

Development of an Aircraft Routing System for an Air Taxi Operator

F.M. van der Zwan¹, K. Wils¹ and S.S.A. Ghijs²

¹*Delft University of Technology*

²*Fly Aeolus*

¹*The Netherlands*

²*Belgium*

1. Introduction

Due to increasing congestion of the road network and the major airports, the established modes of transport like car and scheduled air travel are having difficulty to fulfil the need for efficient business travel. Therefore, recent years have seen a growing demand for business aviation services like business aircraft charter and (fractional) aircraft ownership. These services enable improved time efficiencies, access to a wider range of airports, and the ability to control flight scheduling (Budd & Graham, 2009). However, business aviation is characterized by high prices and due to the economic downturn of 2008 corporate travel departments now seek less costly and more efficient business travel solutions (American Express, 2009; Gall & Hindhaugh, 2009). This has paved the way for the development of a new phenomenon in the air transport industry, namely the on-demand air taxi service. Air taxi operators offer travellers a low fare, on-demand travel service (Budd & Graham, 2009). As such, they provide the flexibility and time efficiency of existing on-demand business aviation, giving them the same time-efficiency advantage over car and scheduled air transport, but at a lower price (Bonnefoy, 2005).

One of the key enablers that allows air taxi operators to offer lower fares and still remain profitable is, apart from their smaller and thus cheaper aircraft, the optimization of resources, particularly during the employment of the aircraft during the aircraft routing phase. This chapter therefore sets out to develop an aircraft routing system which allows a full on-demand per-aircraft air taxi company to generate aircraft routing plans which adhere to its planning objectives, during the first operational years. Firstly, the air taxi business model is discussed. Then the aircraft routing problem and modelling approaches are described, followed by the design, testing and validation of the aircraft routing model, after which conclusions on the approach taken are drawn.

2. The air taxi business model

This section discusses the air taxi business model, the planning process for on-demand air transport (ODAT) and how this differs from scheduled air transport services.

2.1 Air taxi characteristics

Existing air taxi services can be differentiated with respect to their offer. Both per-seat and per-aircraft air taxi services exist. In the former the customer buys a single seat on a flight,

while in the latter the customer charters a whole aircraft. Also 'semi on-demand' and 'full on-demand' air taxi services exist. In the former the air taxi service only operates from a predetermined set of airports, while in the latter the customer can freely choose the departure and arrival airport. Another difference lies in the type of aircraft used; operators either fly very light jets (VLJs) or small piston aircraft. A VLJ typically seats between four and seven passengers, can be flown by a single pilot, has a maximum take-off weight of less than 10,000 lb (4536 kg), and an average range of around 2,000 km (Budd & Graham, 2009, p.289). Piston aircrafts have one or more piston-powered engines connected to the propeller(s). Piston aircraft used for business typically seats one to six passengers and fly relatively short missions of 300-400 miles, using very small general aviation airports that are often without air traffic control towers (NBAA, 2010). The characteristics of the air taxi business model are summarised in Table 1 below.

Focused outside hub airports
Faster than scheduled air carriers for most point-to-point travel
One to six passengers per aircraft (though some take more)
High utilization of aircraft (via optimization of resources)
Cheaper than a charter or a fractionally owned aircraft
On-demand or semi on-demand (no fixed or published schedules)
Sold by the plane or per-seat
Clear, all-in pricing, by distance, time or zone
Using small piston or VLJ aircraft

Table 1. Characteristics of the air taxi business model (after Dyson, 2006)

As mentioned at the beginning of this chapter, air taxi services offer the flexibility and time efficiency of existing on-demand business aviation but at a lower price. Prices for per-aircraft air taxi (PAAT) operators in Europe range from € 2,500 per flight hour (London Executive Aviation, 2010) to € 800 per flight hour (FlyAeolus, 2010). Bonnefoy (2005) indicates that air taxi operators are able to offer these lower fares due to their use of a small aircraft single type fleet and optimization of resources. The small aircraft used by air taxi operators have lower acquisition and operating costs than jets used by existing business aviation (Bonnefoy, 2005). Also, Budd and Graham (2009) and Espinoza et al. (2008a) note that the landing and take-off characteristics of these smaller aircraft enable access to an even greater range of airports, which possibly lie closer to where business travellers want to go. Apart from the smaller aircraft, optimization of resources is also key to the cost feasibility of the air taxi model (Dyson, 2006; Mane & Crossley, 2009). The resource optimization challenges faced by air taxi operators present themselves during the operational planning and management of the day-to-day operations.

2.2 Planning process for air taxi model

The airline planning process for scheduled air transport typically consists of five phases: flight scheduling, fleet assignment, aircraft routing, crew scheduling and crew rostering (Bazargan, 2004). The planning problem for on-demand air travel has a different nature than that of scheduled airlines (Hicks et al., 2005; Yao, 2007). The primary cause for this is the demand mechanism (Yao, 2007): scheduled airlines decide on their flight schedule months in advance, while in contrast, on-demand air transport operators like air taxis and fractional management companies sometimes know their flight requests only several hours in advance

since customers can book up to six or eight hours before departure (Hicks et al., 2005; Yang et al., 2010). In addition, the flight legs flown by on-demand operators are less predictable and differ from day to day and week to week (Ronen, 2000; Yao et al., 2008). However, one can still distinguish the different phases of the scheduled airline's planning process in the planning process of on-demand air travel operators.

The **flight scheduling** phase for on-demand air transportation is customer driven. The customer contacts the air taxi operator a couple of days or hours in advance with his flight request. This process leads to an unpredictable and non-repeatable flight schedule. The customer's flight request specifies a departure location and departure time and an arrival location. Air taxi companies have the option to reject a request. The accept/reject decision depends on a number of factors, for example whether it is possible to fly the request and also whether it is worthwhile flying the flight, both financially and strategically. This assessment should be performed with respect to the flight request itself, with respect to expected demand (Fagerholt et al., 2009), and also future customer demand as a customer denied service may choose not to request future trips (Mane & Crossley, 2007a). During the flight scheduling phase it must also be decided when and where aircraft will undergo scheduled maintenance (Keysan et al., 2010). If the request is accepted, the flight leg(s) of the customer's flight request are added to the flight schedule.

The second phase of the scheduling process considers **fleet assignment**, where fleet types (not specific aircraft) are assigned to each flight leg in the schedule (Bazargan, 2004). For an air taxi company, this phase is virtually non-existent as they mostly operate a single type fleet (Bonnetoy, 2005; Dyson, 2006).

The next phase is **aircraft routing**. During this phase each individual aircraft is assigned a routing, which is a sequence of flight legs, so that each leg is covered exactly once while ensuring that the aircraft visits maintenance stations at regular intervals thereby fulfilling maintenance requirements (Barnhart et al., 2003). The aircraft routing phase for on-demand air transport operators consists of assigning specific aircraft to the customer requested flight legs (Yao, 2007). During aircraft routing the operator thus aims to create aircraft routes that cover all the flight requests while minimizing the operational cost and adhering to the operational constraints, related to the aircraft's maintenance limit, crew regulations and the availability of fuelling facilities at airports (Martin et al., 2003; Yao et al., 2008). As departure and arrival locations can be chosen freely by the customer it is unlikely that an aircraft is always directly available at the departure location. Therefore an ODAT operator will have to conduct repositioning flights, also called 'deadhead' or 'non revenue' flights, to reposition their aircraft to the departure locations. Yao et al. (2008) mentions that deadhead flights may represent over 35% of the total flying conducted by a fractional management company (FMC). Air taxis and FMCs mostly need to bear the cost of repositioning themselves since the customer only pays for his actual flight request (Yang et al., 2008). Therefore, as a major driver during the aircraft routing phase is minimizing operational cost, it is important to minimize repositioning flights, both in number and in length (Bonnetoy, 2005; Mane and Crossley, 2007b; Yao, 2007). It must be noted that per-seat on-demand operators face an additional challenge during the aircraft routing phase. Apart from planning so that all customer requests are served, customer requests should be interleaved where possible such that that different customers can be put on the same aircraft, thereby obtaining a load factor sufficiently high to remain profitable (Espinoza et al., 2008a). This greatly increases the complexity of the underlying aircraft routing problem.

