Destination selection and flight scheduling for regional airlines at slot-constrained airports

Alberto Santini[1,2]

[1] Department of Economics and Business, Universitat Pompeu Fabra, Carrer Ramon Trias-Fargas, 25–57, 08005, Barcelona, Spain
[2] Data Science Centre, Barcelona School of Economics, Carrer Ramon Trias-Fargas, 25–57, 08005, Barcelona, Spain

alberto.santini@upf.edu
0000-0002-0440-0357

August 16, 2024

Abstract

We present the problem of a regional airline based at a slot-constrained airport, which must select a set of destinations to serve, how many flights per day to operate to each destination, at what time the flights take place, and which aircraft operates each flight. Restricting ourselves to the special case of regional airlines, which fly round trips from a central hub, we are able to tackle the four above decisions jointly. By contrast, in the existing literature, these decisions are usually optimised separately for generic airline networks. To solve the proposed problem, we introduce two compact integer formulations: a three-index and a two-index formulation. This latter, however, only solves a relaxation of the original problem because it cannot guarantee that all constraints are respected. Therefore, we embed the two-index formulation into an iterative algorithm that dynamically adds violated constraints. Computational experiments highlight the validity of this approach and provide insights into the characteristics of the solutions.

Keywords: airline optimization; destination selection; fleet assignment; tail assignment.

1 Introduction

We tackle the problem of selecting a set of destinations and operating flights for a regional airline based at a slot-constrained airport. This is a tactical decision that affects the airline's operations in the middle to long term. It consists of choosing which destinations the airline will serve, how many flights per day it will operate to each destination, at what time the flights will take place, and which aircraft will operate each flight. We concentrate on a specific scenario characterised by the following properties:

- The airline operates from a single hub. All the flights are round trips from and back to the hub.
- Operations at the hub are constrained by slot availability, i.e., by the airport capacity. This can be the case if the regional airline shares the hub with many (or larger) carriers. Alternatively, the hub might be a small airport whose facilities only allow a few movements per day. Airport congestion and slot availability hinder airlines' operations and were identified by Barnhart and Cohn (2004) as one of the drivers which push airline planners towards sophisticated optimisation approaches. Indeed, slot assignment is an optimisation problem in itself (see, e.g., Bennell, Mesgarpour, and Potts 2011; Cheung et al. 2021).
- Slot availability is the strongest constraint on the number of flights which the airline can schedule. In other words, there is more demand for flights than available slots to operate them. The airline must then decide which demand it wants to serve and how.

- We assume that the airline fleet is homogeneous, although our approach can easily be adapted to the heterogeneous case (see Section 4). We also assume that the number of available aircraft is fixed, i.e., we do not consider the problem of sizing the fleet by buying or selling aeroplanes.

The planner works with a set of destinations and an ideal number of flights to each of them, depending on the demand. Each destination is associated with a utility. If the airline operates in an unregulated market, the utility coincides with the expected profit of flying one aeroplane to the destination and back. In the case of remote airports, which is of particular interest for our application, the utility can also depend on government subsidies (see, e.g., Pita, Barnhart, and Antunes 2013, for a related problem). For a given destination, the utility changes based on the departure time of the flight: not many people, for example, would like to take off at 04:00 in the morning. The objective of the planner is to create a schedule assigning round trips to the available fleet in order to maximise the total utility.

In the beginning, we will assume that the utility of a flight is given in advance and only depends on the destination and the departure time. This assumption disregards the fact that the demand is affected by the number of scheduled flights, especially if these are close in time. For example, a flight at 09:00 is likely to cannibalise the demand of another flight at 10:00 with the same destination. We mitigate this effect by adding a minimum padding between two consecutive flights to the same destination. Still, in Section 4, we explain how it is possible to extend our model to take advantage of more precise market data, by modelling passenger demand explicitly with a discrete choice model.

The planner is subject to the following constraints:

1. The total flying time of each aeroplane is limited and depends on the planning period. The shortest planning period is one day; for example, the planner might want to assign three destinations to a specific aircraft and repeat the same schedule every day. In this case, the total time required for the flights assigned to the aeroplane cannot exceed twenty-four hours. Indeed, turn-around times and required maintenance will make the actual available time lower. The assumption that the schedule must repeat based on one (or a few) base schedules is justified in the case of a regional airline that flies short-haul flights.

2. The demand limits the maximum number of flights to a destination. Furthermore, flights to the same destination should be spread out during the day because such an arrangement is known to maximise the realised demand (Pita, Barnhart, and Antunes 2013). As such, there is a minimum time between two flights to the same airport.

3. The maximum number of flights that can take off from the hub at nearby times is limited by the number of slots held by the airline. At large airports, it is not uncommon that slots are held in increments of five minutes. For example, if an airline owns two slots at 09:00, it can put two aeroplanes in the air between 09:00 and 09:05. At smaller airports with less capacity, slots can be associated with larger intervals.

4. A company operating aeroplanes incur high capital and fixed costs, such as their purchase or lease costs. The airline pays these costs no matter if the aeroplane is flying or sitting idle on the ground. Airlines operate on tight margins and must make sure that each aircraft is used as much as possible. This often translates into setting a minimum utilisation time for each aeroplane. For example, the planner might want all their aircraft to be busy—flying, turning around, or in maintenance—at least 75% of the time. Such a requirement is a way for the airline to account for strategic decisions—an aircraft can be operated for twenty-five years (Dixon 2006)—during tactical planning (destination selection can affect just one season).

We denote the problem described above as the *Regional Flight Scheduling Problem* (RFSP). In the rest of this section, we highlight the paper's contributions, and we analyse the related literature. In Section 2, we show that the RFSP is $\mathcal{NP}$-complete, and we propose two possible exact approaches to solve the problem. Section 3 presents computational results on a set of realistic instances (with added noise to preserve confidentiality) and reports managerial insights. Section 4 explains how to extend the model to account for further realistic constraints. Section 5 concludes and outlines future research paths.

1.1 Contributions

In this paper, we present a mathematical formulation to tackle the RFSP, which we prove to be $\mathcal{NP}$-complete (see Theorem 1). The RFSP combines the common destination selection, frequency selection, fleet assignment and tail assignment problems from the airline optimisation literature. These problems are usually treated separately. However, leveraging the assumptions of the special case of a regional airline flying round trips from a single hub, we tackle them jointly within a single model.

We start with a three-index formulation, in which each binary variable is associated with an aeroplane, a destination, and a time slot. Next, we introduce a two-index formulation in which we drop the aeroplane index. This formulation, however, only solves a relaxation of the original problem. In particular, because it does not keep track of which destination is assigned to which aircraft, it cannot ensure that the solution respects the minimum utilisation constraint (see point 4 above). For this reason, we embed this formulation into an iterative procedure that dynamically adds violated constraints.

Computational results on realistic instances show that using this cut generation algorithm gives better results in a shorter time compared to solving the three-index formulation. At the same time, the minimum utilisation constraints are usually tight at the optimum, and, therefore, it is not possible to solve the two-index formulation without dynamically separating these cuts.

1.2 Related literature

This work contributes to a rich literature about optimisation methods in the aviation industry and, in particular, to the problem of designing an airline's network. For general surveys on this topic, we refer the reader to the excellent books (Belobaba, A. Odoni, and Barnhart 2016; Barnhart and Smith 2012) and to (Yu and Yang 1998; Barnhart, Belobaba, and A. R. Odoni 2003). In the following, we give an overview of works which include a destination selection component. Some of the below approaches share common characteristics with our work, such as assigning departure times to flights or flights to aircraft. However, there are only a few works which explicitly consider slot availability as a crucial constraint. In fact, we could only find two papers, by Pita, Barnhart, and Antunes (2013) and Vaze and Barnhart (2012), with this characteristic. In (Pita, Barnhart, and Antunes 2013), the authors work with a given network in which hub assignment, destination selection, fleet composition, expected profits, demand patterns, and competition flight frequencies are known. Under the assumption that market share is a piecewise linear function of frequency, they develop a flight scheduling and fleet assignment model which takes into account connections and demand recapture (i.e., the part of the demand spilt and regained in a later period). They apply their Mixed-Integer Programming (MIP) model to the case of TAP Airlines of Portugal and find significant potential for increasing revenue when: slots are flexible, i.e., can be changed within $\pm$ 1 hour of the original slot time; passenger delay costs are explicitly accounted for in the model; there is more padding between flights, thus increasing connection times and decreasing the probability of delays. In (Vaze and Barnhart 2012), the authors attribute the scarcity of slots at congested airports to airline competition. Because market share is roughly proportional to flight frequency, airlines sometimes have the incentive to schedule many flights to the same destination using smaller aeroplanes. This, in turn, leads to an overuse of slots and their scarcity. Furthermore, smaller aeroplanes are both less economically efficient and more polluting than larger ones. The authors propose a game theoretical model in which Nash equilibrium closely corresponds to real-life observed data. Counterintuitively, they find that even a slight *reduction* in the number of assigned slots can cause a large decrease in delays and an increase in airlines' profits.

