Optimizing cost effectiveness and flexibility of air taxis

A case study for optimization of air taxi operations

Author:
Bart de Jong, BSc.
bartdejong@mac.com

Thesis MSc. Business Information Technology
Haarlem, September 2007

Etirc Aviation
Dreef 32
2012 HS Haarlem
The Netherlands

Supervisor
Drs. Teun Mentzel

University of Twente
Drienerlolaan 5
7500 AE Enschede
The Netherlands

First supervisor
Dr. Pascal van Eck

Second supervisor
Prof. Ir. Roel Pieper

Management Summary

In the near future, Etirc Aviation will provide air taxi operators with solutions for cost-effective fleet and network management. This study has researched the requirements for optimization of air taxi operations and techniques for the scheduling of air taxi services. On the basis of this research, this study has developed and tested an algorithm that implements a cost-effective scheduling solution for single operators.

The optimization of air taxi operations requires both cost effectiveness and flexibility of operations. The optimization of cost effectiveness and flexibility is the combination of optimization of the air taxi network, the use of flexible time windows for flight scheduling and the use of an efficient algorithm for dispatching aircraft. The dispatch process is a dynamic problem. The air taxi business requires that the effect of schedule dynamics on operations is minimized. With these requirements, the air taxi optimization problem can be best solved through a dynamic programming approach. A dynamic programming algorithm was developed for this study. In simulations of a real-life case study, this algorithm was able to provide cost optimization, support for time windows and produce a stable operational schedule.

For a selected case study, the algorithm has proven to be more cost effective and flexible than the currently best alternative. Under normal operating conditions, an average 5.5% cost advantage is offered on the cost of empty legs. The addition of a two-hour time window in operations increases the customer acceptance rate with 2-3%. This flexibility can be used to decrease the fleet size, which saves further cost. Most importantly, the developed algorithm is better able to handle unforeseen disruptions and minimizes their negative effects. The algorithm can perform 60% of disrupted flights, instead of the 25% offered by current scheduling solutions.

The developed algorithm is does not guarantee an optimal solution. The algorithm sacrifices 5 to 10% performance in cost optimality. In return for sacrificing some cost optimality, the algorithm offers air taxi operators the operational stability and flexibility that are required in the air taxi market.

Preface

At the time of writing, it is exactly one year ago that I started to create plans for my graduation assignment. I spoke to Mr. Pieper about my wish to graduate in a startup company with innovative services. Through Mr. Pieper, I got into contact with Etirc and Mr. Mentzel. I was lucky to be asked for the project that had caught my eye: Aviation. The assignment required me to move to Haarlem. It is now been half a year since my moving to Haarlem. I will be moving out of Haarlem soon, but I am thinking about returning sometime in the future for a longer period.

I requested Mr. Van Eck and Mr. Pieper to be supervisors for my graduation assignment. I am very grateful that they accepted my request and for all the guidance and advice they have provided me with.

It took me a bit more time to select a focus in my research than had I hoped. The amount of research in the field of transportation and optimization is huge and diverse. With the help of Mr. Mentzel, I was able to focus on the business problem and select a research focus. I would like to thank Mr. Mentzel for introducing me to the air taxi business, for all his critical and practical advise and for the enjoyable time at Etirc.

On a personal level, I wanted the results of my work to be both useful for further research and well as directly practical for Etirc. The combination of a real-life case study and the creation of a prototype solution offered me that opportunity. The process of creating a prototype was time consuming and frustrating at times. However, the direct results made the effort rewarding.

In the end, I am glad to say that I have delivered the deliverables that I planned in the time I originally planned for it. I enjoyed working with my colleagues at Etirc and I have felt challenged during my assignment. There is still a lot to be done in the air taxi business. I hope that the end-results of my work will be a valuable contribution for Etirc and the future of all air taxi businesses.

Bart de Jong, July 26, 2007

Contents

List of figures

1 Introduction

The year 2006 saw the introduction of the first certified Very Light Jet (VLJ). These Very Light Jets have a take-off weight of less than 4540 kilograms. The introduction of these aircraft is creating a new market for quick and affordable point-to-point transport. This has also created a new market for air taxi services. In this market, Etirc will unroll a service network for Air taxi operators. To be an attractive business partner for air taxi operators, Etirc will offer services that will strengthen the operational efficiency of the operators.

This document is the MSc. Thesis of Bart de Jong. The research for this thesis was performed at Etirc in Haarlem, between January 15th and July 30th. The result of this research is an advise on a technique to optimize the cost-effectiveness and flexibility of air taxi operators.

1.1 Air taxi business model

The value proposition of an air taxi operator for its customer is the value of time saved. The door-to-door travel time is shortened by air taxis. Firstly by flying to smaller airports that are closer to the customer's origin and destination, thereby decreasing the travel time to an airport (see Figure 1). Secondly, by offering a shorter boarding time. Due to the smaller volume of travellers and lesser security checks, the boarding process is faster than with traditional airlines. Denial-of-service is kept to a minimum. More than 95% of customer requests will be accepted and serviced.