The two last phases in the scheduled airline's planning process involve the planning of crew resources. The fourth phase is **crew scheduling** or crew pairing. This treats the process of identifying sequences of flight legs that start and end at the same crew base, i.e. trip pairings (Bazargan, 2004). While creating trip pairings one strives for a minimum-cost set of pairings so that every flight leg is covered and that every pairing satisfies the applicable crew work rules (Bazargan, 2004). Finally, in the **crew rostering** phase, crew pairings are combined into monthly schedules which are then assigned to individual crew members (Bazargan, 2004). The crew pairing and crew rostering phase for ODAT operators are different from those used by scheduled airlines. As the definitive flight schedule is not known beforehand, no accurate monthly crew schedules can be created. Therefore crew of ODAT operators work with on-duty periods, during which they must be ready to operate flight legs which may not yet be known at the start of the on-duty period, and off-duty periods (Hicks et al., 2005; Yao et al., 2008). Yao et al. (2008) mention that, as aircraft are not always stationed at crew bases, crews coming on duty may need to travel to the aircraft's location, which counts as on-duty time and (possibly) needs to be paid for by the ODAT operator. Needless to say, ODAT crew schedules need to meet crew regulations as stipulated by the air transport authorities.

With respect to the demand uncertainty of ODAT, Hicks et al. (2005) report, based on the operations of an FMC, that approximately 80% of the trips are requested 48 hours or more prior to departure and 20% with as little as 4 hours notice. In addition, 30% of trips are changed at least once within 48 hours of the requested departure time (Hicks et al., 2005). Historical data analysis on FMC operations shows that on average 5% of demand is unknown for the first day for which the schedule is created, and 10-20% and 25-40% for the second and the third day respectively (Yao et al., 2007, 2008). Despite the uncertainty in demand, ODAT operators need to create advance flight schedules, aircraft routings and crew schedules. There are a number of reasons for this: some airports require arrival information 24 hours in advance, the crews need to be given enough time to relocate to where they are required to come on duty, and maintenance locations need to be booked for scheduled maintenance (Yang et al., 2010). However, these advance schedules need to be adapted as more demand information gets known. For example, Yang et al. (2008) reports that around 25% of the requested trips of an FMC arrived after the original schedule was created, which causes the need to re-optimize the schedule.

Because of this dynamic process of creating and adapting the schedules it is desirable to both have a flexible and persistent schedule (Karaesmen et al., 2007; Yang et al., 2010): a flexible schedule to be able to cost-effectively cater extra demand, and a persistent schedule to avoid that the planning completely changes every time a change in demand occurs. Next to that, the schedule needs to be robust. Robustness denotes the ability of the schedule to cope with and cost-effectively recover from unforeseen changes in supply (e.g. aircraft break-down) or other disruptions (e.g. adverse weather) (Ball et al., 2007). As Bian et al. (2003) state, in a robust schedule enough slack is built in, such that a single disruption does not cause later flights to be delayed. However, a robust schedule comes at a price as schedule slack basically means that aircraft will be standing idle. Finally, the ODAT schedule needs to achieve a sufficient customer service level. This means on the one hand carefully looking at accepting or rejecting of flight requests, but also considering whether to schedule their customer's flights as requested or not (Fagerholt et al., 2009). Delaying the departure time by 30 minutes may enable a more cost-effective aircraft routing but this may

have a negative impact on the service level as perceived by the client. In summary, the ODAT operator has to balance minimising costs whilst keeping a high service level and developing an appropriate schedule (flexible, persistent, and robust).

3. The aircraft routing problem

In this section the aircraft routing problem as it applies to the air taxi operator is laid out. As explained before, the aircraft routing problem arises during the aircraft routing phase, in which the ODAT operator aims to create aircraft routes that cover all flight legs and satisfy operational restrictions while minimizing operational cost. Below, the different aspects of the aircraft routing problem are treated:

- **The scope:** planning horizon
- **The input:** flight legs (trips) and aircraft
- **The main objective:** to create an aircraft routing plan that covers all flight legs while minimizing the operating cost
- **The secondary objectives:** robustness, flexibility and persistence
- **The constraints:** operational requirements

Please note that this work has been undertaken in collaboration with Fly Aeolus, a Belgian full on-demand per-aircraft air taxi operator. This routing problem and subsequent model has been kept as generic as possible. However, in some instances the specifics of the Fly Aeolus business model have been taken as an input for the routing problem. In the text it is clearly stated how the routing problem might be (slightly) different for ODAT operators with a different business model.

An important parameter in the planning process is the **planning horizon** that is used. This determines the size, and thus complexity, of the planning problems that need to be solved during the different phases. If a 24 hour horizon is chosen the aircraft routing system will find a local optimum for that day, but the ending locations of the aircraft will not be optimized towards the start of the next day. Theoretically, the more days that are included in the planning horizon, the closer the solution will be to a global optimum. However, due to the computational complexity and the uncertainty of demand, there is a limit on the number of days that can be incorporated in the planning horizon. Ronen (2000) has tested the effect of the planning horizon by solving aircraft routing problems for 24, 36 and 48 hour horizons for a fractional management company (FMC). For his application a 24 to 36 hour planning horizon proved most effective, as planning decisions taking place beyond that horizon were always changed again later. On the other hand, Martin et al. (2003) state that the operations of the FMC they worked with allowed a two to three day planning horizon. Overall, in the ODAT literature it is generally accepted that planning beyond 72 hours is highly speculative. Therefore, the planning horizon is set to a maximum of 72 hours.

At each instance the operations control centre (OCC) wants to create an aircraft routing plan, the aircraft routing system takes the known flight schedule (output of the flight scheduling phase) and extracts the booked flight requests. Each flight request consists out of at least one **flight leg** (trip). Each of these flight legs has a departure airport, a requested time of departure, an arrival airport and a specified number of passengers for that flight. Customers also have the possibility to make a booking consisting of multiple flight legs. For these requests they are given the option to have the aircraft and crew remaining on standby at the airport in between the coupled flight legs, such that all their separate trips are served by the same aircraft. Apart from the flight requests, the aircraft routing system also retrieves

the current status of all the **aircraft** available within the available fleet. For each aircraft the home base and its lease cost per hour are known. In addition to this, it is known when and where the aircraft will become available, and the aircraft's fuel status and flight hours remaining till maintenance at that time. An aircraft can be unavailable at the start of the planning horizon because it is still in-flight at that time, or because maintenance is performed on the aircraft.

The objective of the aircraft routing system is now to assign to each aircraft a feasible sequence of flight legs in a manner that **minimizes the cost of operating all the flights** while meeting the operational requirements. Of the total operating costs, the variable direct operating costs (DOCs) are the only costs that are controlled by the aircraft routing system. This is because these are the costs that are directly related to the amount of flying that is undertaken (Doganis, 2002), which in turn depends on the aircraft routing plan that is created. The following DOC items incurred by the air taxi operator are directly influenced by the aircraft routing plan: variable aircraft costs, passenger and non-passenger related airport charges, air traffic service charges, fuel costs, and variable flight crew costs. With respect to the fuel costs, these costs depend on the amount of flying that is planned. In addition, they also rely on where and how much the aircraft is planned to be fuelled. Because fuel prices vary significantly between airports (Flyer Forums, 2010), so-called fuel ferrying can in certain cases be a valid strategy to reduce fuel cost. When adopting this approach, excess fuel is carried on a flight leg to decrease the amount of fuel that needs to be bought at the next airport where fuel prices are high. The crew costs consist of the pilot wages which are directly related to how much flying the pilot carries out (as measured by the Hobbs meter of the aircraft, which registers the time that the engine is running and thus also includes taxiing and idling (e.g. when waiting for take-off clearance)). Therefore, this part of crew costs is also controlled during the aircraft routing phase.