S. Yan and C.-R. Wang (2001) study the problem of maximising an airline's profit when determining routes and flight frequencies in a network. The authors model the problem on a graph with a commodity for each origin-destination pair with demand. They use a Lagrangian-based algorithm to solve a multi-commodity network flow problem modelled as a linear programme. Their approach allows the use of heterogeneous fleets and both non-stop and multi-stop flights. Unlike our work, they do not consider flight times or slot availability and only determine the total number of flights per day on each served leg.

A similar problem is tackled by Mashford and Marksjö (2001), who build a schedule for an airline, including destination selection, timetabling, and aircraft assignment. Similar to (S. Yan and C.-R. Wang 2001), the authors work with a network in which vertices are airports, edges represent possible flights, and each pair of vertices is associated with some demand. When building the schedule, the authors consider the possibility that passengers will transfer at an intermediate airport if given enough time. They develop a heuristic algorithm based on simulated annealing and test it on a network comprising five destinations in Australia, one in New Zealand, and one in Japan.

Chang and Lee (2010) study the problem of a new low-cost regional airline which must choose both its hub airport and a network of destinations to serve. The authors take a multi-objective approach involving the minimisation of costs, the maximisation of revenue, and the maximisation of transported passengers. They use a heuristic method to reduce the solution space and then evaluate possible solutions using a Goal Programming approach (Charnes and Cooper 1957). Applying their method to the case of an airline based in Taiwan and flying to southeast Asia, they select 12 out of 17 potential destinations. Because Taipei was the only considered hub airport, the proposed model does not actually perform hub selection in the studied case.

The problem of selecting potential new destinations to fly to from a hub is also relevant for intercontinental travel. In this case, however, the set of potential destinations is usually small, and the number of new routes to open can be as low as one. Chang, Woon, et al. (2017) study the destination selection problem for a Taiwan-based airline wishing to open new routes to the United States. The authors consider a fixed flight frequency (one flight per week) and do not model a time component or slot constraints. They frame the problem as a multi-objective programme, minimising the operating cost and maximising the expected revenue. Unfortunately, the authors did not publish detailed results of their model, and therefore, it is not possible to determine the details of the Pareto frontier. Similarly, Deveci, Demirel, and Ahmetoğlu (2017) study the problem of a Turkish airline which must open a new route to North America among five possible ones. They use a fuzzy multi-criterion decision model and select the best solution using the TOPSIS method (Technique for Order of Preference by Similarity to Ideal Solution, see Hwang and Yoon 1981). Özcan (2018) also seems to use a multi-criterion approach for a route destination selection problem, based on their abstract. However, the rest of the article is in Turkish, and therefore, we cannot provide any further detail on the methodology.

For a qualitative approach to the problem of selecting potential destinations for low-cost carriers, we highlight the work of Chang, Hsu, et al. (2008). The authors use the Delphi method (Dalkey and Helmer 1963) to find potential airports based on the competitiveness of the destination, the airport infrastructure and potential demand. Badi et al. (2023), instead, focus on the problem of choosing a new hub airport rather than new destinations from a given hub.

Among problems which involve only a subset of the decisions considered in the RFSP, we mention aircraft assignment (see, e.g., Hane et al. 1995; Pilla et al. 2008; Fuentes et al. 2021), sometimes tackled in conjunction with timetabling (as in Wei, Vaze, and Jacquillat 2020), or with both frequency determination and timetabling (as in Cadarso and Marín 2011; Cadarso and Marín 2013), or with partial destination selection (in which some destinations are fixed and other are optional, as in C. Yan, Barnhart, and Vaze 2022).

2 Mathematical models and algorithms

Let $\mathcal{D}$ be the set of destinations, $\mathcal{A}$ the set of aircraft and $\mathcal{T}$ the discretised time horizon. For convenience, we use the notation $N = |\mathcal{A}|$ for the number of available aircraft and $T = \max \mathcal{T}$ for the last time instant of the time horizon. Let $m_i \in \mathbb{N}^+$ be the maximum number of flights bound for destination $i \in \mathcal{D}$, which depends on the demand for the given destination. Also, let $t_i \in \mathbb{N}^+$ be the number of time instants after which an aircraft can fly again after starting a return trip to i. Hence, t_i includes the flight time and the turn-around time at both the destination and the hub. We denote with $\sigma \in \mathbb{N}^+$ the minimum number of time instants between two departures for the same destination. Let $s_t \in \mathbb{N}$ be the number of available slots at the hub airport at time $t \in \mathcal{T}$. Scheduling a flight with destination $i \in \mathcal{D}$ at time $t \in \mathcal{T}$ has a utility (or yields a profit) of $p_{it} \in \mathbb{N}^+$ units. Finally, let $\alpha \in [0, 1]$ be the minimum fraction of the time horizon during which an aircraft must be utilised (i.e., flying or performing a turn-around).

In this section, we propose two formulations which can be used to solve the RFSP. The first one, the *three-index formulation*, uses variables indexed on $\mathcal{D}$, $\mathcal{A}$ and $\mathcal{T}$. The second one, the *two-index formulation*, uses variables indexed on $\mathcal{D}$ and $\mathcal{T}$ but only solves a relaxed version of the RFSP. The two-index formulation does not explicitly assign flights to aircraft and, therefore, cannot ensure that the minimum utilisation time requirement is met. This formulation is then paired with an assignment subproblem, determining whether feasible assignments of flights to aircraft exist. In case such an assignment does not exist, we add no-good cuts to the two-index formulation and solve again until the subproblem finds a valid assignment.

Before introducing the formulations, we settle the question of what the computational complexity of the RFSP is.

Theorem 1. *The decision version of the RFSP, which asks if a solution exists with utility at least $U \geq 0$, is $\mathcal{NP}$-complete.*

Proof. Consider an instance of the $\mathcal{NP}$-complete 0–1 Knapsack Problem. Its decision version asks whether it is possible to pack a subset of n items, each with weight $w_i \in \mathbb{N}^+$ and yielding profit $\pi_i \in \mathbb{N}^+$ ($i \in \{1, \dots, n\}$), in a knapsack of capacity $c \in \mathbb{N}^+$ such that the sum of the weights of the packed items does not exceed c and their profit is at least U. We show that it is possible to transform such an instance into an instance of the RFSP and that the transformation takes polynomial time. This shows that the RFSP is $\mathcal{NP}$-complete.

We build the RFSP instance as follows. Let $|\mathcal{D}| = n$, $|\mathcal{A}| = 1$ and $|\mathcal{T}| = c$. For each destination $i \in \mathcal{D}$, we let $m_i = 1$, $t_i = w_i$ and $p_{it} = \pi_i \ \forall t \in \mathcal{T}$. We also let $\sigma = 1$, $s_t = 1 \ \forall t \in \mathcal{T}$ and $\alpha = 0$. We note that such a transformation takes time proportional to $n \cdot c$.

Finding a feasible solution to the RFSP with a utility of at least U corresponds to finding an analogous solution to the 0–1 Knapsack Problem with a total profit of at least U:

(i) The only aircraft of $\mathcal{A}$ acts as the knapsack.

(ii) Because all selected flights must be completed within the end of the time horizon, the capacity of the knapsack equals the length of the time horizon ($c = |\mathcal{T}|$).

(iii) Including a flight to destination $i \in \mathcal{D}$ in the schedule corresponds to packing the i-th item; indeed, irrespective of the departure time, the flight yields utility π_i.

(iv) In the 0–1 Knapsack Problem, each item can be packed only once. This condition is ensured in the RFSP by setting $m_i = 1$ for all destinations $i \in \mathcal{D}$.

(v) Because we only have one aircraft, setting $s_t = 1$ for all $t \in \mathcal{T}$ is not binding. Analogously, because $m_i = 1$ for all destinations $i \in \mathcal{D}$, setting $\sigma = 1$ is not binding.

(vi) Finally, because in the 0–1 Knapsack Problem there is no requirement that at least some capacity is used, we set $\alpha = 0$.

□

2.1 Three-index formulation

Let variable $x_{ijt} \in \{0, 1\}$ take value 1 iff aircraft $j \in \mathcal{A}$ leaves at time $t \in \mathcal{T}$ for a return trip to destination $i \in \mathcal{D}$. Variables x_{ijt} such that either (a) there are no available slots at time t, i.e., $s_t = 0$, or (b) there is not enough time to complete a round trip, i.e., $t + t_i > T$, can be fixed to take value 0.