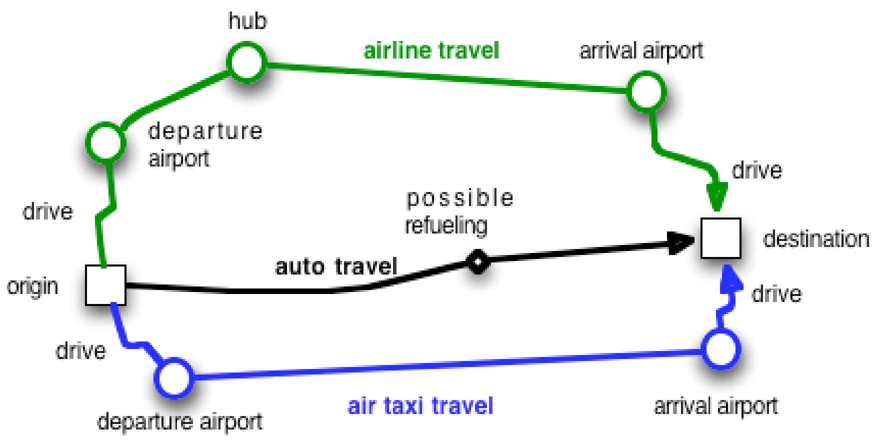

Figure 1 - trip composition for airline, automobile and air taxi travel

Adapted from Mane, 2006

The target customers for the air taxi services are business travellers. Air taxi operators fill a gap between slow auto travel, affordable but rigid business class flights and expensive but flexible charter flights. Air taxis offer full service with an arranged pilot and pickup, while air charters require the customers to arrange these themselves. Overall, the drivers for air taxi services are related to time, cost and travel convenience offered to the traveller.

Travellers can only be serviced if the route that the traveller wishes to travel is in the network that is operated by the operator. The services are distributed to the customer on-demand. In the point-to-point and on-demand business model, the customer requests are considered for acceptance on a First-in-First-Out basis.

The revenue model of air taxi operators is based on leg price. The main factor that will shape air taxis' ability to compete with business aviation, car and train travel will be the cost of airfares. This places a price pressure on the airfares. To be cost-effective, an air taxi operator has to create a flexible and cost-efficient configuration of fleet, crew and covered airports.

The reduced cost of ownership for VLJs drives many initiatives to enter the air taxi market. This will probably lead to a diversification of networks, services and customer segments. In the long term, more competition in the market of personal air-based business travel will probably lead to a further pressure on the price of airfares. Only those air taxi operators that manage their fleet and network effectively will be strong enough to survive the competition.

1.2 Etirc Aviation

Etirc Aviation (Etirc) operates in the domain of personalized business flights in both Europe and Russia. The customer base of Etirc is formed by air taxi operators that are operating Very Light Jets (VLJ). Etirc is developing a business that offers cost-effective fleet and network management for air taxi operators. Current air taxi operations typically use 70 to 80% of their capacity. Etirc will offer fleet and network solutions that improve cost effectiveness. The added value of these services for air operators is the fact that the solutions will lower their operational cost and strengthen the operators in surviving the expected price pressure on airfares.

Etirc has organized its services in four key business divisions:

1. **Asset Management**
   Providing financial services like fleet financing services, aircraft leasing and tax optimization;
2. **Fleet services**
   Offering a real-time scheduling and operations system, enabling rapid pricing;
3. **Crew services**
   Services for handling human resource matters, including contracting pilots and payroll services;
4. **Jet services**
   Offering maintenance facilities for VLJs and pilot training facilities.

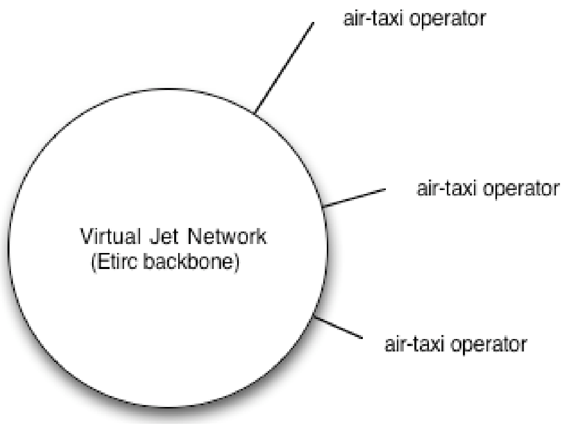

The strategy of Etirc is the creation of a Virtual Jet Network (VJN) through a "Backbone". This Backbone will create a VJN as an ASP service. The primary function of the Backbone is to offer cost-optimizing tools to air taxi operators. Air taxi operators will be connected to the services of this backbone, see Figure 2.

Figure 2 - Etirc backbone

The strategy and execution of both Etirc and its customers are summarized in Figure 3. It is visible that air taxi operators focus on customers, while Etirc focuses on cost optimization of supporting services and on offering these services in a network of operators.

	Etirc	Air taxi operators
strategy	development of a Virtual Jet Network	personalized business travel
execution	employee training, crew pool, aircraft lease/finance aircraft maintenance service backbone	operating daily flights customer contact booking & billing

Figure 3 – activities of Etirc and air taxi operators

1.3 Document structure

This document researches cost and flexibility optimizing solutions for air taxi operators and tests a selected solution technique. This introduction has defined the context. The second chapter define the research problem and research approach. The third chapter focuses on air taxi operation in literature, while the fourth chapter creates a model specifically for air tax operations. The operational model is turned into a cost model through chapters 5 and 6.

After formulating a cost model, the focus is shifted to cost optimization. Chapter 7 and 8 research optimization solutions for the air taxi problem and select the most suitable optimization technique. Chapter 9 describes the creation of an algorithm prototype and defines tests that where performed to measure its performance. Chapter 10 discusses the results of the tests. Finally, chapter 11 writes the conclusions and advises Etirc on optimization of air taxi operations.