As explained in the previous section, achieving **robustness, persistence and flexibility** while creating operational plans are additional objectives to consider during the planning process. These three secondary planning objectives are also to be considered during the design of the aircraft routing system. However, these planning objectives are interrelated and cannot be fully achieved together. For example, requiring a very robust schedule (thus a large amount of slack time), decreases the ability of the air taxi operator to accept future flight requests. During the development of the aircraft routing system it is therefore ensured that the system parameters that control robustness, persistence and flexibility can be altered by the OCC such that they can prioritize these objectives as they wish.

The aircraft routing plans that are created also have to fulfil a number of **operational requirements**. Firstly, aircraft should not exceed their maintenance flight hour limit when flying the assigned flight sequence. Secondly, flight sequences should be fuel feasible. Both because there are airports where it is not possible to fuel and because the number of passengers that is carried influences the amount of fuel that can be taken. For example, an SR22 carrying a pilot and two passengers can only carry 38% of its maximum usable fuel (Cirrus, 2007). In addition, aircraft routing plans must be crew feasible, i.e. they should not violate the crew work rules. For this purpose, it is customary in the ODAT industry to couple an aircraft with its crew for scheduling purposes (Martin et al., 2003; Yao et al., 2007; Yang et al., 2008). Hence, an aircraft is viewed as an operational unit which has to meet both its own restrictions (e.g. maintenance) and crew duty time limits. Because of this, the aircraft routing system will generate routing plans that do not necessarily require crew swaps during the day as a single crew can (legally) stay with a single aircraft throughout the whole day.

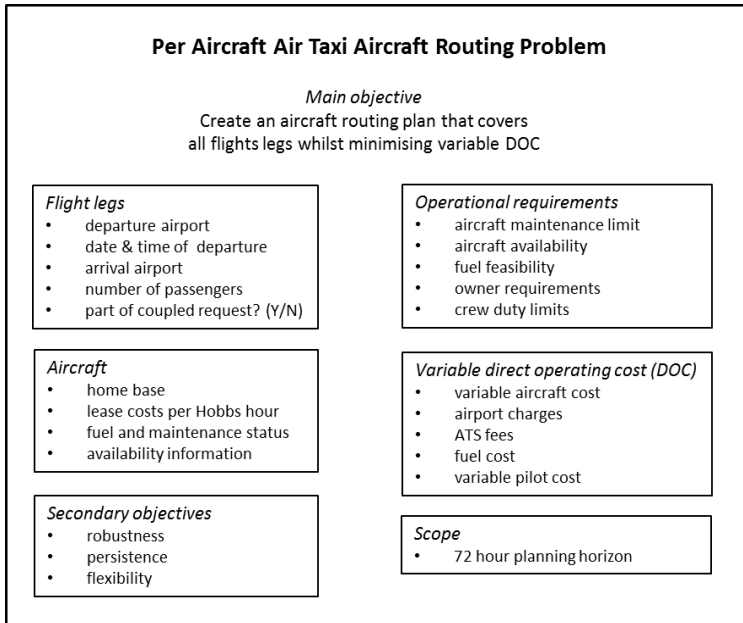


Fig. 1. Full on-demand per-aircraft air taxi routing problem

In summary, during the routing phase the air taxi operator's OCC needs to create aircraft routing plans covering all flight legs taking place during the scheduling horizon with the aim to minimize operating cost while adhering to the aircraft's maintenance flight hour limits, fuel feasibility, aircraft availability and crew duty rules, as shown in Figure 1 above.

4. Choice of modelling approach

In this section the modelling approach for the aircraft routing problem is chosen, after which the chosen methodology is further explained.

4.1 Choosing the appropriate modelling approach

To the authors' best knowledge there exists no literature on the aircraft routing problem faced by a per-aircraft air taxi (PAAT) operator. However, their aircraft routing problem is very similar to that of a FMC or charter. Both aim to minimize the operating cost while serving all the customer requests on time, and satisfying the applicable constraints. However, there are some differences between the aircraft routing problem of an FMC and that of a typical air taxi operator. First of all, since most air taxi operators operate a single type of aircraft, all its aircrafts are in principle compatible with every request. For an FMC, which operates different fleet types, there is a specific aircraft type that fits each request and this requirement must be taken into account during the aircraft routing. In addition, an FMC does not incur the flight hour related aircraft costs as a part of its variable operating costs because they either own or lease their aircraft for a continuous period of time. So their aircraft cost is irrespective of how much they use the aircraft. However, despite these differences the core of the aircraft routing problem for an FMC and a PAAT operator

remains the same. Therefore, this section focuses on studies related to FMCs to develop a system that solves the PAAT aircraft routing problem. These are listed in Table 2 below.

Modelling approach	Author(s)	Maintenance constraints	Crew constraints	Planning horizon [h]	Mathematical formulation*	Solution method#
M1	Keskinocak & Tayur (1998)	N	N	72	0 - 1 IP	CPLEX
	Keskinocak & Tayur (1998)	N	N	72	0 - 1 IP	Heuristic
M2	Ronen (2000)	Y	Y	48	SP	Heuristic
M3	Martin et al. (2003)	Y	Y	48 - 72	0 - 1 IP	CPLEX
M4	Hicks et al. (2005)	Y	Y	24 - 72	Mixed IP	CG, GENCOL
M5	Karaesmen et al. (2005)	Y	N	24	0 - 1 IP	CPLEX
M6	Karaesmen et al. (2005)	N	N	24	Mixed IP	CPLEX
	Yang et al. (2008)	N	N	24	Mixed IP	CPLEX
M7	Karaesmen et al. (2005)	Y	Y	24	0 - 1 IP	Heuristic, CPLEX
	Yang et al. (2008)	Y	Y	24	0 - 1 IP	Heuristic CPLEX
M8	Karaesmen et al. (2005)	N	Y	24	Mixed IP	CPLEX
	Yang et al. (2008)	N	Y	24	Mixed IP	CPLEX
M9	Karaesmen et al. (2005)	Y	Y	24	SP	BP, CPLEX
	Yang et al. (2008)	Y	Y	24	SP	BP, CPLEX
	Yang et al. (2008)	Y	Y	72	SP	BP, CPLEX
	Yang et al. (2008)	Y	Y	96	SP	BP, CPLEX
M10	Yao et al. (2005)	Y	Y	72	SP	CG
	Yao et al. (2008)	Y	Y	72	SP	CG
M11	Yao & Zhao (2006)	Y	Y	72	SP	na
	Yao et al. (2007)	Y	Y	72	SP	CG

* IP = integer programming, SP = set-partitioning # CG = column generation, BP = branch-and-price

Table 2. Modelling and solution approaches for the aircraft routing problem

To choose a modelling approach for a PAAT aircraft routing problem one must first specify the decision criteria that are used during the trade-off. These criteria are derived from the requirements that are posed upon the aircraft routing system and from the nature of the PAAT aircraft routing problem. The modelling approach should allow for both crew feasible and maintenance feasible routings. It should be able to accommodate a multiple day planning horizon and multiple fleet types. It should also have the flexibility to incorporate additional operational rules and finally, the flight leg cost should be dependent on the aircraft route. In Table 2 above the specifics of the different aircraft modelling approaches deployed in FMC literature are given. Since the criteria that schedules must be crew and maintenance feasible need to be fulfilled a number of modelling approaches can already be discarded. These are model 1 proposed by Keskinocak and Tayur (1998) and models 5, 6 and 8 described by Karaesmen et al. (2005) and Yang et al. (2008). Furthermore, model 7 can also be discarded since it is noted by Karaesmen et al. (2005) and Yang et al. (2008) that this modelling approach is not flexible enough to handle planning horizons longer than 24 hours. This leaves us with six models (2, 3, 4, 9, 10 and 11). All these models produce schedules that are crew and maintenance feasible and can be used in multiple day planning horizons.

However, model 10 and 11 do not incorporate the cost of flying customer trips in the operational cost of a schedule. The authors, Yao et al. (2005, 2008); Yao and Zhao (2006), state that the customer of an FMC pays for the fuel and the crew costs of their trip, thus they do not consider these trip costs as a part of operational cost. For the PAAT aircraft routing problem this is not a valid assumption. The cost for one specific flight leg namely differs depending on which aircraft routing the leg is in. Air taxi customers always pay a fixed price per flight hour, which is independent of what total cost the operator has to bear to operate this flight.