The three-index formulation reads as follows:

$$\max \quad \sum_{i\in\mathcal{D}}\sum_{j\in\mathcal{A}}\sum_{t\in\mathcal{T}} p_{it}x_{ijt} \tag{1}$$

$$\text{s.t.} \quad \sum_{i\in\mathcal{D}}\sum_{j\in\mathcal{A}} x_{ijt} \le s_t \qquad \forall t\in\mathcal{T} \tag{2}$$

$$\sum_{i\in\mathcal{I}}\sum_{\substack{t'\in\mathcal{T}\\(t-t_i)^+<t'\le t}} x_{ijt'} \le 1 \qquad \forall j\in\mathcal{A},\ \forall t\in\mathcal{T} \tag{3}$$

$$\sum_{j\in\mathcal{A}}\sum_{\substack{t'\in\mathcal{T}\\t\le t'<t+\sigma}} x_{ijt'} \le 1 \qquad \forall i\in\mathcal{D},\ \forall t\in\mathcal{T} \tag{4}$$

$$\sum_{i\in\mathcal{D}}\sum_{t\in\mathcal{T}} t_i x_{ijt} \ge \alpha T \qquad \forall j\in\mathcal{A} \tag{5}$$

$$\sum_{j\in\mathcal{A}}\sum_{t\in\mathcal{T}} x_{ijt} \le m_i \qquad \forall i\in\mathcal{D} \tag{6}$$

$$x_{ijt}\in\{0,1\} \qquad \forall i\in\mathcal{D},\ \forall j\in\mathcal{A},\ \forall t\in\mathcal{T}. \tag{7}$$

The objective function (1) maximises the expected utility of the assignment. Constraint (2) ensures that slot limits at the hub airport are respected. Constraint (3), in which notation $(t-t_i)^+$ is a shorthand for $\max\{t-t_i, 0\}$, makes sure that flights assigned to the same aircraft do not overlap. Constraint (4) enforces the minimum spacing between departures to the same destination, and constraint (5) forces each aircraft to be operational for at least a fraction α of the time horizon. Finally, constraint (6) ensures that the schedules do not exceed the maximum number of flights to each destination.

2.2 Two-index formulation

Let variable $y_{it}\in\{0,1\}$ take value 1 iff there is a flight scheduled at time $t\in\mathcal{T}$ with destination $i\in\mathcal{D}$. Because $\sigma>0$, there cannot be more than one flight with the same destination leaving at the same time instant, and thus, it is sufficient to consider binary y's. Some variables y_{it} can be fixed to value 0 for reasons analogous to those used to fix variables x_{ijt}.

The two-index formulation is:

$$\max \quad \sum_{i\in\mathcal{D}}\sum_{t\in\mathcal{T}} p_{it}y_{it} \tag{8}$$

$$\text{s.t.} \quad \sum_{i\in\mathcal{D}} y_{it} \le s_t \qquad \forall t\in\mathcal{T} \tag{9}$$

$$\sum_{i\in\mathcal{D}}\sum_{\substack{t'\in\mathcal{T}\\ [t-t_i]<t'\le t}} y_{it'} \le N \qquad \forall t\in\mathcal{T} \tag{10}$$

$$\sum_{\substack{t'\in\mathcal{T}\\t\le t'<t+\sigma}} y_{it'} \le 1 \qquad \forall i\in\mathcal{D},\ \forall t\in\mathcal{T} \tag{11}$$

$$\sum_{i\in\mathcal{D}}\sum_{t\in\mathcal{T}} t_i y_{it} \ge \alpha n T \tag{12}$$

$$\sum_{t\in\mathcal{T}} y_{it} \le m_i \qquad \forall i\in\mathcal{D} \tag{13}$$

$$y_{it}\in\{0,1\} \qquad \forall i\in\mathcal{D},\ \forall t\in\mathcal{T}. \tag{14}$$

The objective function (8) is analogous to (1) and maximises the utility of the assignment. Constraints (9), (11) and (13) are completely analogous to, respectively, constraints (2), (4) and (6) and achieve the same goals. Constraint (10) ensures that no more than N flights are flown at the same time and therefore acts as a no-overlap constraint analogous to (3). The two-index formulation does not explicitly assign flights to aircraft and, thus, does not allow checking that the minimum operation time constraint—enforced by (5)—is satisfied. However, we still add constraint (12), which checks that, on average, each aircraft is used for at least αT units of time, although it does not guarantee that each aircraft meets this utilisation threshold.

Formulation (8)–(14) is a relaxation of the RFSP: there are optimal solutions to this formulation which do not allow any assignment of flights to aircraft such that the minimum operation time constraint is satisfied. To see why this is the case, consider a simple instance with a time horizon of $T = 100$ units, two destinations ($\mathcal{D} = \{1, 2\}$) with $t_1 = 25, t_2 = 100$ and $m_1 = 2, m_2 = 1$, two aircraft ($N = 2$) and $\alpha = 0.75$. Let $p_{it} = 1$ for all indices (i, t), $\sigma = 1$ and $s_t = 2$ for all t. All optimal solutions to the two-index formulation must have $y_{21} = 1$ and $y_{1t} = 1$ for any two values of t, say s_1 and s_2, such that $|s_2 - s_1| \geq 26$ and $\max\{s_1, s_2\} \leq 75$. For example, $y_{1\,1} = y_{1\,26} = 1$ (and $y_{it} = 0$ for all other indices) would provide an optimal solution of utility 3. Such a solution would satisfy inequality (12) because $100 + 25 + 25 = 150 \geq 150 = 0.75 \cdot 2 \cdot 100$. However, because the flight to destination 2 must be assigned to one aircraft, and the flights to destination 1 to the other aircraft, the utilisation rate of the latter aircraft would be 50 time units out of 100. Therefore, inequality (5) would be violated because $25 + 25 = 50 < 75 = 0.75 \cdot 100$.

We propose to use formulation (8)–(14) within a decomposition scheme in which it acts as the master problem. The subproblem, described below, attempts to produce a feasible flight-to-aircraft assignment. If the subproblem succeeds, then the solution to the master problem is optimal for the RFSP, and a concrete schedule for the aircraft can be built using the subproblem solution. Otherwise, if the subproblem proves that no feasible assignment exists, we add a no-good cut, which cuts off the current master problem solution, and we solve the master problem again. The algorithm ends as soon as the subproblem finds a feasible assignment, i.e., an assignment that satisfies all the original RFSP constraints.

2.2.1 Assignment sub-problem

Given a solution y^*_{it} to model (8)–(14), we denote with $\mathcal{I}$ the set of the indices of variables whose value is 1:

$$\mathcal{I} = \left\{(i, t) \in \mathcal{D} \times \mathcal{T} \ : \ y^*_{it} = 1\right\}.$$

To check that solution y^*_{it} allows a feasible assignment of flights to aircraft, one can solve a feasibility problem in which variables $x_{ijt} \in \{0, 1\}$ have the same meaning as in model (1)–(7). In the following problem, we only include variables x_{ijt} if $(i, t) \in \mathcal{I}$.

$$
\begin{aligned}
\max \quad & 0 && && (15)\\
\text{s.t.} \quad & \sum_{j \in \mathcal{A}} x_{ijt} = 1 && \forall (i, t) \in \mathcal{I} && (16)\\
& \sum_{\substack{i \in \mathcal{I} \\ (i,t) \in \mathcal{I}}} \ \sum_{\substack{t' \in \mathcal{T} \\ (i,t') \in \mathcal{I} \\ (t - t_i)^+ < t' \leq t}} x_{ijt'} \leq 1 && \forall j \in \mathcal{A},\ \forall t \in \mathcal{T} && (17)\\
& \sum_{(i,t) \in \mathcal{I}} t_i x_{ijt} \geq \alpha T && \forall j \in \mathcal{A} && (18)\\
& x_{ijt} \in \{0, 1\} && \forall (i, t) \in \mathcal{I},\ \forall j \in \mathcal{A}. && (19)
\end{aligned}
$$

Constraint (16) ensures that, for each destination-time pair identified by the master problem, exactly one aircraft operates the corresponding flight. Constraints (17) and (18) are no-overlap and minimum operation time constraints.

If subproblem (16)–(19) is feasible, solution y^*_{it} is valid for the original RFSP, and the schedule of each aircraft can be built using the value of variables x_{ijt} in any sub-problem solution (all subproblem solutions give overall solutions of the same utility). Otherwise, we add a no-good cut to model (8)–(14) to cut off the infeasible solution. Next, we detail the cut generation procedure.

2.2.2 Cut generation

The most straightforward way to cut off an RFSP-infeasible solution y^*_{it} in the master problem is by adding the following no-good cut:

$$\sum_{(i,t)\in\mathcal{I}} (1 - y_{it}) + \sum_{(i,t)\notin\mathcal{I}} y_{it} \geq 1, \tag{20}$$

which forces at least one variable y_{it} to change value. This cut is quite weak, excluding only the particular combination of destinations and departure times identified by variables y^*_{it}. At the next iteration, the master problem could return a different schedule in which only one flight is modified (e.g., by postponing its take-off time by one time unit), and such a solution would most likely be again infeasible for the RFSP and have the same or a very similar utility. In our algorithm, then, we first attempt to generate a stronger cut and, only if we do not succeed, we add cut (20).

