$$\left(\sum_{t\in T} x_t^m t - (\tau_m - q_m)\right) \geq \frac{1}{2} - y_m T, \qquad m \in M \tag{10}$$

$$(\tau_m + q_m) - \sum_{t\in T} x_t^m t \geq -y_m T, \qquad m \in M \tag{11}$$

$$(\tau_m + q_m) - \sum_{t\in T} x_t^m t \leq -\frac{1}{2} + (1 - w_m)\mathrm{T} \tag{12}$$

$$\sum_{t\in T} x_t^m t - (\tau_m - q_m) \leq \frac{1}{2} + (1 - v_m)\mathrm{T} \tag{13}$$

$$y_m \leq w_m + v_m \qquad m \in M \tag{14}$$

Objective functions (8) and (9) express the minimization of the total schedule displacement (total 'cost' of schedule delay) and the minimization of the number of violated slot assignments respectively. Constraints (10)-(14) are used to define variables y_m. In particular, constraint (10) implies that if the slot allocated to request m is earlier than ($\tau_m - q_m$) (i.e., the left part of the inequality is negative), then variable y_m is set equal to 1 (i.e., the slot request assignment is violated). Similarly, constraint (11) implies that if the slot allocated to request m is later than ($\tau_m + q_m$) (i.e., the left part of the inequality is negative) then variable y_m is set equal to 1. On the other hand, constraint (12) implies that if the slot allocated to movement m is earlier than ($\tau_m + q_m$), then the variable w_m is set to 0, while constraint (13) implies that if the slot allocated to movement m is later than ($\tau_m - q_m$) then the variable v_m is set to 0. Based on (10)-(13), constraint (14) implies that if the slot allocated to movement m is earlier than (or equal to) ($\tau_m + q_m$) and later than (or equal to) ($\tau_m - q_m$), then variable y_m is set equal to 0, i.e., the slot assigned to movement m lies within the tolerance limits of the airline requesting a slot. The Grandfather rights can be included in the model described above by customising sets T_m, $m \in M_G$ as in SAM-I.

Interesting insights may emerge by considering the trade-off between total schedule displacement and the number of violated slot assignments. Consider example 2 where there are six slot requests m_1-m_6 with requested time intervals: τ_{m_1}, τ_{m_2}, equal to 7, and τ_{m_4}, τ_{m_5}, equal to 10. No turnaround time is assumed. The acceptable displacement per request is equal to $\pm$ time intervals. The rolling capacity constraint is 2 movements per three time intervals, (i.e., it is forbidden to schedule more than two movements within any time period of length 3). Figure 3 illustrates the solution of the problem that optimises total schedule displacement. The solution representation scheme is in line with the one used in Example 1 (Section 3.1).

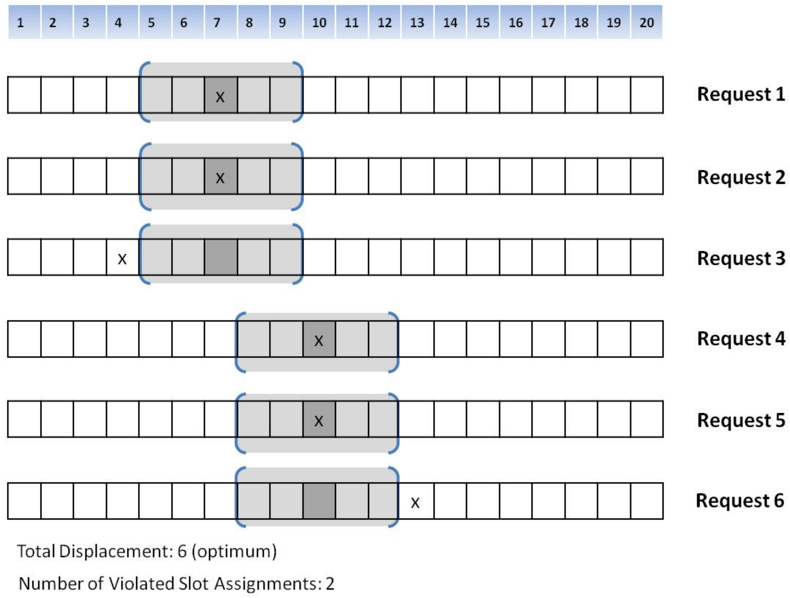

Figure 3: Optimal Solution of the Problem in Example 2
(optimising total schedule displacement)

Figure 4 represents an alternative solution to the same problem, which however minimises the number of violated slot assignments. In addition, it presents the changes that have been performed in solution 1 in order to get solution 2. It is evident that reducing the number of violated slot assignments by 1 results in an increase of total displacement by 1. As illustrated in Figure 4, this result is due to the fact that decreasing the distance by 1 of the scheduled time of m_6 from τ_6 leads to increasing the distance by 1 for two requests: m_1 and m_4 (due to capacity constraints). This example indicates that the number of violated slot assignments and the total schedule displacements are conflicting scheduling criteria. This trade-off between objective functions (8) and (9) is further verified by the computational results presented later in the paper.

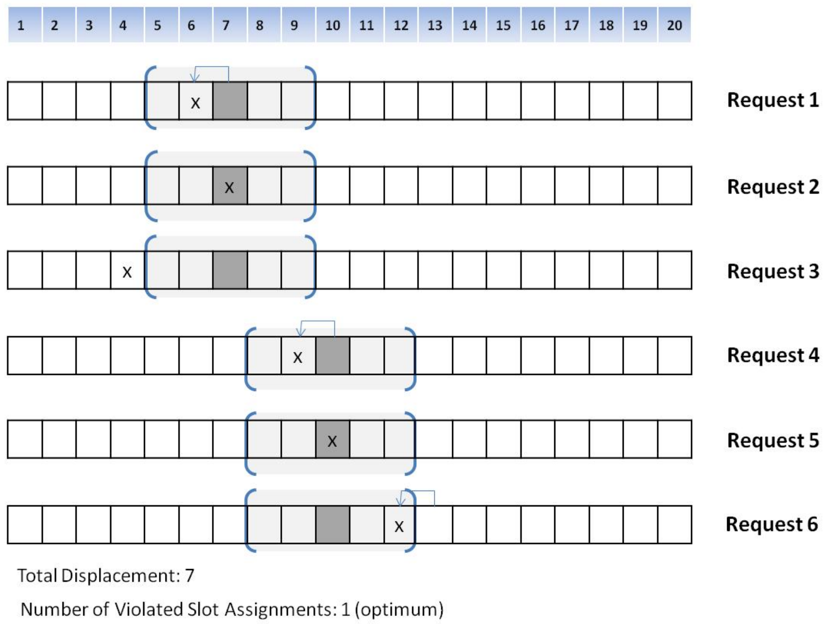

Figure 4: Optimal Solution of the Problem in Example 2

(optimising the number of violated slot assignments)

4. Solution Approach

Both bi-objective slot allocation models aim to determine the alternative non-dominated solutions for an airport within a given time horizon and capacity level. One of the methods that can be used to approximate the efficient frontier is the ε-Constraint Method (Haimes 1973). According to this method, a given bi-objective problem is converted into a parametric single-objective problem, which considers the optimisation of only one of the two objective functions of the bi-objective problem, the minimisation of the total schedule displacement cost in our case, while the second objective is converted into a parametric constraint. The steps followed for the solution of both models are similar. A brief description of the steps followed for SAM-I and SAM-II is presented below:

i) A single-objective slot allocation model is built and solved using objective function (1) (objective function 8 for SAM-II), and constraints (3-7) (constraints 3-5 and 10-14 for SAM-II). The optimal solution of both models corresponds to the minimum total schedule displacement.

ii) A single-objective slot allocation model is built and solved using objective function (2) representing the maximum schedule displacement (objective function 9 for SAM-II), and constraints (3-7) (constraints 3-5 and 10-15 for SAM-II). The optimal solution of this model corresponds to the min-max schedule displacement (the minimum number of violated slot assignments for SAM-II).

iii) A new parametric slot allocation model is built using objective function (1), constraints (3)-(7) plus (15) (objective function 8 and constraints 3-5, 10-14 plus 16 for SAM-II) presented below.

$$z \leq \varepsilon \tag{15}$$

$$\sum_{m \in M} y_m \leq \varepsilon \tag{16}$$

This model is solved iteratively for a series of values of starting with the optimum value determined in (ii) and increasing by one unit until the optimum value of the total displacement becomes equal to the optimum value of the model in step (i).