Models 2 and 9 utilize a set-partitioning formulation. In an overview of vehicle routing models, Bunte and Kliewer (2010) state that an advantage of the set-partitioning formulation, in contrast to other formulations, is that additional constraints or operational rules can be easily incorporated. Also, Ronen (2000) and Karaesmen et al. (2005) note that the set-partitioning approach is very appropriate to problems where costs are nonlinear and discrete, and complicated rules are imposed, as is the case in the PAAT aircraft routing problem. A direct comparison with a 0-1 and mixed IP formulation is made by Karaesmen et al. (2005) and Yang et al. (2008). They conclude that the set-partitioning formulation is the only one that has adequate flexibility to cover long planning horizons and to incorporate complex operational rules (Karaesmen et al., 2005). Therefore, in order to model the PAAT aircraft routing problem the set-partitioning formulation is chosen.

4.2 The set-partitioning model

For the set-partitioning formulation of the PAAT aircraft routing problem, which is adopted from Yang et al. (2008), the following notation is used:

$$\begin{aligned}
 \Omega_i &= \text{the set of feasible routes for aircraft } i, & i \in N \\
 p &= \text{a feasible route}, & p \in \Omega_i \\
 c_{ip} &= \text{the total cost of taking route } p \text{ by aircraft } i, & i \in N \text{ and } p \in \Omega_i \\
 a_{jp}^i &= 1 \text{ if trip } j \text{ is on route } p; \quad 0 \text{ otherwise}, & i \in N, p \in \Omega_i \text{ and } j \in M \\
 \theta_{ip} &= 1 \text{ if aircraft } i \text{ takes route } p; \quad 0 \text{ otherwise}, & p \in \Omega_i \text{ and } i \in N
 \end{aligned}$$

Where $N = \{1, 2, \dots, n + m\}$ is the set of all aircraft such that $\{1, \dots, n\}$ represents the own fleet, and $\{n+1, \dots, m\}$ represents any charter aircraft for subcontracting the trips (in case the own fleet proves insufficient to operate all the flights). For each $i \in (n + 1, \dots, m)$, Ω_i consists of a single route that only takes trip j with a cost b_j (the cost that is incurred when subcontracting this trip). Furthermore, M denotes the set of all trips. Using this notation, the aircraft routing problem of a PAAT operator can be formulated as the following set-partitioning problem (SP):

$$\text{(SP)} \quad \min \sum_{i \in N} \sum_{p \in \Omega_i} c_{ip} \theta_{ip} \quad (1)$$

$$\text{s.t.} \quad \sum_{i \in N} \sum_{p \in \Omega_i} a_{jp}^i \theta_{ip} = 1, \quad j \in M \quad (2)$$

$$\sum_{p \in \Omega_i} \theta_{ip} \leq 1, \quad i \in N \quad (3)$$

$$\theta_{ip} \in \{0, 1\}, \quad i \in N \text{ and } p \in \Omega_i \quad (4)$$

This way aircraft are assigned to feasible routes while minimizing the total cost of serving all trips (1). Constraints (2) ensure that each trip is covered exactly once in the solution. Constraints (3) make sure every aircraft is assigned at most one route. Also, either a route is part of the solution or it is not, therefore parameters θ_p must be either 0 or 1 (4).

To solve the set-partitioning model of the aircraft routing problem it is possible to use either an exact solution method or a heuristic (Hillier and Lieberman, 2008). As Hillier and Lieberman (2008) and Silver (2004) note heuristics are often used when the time required to find an optimal solution for an accurate model of the problem would be very large. For the aircraft routing system under consideration computational time is indeed an important requirement. However, the FMC studies mentioned in Table 2 that use exact solution methods to solve a set-partitioning model already achieve solution times in the order of seconds. The second performance measure for a solution method, apart from computational speed, is the degree of optimality of the solution (Silver, 2004). In practice, and even for exact solution methods, there is a maximum optimality gap given which is the maximum percentage that the obtained solution differs from the optimal solution. For the exact methods computational time or memory limits can cause the model not to be solved to optimality. The maximum optimality gap for the exact solution methods for the set-partitioning models has been found to be smaller than half a percent. Also, in the considered literature on the aircraft routing problem for on-demand air transport no heuristics have been used to solve the set-partitioning model. In addition, when designing a heuristic or adapting an existing one, both computational performance and solution quality can only be assessed after implementation. So when opting for a heuristic solution method for this problem no performance guarantees can be given beforehand. Because exact solution methods have proven to yield an adequate computational performance and solution quality for solving the set-partitioning model of the aircraft routing problem for on-demand air transport, the authors have chosen to use an exact solution method to solve the set-partitioning model that is created by the model creator.

To implement an exact solution method for the solver of the aircraft routing system, there are two options. First of all, one can program and implement an exact solution method, like branch-and-price or column generation, themselves. However, as Feillet (2010) has recently noted, carrying out an accurate implementation of an exact solution algorithm is a long and difficult task due to the inherent complexity of these methods and the lack of simple and comprehensive descriptions of these methods. The second method is to use a solver software package that embeds one or more of the exact solution methods needed to solve the set-partitioning problem. There already exist a number of solver packages that are applicable for the problem at hand, as can be seen in Table 2. Therefore, in this project a solver software package is used to solve the set-partitioning formulation of the PAAT aircraft routing problem. The solver package that is used is the IBM ILOG CPLEX Optimizer, since it has already proven its worth in ODAT applications. To solve the set-partitioning problem with CPLEX use is made of the built-in "mipopt" optimizer module. This optimizer uses a branch-and-cut algorithm, which is a hybrid method of branch-and-bound and cutting plane methods (Smith and Taşkin, 2008).

5. Designing the aircraft routing system

To construct the set-partitioning model of the aircraft routing problem for a PAAT operator, the model creator receives the following input data via the I/O interface of the aircraft

routing system: flight leg (trip) data of customer and dummy flight requests, parameters of the real and dummy aircraft, airport parameters, a list of persistent aircraft-trip assignments, and operational parameters. This is shown in Figure 2 below.

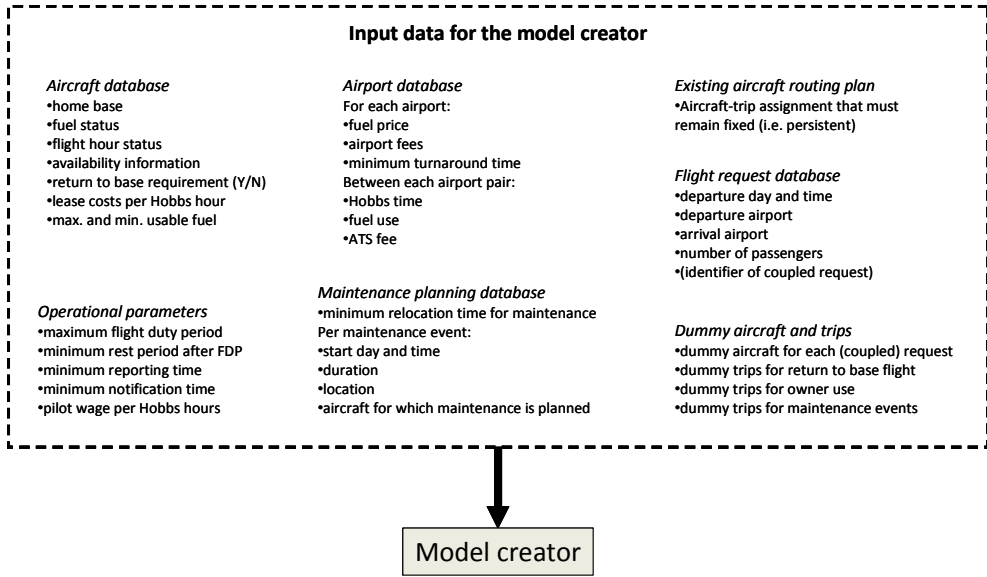


Fig. 2. Input data of the aircraft routing system.