To generate a stronger cut, we check whether there exists any feasible RFSP solution which includes flights to the same airports (with their given multiplicity, i.e., how many flights there are for each destination) as the solution from the master problem, but possibly at different times. For example, given a solution with $y^*_{1,12} = y^*_{1,80} = y^*_{2,24} = y^*_{3,98} = y^*_{7,130} = 1$ and all other variables equal to 0, we ask whether there is a feasible RFSP solution which visits destination 1 twice and destinations $2, 3, 7$ once, without imposing any limitation on the flight times. If such a solution exists, then (a) this solution is feasible for the RFSP and therefore provides a primal bound for the original problem (in this sense, we are performing a local search attempting to *repair* the infeasible master problem solution), and (b) we can only conclude that the particular choice of destination-time pairs in y^*_{it} is unfeasible for the RFSP, and thus we must add no-good cut (20). If, on the other hand, we can prove that no such solution exists, we reach a much stronger conclusion: no master problem solution visiting the given subset of destinations will be feasible for the RFSP, no matter what the flight departure times are.

We need, then, two ingredients. First, a mechanism to determine whether there is an RFSP-feasible solution visiting the same airports (but at possibly different times) as the master problem solution. Next, in case there is no such solution, a mechanism to generate a stronger cut which excludes the given destination choice.

Checking whether an RFSP-feasible solution exists. To this end, we can use a reduced three-index formulation. The difference between this reduced formulation and the original (1)–(7) is that (a) we fix to 0 variables x_{ijt} corresponding to indices $i \in \mathcal{D}$ for which no $y^*_{it'} = 1$ (for any value of t'); (b) we change constraint (6) into an equality and replace its right-hand side with $M_i := |\{t \in \mathcal{T} \; : \; y^*_{it} = 1\}|$, i.e., by the number of flights with destination i selected in the master problem solution; (c) we set the MIP solver focus to finding primal solutions quickly or proving that none exist.

Adding a stronger cut. If there is no RFSP-feasible solution visiting the same airports as the master problem solution, we want to add a cut which excludes the destinations (with their multiplicity) chosen by the master problem. To this end, we add to the master problem two-index formulation (8)–(14) the following binary variables: $z_{i\ell} \in \{0, 1\}$ for each $i \in \mathcal{D}$ and $\ell \in \{0, \ldots, m_i\}$ (i.e., for each possible number of flights to i). Variable $z_{i\ell}$ will take value 1 iff there are exactly ℓ flights with destination i in the solution. We can set these variables by adding to the model the following indicator (Bonami et al. 2015) and linear constraints:

$$\big(z_{i\ell} = 1\big) \rightarrow \big(\sum_{t\in\mathcal{T}} y_{it} = \ell\big) \qquad \forall i \in \mathcal{D} \tag{21}$$

$$\sum_{\ell=0}^{m_i} z_{i\ell} = 1 \qquad \forall i \in \mathcal{D}. \tag{22}$$

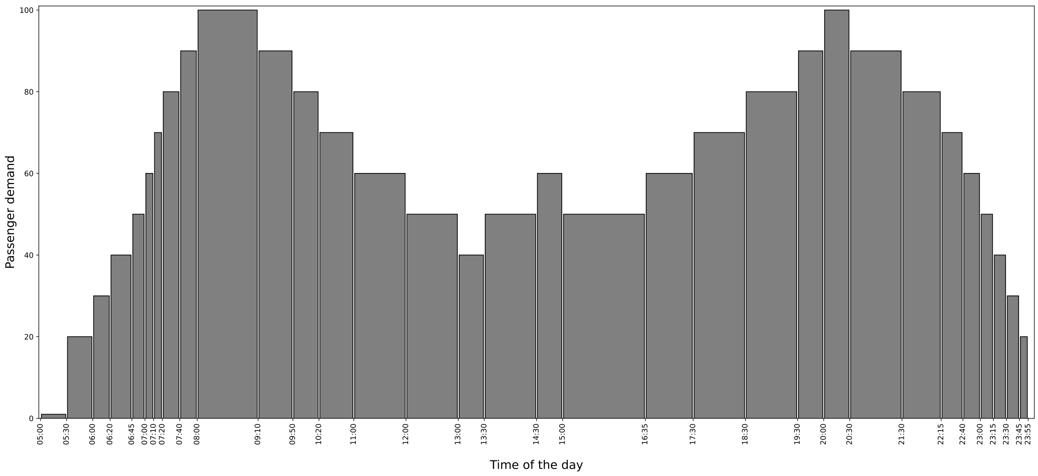

Figure 1: Example demand between two destinations, from which the utilities have been estimated.

Then, the stronger cut is:

$$\sum_{i \in \mathcal{D}} \Big(1 - z_{iM_i} + \sum_{\substack{\ell=0 \\ \ell \neq M_i}}^{m_i} z_{i\ell}\Big) \geq 1, \tag{23}$$

which forces changing the number of flights to at least one airport.

3 Computational results

In this section, we present insights from computational experiments. We used a set of fifteen noisy instances obtained from realistic data relative to a small regional airline. The utilities are generated based on demand patterns, such as the one shown in Figure 1. The figure shows the passenger demand on a short-range route during a weekday. Instance characteristics are summarised as follows: the number of destinations is either 10, 15, or 20; the fleet size is either 3, 5, or 10; the minimum utilisation threshold is either 0.75, 0.85, or 0.95; the time horizon consists of 228 time intervals (19 hours with a resolution of 5 minutes); the minimum padding time between two flights to the same destination is 1 hour. The instances are available on GitHub (Santini 2023).

Table 1 reports the performance of the two solution approaches on the fifteen instances. All computational tests used a time limit of 1800 seconds on a machine equipped with a 4-core Intel i7 processor running at 2.8GHz and 8GB of memory. The solver used was `Gurobi` version 9.5.2.

Columns "Gap%" report the percentage optimality gap computed as

$$\text{Gap\%} = \frac{\text{UB} - \text{LB}}{\text{UB}},$$

where UB is the best dual bound and LB is the best primal bound found within the time limit. Columns "Time (s)" are the runtimes in seconds. A horizontal line in the gap columns denotes instances for which the solver does not produce any primal feasible solution before timing out.

The decomposition approach applied to the two-index formulation solves to optimality 12 out of 15 instances. The three open instances all have gaps under 2%. The three-index formulation solves seven instances to optimality, produces primal solutions for 5, and cannot find any primal feasible solution for the remaining 3. There are four instances (2, 6, 8, 9) which are solved to optimality by the two-index formulation, while the 3-index formulation provides a primal bound but does not close the gap. Inspecting the model results, we noticed that in three cases (instances 2, 6, 9), the three-index formulation found the optimal solution, but it could not prove its optimality. In the remaining case

				3-index formulation		2-index formulation	
Instance	$\lvert\mathcal{D}\rvert$	$\lvert\mathcal{A}\rvert$	α	Gap%	Time (s)	Gap%	Time (s)
1	10	3	0.75	0.00	1.10	0.00	0.39
2	10	3	0.85	2.76	1800.00	0.00	5.21
3	10	3	0.95	0.00	5.21	0.00	0.73
4	10	5	0.75	0.00	10.83	0.00	1.41
5	10	5	0.85	—	1800.00	0.00	684.17
6	10	5	0.95	8.13	1800.00	0.00	1688.39
7	15	3	0.75	0.00	1.99	0.00	0.60
8	15	3	0.85	2.88	1800.00	0.00	22.42
9	15	3	0.95	0.15	1800.00	0.00	2.58
10	15	5	0.75	0.00	3.50	0.00	0.59
11	15	5	0.85	—	1800.00	0.51	1800.00
12	20	5	0.75	0.00	9.41	0.00	1.46
13	20	5	0.85	5.39	1800.00	1.47	1800.00
14	20	5	0.95	—	1800.00	1.98	1800.00
15	20	10	0.75	0.00	48.69	0.00	4.13

Table 1: Computational results on the 15 instances. Column "Gap%" is the optimality gap, and column "Time (s)" is the solution time in seconds. A horizontal line for the gap means that no primal feasible solution was found within the time limit.

(instance 8), the best primal solution returned by the 3-index formulation was sub-optimal. Regarding solution times, even when both approaches find a provably optimal solution, the 2-index formulation is quicker.

In conclusion, the decomposition approach with the two-index formulation dominates the three-index formulation. For each instance, the two-index model either gives lower gaps or, if both models solve the instance to optimality, the two-index approach uses less time.

Table 2 gives a summary of some important characteristics of the instances and the solutions of the RFSP. We use five indicators,—available flights per aircraft (AV), assigned flights per aircraft (AS), percentage of flights assigned (FA%), aircraft utilisation (U%), and flights on the best slot (FB%)—as explained below. For clarity, we represent a solution in terms of the three-index x variables: we use the indicator $\hat{x}_{ijt} \in \{0, 1\}$, taking value 1 iff the solution assigns a flight to $i \in \mathcal{D}$ at time $t \in \mathcal{T}$ to aircraft $j \in \mathcal{A}$.