The models emerging in the above process are solved by ILOG Cplex 12.6. The solutions determined throughout steps (i)-(iii) constitute non-dominated solutions of the relevant problem.

5. Model Application and Results

The application of slot allocation models SAM-I and SAM-II may provide valuable insights regarding: i) the trade-off between the maximum and the total schedule displacement at a slot-controlled airport (SAM-I Model results), ii) the trade-off between the number of violated slot assignments and the associated total schedule displacement (SAM-II Model results) and iii) the effect of various slot scheduling constraints (e.g., grandfather rights - GFR, declared capacity levels) on the relationships stated above. The latter is applicable to both SAM-I and II Models.

Here, it is important to underline the fact that a primary aim of our analysis is to demonstrate the potential benefit of compromising scheduling efficiency in favour of airlines' acceptability examined at different tolerance limits. The understanding of airlines' slot strategies and the way that these are reflected onto preferred slot requests and acceptability threshold values requires an in-depth analysis of airlines' behavioural patterns (e.g., slot strategies, actual slot preferences, maximum displacement thresholds, timing flexibility) that goes beyond the scope of our research.

The proposed models were applied to real-world airport scheduling data. The selected airport is a small, yet coordinated during the summer scheduling season, regional airport serving roughly 2 million passengers and 16 thousand flights annually by accommodating around 500 requests for slot series during a typical summer scheduling season. The tested data pertain to the typical scheduling practice applied at the airport under consideration. Section 5.1 presents the data/input and problem parameters for the various tested scenarios, while Sections 5.2 and 5.3 discuss the results from the application of SAM-I and SAM-II models, respectively.

5.1 Input Data and Problem Parameters

Two alternative scenarios are introduced to investigate the impact of grandfathered right constraints on both objectives: i) no prioritisation for grandfathered slot requests is considered (NGFR scenario) and ii) scheduling takes into account the IATA Worldwide Scheduling Guidelines requiring the satisfaction of grandfathered requests (GFR scenario). Each scenario (i.e., NGFR, GFR) is examined at different declared capacity levels (or "capacity cases") for both SAM-I and II Models. Starting from the current declared capacity levels (4 arrivals/ARR - 6 departures/DEP - 10 total/TOT movements per hour) typically applied by the coordinator at the given airport, we develop a range of capacity cases to investigate the impact of relaxing or tightening capacity constraints on scheduling efficiency and acceptability objectives, respectively. It is worth noting here that the NGFR scenario, being less constrained, can be solved at lower declared capacity levels (starting from 4 ARR - 4 DEP - 8 TOT) than the GFR scenario. However, the starting capacity level for SAM-II model (NGFR) was 4 ARR - 5 DEP - 9 TOT since the problem could not be solved to optimality (using CPLEX 12.6 on a PC with RAM 16GB) for the capacity level 4 ARR - 4 DEP - 8 TOT. The maximum capacity case examined for the NGFR scenario is 7 ARR – 7 DEP – 14 TOT; this represents the capacity level above which there is a single optimal solution optimising both objective functions (i.e., the trade-off among the two scheduling objectives is eliminated from instances of the problem involving higher capacity levels). On the other hand, the initial capacity case examined for the GFR problem (7 ARR - 7 DEP - 14 TOT) represents the minimum capacity levels for which the problem can be solved. Similar to the NGFR problem, the maximum capacity case examined for the GFR scenario is 8 ARR – 9 DEP – 17 TOT, representing the capacity levels above which there is a single optimal solution optimising both objective functions (no trade-off among objectives).

Thus, the capacity levels examined for each scenario (i.e., NGFR, GFR) are not identical. However, the NGFR and GFR scenarios are compared for a given capacity level (7 ARR – 7 DEP – 14 TOT), for which both problems are feasible and solved with view to examining the effect of grandfathered slot priorities on the various scheduling performance metrics. Results derived from the GFR/NGFR scenarios and the various capacity cases are meant to enable the decision maker to have a clear view of the benefits and costs (on the basis of scheduling performance) that can be realised by experimenting with different declared capacity levels or loosening grandfather right constraints. Furthermore, the SAM-II Model examines the impact of different tolerance limits (q_m) applied by airlines with respect to maximum acceptable schedule displacement.

Table 1 provides an overview of the test data used for both models SAM-I and SAM-II, as well as the associated main scenarios (i.e., GFR, NGFR), the capacity cases (i.e., various declared capacity levels in terms of hourly movements for Arrivals, Departures and Total movements) and different tolerance limits (q_m) of airlines with respect to maximum acceptable schedule displacement measured as the number of coordination time intervals. The coordination time interval is the typical unit of measurement applied to all schedule displacement metrics. The length of the coordination time interval used for slot coordination purposes is 5 minutes. For instance, a maximum schedule displacement of 40 intervals corresponds to a 200-minute displacement from initially requested slot time. Consequently, a tolerance limit q_m = 9 implies that the maximum schedule displacement (before or after the requested time interval) considered acceptable for airlines is 9 coordination time intervals, that is, 45 minutes. The minimum turnaround time was set equal to 60 minutes (i.e., 12 coordination time intervals). Minimum turnaround times signify the minimum time separation required between the arrival and the departure leg of the associated paired flight.

		SAM-I Model	**SAM-II Model**
NGFR	Slot Requests	449 requests (no GFR priorities applied)	
	Capacity Cases / Levels	4 ARR – 4 DEP – 8 TOT 4 ARR – 5 DEP – 9 TOT 5 ARR – 5 DEP – 10 TOT 5 ARR – 6 DEP – 11 TOT 6 ARR – 6 DEP – 12 TOT 7 ARR – 7 DEP – 14 TOT* *(used for comparison with GFR)	4 ARR – 5 DEP – 9 TOT 5 ARR – 5 DEP – 10 TOT 5 ARR – 6 DEP – 11 TOT 6 ARR – 6 DEP – 12 TOT 7 ARR – 7 DEP – 14 TOT* *(used for comparison with GFR)
	Tolerance Limit (q_m)		q_m = 3, 6, 9, 12 intervals
	Coordination time interval	5 minutes	
	Minimum turnaround time	60 minutes	
GFR	Slot Requests	203 grandfathered, 24 new entrant, 222 other 449 requests in total	
	Capacity Cases / Levels	7 ARR – 7 DEP – 14 TOT* 7 ARR – 8 DEP – 15 TOT 8 ARR – 8 DEP – 16 TOT 8 ARR – 9 DEP – 17 TOT *(used for comparison with NGFR)	7 ARR – 7 DEP – 14 TOT* 7 ARR – 8 DEP – 15 TOT 8 ARR – 8 DEP – 16 TOT 8 ARR – 9 DEP – 17 TOT *(used for comparison with NGFR)
	Tolerance Limit (q_m)		q_m = 3, 6, 9, 12 intervals
	Coordination time interval	5 minutes	
	Minimum turnaround time	60 minutes (12 intervals)	

Table 1: Test Data for each Model and Scenario

5.2 Results of Model SAM-I

The SAM-I Model examines the trade-off between total and maximum schedule displacement (both measured in 5-min. coordination time intervals). Initially, SAM-I was used to examine the effect of prioritisation of grandfathered slots over new entrant and other slot requests on the overall slot scheduling performance (in terms of both total and maximum schedule displacement) for a given declared capacity level. Figure 5 illustrates the results (efficient frontier) of SAM-I both with (GFR scenario) and without GFR constraints (NGFR scenario). Slot demand is expressed in the form of actual airlines' slot requests, while slot availability is based on the declared capacity level of 7-7-14 (i.e., 7 arrivals, 7 departures, 14 total movements per hour) for both scenarios. According to the NGFR scenario, no distinction is made between different categories of slot requests (i.e., no priority for grandfathered slots). For the GFR scenario, the Model SAM-I is solved hierarchically satisfying first the requests for grandfathered slots, followed by new entrant and other slot requests in line with the current slot allocation regulatory framework (IATA 2014). Due to the application of the GFR constraints, a single solution is obtained for the GFR scenario as compared to the frontier obtained for the less constrained NGFR scenario (Figure 5).