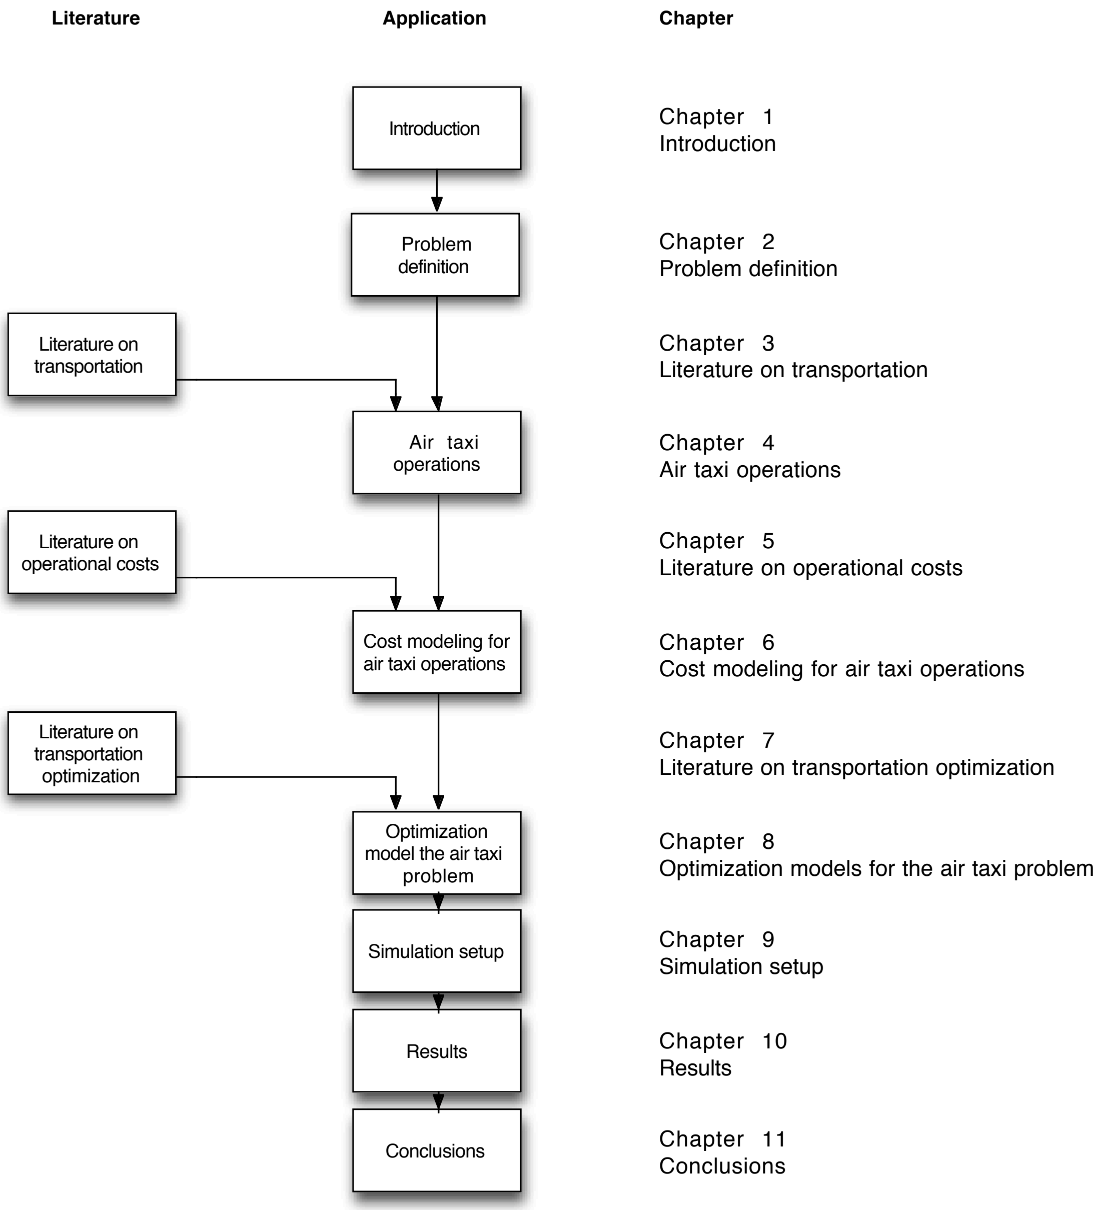

2 Problem definition

2.1 Project context

A demand for services that optimize the cost-effectiveness of operations for air taxi operators is expected. Etirc plans to offer these services to air taxi operators in Europe and Russia and needs to know what the cost factors in air taxi operations are and how these cost factors can be optimized.

2.2 Research objective

The objective of this research is to find one or more optimization techniques that improve the cost-effectiveness and flexibility of air taxi operators and to prove the effectiveness of those techniques in a case study.

2.3 Problem statement

Air taxis have existed for some years, but in many different forms. For the purpose of this study, an air taxi is defined as a personal means of business travel that uses Very Light Jets (VLJ) to provide on-demand flight services. An air taxi is operated by an air taxi operator, which owns and manages several VLJs over a network of airports. An operator is driven by the demand in the demand area that is covered by its airport network. The system of an air taxi operator is depicted as a black box in Figure 4.