- AV is the number of flights for which there is demand, divided by the fleet size. This is a property of the instance, and it does not depend on the solution. Formally,

$$\text{AV} = \frac{1}{N} \sum_{i \in \mathcal{D}} m_i.$$

- AS is the number of scheduled flights divided by the fleet size, i.e., the average number of flights operated by an aircraft. Formally,

$$\text{AS} = \frac{1}{N} \sum_{i \in \mathcal{D}} \sum_{j \in \mathcal{A}} \sum_{t \in \mathcal{T}} \hat{x}_{ijt}.$$

- FA% is the percentage of flights included in some aircraft's schedule over all possible flights. Formally,

$$\text{FA\%} = 100 \cdot \frac{\sum_{i \in \mathcal{D}} \sum_{j \in \mathcal{A}} \sum_{t \in \mathcal{T}} \hat{x}_{ijt}}{\sum_{i \in \mathcal{D}} m_i} = 100 \cdot \frac{\text{AS}}{\text{AV}}.$$

Instance	$\|\mathcal{D}\|$	$\|\mathcal{A}\|$	α	AV	AS	FA%	U%	FB%
1	10	3	0.75	4.67	3.33	71.31	82.31	20.00
2	10	3	0.85	4.33	3.33	76.91	89.33	0.00
3	10	3	0.95	5.67	4.33	76.37	98.68	23.08
4	10	5	0.75	3.20	3.20	100.00	83.51	31.25
5	10	5	0.85	3.60	3.60	100.00	91.23	11.11
6	10	5	0.95	3.60	3.60	100.00	95.70	11.11
7	15	3	0.75	6.67	3.00	44.98	83.33	33.33
8	15	3	0.85	6.67	4.00	59.97	94.74	8.33
9	15	3	0.95	7.00	4.00	57.14	96.78	33.33
10	15	5	0.75	4.80	3.60	75.00	88.95	27.78
11	15	5	0.85	3.80	3.00	78.95	85.61	6.67
12	20	5	0.75	5.00	3.20	88.57	81.32	43.75
13	20	5	0.85	5.20	3.60	64.00	92.89	16.67
14	20	5	0.95	5.20	3.40	69.23	96.84	17.65
15	20	10	0.75	3.50	3.10	65.38	83.51	35.48

Table 2: Summary of solution characteristics for the best primal solutions returned by the two-index approach.

- U% is the percentage of time that an aircraft was utilised, averaged over all aircraft. Formally,

$$\text{U\%} = 100 \cdot \frac{1}{N} \cdot \sum_{j \in \mathcal{A}} \frac{\sum_{i \in \mathcal{D}} \sum_{t \in \mathcal{T}} t_i \hat{x}_{ijt}}{T}.$$

- FB% is the percentage of operated flights which depart on the slot with the best possible utility (for the given destination) over all operated flights. For example, consider a destination $i \in \mathcal{D}$, let $\mathcal{T} = \{1, \dots, 5\}$ and assume the utilities of flying to i are $2, 1, 3, 2, 1$ for the five time intervals. If a solution includes three flights to i at times $1, 3, 5$, then only two out of three are in slots with the best possible utilities: the flight at time 3 (utility of 3) and the flight at time 1 (utility of 2). The next-best slot would be at time interval 4 (utility of 2), but it was not used; instead, the third flight departed at time interval 5.

  To formally define this metric, let $\hat{T}_i$ be the set of departure times of flights to i in solution $\hat{x}$,

$$\hat{T}_i = \big\{t \in \mathcal{T} \text{ s.t. } \exists j \in \mathcal{A} \ : \ \hat{x}_{ijt} = 1\big\}.$$

  Next, let $(\pi^i_1, \dots, \pi^i_{\kappa_i})$ be a sequence of time intervals such that: the corresponding profits are in decreasing order, i.e., $p_{i\pi^i_1} \geq p_{i\pi^i_2} \geq \dots \geq p_{i\pi^i_{\kappa_i}}$; time instants of $\hat{T}_i$ are removed from the sequence. In the previous example, we would have $\hat{T}_i = \{1, 3, 5\}$ and the sequence would only have the two time instants $\pi^i_1 = 4$ (utility 2), and $\pi^i_2 = 2$ (utility 1) during which no flight to i departs. The number of flights departing on the slots with the best possible utility is simply the number of elements of $\hat{T}_i$ with corresponding utility not less than $p_{i\pi^i_1}$ (the highest utility among unused slots). Therefore, formally,

$$\text{FB\%} = 100 \cdot \frac{\sum_{i \in \mathcal{D}} \left| \big\{t \in \hat{T}_i \ : \ p_{it} \geq p_{i\pi^i_1}\big\} \right|}{\sum_{i \in \mathcal{D}} \sum_{j \in \mathcal{A}} \sum_{t \in \mathcal{T}} \hat{x}_{ijt}}.$$

Comparing columns AV and AS in Table 2 highlights a strong destination selection component in the majority of RFSP instances, as the operated flights per aircraft are fewer than the available ones. This is further confirmed by metric FA% (which, indeed, can be written as the ratio of AS over AV). In only three instances, it is feasible to operate all flights for which there is demand, and the percentage of operated flights can be even less than half (instance 7).

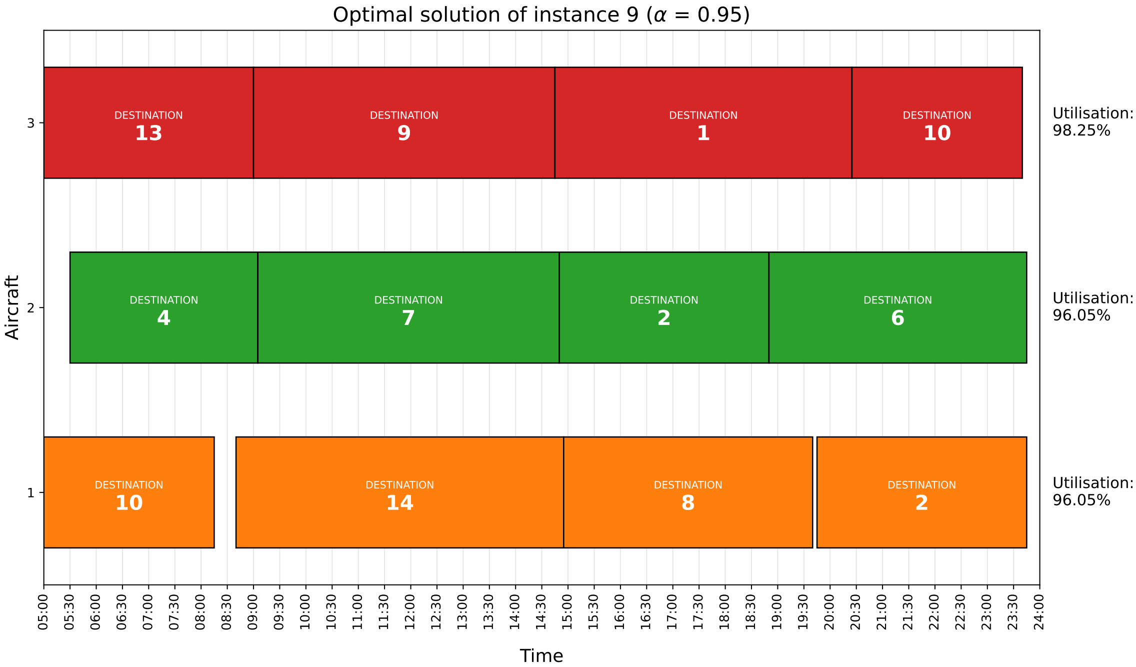

Figure 2: Optimal solution of instance 9, with the minimum utilisation constraints ($\alpha = 0.95$).

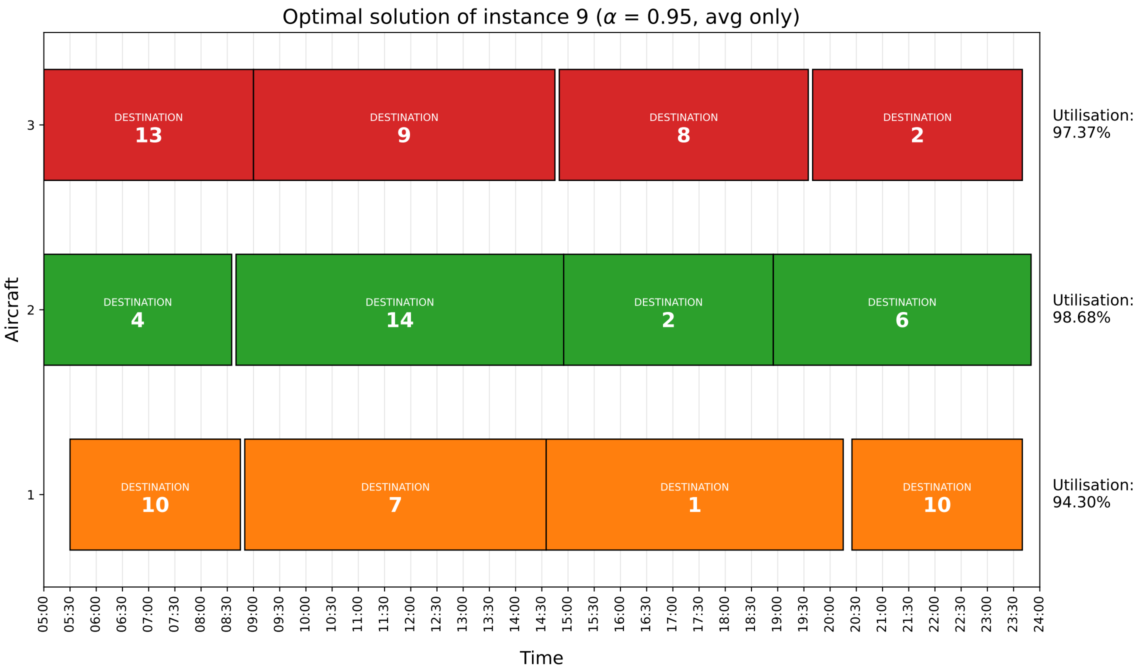

Figure 3: Optimal solution of instance 9 when we do not require that each aircraft complies with the minimum utilisation constraint, but only that the average utilisation across the entire fleet is at least α.