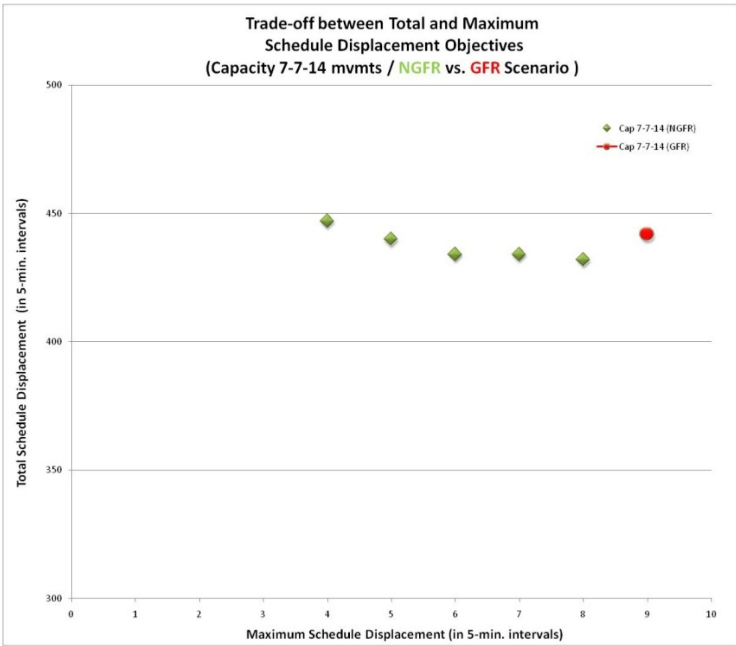

Figure 5: SAM-I Model Results for a Given Declared Capacity Level (7-7-14)
(NGFR vs. GFR Scenario)

Some interesting results emerge from the application of SAM-I Model to the GFR and NGFR scenarios (Figure 5). The range of values for total schedule displacement under both scenarios is between 432 and 447 intervals, while the maximum schedule displacement ranges between 4 and 9 intervals. It is, however, interesting to observe that, for the same level of total schedule displacement (i.e., 440 intervals), the maximum schedule displacement almost doubles from 5 to 9 intervals (or from 25 to 45 minutes, corresponding to an increase in the order of 80%) when the satisfaction of the GFR rights is assumed. Furthermore, the solution to the GFR scenario is dominated by all solutions to the NGFR problem, except the one (upper left point in NGFR curve), which minimises the maximum schedule displacement (at 4 intervals). Overall, the comparison between NGFR and GFR scenarios demonstrates clearly the strong impact of GFR on potential schedule acceptability for airlines. Although not modelled, the influence on schedule acceptability is expected to be even more prominent for those airlines (e.g., new entrants) that have not already established a GFR foothold at congested (coordinated) airports. In particular, given the fact that new entrants represent a lower priority slot class (as compared to grandfathered slots) under a GFR regime, their slot requests will be more severely displaced so that grandfathered requests can be accommodated.

Previous studies (Zografos et al. 2012) have demonstrated the strong effect of declared capacity on total schedule displacement. However, the effect of declared capacity on the trade-off between total and maximum schedule displacement has not been studied so far. The SAM-I Model was also used to examine the impact of declared capacity on both schedule displacement metrics. Figure 6 presents the efficient frontiers for different declared capacity levels (8-14 total movement per hour) for the NGFR scenario.

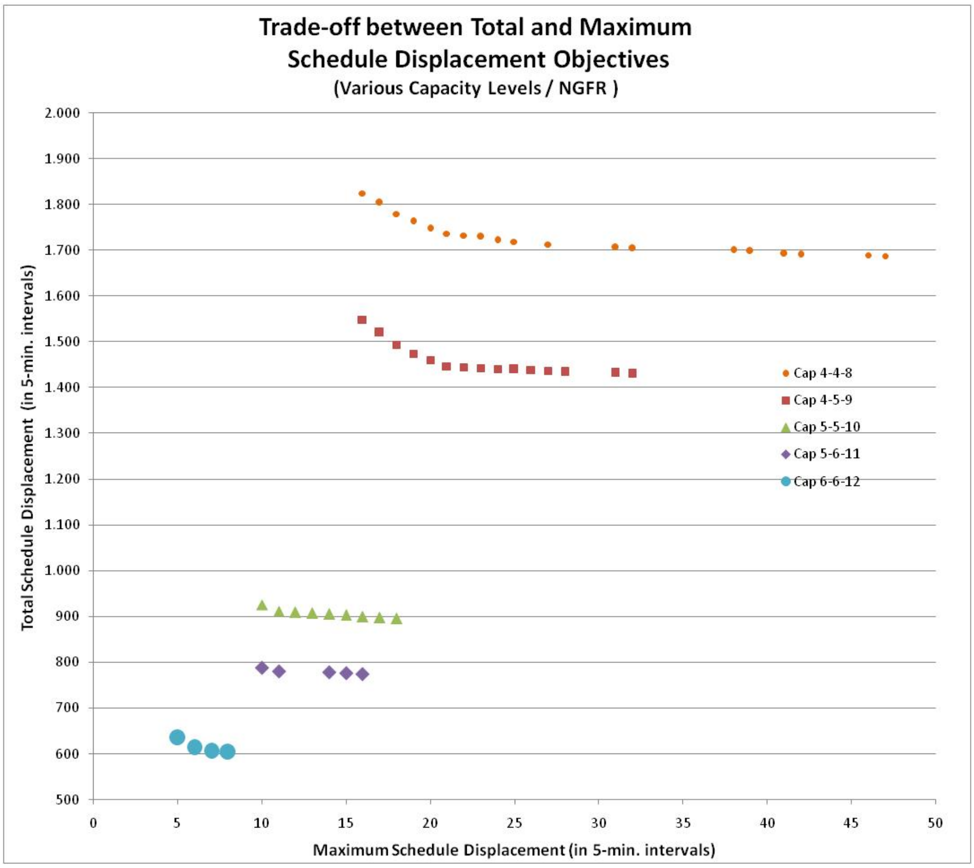

Figure 6: SAM-I Model Results for Different Capacity Levels (NGFR Scenario)

The first result emerging from the analysis presented in Figure 6 is that each frontier is clearly dominated by a higher capacity frontier. Furthermore, there is a strong and largely varying effect of declared capacity on both maximum and total schedule displacement. Notably, moving from the most constrained case (8 total movements per hour) to higher capacity cases (9, 10, 11 and 12 total movements per hour), the rate of decrease of the optimal point of maximum schedule displacement (extreme left in each curve) amounts to 0%, 37.5%, 0%, and 50%. On the other hand, the rate of decrease of the optimal value of total schedule displacement (extreme right point in each curve) is 15.3%, 37.4%, 13.5%, and 21.9%, respectively. More specifically, it can be observed in Figure 6 that when declared capacity increases by 1 movement (from 8 to 9 hourly movements), the optimal point of maximum schedule displacement remains unchanged with the optimal point of total schedule displacement being reduced by 15.3%. However, further increase in declared capacity by 1 movement (from 9 to 10 hourly movements) results roughly in a 37.5% decrease in both total and maximum schedule displacement simultaneously. The latter finding is clearly illustrated by the abrupt shift of the frontier while increasing declared capacity from 9 to 10 hourly movements.