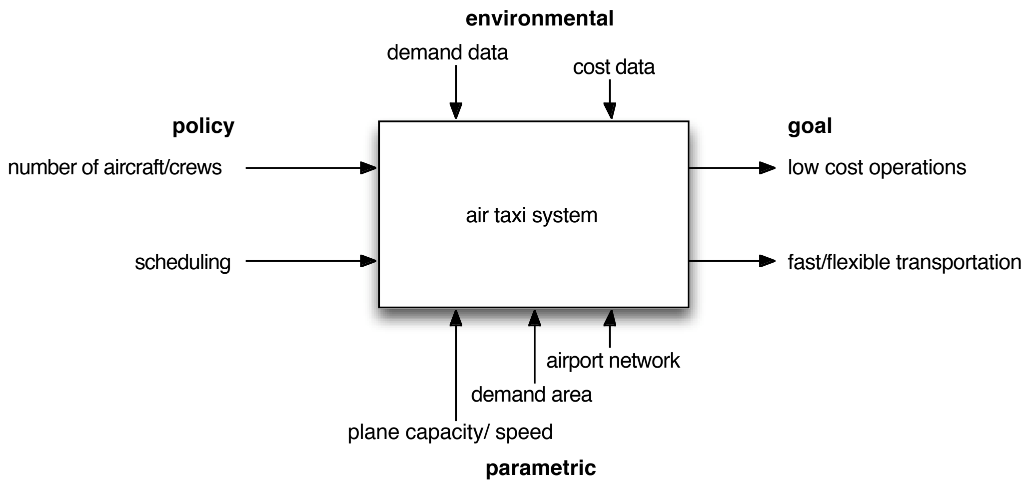

Figure 4 - black box of an air taxi operator

Adapted from Bailey, Clark, 1992

The costs of an air taxi operator consists of a fixed overhead and incremental operational costs. The number of revenue flights determines the revenue of an air taxi operator. There are two basic pricing methods for a revenue flight: per flight leg, or per seat. In the early stages of air taxi services, it is assumed that flights are sold on a per leg basis. The demand is considered exogenous and fixed. With a fixed pricing system and a fixed demand volume, the revenue is also a fixed value. Cost optimization has to be performed purely on the cost side.

The scheduling policy is considered the only variable in the air taxi system. The scheduling solution of an air taxi operator determines its ability to offer cost-efficient and flexible services. It is not known what kind of scheduling solutions are available for air taxi systems and what scheduling solution is the most suitable.

2.4 Research questions

The central research question for this master-thesis is:

What are the operational costs of air taxi operations, which factors determine these costs and what is an effective scheduling solution for the minimization of these costs through services oriented at air taxi operators?

This central research question can be split into four sub-questions:

1. What factors form the variable operational costs of a single air taxi operator?
2. How can the cost minimization problem be modelled?
3. What optimization techniques can be used to create optimal solutions?
4. Which solutions can Etirc implement for these optimizations and how can these solutions be customized?

2.5 Research method

The research objective is maximizing the cost-effectiveness of air taxi operators. The nature of the research is diagnostic practice-oriented. The effectiveness of the selected solution is tested through simulations. The research focuses on the cost structure of air taxi operations and on optimization techniques for operations scheduling. The resulting research framework is shown in Figure 5.

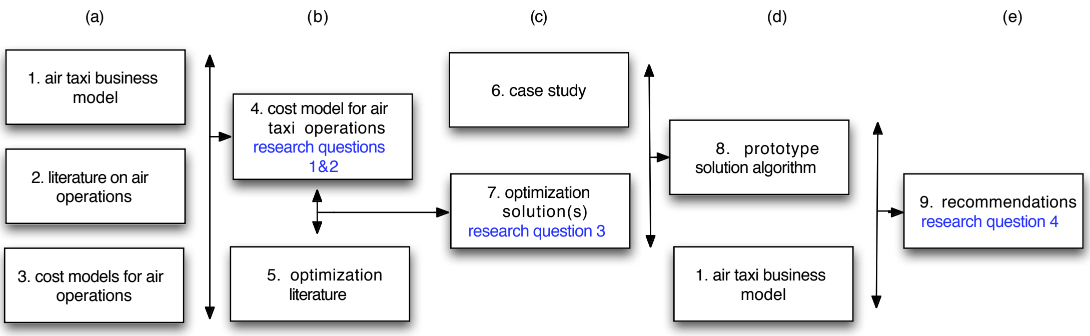

Figure 5 - research tree

(a) An analysis of the air taxi business model, air operations and costs of air taxi operations results in (b) a cost model for air taxi operations. This cost model is confronted with research on transportation optimization. This will results in (c) a selection of the most suitable optimization technique. The confrontation of a prototype based on the selected optimization technique with a case model in a simulation will result in prototype solution for the utilized case study. The performance of the prototype is confronted with the requirement from the air taxi business model. This results in (e) recommendations for the optimization of air taxi operations.

3 Literature on transportation

The introduction chapter has described the business model of air taxi operators. Cost effectiveness and flexibility are important qualities in the air taxi business. Optimizing of cost and flexibility requires an optimization model. However, at this stage, Etirc has no detailed operational models or cost models for air taxis. These models have to be created. This chapter surveys literature on demand for transportation, air operations and methods of modelling demand and operations.

3.1 Definitions

The field of transportation has several terms that are used throughout literature and practice. The term *trip* is frequently used in the field of transportation. A trip is defined as “a single period of non-stop travel from one point to another”. In the field of aviation, a trip is often called a *leg*. A trip (or leg) is often a paid service towards a customer. When a trip is paid for by a customer it is called a *revenue trip*. Likewise, trip that are not paid for by a customer are *non-revenue trips*. In the field of aviation, a non-revenue trip is also called an *empty leg*. An empty leg is an important concept for cost optimization.