Dummy requests are created for each return to base trip, for owner use and for maintenance events. Dummy aircraft on the other hand are added to represent the aircraft of another company when a trip is subcontracted. The model creator uses the above input data to generate the set of feasible routes for each aircraft (Ω_i), the total cost of operating each of these routes (c_{ip}) and information as to which trips are contained in each route (a_{ijp}). With these parameters the set-partitioning formulation of PAAT aircraft routing problem is constructed and written to an LP file. This file is read by CPLEX which then solves the set-partitioning model. The obtained solution specifies which routes are part of the aircraft routing plan (i.e. which θ_{ip} 's are 1). This information is processed by the solver module which uses it to generate the aircraft routing plan that forms the solution of the aircraft routing problem. The model creator and the solver module of the aircraft routing system described in this chapter are programmed in MATLAB. The model creator uses the input data provided by the I/O interface to generate an LP file of the set-partitioning model that represents the aircraft routing problem. This LP file serves as input for the solver module. Figure 3 provides a detailed overview of this process.

First, the preprocessor module of the model creator uses the input data to create matrices in which it is specified whether aircraft i is initially compatible with trip j and whether aircraft i can fly trip j_2 after trip j_1 . Together with the input data these two matrices are passed to the next module of the model creator, namely the route generator. The route generator module uses the information stored in the matrices to create the set of possible routes for each aircraft.

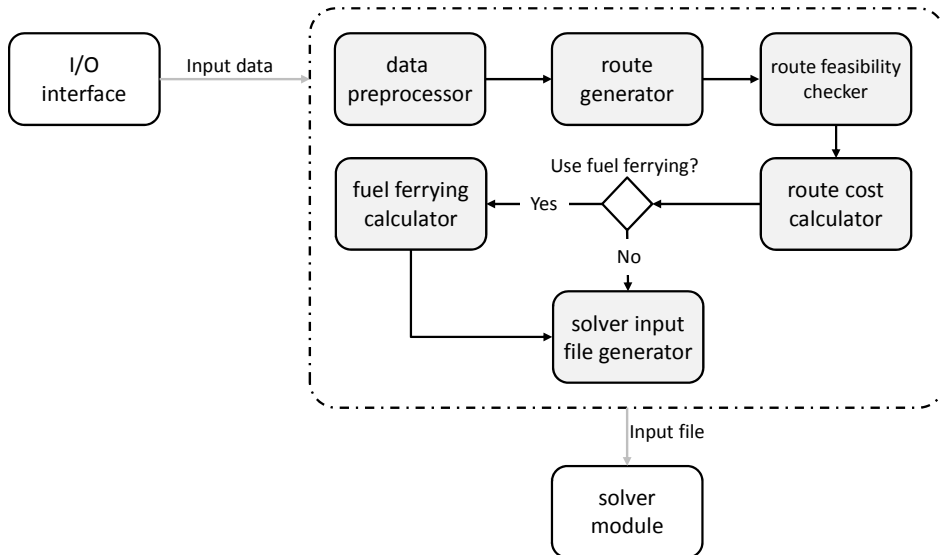


Fig. 3. The model creator module of the aircraft routing system.

Each of these routes consists of a sequence of flight legs. This set of routes for each aircraft is passed to the route feasibility checker. The route feasibility checker processes the route list for each aircraft generated by the route generator module and removes the routes that are infeasible. The list of feasible routes is then forwarded to the route cost calculator. The route cost calculator processes each of the feasible routes and calculates the total route cost for each of the feasible routes. The total route cost consists of the aforementioned variable aircraft cost, airport charges, ATS charges, fuel costs, and variable crew cost. If the fuel ferrying strategy is used, the fuel cost is calculated using the fuel ferrying calculator. This module uses the assumed percentage burn-off fuel ferrying calculation method. The final module of the model creator is the solver input file generator. This module writes an LP file of the set-partitioning model that is created using the previous modules. This LP file can be directly read and solved by CPLEX. The solver module of the aircraft routing system consists out of a MATLAB script that solves the created set-partitioning formulation, as written to the LP file, using CPLEX and retrieves the solution, i.e the aircraft routing plan.

6. Testing and validating the model

In this section the aircraft routing model is both tested and validated.

6.1 Testing the model

To test and validate the developed aircraft routing system operational data is needed as input. However, at the time of writing Fly Aeolus is not yet fully operational and as such the system can not be tested in the environment where it will eventually be used. Therefore a virtual test environment is created. This environment contains a virtual airport network in which a per-aircraft air taxi operates. The virtual environment is created by an input data generator which creates all the input data needed by the model creator. The input data generator utilizes operational parameters that reflect the (expected) PAAT operations of Fly Aeolus, see Table 3.

Parameter	Value	Unit
Parameters used to generate airport test data		
Farthest point-to-point distance in the network	800	km
Number of airports	225	-
Average cruise speed	310	km/h
Average fuel burn	1.12	litre per minute
ATS fee	0.00	€ per km
Passenger arrival fee	0 - 10	€ per pax
Passenger departure fee	0 - 30	€ per pax
Aircraft arrival fee	10 - 150	€
Aircraft departure fee	0 - 60	€
Turnaround time	45	minutes
Avgas price	0.90 - 2.70	€ per liter
Number of maintenance bases	6	-
Parameters used to generate trip test data		
Planning horizon	72	hours
Minimum Hobbs time of flight request	50	minutes
Average Hobbs time of flight request	90	minutes
Departure time distribution	-	xx
Average number of passengers	1.3	-
Percentage of coupled flight requests	10	% of flight requests
Number of trips in a coupled request	2	-
Waiting time between each trip of a coupled request	6 - 8	hours
Aircraft that have maintenance events	10	% of aircraft
Planned maintenance duration	6	hours
Minimum relocation for maintenance	120	minutes
Aircraft that have owner use	30	% of aircraft
Start time of owner use, between	5 - 20	hour
Aircraft that have return to base requirement	30	% of aircraft
Aircraft has to be back at base at	22	hour
Parameters used to generate aircraft test data		
Time when aircraft becomes available, between	0 - 17	hours of the 1 st day
Initial fuel status, between	52 - 87	litres
Maximum useable fuel for 0, 1, and 2 pax	348; 272; 132	litres
Minimum reserve fuel	52	litres
Flight hours remaining till maintenance, between	6 - 50	flight hours
Lease costs per Hobbs hour, between	n/a	€
Charter cost per Hobbs hour, between	n/a	€
Parameters used to in the aircraft routing system		
Maximum flight duty period	13	hours
Minimum rest period after an FDP	12	hours
Time pilots need to report for duty	60	minutes
Pilot wage per Hobs hour	50	€ per hour
Surplus fuel burn (sf_{hour})	4	% per hour
Allowance for unknown costs (a_{ii})	0.05	€ per litre
Number of closest routes to append (k)	10	-

Table 3. Parameters used to generate input test data

In addition, to test the use of persistent aircraft-trip assignments a demand uncertainty simulator is programmed. The input generator that is created to generate the test environment and corresponding data consists of two MATLAB scripts: an 'airport data generator' and a 'trip and aircraft data generator'. The airport data generator creates a virtual world containing a number of airports, i.e. the network in which a virtual per-aircraft air taxi service operates. The trip and aircraft data generator in turn creates flight requests that need to be served by that virtual operator and the data of the virtual operator's aircraft. Both the virtual world and the virtual PAAT operator are modelled on our Belgian air taxi company and its operations to ensure that the order of magnitude of the generated input data is the same as that of the environment where the system will eventually be employed.

The final parameter, k , determines the number of closest trips that are appended to a route during the route generation process. To choose the value for k , tests are conducted with the aircraft routing system with k ranging from 3 to 15. The results of these experiments are given in Tables 4 and 5. Note that each line of results that is reported in the results tables represents the average result of ten tests with independent random data sets. Furthermore the number of flight requests per day in the planning horizon is indicated in the results by j and the number of available aircraft by i . The data sets are generated using the virtual environment and virtual air taxi operator as described in this section.