Another interesting result emerges when one looks on the effect of increase in declared capacity for the same level of maximum schedule displacement. In particular, when declared capacity increases by 1 movement (from 8 to 9 hourly movements), for the same level of maximum schedule displacement (e.g., 18 time intervals), the total schedule displacement improves by 16.1%. However, for the same level of maximum schedule displacement (e.g., 18 time intervals), further increase in declared capacity by 1 movement (from 9 to 10 hourly movements) results in improvements in total schedule displacement in the order of 40%. This provides further evidence on the strong effect of declared capacity on total schedule displacement as a central measure of slot scheduling efficiency. Furthermore, the magnitude of the effect of declared capacity on both scheduling objectives varies substantially with the capacity levels. This further implies that any consideration towards modifying an airport's declared capacity should be based on the investigation of alternative capacity levels potentially capable of maximising the marginal improvement in both scheduling objectives.

From a different viewpoint, the trade-off among the two objectives for various capacity levels is also evident in Figure 6. As it can be concluded from Figure 6, for the same level of declared capacity, substantial gains in maximum schedule displacement can be achieved for quite small sacrifices in total schedule displacement. For instance, for the most capacity-constrained case (uppermost curve of Figure 6), a 2.9% increase in total schedule displacement (from 1,688 to 1,737 intervals) may yield a 55.3% improvement in maximum schedule displacement (from 47 to 21 intervals). Furthermore, the maximum schedule displacement objective can be minimised (16 intervals) at a "cost" of deviating 8.1% from the optimal total schedule displacement objective (from 1,688 to 1,825 intervals). Notably, this "low-cost" gain in maximum schedule displacement is present for all declared capacity cases. As a matter of fact, there is a point in all curves presented in Figure 6 at which substantial gains in maximum schedule displacement can be achieved at practically "no cost" with respect to total schedule displacement objective. For example, for a less than 1% deviation from optimal total schedule displacement objective, the improvement in maximum schedule displacement objective ranges between 34%, 28%, 20%, 31%, and 25% for the examined declared capacity levels, respectively. This finding is very useful in making slot scheduling decisions as a very small sacrifice in total schedule displacement (as a measure of scheduling efficiency) may increase the capability of the airport coordinator to allocate slots closer to their requested times for all slot requests. This, in turn, brings promises to produce more acceptable schedules from the airlines' perspective, since the maximum displacement from airlines' initial slot requests becomes significantly smaller. In addition, this threshold point may provide useful insights to airlines in setting up their maximum acceptable displacement levels.

Figure 7 presents the corresponding frontiers for different levels of declared capacity for the GFR scenario. A practical confirmation of the strong effect of GFR on the feasibility *per se* of the scheduling problem stems from the fact that the GFR scenario can be solved for declared capacity levels starting from 7-7-14 hourly movements as compared to the NGFR scenario obtaining feasible solutions at capacity levels starting from 4-4-8 movements per hour. This was reasonably expected since the GFR problem is much more constrained such that rolling capacity constraints for arrivals, departures or total movements are often violated.

A number of interesting conclusions can be drawn from Figure 7. First, both total and maximum schedule displacement are kept at low levels, while simultaneously exhibiting rather low variability, namely 360-440 intervals and 5-9 intervals, at different capacity levels. This is basically attributed to the fact that the GFR scenario can only be solved at high capacity levels (starting from 14 hourly movements). Furthermore, although each frontier is dominated by a higher capacity frontier, the impact of increased capacity levels on both objectives significantly weakens as compared to the NGFR scenario. For instance, moving from the most constrained case (14 total movements per hour) to higher capacity cases (15, 16 and 17 total movements per hour), the rate of decrease of the optimal point of total schedule displacement (extreme left in each curve) ranges between 2% and 10%. On the other hand, the rate of decrease of the optimal point of maximum schedule displacement (extreme right point in each curve) amounts to 22%, 0% and 28.5%, respectively but it is already at substantially lower levels (below 10 intervals in all cases). This is also schematically illustrated by the small to negligible shift (from upper right to lower left) of curves at different declared capacity levels. Overall, the strong effect of declared capacity on both maximum and total schedule displacement is re-confirmed. Furthermore, there is clear evidence on the "trimming" impact of GFR constraints on declared capacity's potential of improving allocation efficiency and acceptability.

Another interesting observation stems from the examination of the trade-off relationship among the two objectives. As it is demonstrated in Figure 7, the trade-off between total and maximum schedule displacement is extremely weak (if any at all), as evidenced by the rather flat profile of all curves, at all capacity levels. Interestingly, the capacity case 7-7-14 makes the GFR problem loose enough so that a single schedule (single data point rather than frontier in Figure 7) optimises both objective functions (no trade-off exists at this capacity level). As a general observation, the maximum schedule displacement can be minimised at a very "low-cost" in terms of total schedule displacement (less than 4% deviation from optimal total schedule displacement objective) for all declared capacity cases. However, it should be underlined that the gradual elimination of the trade-off should be viewed as a direct result of increasing capacity levels rather than the imposition of GFR priorities. This practically means that airlines operating at non overly congested or over-capacitated (still slot-controlled) airports can have substantial schedule acceptability gains without impacting the overall schedule efficiency for the airport users' community.

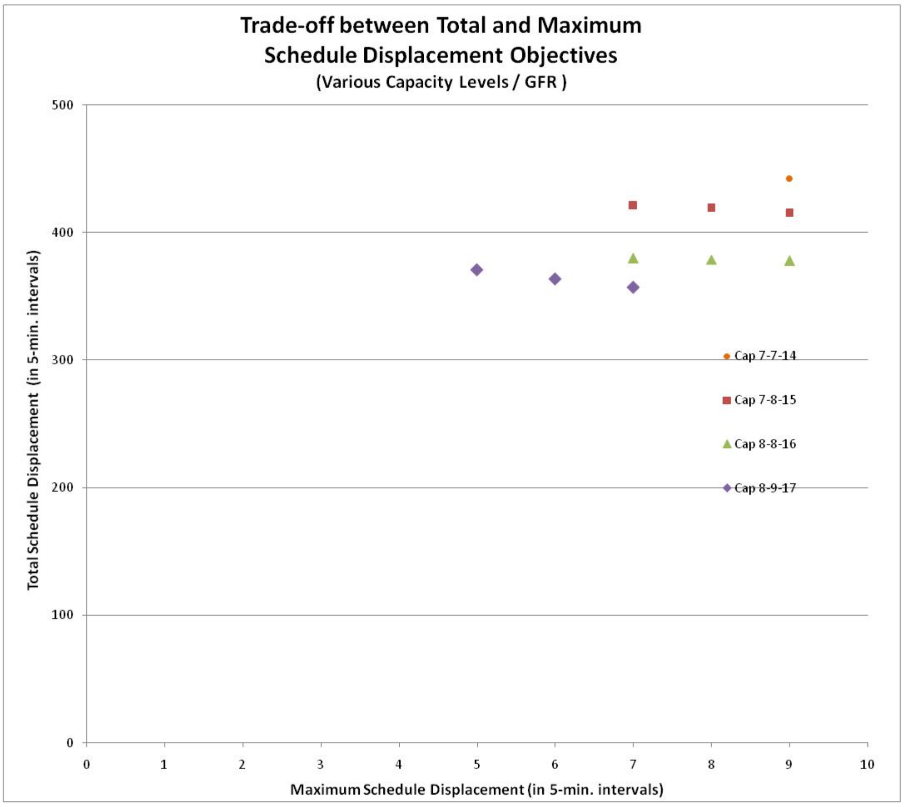

Figure 7: SAM-I Model Results for Different Capacity Levels (GFR Scenario)

5.3 Results of Model SAM-II

The SAM-II Model aims to investigate the relationship between the number of violated slot assignments and the total schedule displacement. Initially, a comparison between the GFR and the NGFR scenario is performed at a certain declared capacity level (i.e., 7 arrivals, 7 departures, 14 total movements per hour). Figure 8 presents the results from the application of SAM-II Model to the GFR and NGFR scenarios.