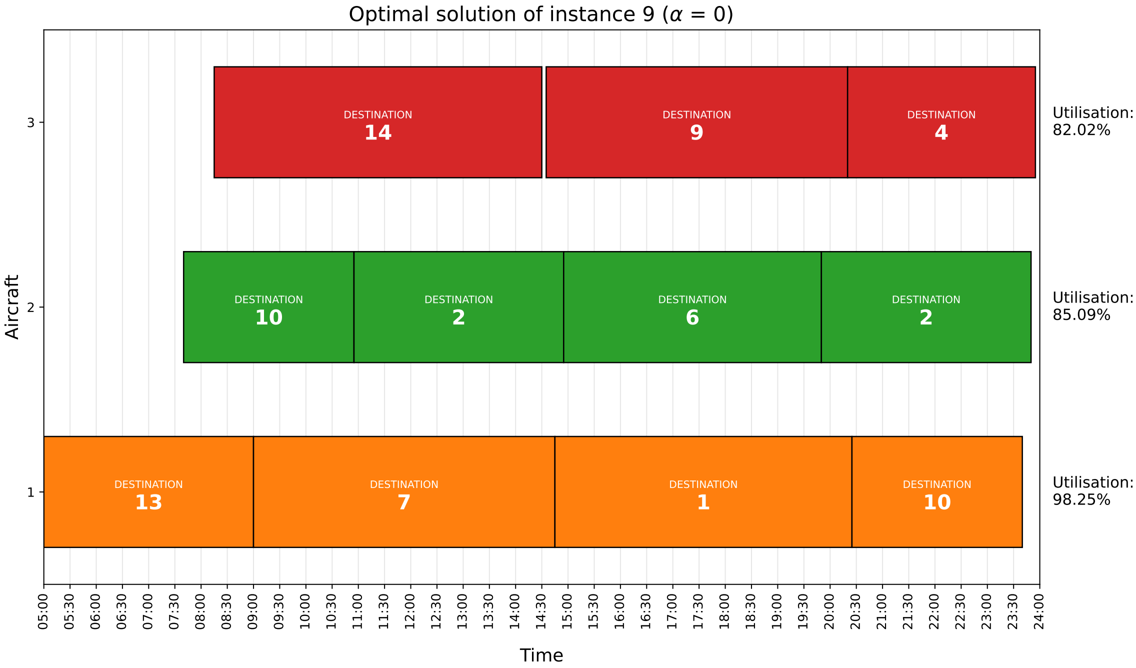

Figure 4: Optimal solution of a modified version of instance 9 in which we set $\alpha = 0.0$.

Aircraft utilisation is high, and it is always above 80%, including in instances with $\alpha = 0.75$. The minimum utilisation constraint is the "complicating constraint" of the RFSP, which makes solutions of the two-index formulation (8)–(14) potentially infeasible. Therefore, in light of the results presented in Table 2, it makes sense to question whether this constraint is necessary or if it can be removed while still obtaining solutions which comply with the utilisation constraint.

To this end, we ran the two-index formulation on the same 15 instances but tried two further approaches. In the first approach, we no longer enforce that each aircraft complies with the minimum utilisation constraint; rather, we only require that the average utilisation across all aircraft is at least α. In the second approach, we set $\alpha = 0$, completely removing any minimum utilisation constraint.

To provide a visual example, compare Figures 2 to 4. Figure 2 displays the optimal solution of instance 9, in which the minimum utilisation parameter is $\alpha = 0.95$, the average aircraft utilisation is 96.78%, and the achieved utility is 957.20. Figure 3 depicts the optimal solution of the same instance when only requiring that the average utilisation is at least 0.95. In the solution, this average is 96.78%, but aircraft 1 does not comply with the original constraint because its utilisation is 94.30%. The utility of this solution is 958.6. Finally, Figure 4 shows the optimal solution of instance 9, but with the parameter α set to 0. In this case, the average aircraft utilisation is 88.45%, and the achieved utility is 965.60. Indeed, the solution depicted in Figure 4 uses aircraft 2 and 3 much less than required in the original problem. However, it schedules the first flight of each of these aircraft between 07:30 and 08:30, which are periods corresponding with high utility. The solutions presented in Figures 2 and 3, on the other hand, must schedule the first flights of these aircraft earlier in order to ensure a higher utilisation.

Table 3 shows a complete comparison of the three approaches discussed above over all instances. The AC columns indicate how many aircraft comply with the per-aircraft minimum utilisation constraint. The U columns indicate the percentage increase in utility achieved by relaxing the constraint. Subscript 1 refers to solutions of the RFSP (therefore, AC_1 is always equal to $|\mathcal{A}|$). Subscript 2 refers to the model in which the minimum utilisation constraint is only enforced on average. Subscript 3 refers to the model in which no minimum utilisation constraint is enforced (equivalent to setting $\alpha = 0$).

In 10 out of 15 instances AC_2 is strictly less than AC_1 and in 11 instances AC_3 is strictly less than

Instance	$\|\mathcal{D}\|$	$\|\mathcal{A}\|$	α	AC_1	AC_2	AC_3	U_2	U_3
1	10	3	0.75	3	3	2	+0.00%	+0.00%
2	10	3	0.85	3	1	0	+0.05%	+0.06%
3	10	3	0.95	3	2	0	+0.00%	+0.03%
4	10	5	0.75	5	5	5	+0.00%	+0.00%
5	10	5	0.85	5	2	0	+0.04%	+0.05%
6	10	5	0.95	5	4	0	+0.08%	+0.17%
7	15	3	0.75	3	3	3	+0.00%	+0.00%
8	15	3	0.85	3	2	1	+0.02%	+0.03%
9	15	3	0.95	3	2	1	+0.00%	+0.01%
10	15	5	0.75	5	5	5	+0.00%	+0.00%
11	15	5	0.85	5	3	0	+0.03%	+0.05%
12	20	5	0.75	5	4	4	+0.00%	+0.00%
13	20	5	0.85	5	2	2	+0.05%	+0.16%
14	20	5	0.95	5	4	0	+0.03%	+0.05%
15	20	10	0.75	10	10	10	+0.00%	+0.00%

Table 3: Comparison of solutions with the original and the relaxed versions of the minimum utilisation constraint. The AC columns indicate how many aircraft comply with the per-aircraft constraint. The U columns indicate the percentage increase in utility achieved by relaxing the constraint. Subscript 1 refers to solutions of the RFSP. Subscript 2 refers to the model in which the minimum utilisation constraint is only enforced on average. Subscript 3 refers to the model in which no minimum utilisation constraint is enforced (equivalent to setting $\alpha = 0$).

AC_1. We remark that when AC_2 is equal to AC_1, the cut generation algorithm terminates at the first iteration because the assignment subproblem finds a feasible assignment. Therefore, an instance which would be amenable to be solved with a simplified constraint (the one which only enforces average minimum utilisation) because such a solution would happen to respect the per-aircraft utilisation constraint is precisely an instance which is solved quickly using the original constraint and the cut separation algorithm. This fact negates any potential computational advantage of simplifying the minimum utilisation constraints.

At the same time, columns U_2 and U_3 show that the increase in utility achieved by relaxing the constraints is very modest. This means that if there are strategic reasons to maintain aircraft utilisation high, the corresponding constraints should be taken into account. The maximum utility increase obtained by setting $\alpha = 0$, in fact, is only 0.17%.

In summary, the results presented in Table 3 justify the use of the minimum utilisation constraints and, therefore, the dynamic cut generation approach presented in Section 2.2.

Finally, we remark that the percentage of flights assigned to the best possible slot (column FB% in Table 2) varies considerably from instance to instance, spanning from 0% up to 43.75%. Considering that slots with high utility are usually close to each other (see Figure 1), it is not surprising that flights to the same destination cannot all be assigned to high-utility slots, as this would violate the minimum spacing constraints (4) and (11).