1. k	2. (j, i)	3. Routes	4. ct_{MC}	5. ct_{CPLEX}	6. Solution cost
3	(7, 19)	1,729	5.3	0.01	18,649
5	(7, 19)	3,427	8.0	0.03	18,491
10	(7, 19)	8,090	15.6	0.08	18,401
15	(7, 19)	10,799	20.3	0.12	18,395
3	(15, 25)	6,086	17.7	0.13	38,469
5	(15, 25)	16,631	37.2	0.22	38,149
10	(15, 25)	72,260	177.9	1.21	37,849
15	(15, 25)	180,389	608.9	5.05	37,832

Table 4. The effect of k on the aircraft routing solution

The first column in Table 4 indicates the value of parameter k that is used for the experiments. Column 2 gives the number of flight requests per day that are generated and the number of aircraft in the fleet (excluding subcontractor aircraft). In the next column the total number of feasible routes present in the set-partitioning model is stated (thus for all the aircraft). Column 4 and 5 respectively denote the computational time in seconds of the model creator and CPLEX. Finally, in the last column the total cost of the aircraft routing plan is given. Table 5 below gives the results of the same experiments as shown in Table 4, but this time relative to the results for $k = 15$.

1. k	2. (j, i)	3. Routes	4. ct_{MC}	5. ct_{CPLEX}	6. Solution cost
3	(7, 19)	16%	27%	13%	101.4%
5	(7, 19)	33%	40%	33%	100.5%
10	(7, 19)	77%	78%	71%	100.0%
15	(7, 19)	100%	100%	100%	100.0%
3	(15, 25)	4%	3%	4%	101.8%
5	(15, 25)	10%	7%	5%	100.8%
10	(15, 25)	40%	30%	24%	100.0%
15	(15, 25)	100%	100%	100%	100.0%

Table 5. The effect of k on the aircraft routing solution (relative to $k = 15$)

It can be seen in column 3 of Table 4 that decreasing k causes the number of feasible routes that is generated by the model creator to decrease. This is because fewer of the closest trips are appended during the route generator loop. Due to the lower number of routes that are generated, both the model creator and CPLEX need less time to create and solve the set-partitioning model. This effect can be observed in column 4 and 5 of Table 5. While the number of created routes and the computational time for $k = 3$ can be as little as 4% of those for $k = 15$, the average increase in total cost of the aircraft routing solution is maximum 1.8%. When changing k from 15 to 10, the increase in solution cost is less than 0.1% while the computational time is decreased by at least 22% and in the (15, 25) case even by 70%. The latter can be explained by the fact that the number of routes that is created grows exponentially with the number of flight requests per day. A change in parameter k therefore has an increasingly larger effect on the change in the number of generated routes, and thus in computational time, when increasing the flight requests per day. In summary, varying parameter k from 15 to 3 has a large effect on the number of feasible routes that are created and, consequently, on the computational time needed by the aircraft routing system. However, the conducted experiments show that the effect on the solution cost is limited to an increase by maximum 1.8%. These findings are similar to the ones reported in the research of Ronen (2000) who varied k from 11 to 5 using 47 aircraft and 50 trips and reported a maximum cost increase of less than 0.1%. It can therefore be concluded that while having a high value of k causes the system to generate a larger number of feasible routes and thus possible solutions, the chosen aircraft routing solution mostly contains those routes that minimize the repositioning cost. In other words, when creating an aircraft routing it is a good rule of thumb to only consider those trips that are close to the arrival airport of the previous trip. As shown by the results in Table 5 only considering the five trips for which the departure airport is closest to the arrival airport of the previous trip, causes the routing solution to be at most 0.8% more costly compared to when considering the fifteen closest ones. Because changing k from 15 to 10 only increases the solution cost by less than 0.1% while reducing the solution time sometimes by as much as 70%, a k value of 10 is adopted as the standard value.

When creating the virtual airport network, demand and PAAT operator it is attempted to create input data for the aircraft routing system that has the same order of magnitude of the expected real input data. However, there are a number of differences between the virtual test environment and the air taxi operator expected operations. First of all, the virtual test environment uses a square airport grid in which airports are uniformly distributed instead of the actual locations of the airports in Western Europe. In addition, the airport charges the virtual environment uses are therefore not the ones encountered in reality. Also related to the airports is the fact that the turnaround time is taken to be 45 minutes for all airports. In reality the turnaround time depends on the airport.

The trip data generator randomly chooses a departure and arrival airport for all the trips. It therefore does not simulate demand hot spots, i.e. airports where a lot of demand originates or arrives. In reality, certain airports do attract significantly more business flight traffic than others. For example, Geneva Cointrin airport, which is the second busiest business airport in Europe, has almost 50% more business flight departures than the fifth busiest airport, Nice (Eurocontrol, 2010). In addition, the trip data generator only uses the hourly business flight departure pattern of France. In reality, however, the hourly departure pattern is dependent on the country the flight departs from and on the month and the day of the week (Eurocontrol, 2010). Also the trip data generator creates few day returns, i.e. customers that

fly from A to B in the morning and return from B to A in the evening. Furthermore, the data generator only earmarks aircraft unavailable for a maximum of a single day. However, it is possible that the aircraft is unavailable for a longer period of time. Yao et al. (2008) report that a FMC with 35 aircraft encountered 49 mid-day unscheduled maintenance events during a one-month period. It is therefore probable that unplanned maintenance will also occur during air taxi operator's operations. However, this virtual environment does not account for this.

It is apparent from the above that there exist differences between the virtual test environment that is created and the expected operational environment of the air taxi operator. However, as explained before, the virtual test environment is only used in this project to test and validate the workings of the aircraft routing system, and not for example to find out what the exact effect of the fleet size is on the air taxi operations. Therefore, it is of main importance that all the necessary input elements (like airport turnaround times, owner use trips and coupled requests) are specified in the virtual environment and that these have the same order of magnitude as in reality. As such it can be assessed whether the aircraft routing system is capable of handling all these inputs correctly. It is in this stage not important that the virtual test environment does exactly represent reality. The main aim is the development of the aircraft routing system and therefore the virtual test environment described in this section is adequate for its purpose, namely validating the aircraft routing system.

6.2 Validating the model

Validation is an important step in model development because through validation the level of confidence in and the credibility of the aircraft routing system is established. To carry out the validation in this project five types of validation are conducted, namely conceptual model validation, computerized model verification, data validation, experimental validation and operational validation (Landry et al., 1983; Sargent, 2007).

In the **conceptual validation** phase, in which it is determined whether the underlying theories and assumptions of the system are correct, face validation and desk checking are the used validation techniques. It is argued that the theory of the aircraft routing system is indeed correct. However, there are some assumptions of which the exact effect on the accuracy of the conceptual model could not be determined as no real operational data is available, namely the assumptions that a fixed percentage of ferried fuel is burned per hour, that the Hobbs time and fuel use between airport pairs is specified for the one passenger case, and that it is acceptable to impose crew constraints on the aircraft. It is therefore advised that these assumptions and their effect are checked as soon as real operational data is available. Next is the **computerized model verification**. This process ensures that the implementation of the conceptual model in programming code is correct. To make sure this holds true for the MATLAB implementation of the aircraft routing system the authors used bottom-up development and testing, debugging and desk checking. In addition, experiments show that the results produced by the routing system are consistent. Next, **data validation** is performed to ensure that the data used to build and evaluate the model is adequate and correct. For building the model the mental and written databases provided by the air taxi operator and found in ODAT literature are considered appropriate. In addition, the virtual test environment and corresponding numerical input data created to support the evaluation of the aircraft routing system is

deemed adequate for its purpose. **Experimental validation** aims to check whether or not the model contradicts qualitative, expert knowledge. To this end both the convergent validation and face validation techniques are utilized. For the first technique the results produced by the aircraft routing system are checked against six expert statements. In all cases the results of the routing system adhered to the expert knowledge. In addition face validation is carried out with Fly Aeolus' director. He is given an example aircraft routing problem to assess whether the solution produced by the routing system is correct, which he indeed found to be true. Finally, **operational validation** assesses the usefulness and timeliness of the solutions and cost of implementing the system. To carry out the routing plans produced by the aircraft routing system the air taxi operator must first decide on and then apply the waiting strategy for the routing plans. Furthermore, as mentioned in the conceptual validation, it must be assessed what the exact effect on the accuracy of the routing plans is of the fuel ferrying, and Hobbs time and fuel use assumptions before implementing the routing system. With respect to the timeliness it is found that the aircraft routing system is capable of solving routing problems with sizes that correspond to air taxi operator's first operational years in under five minutes. Noteworthy in this respect is that the routing system is capable of producing a solution to a routing problem in little over one second that takes a human dispatcher 30 minutes to solve.