It is interesting to observe that for tight acceptability thresholds (q_m = 3 or 6), both objectives are simultaneously reduced as depicted by the sharp left shift of curves from the GFR to the NGFR scenario. For instance, for q_m = 3, the minimum total schedule displacement decreases from 442 intervals (GFR) to 432 intervals (NGFR) and the optimal value for the violated slot assignment objective drops from 13 (GFR) to 4 (NGFR) violated assignments. Most importantly, for the same level of total schedule displacement (e.g., 440 intervals) and tolerance level (e.g., q_m = 3), the number of violated slot assignments decreases from 20 (GFR) to 4 (NGFR), corresponding to an improvement in the order of 80%. This finding highlights the influence of GFR requirements on the overall slot scheduling efficiency (the imposition of GFR leads to an upward shift of total schedule displacement) but, most importantly, it unfolds the strong impact of GFR on maximum displacement. Larger values of maximum displacement impact, in turn, airline acceptability and consequently the potential utilisation of slots by airlines that are forced to drastically modify their scheduling preferences during high demand periods.

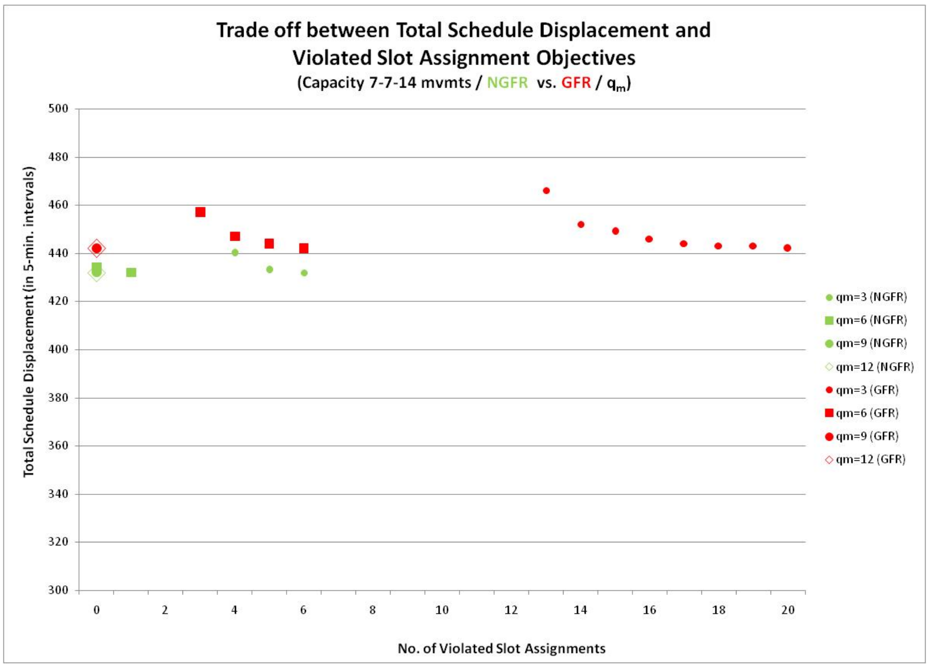

Figure 8: SAM-II Model Results for a Given Declared Capacity Level (7-7-14) (NGFR vs. GFR Scenario)

The SAM-II Model was also used to examine the relationship between the number of violated slot assignments and the total schedule displacement at different values of maximum acceptable displacement (tolerance limits) and different levels of declared capacity. For the NGFR scenario, four different tolerance limits (i.e., q_m = 3, 6, 9 and 12) are examined, namely 15, 30, 45 and 60 minutes, and six different levels of declared capacity, that is, 9, 10, 11 and 12 hourly movements, respectively. Figures 9a-d present the results (efficient frontiers) of the NGFR scenario for different declared capacity levels and tolerance limits. For all analysed scenarios, the number of violated slot assignments decreases as the maximum acceptable schedule displacement increases. This result is attributed to the fact that the increase of maximum acceptable displacement leaves a smaller number of slot requests (non-accommodated) displaced beyond the acceptable level.

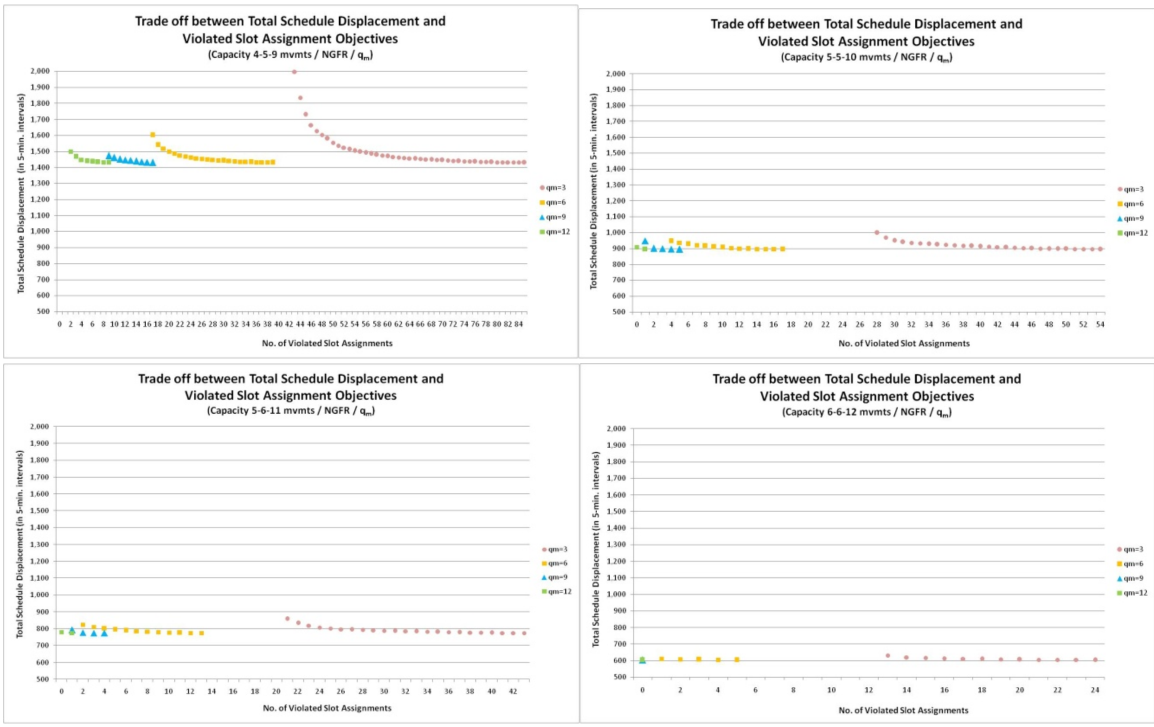

Figure 9a-d: SAM-II Model Results for Different Capacity Levels and Tolerance Limits (NGFR Scenario)

Another intuitive result emerging from the application of the SAM-II Model is that for the same value of maximum acceptable displacement (e.g., 15 minutes), the number of violated slot assignments and the associated total schedule displacement are decreasing as the airport declared capacity increases. For instance, at 4-5-9 capacity level, when the maximum acceptable displacement value is set to 15 minutes, the minimum total schedule displacement value is 1,430 intervals, which corresponds to a maximum number of 85 violated slot assignments (Figure 9a). The corresponding values of minimum total schedule displacement and violated slot assignments for declared capacity levels of 10, 11, and 12 hourly movements are 895 intervals and 54 violated slot assignments (Figure 9b), 774 intervals and 43 violated slot assignments (Figure 9c), and 604 intervals and 24 violated slot assignments (Figure 9d), respectively. A similar pattern exists when the maximum acceptable displacement increases. It is also worth noting that, as the maximum acceptable displacement increases, both the total schedule displacement and the associated number of violated slot assignments decrease simultaneously. This means that as the airport's declared capacity and the airlines' tolerance levels increase, there is less need to displace slots to periods of the day experiencing lower demand and thus the total schedule displacement decreases. At the same time, a higher tolerance level of schedule displacement generates a smaller number of associated violated slot assignments as more slots are available closer (in terms of time) to airlines' requests.