3.2 Demand for transport

The size of transport operations is based on the demand for transport. The specific demand for flight services can be determined from the total demand for transportation. The total demand for transport is mostly dependent on socio-economical factors like income and the distance between home and work. The total demand for transport is often estimated as a total number of trips starting from a single region. The total number of trips can be distributed to trips between individual zones. For each of these zone-trips, the number of trips using an individual transportation mode is estimated. This is depicted in Figure 6 (adapted from Teodorovic, 1988).

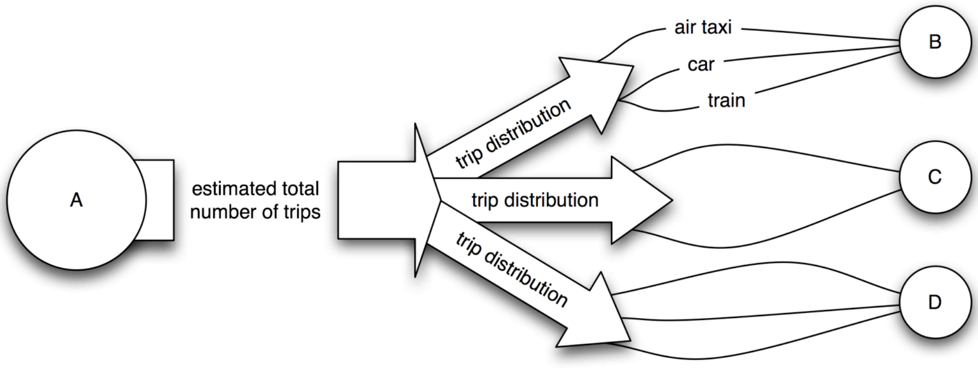

Figure 6 - demand forecast methodology

The trip-distribution model has suffered criticism. The model can be computationally expensive. The model does not include daily and seasonal demand fluctuations (Isikveren, 2002, Bonnefoy, 2005, Lee et al. 2006) and is based on the assumptions that each trip can be analyzed independently from other trips and that the time-of-day dimension can be ignored (Kitamura et al, 2000).

3.2.1 Customer request

The previous paragraph showed that total demand can be split into individual trips or 'requests for transport'. In the air taxi domain, such a trip is called a *customer request*. Each customer request consists of (Bonnefoy 2005):

- An origin airport;
- a destination airport;
- a requested departure time.

A customer request is sent in at a specific time. The request is placed hours, days or months ahead of the requested departure time. The period between the request being entered and its requested departure time is called the *schedule window*.

3.2.2 Time windows

Often, a request for transport has no fixed departure time but is set for a certain time window with an earliest acceptable departure and latest acceptable arrival time (Teodorovic 1988, Cordeau 2004, Savelsbergh 2005). For example: a customer may express the wish that his departure time is somewhere between 1 PM and 3 PM. This constitutes a two-hour time window. Time windows create more flexibility when flights are scheduled, but increase the complexity of the flight scheduling process.

3.2.3 Cyclic demand influences

The demand for transport suffers from seasonal influences. Demand also fluctuates during the day. In airline traffic, the demand for transport can be modelled as two peaks a day (Teodorovic, 1988). There is a demand-peak in the morning and a second demand-peak in the early evening. The demand as a function of time of day is depicted in Figure 7.

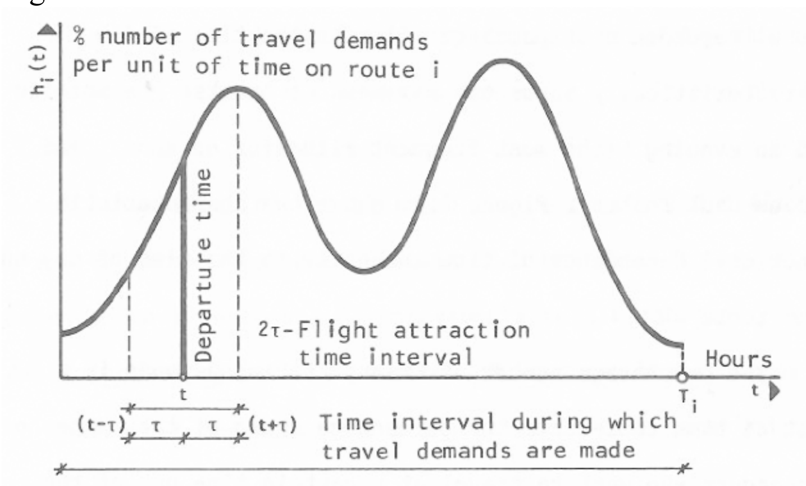

Figure 7 - demand as a function of the time of day

Source: Teodorovic, 1988

3.2.4 Demand simulation models

Due to the complexity of socio-economical demand estimations, very few researchers (Wong et al. 2003) use a socio-economical demand model. Simplified demand models are often applied. In situations where demand is highest just before the moment of transportation, as is the case with urban taxis, an exponential distribution is useful to model demand (Bailey and Clark, 1992). For more complex demand distributions, a Poisson function can be used (Schroeter, 1983).