7. Discussion & conclusions

In this section important limitations of the work are discussed and conclusions are drawn.

7.1 Discussion

As this aircraft routing system has been developed for an air taxi operator who – at time of writing – is not yet fully operational, no real operational input data can be used for system validation and testing and instead a virtual test environment that approximates the air taxi operator's operations is utilized. Therefore, this research has some limitations.

The air taxi operator needs to check some assumptions underlying the routing system before putting it to use. The exact effect of these assumptions on the routing plans could not be assessed in this project because no real operational data is available. First, it needs to be assessed whether the increase in routing costs caused by imposing crew constraints upon the aircraft is acceptable. In addition, the air taxi operator must check whether specifying Hobbs time and fuel use between airport pairs for the one passenger load case is accurate enough. Finally, the relation to reality of the assumption that a fixed percentage of ferried fuel is burned per hour must be checked. Once operational data becomes available, the air taxi operator is advised to check these assumptions. This is deemed a necessary step in the implementation of the aircraft routing system. However do note that with respect to the Hobbs time and fuel use assumption, the air taxi operator can also opt to remove this assumption from the system altogether by specifying these input values for each passenger load case.

The aircraft routing plans created with the aircraft routing system do not contain a waiting strategy for repositioning flights, only the earliest and latest start time for deadheads are indicated. If the air taxi operator wants to adopt a 'fly-first' or 'wait-first' waiting strategy, this can be directly applied to the created routing plans and as such this limitation is easily removed from the system. However, the authors suggest that the aircraft routing system is

augmented with a module which calculates the optimal waiting strategy per aircraft route instead of adopting the 'fly-first' or 'wait-first' approach as the latter two are suboptimal approaches. Once the waiting strategy is determined, also aircraft parking fees, crew waiting wages, and airport opening hours can be incorporated into the aircraft routing system, as these all depend on the strategy that is adopted.

With respect to the secondary planning objectives, the degree of robustness, flexibility and persistence of the aircraft routing plans are not measured. The developed system does allow the OCC to control these parameters, but only to a limited extent. Robustness is solely controlled by slack time, and no other means of enforcing robustness are currently provided. The same holds for persistence and flexibility, whose balance is only altered by using the persistent aircraft-trip assignments. But the aircraft routing system does not provide information as to exactly how robust, flexible or persistent the created routing plans are. Though the dispatcher knows he is increasing/decreasing robustness, flexibility or persistence when he is adapting the corresponding control parameters. The authors would like to further explore and implement other control parameters to enforce these three secondary planning objectives of an aircraft routing plan in an on-demand air transport (ODAT) context.

7.2 Conclusions

This paper has aimed to be the first study that provides a detailed description of the aircraft routing problem as faced by a per-aircraft air taxi (PAAT) operator and also the first study that treats the development of an aircraft routing system for a per-aircraft air taxi operator. The developed aircraft routing system, which consists out of the model creator, CPLEX and the solution module, is capable of solving the aircraft routing problem faced by the air taxi operator in under five minutes for the first operational years. This will allow the air taxi operator to achieve an estimated cost reduction of 12% on their routing plans with respect to using a human dispatcher. This amounts to estimated savings of 10 % on variable direct operating costs during the air taxi's first operational year. By altering the slack time parameter the OCC can control the robustness, i.e. the capability to cope with and recover from external disruptions, of the aircraft routing plans it creates. Furthermore, via the persistent aircraft-trip assignments input, the system allows the OCC to control the balance between persistence, the degree to which the aircraft routing plan deviates from the previous one when it is re-optimized, and flexibility, the ability of the schedule to cost-effectively deal with future demand, of the generated aircraft routing plans. In addition, by setting aircraft-trip assignments as input the OCC can specify 'hard' input decisions. The aircraft routing system also allows the air taxi operator to incorporate fuel ferrying in the aircraft routing phase as a strategy to reduce the operational cost of the aircraft routing plan. The next stage for our project is for our Belgian air taxi operator to become fully operational, further develop the routing system and eventually implement it at their OCC. Then the various databases in the aircraft routing system can be filled with real life operational data and help the air taxi operator to provide affordable and reliable full on-demand personal air transport.

8. References

- Abdelghany, A. & Abdelghany, K. (2009). *Modeling Applications in the Airline Industry*, Ashgate, ISBN 978-0-7546-7874-8, Farnham, United Kingdom

- American Express (2009). 2010 Global business travel forecast. In: *American Express*, 02.11.2011, Available at:
<http://corp.americanexpress.com/gcs/travel/news/pressreleases.aspx>
- Ball, M., Barnhart, C., Nemhauser, G. & Odoni, A. (2007). Chapter 1 Air Transportation: Irregular operations and control, In: *Transportation. Vol. 14 of Handbooks in Operations Research and Management Science*, Barnhart, C. & Laporte, G. (Eds.), pp. 1 - 67, Elsevier, ISBN: 978-0-444-51346-5, Amsterdam, the Netherlands
- Barnhart, C., Belobaba, P. & Odoni, A. (2003). Applications of operations research in the air transport industry, *Transportation Science*, Vol. 37, No. 4, pp. 368 - 391, ISSN: 0041-1655
- Bazargan, M. (2004). *Airline Operations And Scheduling*, Ashgate Publishing, ISBN 978-0-7546-9772-5, Aldershot, United Kingdom
- Bian, F., Burke, E., Jain, S., Kendall, G., Koole, G., Landa Silva, J., Mulder, J., Paelinck, M., Reeves, C., Rusdi, I., Suleman, M. (2003). Making airline schedules more robust. In: *Proceedings of the 1st Multidisciplinary International Conference on Scheduling: Theory and Applications (MISTA 2003)*, pp. 678-693
- Bonnefoy, P. (2005). Simulating air taxi networks. *Proceedings of 37th Winter Simulation Conference*, pp. 1586 - 1595, ISBN: 0-7803-9519-0
- Budd, L. & Graham, B. (2009). Unintended trajectories: liberalization and the geographies of private business flight. *Journal of Transport Geography*, Vol. 17, No. 4 (July 2009), pp. 285 - 292, ISSN: 0966-6923
- Bunte, S. & Kliewer, N. (2010). An overview on vehicle scheduling models. *Public Transport*, Vol. 1, No. 4, (March 2010), pp. 299 - 317
- Cirrus (2007). Cirrus SR22 weight and balance loading form, In: *Airplane information manual for the Cirrus Design SR Cirrus, 04.05.2011*, Available on:
<http://servicecenters.cirrusdesign.com/TechPubs/pdf/POH/sr22/pdf/20880-001InfoManual.pdf>
- Doganis, R. (2002). *Flying Off Course*, 3rd Edition, Routledge, ISBN: 0-415-21323-1, New York
- Dyson, E. (2006). Visible demand: The new air-taxi market, In: *Fly Your Jet website*, 17.11.2010, Available at:
<http://www.flyyourjet.com/Portals/0/DysonRelease.pdf>
- Espinoza, D., Garcia, R., Goycoolea, M., Nemhauser, G. & Savelsbergh, M. (2008a). Per-seat, on-demand air transportation part i: problem description and an integer multicommodity ow model. *Transportation Science*, Vol. 42, No. 3, (August 2008), pp. 263 - 278, ISSN: 0041-1655
- Espinoza, D., Garcia, R., Goycoolea, M., Nemhauser, G. & Savelsbergh, M. (2008b). Per-seat, on-demand air transportation part ii: Parallel local search. *Transportation Science*, Vol. 42, No. 3, (August 2008), pp. 279 - 291, ISSN: 0041-1655
- Eurocontrol (2010). Business aviation in Europe 2009, In: *Eurocontrol website*, 04.05.2010, Available at: <http://www.eurocontrol.int/statfor/public/subsitehomepage/homepage.html>