It is worth noting that for the most capacity-constrained case (i.e., 9 movements) and the minimum acceptable displacement value (15 minutes), the rate of change (decrease) of violated slot assignments is much higher as compared to the cases of 30, 45 and 60 minutes of maximum acceptable schedule displacement. More specifically, for the 15-minute maximum acceptable displacement, a 50% reduction of the number of violated slot assignments (moving from the rightmost point of the curve, that is, 85 to 43) results to a 40% increase of the associated minimum total schedule displacement (from 1,430 to 1,997 intervals). This rate of change decreases while moving progressively to higher levels of maximum acceptable displacement (from the rightmost to leftmost curves of Figures 9a-d). Therefore, in the case of 30-minute maximum acceptable schedule displacement, halving the maximum number of violated slot assignments (from 39 to 19) results to only a 6% increase of the minimum total schedule displacement (from 1,430 to 1,517 intervals). For the cases of 45 and 60 minutes, halving the maximum number of violated slot assignments would improve total schedule displacement by only 3% and 1%, respectively. This practically means that the stricter the acceptability thresholds for airlines, the more "expensive" (in terms of total schedule displacement) for the airport users as a whole to satisfy their slot requests within tolerable limits. This finding might be essential for triggering some shift in airlines' behaviour on the grounds that, by imposing tight acceptable thresholds, airlines eventually "penalise" themselves both in terms of more violated slot assignments and higher schedule displacement for the entire airport user community.

Another interesting result relates to the observation that, for the same level of declared capacity, the relationship between the number of violated slot assignments and the total schedule displacement is relatively flat, and that the rate of change of total schedule displacement as a function of the number of violated assignments increase fast beyond a threshold value of the number of violated assignments. This means that substantial gains in the reduction of the number of violated assignments may be realised through small or even negligible total schedule displacement "sacrifices". For the most capacity-constrained case (9 total movements) and the tightest acceptable schedule displacement (15 minutes) (Figures 9a), a 41% reduction in the number of violated assignments (from 85 to 50) can be achieved at less than 9% increase in total schedule displacement (from 1,430 to 1,554 intervals). Interestingly, for the same capacity case (9 total movements), comparable savings (40%) in the number of violated assignments can be achieved at a total schedule displacement "cost" that does not exceed 3% for less strict tolerance levels (i.e., 30, 45 and 60 minutes). Furthermore, it can be observed that the same level of total schedule displacement may lead to significantly different number of violated assignments depending on the maximum acceptable schedule displacement value. Again, in the most capacity-constrained case (9 total movements) (Figures 9a), for every 15-minute increase of the maximum acceptable schedule displacement (q_m = 3, 6, 9, 12), there is approximately a 50% reduction in the number of violated assignments (85, 35, 16, and 9 violated assignments, respectively) for the same level of total schedule displacement (around 1,430 intervals). Moreover, it is worth noting that, as declared capacity increases, these gains in the number of violated assignments require lower sacrifices in terms of the associated total schedule displacement.

Overall, it can be concluded that substantial improvements in the violated assignment objective can be "easily" (close to freely up to a threshold point) achieved without sacrificing much on the total schedule displacement objective for any value of q_m. Improvements in the violated assignment objective (as a function of the total schedule displacement) increase substantially with capacity levels (they practically come at no cost for a declared capacity level above 10 total movements per hour). The selection of the highest possible value for q_m (e.g., 60 minutes) is clearly preferable, since it practically eliminates violated assignments at only a negligible deviation from the optimum total schedule displacement. On the other hand, it may result in unacceptable level of service for airlines, which, in turn, may give rise to slot misuse patterns and hence under-utilisation of scarce airport capacity (e.g., late return of unwanted slots, flights operated at different than allocated slot times / "off slot", failure to operate allocated slots / "no shows"). The above findings are shedding light to the relationship between slot displacement (a measure of airline disutility) and the optimum utilisation of declared airport capacity, as they can help airlines understand that the acceptance of a reasonable level of maximum displacement (e.g., 15-30 minutes) can reduce substantially the number of violated slot assignments (unsatisfied demand), while at the same time the overall slot scheduling efficiency is not significantly degraded, and the overall airport capacity utilisation is improved. It should be underlined here that in policy terms, the concept of maximum acceptable schedule displacement can perfectly fit with a slight adaptation and stricter enforcement of the "Timing Flexibility" option (i.e., time range within which slots are considered acceptable) that has been recently introduced in the IATA-based slot allocation scheme (IATA 2014).

Figure 10a-d present the corresponding frontiers for different levels of declared capacity and tolerance limits for the GFR scenario. Similarly to the results of the SAM-I Model, the GFR scenario can be solved for higher declared capacity levels starting from 7-7-14 hourly movements as compared to the NGFR scenario that obtains feasible solutions for capacity levels starting from 4-5-9 movements per hour.

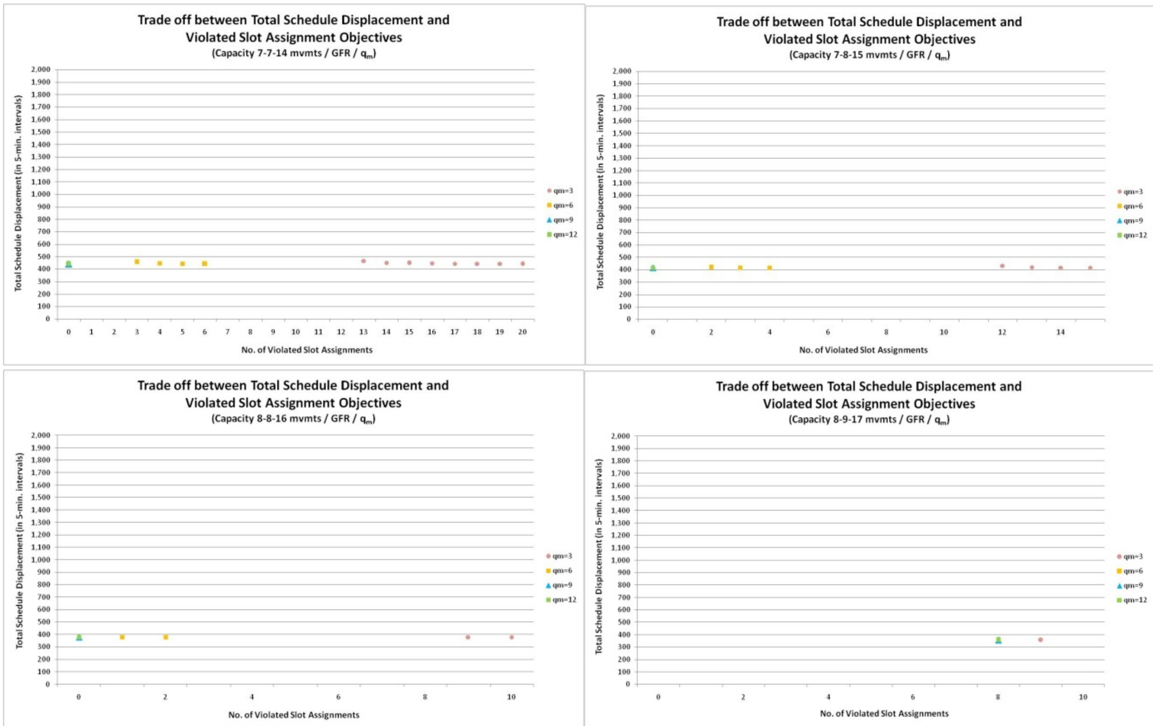

Figure 10a-d: SAM-II Model Results for Different Capacity Levels and Tolerance Limits (GFR Scenario)