Most research into transportation optimization utilizes a discrete event model to simulate reality. A discrete event model handles demand on a per request basis. This creates a realistic simulation, but it has the disadvantage that it is computationally complex for larger problems. A discrete event model can be simplified to form an aggregate model that aggregates the effects of all events in a period (Lee et al., 2006). The type of demand problem determines the type of simulation model that is best. Table 1 list the characteristics of demand problems and the applicable simulation techniques.

	Deterministic	**Stochastic**
Static	- logics - algebraic	Monte Carlo methods
Dynamic: discrete event	- mechanistic algorithms	- queuing theory - discrete event simulation
Dynamic: continuous event	differential equations	stochastic differential equations

Table 1 - study methods for different demand types

Source: Jurkat, M.P., Discrete Event Simulation Systems

3.3 Modelling and optimizing air operations

Air taxi operations resemble airline operations. Both perform cost optimizations. However, the primary focus of airliners is load maximization, while air taxis are focused on minimization of empty legs. Before describing the specific characteristics of air taxis, this paragraph will describe the basics of optimizing air operations.

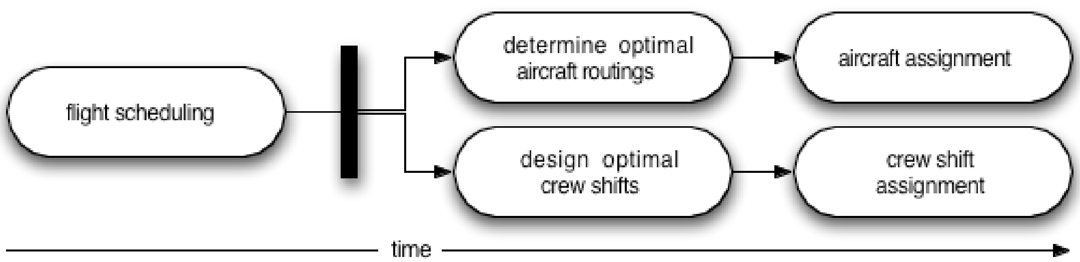

Figure 8 – sequential steps in the air scheduling process

Optimization of airline operations is performed in three steps (see Figure 8). The first step in the air scheduling process is flight scheduling. Airline operators use a fixed schedule that is based on general demand distribution (see paragraph 3.2). The airline operator creates flight schedules that maximize the number of passengers that can be serviced. This scheduling step is performed weeks or months ahead of time. After the flight schedules are created, the operator performs a second scheduling step: the scheduling of aircraft and crews. The aircraft routing process and the crew routing process are usually separated in order to facilitate problem solving (Teodorovic, 1988). However, it is often significantly more cost-efficient to integrate the aircraft and crew routing processes (Freling et al., 2003, Mercier, 2005). After creating routes, the third step is the assignment of specific crews and aircraft to the created routes.

3.3.1 Aircraft routing

An aircraft route is a path that a single aircraft performs to execute any number of scheduled flights. The routing problem is often modelled as nodes on a directed graph (Desrochers, 1992, Ombuki, 2004). Figure 9 shows an example with two routes. Every node in the graph denotes a customer request. Every arc between nodes denotes that those two customer requests can be sequenced. Not all flights can be sequenced, due to constraints on departure time and the fact that an aircraft can only be at one location at any time.

Every arc can have a cost involved. This is the cost of sequencing its two flights. These costs are either monetary (based on the cost of performing required flights), a time duration (flight duration) or both. With this modeling technique, the cost optimization of operations is effectively turned into an optimizable *shortest path problem*, that can be solved with existing optimization techniques.

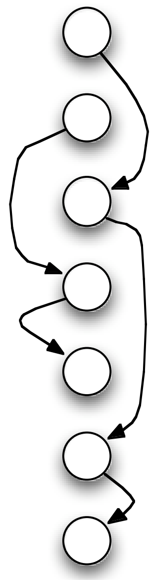

Figure 9 - two routes modelled as directed graph

3.3.2 Crew rotation, crew rostering and duty period regulations

A crew rotation is the path that a single crew performs to execute any number of scheduled flights. Crew rotations are created to minimize the number of crews needed and can be modelled in a similar fashion as aircraft routings. After crew rotations are created, the rotations can be assigned to specific crews. This process is called *rostering*. The rostering process has to take strict regulation on duty periods into account. The duty periods of crews are regulated both by (inter-)national regulations and company policies (Bonnefoy, 2005, appendix I). A schematic representation of a crew rotation is depicted in Figure 10. During a single work duty, a crew can perform multiple legs. Between these legs, a crew often spends time on the ground. This ground time is considered resting time if the time spend on the ground is greater or equal to the prescribed resting time. Additionally, the maximum working time is regulated by flight provisions. A series of work duties are separated by a rest after returning to the crew's home base.