- Fagerholt, K., Foss, B. A. & Horgen, O. J. (2009). A decision support model for establishing an air taxi service: A case study. *Journal of the Operational Research Society*, Vol. 60, No. 9, (September 2009), pp. 1173 – 1182, ISSN: 0160-5682
- Feillet, D. (2010). A tutorial on column generation and branch-and-price for vehicle routing problems. *4OR, A Quarterly Journal of Operations Research*, Vol. 8, No. 4, (June 2010), pp. 407 – 424, ISSN: 1619-4500
- Fly Aeolus (2010). Private Jet Price Private Plane Charter Prices, In: *Fly Aeolus website*, 03.05.2011, Available at <http://www.flyaeolus.com/affordable-and-low-price/>
- Flyer Forums (2010). Avgas price table, In: *Flyer Forums website*, 25.10.2010, Available at: <http://forums.yer.co.uk/viewtopic.php?t=51535>
- Gall, C. & Hindhaugh, J. (2009). Communicating change to sr. management economic & industry outlook 2010, In: *National Business Travel Association website*, 12.07.2010, Available at: <http://www2.nbta.org/Lists/Resource%20Library/F103%20Communicating%20Change%20-%20Industry%20Outlook.pdf>
- Hicks, R., Madrid, R., Milligan, C., Pruneau, R., Kanaley, M., Dumas, Y., Lacroix, B., Desrosiers, J. & Soumis, F. (2005). Bombardier flexjet significantly improves its fractional aircraft ownership operations. *Interfaces*, Vol. 35, No. 1, (January – February 2005), pp. 49 – 60, ISSN: 0092-2102
- Hillier, F. S. & Lieberman, G. (2008). *Introduction to Operations Research*. 8th Edition, McGraw Hill, ISBN: 007-123828-X, New York
- Karaesmen, I., Keskinocak, P., Tayur, S. & Yang, W. (2005). Scheduling multiple types of time-shared aircraft: Models and methods for practice. In: *Proceedings of the 2nd Multidisciplinary International Conference on Scheduling: Theory and Applications*. Kendall, G., Lei, L. & Pinedo, M. (Eds.), New York (July 2005), Vol. 1, pp. 19 – 38
- Karaesmen, I., Yang, W. & Keskinocak, P. (2007). Optimization vs. persistence in scheduling: heuristics to manage uncertainty in on-demand air travel. In: *Proceedings of the 3rd Multidisciplinary International Conference on Scheduling: Theory and Applications*, Paris (August 2007), Vol. 1., pp. 242 - 250
- Keskinocak, P. & Tayur, S. (1998). Scheduling of time-shared jet aircraft, *Transportation Science*, Vol. 32, No. 3, pp. 277 – 294, ISSN: 1526-5447
- Keysan, G., Nemhauser, G. & Savelsbergh, M. (2010). Tactical and operational planning of scheduled maintenance for per-seat, on-demand air transportation. *Transportation Science*, Vol. 44, No. 3, (August 2010), pp. 291 – 306, ISSN: 0041-1655
- Landry, M. & Muhittin, O. (1993). In search of a valid view of model validation for operations research. *European Journal of Operational Research*, Vol. 66, No. 2, (April 1993), pp. 161 – 167, ISSN 0377-2217
- London Executive Aviation (2010). In: *London Executive Aviation website*, 23.11.2010, Available at: <http://www.flylea.com/>
- Mane, M. & Crossley, W. A. (2007a). An approach to predict impact of demand acceptance on air taxi operations. In: *7th AIAA Aviation Technology, Integration, and Operations Conference*. Vol. 1, Belfast, Northern Ireland, pp. 954 – 961

- Mane, M., & Crossley, W. A. (2007b). Probabilistic approach for selection of maintenance facilities for on-demand operations. In: *7th AIAA Aviation Technology, Integration, and Operations Conference*. Vol. 1. Belfast, Northern Ireland, pp. 944 - 953
- Mane, M. & Crossley, W. A. (2009). Importance of aircraft type and operational factors for air taxi cost feasibility. *Journal of Aircraft*, Vol. 46, No. 4, pp. 1222 - 1230, ISSN: 0021-8669
- Martin, C., Jones, D. & Keskinocak, P. (2003). Optimizing on-demand aircraft schedules for fractional aircraft operators. *Interfaces*, Vol. 33, No. 5, (September - October 2003), pp. 22 - 35, ISSN: 0092-2102
- NBAA (2010), Piston Engine Aircraft, In: *National Business Aviation Association website*, 03.05.2011, Available at:
<http://www.nbaa.org/businessaviation/aircraft/pistons/>
- Ronen, D. (2000). Scheduling charter aircraft. *The Journal of the Operational Research Society*, Vol. 51, No. 3, (March 2000), pp. 258 - 262, ISSN: 0160-5682
- Sargent, R. (2007). Verification and validation of simulation models. *WSC '07 Proceedings of the 39th conference on Winter simulation: 40 years! The best is yet to come*, IEEE Press Piscataway, NJ, pp. 124 - 137, ISBN: 1-4244-1306-0
- Silver, E.A. (2004). An overview of heuristic solution methods. *The Journal of the Operational Research Society*, Vol. 55, No. 9, (May 2004), pp. 936 - 956, ISSN: 0160-5682
- Smith, J. C. & Taşkin, Z. C. (2008). A tutorial guide to mixed-integer programming models and solution techniques. In: *Optimization in Medicine and Biology*. Lim, G. J., Lee, E. K. (Eds.), pp. 521 - 548, Taylor and Francis, Auerbach Publications, ISBN: 978-0-849-305-634. New York
- Yang, W., Karaesmen, I., Keskinocak, P. & Tayur, S. (2008). Aircraft and crew scheduling for fractional ownership programs. *Annals of Operations Research*, Vol 159, (December 2007), pp. 415 - 431, ISSN: 1572-9338
- Yang, W., Karaesmen, I. Z. & Keskinocak, P. (2010). Managing uncertainty in on-demand air travel. *Transportation Research Part E: Logistics and Transportation Review*, Vol. 46, No. 6, (November 2010) p. 1169 - 1179, ISSN: 1366-5545
- Yao, Y., Zhao, W., Ergun, O. & Johnson, E. (2005). Crew Pairing and Aircraft Routing for On-Demand Aviation with Time Window. In: *Social Sciences Research Network eLibrary*, 04.05.2011, Available from: <http://ssrn.com/abstract=822265>
- Yao, Y. & Zhao, W. (2006). Optimization problem in fractional ownership operations. *Collection of Technical Papers - 6th AIAA Aviation Technology, Integration, and Operations Conference*, Vol. 1, pp. 319 - 323
- Yao, Y. (2007). *Topics in fractional airlines*. Georgia Institute of Technology, School of Industrial and Systems Engineering, PhD thesis, May 2007
- Yao, Y., Ergun, O., Johnson, E., (2007). Integrated model for the dynamic on-demand air transportation operations. In: *Dynamic Fleet Management*. Sharda, R., Vo, S., Zeimpekis, V., Tarantilis, C. D., Giaglis, G. M. & Minis, I. (Eds.), Vol. 38 of Operations Research/Computer Science Interfaces Series, pp. 95 - 111, Springer US, New York, ISBN 978-0-387-71721-0.

Yao, Y., Ergun, O., Johnson, E., Schultz, W. & Singleton, J. M. (2008). Strategic planning in fractional aircraft ownership programs. *European Journal of Operational Research*, Vol. 189, No 2, (September 2008), pp. 526 - 539, ISSN: 0377-2217