Unlike to the NGFR scenario, the trade-off between total schedule displacement and the number of violated slot assignments is practically eliminated at all capacity levels and tolerance limits for the GFR scenario. Interestingly, both objectives are kept at low levels without much variability with capacity levels or tolerance limits. In particular, the total schedule displacement ranges between 357 and 466 intervals, while the number of violated slot assignments is kept at less than 13 slot requests at all cases. Furthermore, for the same level of maximum acceptable displacement (e.g., 15 minutes), the number of violated slot assignments and the associated total schedule displacement are decreasing as the airport declared capacity increases. For example, setting the maximum acceptable displacement value at 15 minutes and shifting from 7-7-14 to 8-9-17 capacity levels, the minimum total schedule displacement takes values 442, 415, 378 and 357, with the corresponding values of the violated slot assignments being 20, 15, 10 and 9 violations, respectively (Figure 10a). On the other hand, there seems to be a sharp decrease in the number of violated slot assignments when shifting from 15 to 30 minutes of maximum acceptable displacement (please refer to an abrupt shift of the frontiers between q_m = 3 and 6) at all capacity levels. At the same time, a relaxation of acceptability thresholds is still accompanied by small, incremental improvements in the total schedule displacement, which are, however, mostly evident in lower capacity levels (e.g., 7-7-14 movements).

Another interesting finding stems from the observation that, in the presence of GFR constraints, when relaxing tolerance levels above 30 minutes (e.g., q_m = 9 and 12), there is a single optimal solution at all capacity levels. This solution corresponds to a total schedule displacement ranging between 357 and 442 intervals, while the number of violated assignments drops to zero. This practically means that while increasing capacity levels and relaxing tolerance limits, the number of violated slot assignments can be practically eliminated, but there seems to be an "irreducible minimum" for the overall slot scheduling efficiency in terms of total schedule displacement.

6. Concluding Remarks

In this paper, we investigated alternative metrics expressing the acceptable schedule displacement and airlines' tolerance levels against schedule displacements. We introduced the concept of violated slot assignments, i.e., slots allocated to airlines beyond a maximum acceptable displacement as compared to the slot requested by airlines. The concept of violated slot assignment is useful for approximating the actual utility derived by a given airline from a specific slot. It essentially implies that, for a given airline, the commercial value of an assigned slot drops to zero, and hence does not make sense to operate it, if it is assigned beyond a threshold value of time distance from the requested slot. Even at congested airports, a significant number of allocated slots are not eventually used (Steer Davies Gleave 2011) by the airlines due to the significant displacement of the allocated slot from their originally requested slot. Therefore, the concept of violated slot assignments can be used as a proxy for increasing the acceptability of the allocated slots, hence increasing the potential utilisation of scarce airport resources.

Two bi-objective models were introduced to examine the relationship between slot acceptability (expressed through different metrics of acceptable schedule displacement) and slot scheduling efficiency (expressed in the form of total schedule displacement). For both models, two scenarios were developed to examine the effect of the IATA Grandfather Rights (GFR) on total schedule displacement and acceptability. The models were implemented and solved for a small regional airport using test data regarding declared airport capacity, airlines' slot requests, as well as other coordination parameters (e.g., coordination time interval, minimum turnaround time) specified by the respective slot coordination authority. Both models were solved for a range of capacity values ("capacity cases") starting from current declared capacity levels and then varying (loosening or tightening) capacity constraints accordingly. Despite the fact that certain capacity levels go beyond the actual operational capability of the given airport, these were examined with the aim to investigate the impact of capacity constraints on scheduling efficiency and acceptability objectives.

The first model (Model SAM-I) was used to examine the trade-off between total and maximum schedule displacement for an airport of a given declared capacity and profile of slot requests. The SAM-I Model provided the efficient frontier between total schedule displacement, as a central measure of schedule efficiency, and maximum schedule displacement, expressing a metric of schedule acceptability and potential slot utilisation. The maximum schedule displacement provides a measure of the worst-case service level provided to airlines at the airport under consideration. Therefore, it can be implicitly viewed as a “guaranteed service level” on the grounds that airlines can be reasonably ensured in advance of the scheduling season that slots will be either granted as requested or rescheduled within a time range that will not definitely exceed this worst-case service level (i.e., maximum displacement). This information can facilitate and promote Airport Collaborative Decision Making (A-CDM) practices at strategic airport level. Moreover, it may serve as a marketing tool for airports and coordination authorities aiming to promote connections from/to a given airport by committing in advance to a certain service level to the airport user community (e.g., airlines, ground handling agents, passengers).

The efficient frontier developed for the SAM-I Model can help decision makers, airport operators and airlines better explore the capacity capabilities of a given airport and quantify what should be “sacrificed” in terms of total schedule displacement in order to increase slot/capacity utilisation. Interestingly, for the most capacity-constrained case of the NGFR scenario, substantial improvements in maximum schedule displacement can be achieved at rather low “cost” in terms of total schedule displacement. Notably, this "low-cost" gain in maximum schedule displacement was present for all declared capacity cases. This finding is intuitively appealing since very small sacrifices in total schedule displacement (as a measure of scheduling efficiency) may increase the capability of the airport coordinator to allocate slots closer to their requested times for all slot requests. In effect, this would reasonably result in more acceptable slots allocated and hence more intensively used.

The second model (Model SAM-II) was used to examine the effect of various threshold values of acceptable schedule displacement and declared capacity levels on total schedule displacement and the associated number of violated slot assignments. The results of the SAM-II Model can provide decision support to both airlines and airport operators in order to determine the maximum acceptable level of schedule displacement for the airlines. From the analysis of SAM-II Model results, it was concluded that substantial improvements in the violated assignment objective can be “easily” achieved without sacrificing much on total schedule displacement objective for any value of maximum tolerance limits (q_1m) . The selection of large values of q_m (e.g., 60 minutes) is clearly preferable, since it practically eliminates violated assignments at only a negligible deviation from the optimum total schedule displacement. On the other hand, a large value for q_m may result in unacceptable level of service for airlines, hence raising again slot misuse (e.g., no shows, off slot, late slot returns) patterns and under-utilisation of scarce airport capacity. The findings of the SAM-II Model suggested that if airlines are ready to formally accept a reasonable level of maximum displacement (e.g., 15-30 minutes), substantial benefits can be achieved in terms of both violated slot assignments and better capacity utilisation rates without “sacrificing” a lot in terms of the overall scheduling efficiency expressed through total schedule displacement.

An interesting direction for future research stems from the fact that the current slot allocation process leaves room for “strategic manipulation” by airlines. In particular, it does not ensure that airlines will submit their schedules truthfully in terms of the desired number/quantity of slots, actually preferred slot times or actual tolerance limits (acceptability thresholds). For example, airlines may intentionally "inflate" their slot requests in order to secure their actual slot preferences. However, this may result in a rather unrealistic demand profile that eventually jeopardises the inherent ability of scheduling methods/models to allocate capacity in the presence of scarce resources. From an operations research (scheduling) point of view, recent research has explored modelling formulations with fairness considerations (Zografos and Jiang 2016) through which airlines are "penalised" (in terms of slot displacements) in direct relationship with the number of requested slots. In other words, the more they compete for a scarce resource, the more the total displacement that is assigned to the airline as a disincentive against unrealistic or inflated slot requests. A different stream of research lies on market-driven mechanisms aiming to reveal the actual preferences of airlines. One method successfully used in previous research was strategic simulations in which actual decision makers and industry experts are invited to gradually unfold their preferences through an iterative simulation game procedure (Ball et al. 2007b). In a similar context, auctions provide a simple and transparent price discovery mechanism urging all airlines to reveal their preferences and economic valuations of scarce airport resources (NERA 2004; Ball et al. 2007b). More recently, Gillen et al. (2016) proposed an integrated framework pursuing synergies between operations research (scheduling) methods and economic schemes in order to better deal with airline preferences and improve the efficiency of the slot allocation outcome.