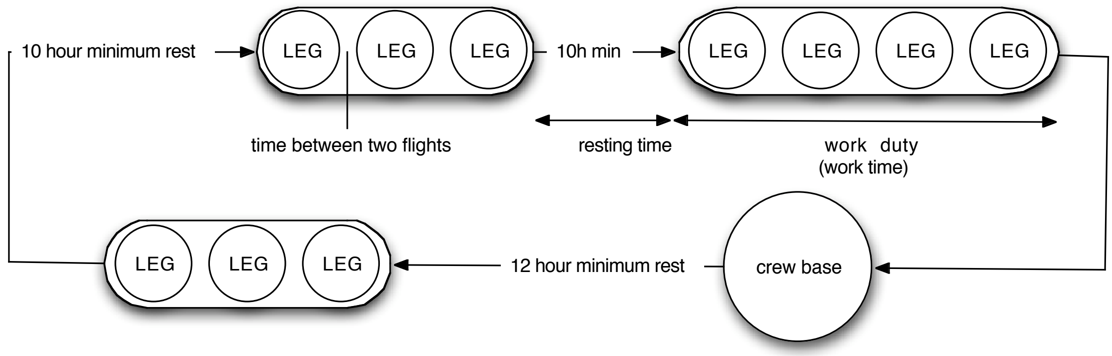

Figure 10 - schematic of a crew rotation

Adapted from Ball and Roberts, 1985

3.3.3 Aircraft maintenance and maintenance scheduling

Aircraft maintenance is a constant that can be included in operator schedules. Aircraft are required to be inspected at periodically (Teodorovic, 1988, Gopalan, 1998). The frequency of maintenance checks depends on the combination of number of flight hours and the number of take-off and landing cycles (Feo, Bard, 1989). The number of flight hours and take-off landing cycles is quite constant in airline operations, as are the flight routes. As a result, maintenance can be pre-scheduled. In practice, maintenance is often not implemented in the first design of a flight schedule. Most operators send the design of a flight schedule to their maintenance department for consideration. The maintenance requirements are then added to the aircraft routing.

3.3.4 Disruptions

During the execution of their flight schedule, aircraft and crew suffer schedule *disruptions*. Disruptions are a form of constraints on operations that cannot be fully implemented into operational plans. Disruptions in operations are caused by tardy customers, weather conditions, shortage of resources or airfield delays. The cost of disruptions for airlines range from 2% to 3% of annual revenue. (Barnhart, 2003).

There are two types of disruptions. A short disruption is a *delay*. A delay is a disruption that prolongs a dispatch. The *delay duration* is measured as the difference between scheduled time and realized time. Typically, a delay lasts 15 to 45 minutes. A second type of disruption is a *disturbance*. A disturbance is encountered at a specific time. Typically a disturbance lasts between 1.5 and 4.5 hours, but in few instances it may last a number of days. Figure 11 depicts the frequency and duration of flight delays and disruptions for airliners.

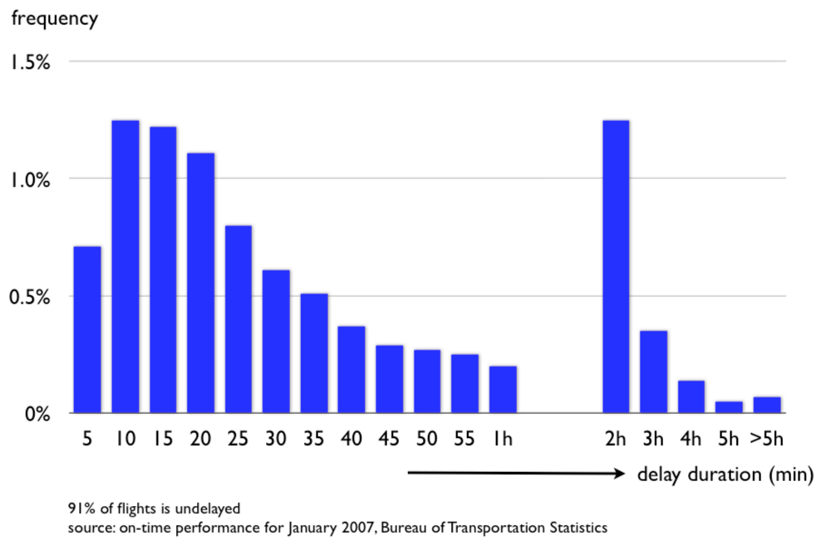

Figure 11 – frequency and duration of delays in airline operations

Figure 12 shows how a prior delay of a flight can cause the further delay of the crew and the aircraft. Most delays are between 15 and 45 minutes (FAA, 2007). The Bureau of Transportation Statistics considers a flight *on time* if it operated less than 15 minutes later than the scheduled time shown in the operator's schedule. A delay that is shorter than 15 minutes will not be noticeable to the customer. However, if the duration of a delay/disturbance is longer, it can cause scheduled flights to be cancelled.

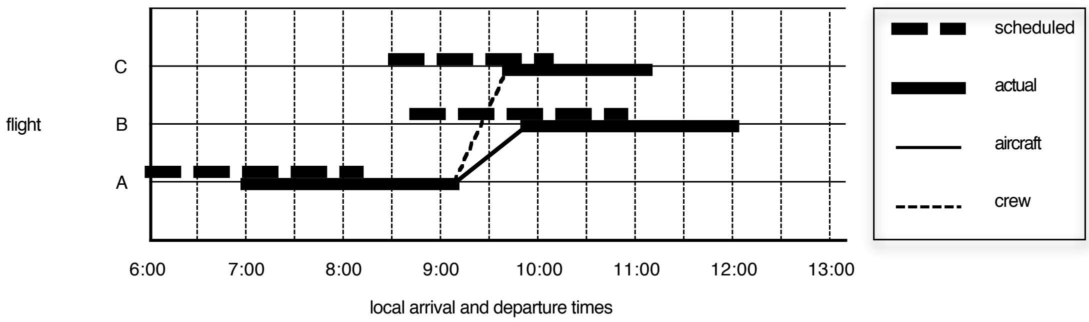

Figure 12 - delay chain

Adapted from Barnhart, 2003

3.3.5 Disruption prevention

Schedule disruptions can be prevented by increasing flexibility of the flight schedule or by predicting disruptions and including these predictions into the flight schedule. Aircraft failure for example can be modelled as a function of the time since the last maintenance check (Bonnefoy, 2005). Airline operators usually include some room for unexpected maintenance requirements in their flight schedules.
Specific airports can cause disturbances due to congestion, security issues or weather conditions. During winter season, some airports are frequently closed due to meteorological reasons. Airline operators include such conditions into their flight schedules. To do so, operators can use flexible departure times (time windows) or schedule multiple alternative routes.