4 Extensions

In this section, we describe how our approach can be extended to take into account more constraints and relax some of the assumptions laid out in Section 1.

4.1 Heterogeneous fleet

In the case of a heterogeneous fleet, we would have to add index j to parameters p and t, i.e., they would become p_{ijt} (utility of flying to i at time t with aircraft j) and t_{ij} (round-trip time to i when using aircraft j).

4.2 Avoid or penalise destination selection

Next, the model allows the planner to exclude some destinations from the schedule. This might be acceptable for a start-up regional airline that is planning its first network. However, established airlines wanting to restructure their existing operations might lose market or a valuable slot at a constrained airport if they stop service to an existing destination. To this end, we propose two model extensions.

The first ensures that each destination $i \in \mathcal{D}$ is served with at least h_i flights. Eventually, h_i can be zero for destinations with no such restriction. Taking the two-index formulation as the reference model, we can then add the following constraint to the master problem.

$$\sum_{t \in \mathcal{T}} y_{it} \geq h_i \quad \forall i \in \mathcal{D}.$$

The second extension, which can be used in combination with the first one, adds a penalty for not serving destinations. This penalty can correspond, for example, to the value of a lost slot. Let δ_i be the penalty for not serving destination $i \in \mathcal{D}$. We introduce variables $w_i \in \{0, 1\}$ for $i \in \mathcal{D}$, taking value 1 iff destination i is served. Then, we can modify objective function (8) as follows:

$$\sum_{i \in \mathcal{D}} \left(\sum_{t \in \mathcal{T}} p_{it} y_{it} - \delta_i (1 - w_i) \right),$$

and constraint (13) as follows:

$$\sum_{t \in \mathcal{T}} y_{it} \leq m_i w_i \quad \forall i \in \mathcal{D}.$$

4.3 Schedule impact on the utility

As mentioned in Section 1, we assume that flight utilities are deterministic and given in advance. As a consequence, in our formulations, scheduling a flight does not change the utilities of other flights. This is an important limitation of our models because, e.g., it is not realistic to consistently obtain high utilities if we schedule so many flies that a destination becomes oversupplied. We propose three ways to overcome this limitation. The third is more realistic than the first two but also requires more market knowledge, which could be a challenge for a start-up regional airline with little access to passenger data. In the following, we again take the two-index formulation as the reference to amend our proposed models.

In the first case, we assume that the utility of a flight to destination i only depends on the number of flights scheduled to i. I.e., the utility can be expressed as $p_{i\ell}$, where ℓ is the number of daily flights to i. In this case, the utility no longer depends on the time of the day but on other decisions made by the model, namely, the number of flights scheduled to each destination. The assumption of independence on time might be justified if the planning horizon only includes times that are "reasonable" (i.e., 04:00 is not included in the planning horizon) and especially for low-cost airlines that cater to a price-sensitive but time-insensitive market. If this assumption holds, then it is possible to model the impact of scheduled flights on the utilities by changing objective function (8) as follows:

$$\sum_{i \in \mathcal{D}} p_{i\ell} z_{i\ell},$$

where we recall that $z_{i\ell}$ is a binary decision variable introduced in Section 2.2.2 and taking value 1 iff there are ℓ flights scheduled to i.

The second approach consists of assuming that the utility of a flight to destination i depends both on the departure time t and on the number of flights ℓ scheduled to i. In this case, we express the utility as $p_{it\ell}$ and we introduce variables $\eta_{it\ell} \in \{0, 1\}$, taking value 1 iff we schedule a flight to i at time t, and there are overall ℓ scheduled flights to i. We can then modify the objective function (8) as follows:

$$\sum_{i \in \mathcal{D}} \sum_{t \in \mathcal{T}} \sum_{\ell=1}^{m_i} p_{it\ell} \eta_{it\ell},$$

and add the following constraints:

$$\eta_{it\ell} \leq y_{it} \qquad \forall i \in \mathcal{D},\ \forall t \in \mathcal{T}$$
$$\eta_{it\ell} \leq z_{i\ell} \qquad \forall i \in \mathcal{D},\ \forall \ell \in \{1, \dots, m_i\}.$$

The third approach relies on estimating the number of passengers carried on each scheduled flight and using this information to compute the utilities. The procedure uses a Generalised Assignment Model (GAM) similar to the one proposed by Wei, Vaze, and Jacquillat (2020). This model assumes that the planner can estimate: (i) the market size μ_i of each destination, i.e., the number of people who want to travel from the hub to i (and vice-versa) on a given day; (ii) the attractiveness A_{it} of a flight to i departing at time t; (iii) the attractiveness B_i of the "outside option", i.e., of travelling with a different means or not at all. In our case, the attractiveness A_{it} could be approximated with the original utilities p_{it}; still, more data is needed to estimate the market size and the outside option attractiveness.

A more general GAM would model subsets of passengers travelling to i separately. For example, when the flight from the hub to i is part of a longer itinerary, the attractiveness of a flight depends on the entire itinerary because different departure times might allow or forbid some connections or increase the total travel time. Given our assumption of a small regional airline that travels point-to-point, however, it is appropriate to consider a single passenger market.

Under the GAM, the probability that a passenger who wants to travel to destination i chooses a flight taking off at time t is

$$\frac{A_{it}}{B_i + \sum_{t' \in \mathcal{T}_i^*} A_{it'}}, \tag{24}$$

where $\mathcal{T}_i^*$ is the set of departure times of flights to i. Because sets $\mathcal{T}_i^*$ depend on the planner decisions (more specifically, on the values of variables y), the GAM model captures the influence of the entire flight offer by setting the probability that a flight is used as proportional to its relative attractiveness compared to the alternatives. Using the GAM model, we can estimate the number of passengers on each flight, depending on all the flights we schedule to a destination. To transform this number into an economic utility, the planner should introduce three further parameters. The first, φ_{it}, is the unit fare earned by a passenger flying to i at time t. The second, γ_{it}, is the cost of operating a flight to i at time t. The third, Q, is the aircraft capacity.

To incorporate the GAM into our two-index formulation, we use the technique introduced by D. D. Wang, Klabjan, and Shebalov (2014). To this end, we introduce variables $\pi_{it} \geq 0$ that will hold the number of passengers flying to i at time t, and variables $\pi_{i0} \geq 0$ that will hold the number of passengers taking the outside option for destination i. The objective value (8) is then replaced by

$$\sum_{i \in \mathcal{D}} \sum_{t \in \mathcal{T}} \Big(\varphi_{it} \pi_{it} - \gamma_{it} y_{it} \Big).$$

The following constraints ensure that variables π_{it} take on the correct values:

$$\pi_{i0} + \sum_{t \in t} \pi_{it} = \mu_i \qquad \forall i \in \mathcal{D} \tag{25}$$

$$\pi_{it} \leq Q y_{it} \qquad \forall i \in \mathcal{D},\ \forall t \in \mathcal{T} \tag{26}$$

$$B_i \pi_{it} \leq A_{it} \pi_{i0} \qquad \forall i \in \mathcal{D},\ \forall t \in \mathcal{T}. \tag{27}$$

Constraint (25) ensures that the entire market is accounted for, either as a customer of our airline or as a passenger taking the outside option. Constraint (26) ensures that capacities are respected on all operated flights. Finally, constraint (27) sets the passenger values as proportional to the probabilities defined above. This constraint is the main methodological improvement proposed by D. D. Wang, Klabjan, and Shebalov (2014) because it allows modelling the GAM using linear constraints.

5 Conclusions

This paper introduces the Regional Flight Scheduling Problem, which combines destination selection, frequency selection, fleet assignment and tail assignment under a unified approach. We are able to tackle these problems jointly because we restrict ourselves to the setting of a small regional airline which operates round trips from a slot-constrained hub airport. Typical application scenarios for this problem include feeder airlines, which are the sole operators of regional routes from busy hubs under agreements with larger carriers, and airlines that fly from a capital city to smaller airports in countries with basic infrastructure. We propose a three-index and a two-index formulation; the latter uses fewer variables but only solves a relaxation of the original problem, with optimal solutions possibly violating some constraints. Therefore, we embed the two-index formulation within a cut separation algorithm, which dynamically adds violated constraints. Computational experiments show the potential of this approach and, at the same time, confirm that the violated constraints are tight and cannot be disregarded.

In the experiments mentioned above, we use 15 instances derived from perturbed realistic data, which contain up to 20 destinations and 10 aircraft. We solve to optimality 12 instances and report gaps under 2% for the remaining three. Although the analysis of larger instances is beyond the scope of this paper, the results hint that, under certain conditions, the proposed approach could scale further. For example, the largest instance (with 20 destinations and 10 aircraft) solves in less than five seconds, but, in this instance, the minimum utilisation bound is low ($\alpha = 0.75$). On the other hand, when this bound is high (e.g., $\alpha = 0.95$), there are smaller instances (e.g., with 15 destinations and 5 aircraft) that remain open. Future research could investigate alternative models, such as extended formulations, that can result in smaller gaps for larger instances.