At the outset, the allocative efficiency of any scheduling model or other allocation scheme is largely determined by the actual behavioural patterns and preferences of airlines. The latter need to be viewed under the prism of the entire process of developing airline slot strategies in the absence of information about competitors' strategies. This may pave the way to the development of game-theoretical models aiming at determining Nash equilibrium slot scheduling strategies (Vaze and Barnhart 2012). In this paper, we demonstrated *inter alia* the potential benefit of compromising scheduling efficiency in favour of airlines' acceptability tested at different tolerance limits. The explicit analysis and modelling of airlines' behavioural patterns (e.g., slot strategies, actual slot preferences, maximum displacement thresholds, timing flexibility) does not fall within the scope of this paper and constitutes a promising research area that merits further investigation both from the operations research and economic standpoint.

A number of issues identified throughout the course of this research have been already brought forward for further research that is currently underway within the framework of the OR-MASTER programme grant (OR-MASTER 2016). These issues relate to the modelling of the utility function for different types of airlines (e.g., low-cost carriers vs. traditional airlines, long haul hub-to-hub flights vs. feeder flights), the consideration of multiple objectives (including fairness considerations) for the scheduling of slots, as well as the development of efficient heuristics for effectively solving large-scale scheduling problems at large hub airports.

Acknowledgments

The work reported in this paper has been partially supported by the UK's Engineering and Physical Sciences Research Council (EPSRC) through the Programme Grant EP/MO20258/1 "Mathematical models and algorithms for allocating scarce airport resources (OR-MASTER)".

List of Acronyms

A-CDM	Airport Collaborative Decision Making
ACI	Airports Council International
ARR	Arrival(s)
ATFM	Air Traffic Flow Management
ATM	Air Traffic Management
DEP	Departure(s)
EU	European Union
EWR	Newark Liberty International Airport
GFR	Grandfathered Rights (or slots)
HDR	High Density Rule
IATA	International Air Transport Association
JFK	John F. Kennedy International Airport
SAM	Slot Allocation Model
TOT	Total Number of Movements (arrivals and departures)

List of References

A. Jacquillat and A.R. Odoni. An Integrated Scheduling and Operations Approach to Airport Congestion Mitigation. *Operations Research*, 63(6):1390-1410, Published Online December 31, 2015, http://dx.doi.org/10.1287/opre.2015.1428.

A.R. Odoni and T. Morisset. Performance Comparisons between US and European Airports. *Proceedings of the 12th World Conference on Transport Research (WCTR)*, July 11-15, Lisbon, Portugal, 2010.

Airport Council International (ACI) Europe. *ACI Europe position on the proposed revision of the Council Regulation (EEC) No 95/93 on common rules for the allocation of slots at Community airports*, Presentation at the TRAN Meeting at the European Parliament, March 25, Strasbourg, France, 2009.

D. Bertsimas and S. Stock Patterson. The Air Traffic Flow Management Problem with Enroute Capacities. *Operations Research*, 46(3):406–422, 1998.

D. Gillen, A. Jacquillat and A.R. Odoni. Airport demand management: The operations research and economic perspectives and potential synergies. *Transportation Research Part A – Policy and Practice*, 94:495-513, 2016.

D. Koesters. Airport Scheduling Performance – An Approach to Evaluate the Airport Scheduling Process by Using Scheduled Delays as Quality Criterion. *Proceedings of the Air Transport Research Society Annual World Conference*, June 21-23, Berkeley, U.S., 2007.

Eurocontrol. *Challenges of Growth 2013*, Task 4: European Air Traffic in 2035, Brussels, Belgium, 2013.

European Commission. *European Council Regulation No. 95/93 of January 1993 on Common Rules for the Allocation of Slots at Community Airports*, Official Journal of the European Union, L014, pp. 0001-0006, Brussels, Belgium, 1993.

European Commission. *European Council Regulation No 793/2004 of April 2004 amending Council Regulation No 95/93 on common rules for the allocation of slots at Community airports*, Official Journal of the European Union, L138, pp. 50-60, Brussels, Belgium, 2004.

European Commission. *European Council Regulation No 545/2009 of June 2009 amending Council Regulation No 95/93 on common rules for the allocation of slots at Community airports*, Official Journal of the European Union, L167, pp. 24-25, Brussels, Belgium, 2009.

International Air Transport Association (IATA). *Worldwide Slot Guidelines*, 6th Edition, Montreal, Canada, 2014.

J. Bottcher, A. Drexl, R. Kolisch and F. Salewski. Project Scheduling Under Partially Renewable Resource Constraints. *Management Science*, 45(1):543-599, 1999.

K.G. Zografos, Y. Salouras and M.A. Madas. Dealing with the Efficient Allocation of Scarce Resources at Congested Airports. *Transportation Research Part C – Emerging Technologies*, 21(1):244-256, 2012.

K.G. Zografos and Y. Jiang. Modelling and Solving the Airport Slot Scheduling Problem with Efficiency, Fairness, and Accessibility Considerations. *TRISTAN Symposium 2016*, June 13-17, 2016, Oranjestad, Aruba.

K.G. Zografos, M.A. Madas and K.N. Androutsopoulos. Increasing Airport Capacity Utilisation through Optimum Slot Scheduling: Review of Current Developments and Identification of Future Needs. *Journal of Scheduling*, 20(1): 3-24, 2017.

L. Castelli, P. Pellegrini and R. Pesenti. Ant Colony Optimization for Allocating Airport Slots. *2nd International Conference on Models and Technologies for ITS (MT-ITS)*, Leuven, Belgium, June 22-24, 2011.

L. Castelli, P. Pellegrini and R. Pesenti. Airport Slot Allocation in Europe: Economic Efficiency and Fairness. *International Journal of Revenue Management*, 6(1/2):28-44, 2012.

L. Corolli, G. Lulli and L. Ntaimo. The Time Slot Allocation Problem under Uncertain Capacity. *Transportation Research Part C - Emerging Technologies*, 46:16-29, 2014.

M.A. Madas and K.G. Zografos. Airport Slot Allocation: A Time for Change? *Transport Policy*, 17(4):274-285, 2010.

M. Ball, C. Barnhart, G. Nemhauser, and A. Odoni. Air Transportation: Irregular Operations and Control. *Handbook in Operations Research & Management Science*, 14:1-67, 2007a.

M. Ball, L. Ausubel, F. Berardino, P. Cramton, G. Donohue, M. Hansen, and K. Hoffman. Market-Based Alternatives for Managing Congestion at New York's LaGuardia Airport. *Proceedings of the AIRNETH/GARS Workshop*, April 2007b, Netherlands.

N. Pyrgiotis and A. Odoni. On the Impact of Scheduling Limits: A Case Study at Newark Liberty International Airport. *Transportation Science*, 50(1):150-165, Published Online January 12, 2015, http://dx.doi.org/10.1287/trsc.2014.0564.

National Economic Research Associates (NERA). *Study to Assess the Effects of Different Slot Allocation Schemes*, Technical Report prepared for the European Commission (DG TREN), London, UK, 2004.

OR-MASTER. *Mathematical Models and Algorithms for Allocating Scarce Airport Resources (OR-MASTER) Programme Grant*, Available online at: http://www.lancaster.ac.uk/news/articles/2015/uk-experts-to-lead-new-drive-to-unlock-airport-capacity-and-tackle-congestion/ (accessed May 8, 2016).

SESAR Joint Undertaking. *European ATM Master Plan*, Edition 2015, Belgium, 2015.

Steer Davies Gleave. *Impact Assessment of Revisions to Regulation 95/93*, Study prepared for the European Commission (DG MOVE), London, UK, 2011.

V. Vaze and C. Barnhart. Modeling Airline Frequency Competition for Airport Congestion Mitigation. *Transportation Science*, 46(4):512–535, 2012.

Y. Haimes. Integrated System Identification and Optimization. In *Control and Dynamic Systems: Advances in Theory and Applications*, (C. Leondes, ed.). Vol. 9, pp. 435-518. Academic Press, New York, 1973.