Our models rely on a set of assumptions that are reflected in our available data. Still, in Section 4, we propose several model extensions that generalise our methodology. In particular, we address the cases of heterogeneous fleets, imposing a minimum frequency at some destination, penalising discarded destinations, making flight utility depend on the number of flights on a leg and, finally, fully modelling passenger choice under a discrete choice model. Because our dataset does not contain enough information to test these extensions, a further avenue of research is to partner with industrial actors that can provide more detailed data. This would allow us to test the business impact of the extensions, how the model performs computationally when extended, and to what extent the methodology would scale to larger instances.

Acknowledgements

The work of Alberto Santini has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska Curie grant agreement number 945380. The author is grateful to Vikrant Vaze for fruitful discussions on this topic.

References

Badi, Ibrahim, Abdulaziz Alosta, Omar Elmansouri, Ali Abdulshahed, and Salem Elsharief (2023). "An application of a novel grey-CODAS method to the selection of hub airport in North Africa". In: *Decision Making: Applications in Management and Engineering* 6.1, pp. 18–33. DOI: 10.31181/dmame0313052022i.

Barnhart, Cynthia, Peter Belobaba, and Amedeo R. Odoni (2003). "Applications of Operations Research in the Air Transport Industry". In: *Transportation Science* 37.4, pp. 368–391. DOI: 10.1287/trsc.37.4.368.23276.

Barnhart, Cynthia and Amy Cohn (2004). "Airline Schedule Planning: Accomplishments and Opportunities". In: *Manufacturing & Service Operations Management* 6.1, pp. 3–22. DOI: 10.1287/msom.1030.0018.

Barnhart, Cynthia and Barry Smith, eds. (2012). *Quantitative Problem Solving Methods in the Airline Industry: A Modeling Methodology Handbook*. Vol. 169. International Series in Operations Research & Management Science. Springer. DOI: 10.1007/978-1-4614-1608-1.

Belobaba, Peter, Amedeo Odoni, and Cynthia Barnhart (2016). *The global airline industry*. Second Edition. Aerospace series. Wiley. 512 pp. ISBN: 978-1-118-88117-0.

Bennell, Julia, Mohammad Mesgarpour, and Chris Potts (2011). "Airport runway scheduling". In: *4OR* 9, pp. 115–138. DOI: 10.1007/s10288-011-0172-x.

Bonami, Pierre, Andrea Lodi, Andrea Tramontani, and Sven Wiese (2015). "On mathematical programming with indicator constraints". In: *Mathematical Programming* 151, pp. 191–223. DOI: 10.1007/s10107-015-0891-4.

Cadarso, Luis and Ángel Marín (2011). "Integrated Robust Airline Schedule Development". In: *Procedia — Social and Behavioral Sciences* 20, pp. 1041–1050. DOI: 10.1016/j.sbspro.2011.08.113.

– (2013). "Robust passenger oriented timetable and fleet assignment integration in airline planning". In: *Journal of Air Transport Management* 26, pp. 44–49. DOI: 10.1016/j.jairtraman.2012.10.004.

Chang, Yu-Chun, Chia-Jui Hsu, George Williams, and Mei-Ling Pan (2008). "Low cost carriers' destination selection using a Delphi method". In: *Tourism Management* 29.5, pp. 898–908. DOI: 10.1016/j.tourman.2007.11.004.

Chang, Yu-Chun and Ning Lee (2010). "A Multi-Objective Goal Programming airport selection model for low-cost carriers' networks". In: *Transportation Research Part E: Logistics and Transportation Review* 46.5, pp. 709–718. DOI: 10.1016/j.tre.2010.01.004.

Chang, Yu-Chun, Hui Kee Woon, Jin-Ru Yen, and Chia-Jui Hsu (2017). "A destination selection model for long haul routes — The case of Taiwan-US". In: *Journal of Air Transport Management* 64, pp. 60–67. DOI: 10.1016/j.jairtraman.2017.06.026.

Charnes, Abraham and William Wager Cooper (1957). "Management Models and Industrial Applications of Linear Programming". In: *Management Science* 4 (1), pp. 38–91. DOI: 10.1287/mnsc.4.1.38.

Cheung, Wai Lun, Rajesh Piplani, Sameer Alam, and Lionel Bernard-Peyre (2021). "Dynamic capacity and variable runway configurations in airport slot allocation". In: *Computers & Industrial Engineering* 159. DOI: 10.1016/j.cie.2021.107480.

Dalkey, Norman and Olaf Helmer (1963). "An Experimental Application of the DELPHI Method to the Use of Experts". In: *Management Science* 9.3, pp. 458–467. DOI: 10.1287/mnsc.9.3.458.

Deveci, Muhammet, Nihan Çetin Demirel, and Emine Ahmetoğlu (2017). "Airline new route selection based on interval type-2 fuzzy MCDM: A case study of new route between Turkey-North American region destinations". In: *Journal of Air Transport Management* 59, pp. 83–99. DOI: 10.1016/j.jairtraman.2016.11.013.

Dixon, Matthew (2006). *The maintenance cost of aging aircraft*. Tech. rep. RAND Corporation. URL: https://apps.dtic.mil/sti/pdfs/ADA459880.pdf.

Fuentes, Manuel, Luis Cadarso, Vikrant Vaze, and Cynthia Barnhart (2021). "The Tail Assignment Problem: A Case Study at Vueling Airlines". In: *Transportation Research Procedia* 52, pp. 445–452. DOI: 10.1016/j.trpro.2021.01.052.

Hane, Christopher, Cynthia Barnhart, Ellis Johnson, Roy Marsten, George Nemhauser, and Gabriele Sigismondi (1995). "The fleet assignment problem: Solving a large-scale integer program". In: *Mathematical Programming* 70.1, pp. 211–232. DOI: 10.1007/BF01585938.

Hwang, Ching-Lai and Kwangsun Yoon (1981). *Multiple Attribute Decision Making. Methods and Applications — A state-of-the-art survey*. Lecture notes in economics and mathematical systems. Springer Berlin. 269 pp. ISBN: 978-3-642-48318-9. DOI: 10.1007/978-3-642-48318-9.

Mashford, John and Bertil Marksjö (2001). "Airline Base Schedule Optimisation by Flight Network Annealing". In: *Annals of Operations Research* 108, pp. 293–313. DOI: 10.1023/A:1016027516013.

Özcan, Tuncay (2018). "A Hybrid Multi-Criteria Decision Making Approach for New Destination Selection in Aviation Industry". (In Turkish). In: *Alphanumeric Journal* 6 (1), pp. 1–12. DOI: 10.17093/alphanumeric.369758.

Pilla, Venkata, Jay Rosenberger, Victoria Chen, and Barry Smith (2008). "A statistical computer experiments approach to airline fleet assignment". In: *IIE Transactions* 40.5, pp. 524–537. DOI: 10.1080/07408170701759734.

Pita, João Pedro, Cynthia Barnhart, and António Pais Antunes (2013). "Integrated Flight Scheduling and Fleet Assignment Under Airport Congestion". In: *Transportation Science* 47.4, pp. 477–492. DOI: 10.1287/trsc.1120.0442.

Santini, Alberto (Feb. 22, 2023). *Instances for the Regional Flight Scheduling Problem*. Github repository. DOI: 10.5281/zenodo.7664132. URL: https://github.com/alberto-santini/rfsp-instances/.

Vaze, Vikrant and Cynthia Barnhart (2012). "Modeling Airline Frequency Competition for Airport Congestion Mitigation". In: *Transportation Science* 46.4, pp. 512–535. DOI: 10.1287/trsc.1120.0412.

Wang, Desmond Di, Diego Klabjan, and Sergey Shebalov (2014). "Attractiveness-Based Airline Network Models with Embedded Spill and Recapture". In: *Journal of Airline and Airport Management* 4.1, pp. 1–25. DOI: 10.3926/jairm.20.

Wei, Keji, Vikrant Vaze, and Alexandre Jacquillat (2020). "Airline Timetable Development and Fleet Assignment Incorporating Passenger Choice". In: *Transportation Science* 54.1, pp. 139–163. DOI: 10.1287/trsc.2019.0924.

Yan, Chiwei, Cynthia Barnhart, and Vikrant Vaze (2022). "Choice-Based Airline Schedule Design and Fleet Assignment: A Decomposition Approach". In: *Transportation Science* 56.6, pp. 1410–1431. DOI: 10.1287/trsc.2022.1141.

Yan, Shangyao and Chung-Rey Wang (2001). "The planning of aircraft routes and flight frequencies in an airline network operations". In: *Journal of Advanced Transportation* 35.1, pp. 33–46. DOI: 10.1002/atr.5670350104.

Yu, Gang and Jian Yang (1998). "Optimization Applications in the Airline Industry". In: *Handbook of Combinatorial Optimization*. Ed. by Ding-Zhu Du and Panos Pardalos. Springer, pp. 1381–1472. DOI: 10.1007/978-1-4613-0303-9_23.