3.3.6 Schedule recovery

The previous paragraph on disruptions has pointed out that some disruption may be prevented. For those disruptions that could not be prevented, there is a need for recovery methods. When aircraft or crews are taken out of operations or when flights are delayed, an operator can use one of four recovery strategies to remove the disruption (Teodorovic, 1988, Barnhart, 2003):

a) Introduce more aircraft into the operations;
b) design a new airline schedule so that no planned flights are cancelled, but with a number of flights being delayed;
c) cancel a number of planned flights with a certain number of remaining flights being delayed;
d) cancel a number of planned flights, with no remaining flight being delayed;

From a customer's perspective, the first strategy is the most attractive, as it prevents any negative effects for the customer. However, extra aircraft can be costly. Cancelling flights will safeguard any future scheduled flight, but will also disappoint customers and potentially result in a loss of customers. Delaying flight prevent flight cancellation. However, flight delays can upset an entire flight schedule. A delay might not always be acceptable to the customer.

What recovery strategy an operator prefers depends on how the operator weighs cost versus service to passengers.

3.4 Conclusion

Transportation operations are scheduled on the basis of demand. Demand can be predicted or simulated. On the basis of demand, operations can be scheduled. With the use of special modelling techniques, these schedules can be optimized. However, when executed, such optimized schedules will suffer from disruptions. Any model to describe air taxi operations will have to include considerations like demand modelling, crew rotations, aircraft maintenance and schedule disruptions.

4 Air taxi operations

The previous chapter has described the operations in transportation and air industry. The operations of air taxi operators can be characterized as a combination of the flexibility of urban taxis and the cost optimization of airliners. However, the operations of air taxis are also very different from operations of airliners and urban taxis. This chapter models the operations of air taxi operations. This model will serve as the basis for cost modelling in chapter 6.

4.1 Demand for air taxi services

The distribution of demand for air taxi services over time can be described by the standard demand curve for transportation (see Figure 7). There is a high demand for transport during the morning and late afternoon and low demand during the night hours. Demand for air taxi services is dynamic and stochastic. Unlike airliners, air taxis cannot create flight schedules before actual demand is apparent. This requires the demand to be handled on a per-request basis, where an demand for transport is entered into an existing schedule. Scheduling a particular aircraft to service a particular request is called a *dispatch*.

4.2 Modelling air taxi operations

All air taxi operations can be modelled on the basis of *dispatches*, *aircraft routes* and an *end-of-day procedure*. The next paragraphs will define these concepts.

4.2.1 Dispatch definition

A customer request is accepted by scheduling that request as a *dispatch*. A dispatch is defined as *a time scheduled trip that is part of a single aircraft's flight route*. A dispatch has an origin and a destination and a pre-determined departure time. This departure time may still have some time-flexibility (a time window). A dispatch can be scheduled to service a particular customer request, but it can also be scheduled for any other reason that requires an aircraft to be relocated. A schematical display of a dispatch is given in Figure 13.

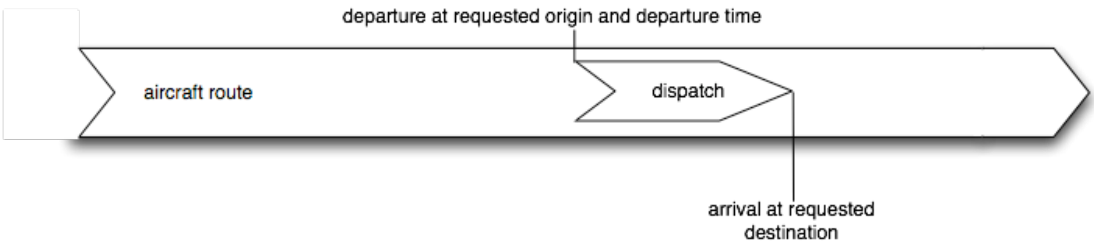

Figure 13 – a dispatch is a scheduled trip that is part of a single aircraft route

A dispatch is scheduled on the basis of the departure time that is requested by the customer. An operator is not always able to service at the requested time. Customers can be given the option to define an earliest acceptable departure and latest acceptable arrival time. This increases the time window wherein the dispatch can be performed, and as a result, this increases the chance that the request can be serviced.

The Etirc business model utilizes a two-hour time window. A customer can request a specific departure time. Etirc will then try to schedule the request between one hour earlier and one hour later. If the request cannot be scheduled, the request is denied service.

4.2.2 Flight route

Every dispatch is part of a flight route. A flight route is actually a set of subsequent dispatches for a single aircraft. Unlike airlines, air taxi routes are dynamic and may change during the day. The duration of a route is arbitrary. Dispatches can be constantly added. A flight route consists of:

- The aircraft tail number;
- a set of dispatches assigned to the aircraft;
- the crew assigned to a sequence of dispatches.

An example of a 24-hour flight route is depicted in Figure 14. Dispatches are depicted in the figure by an arrow.

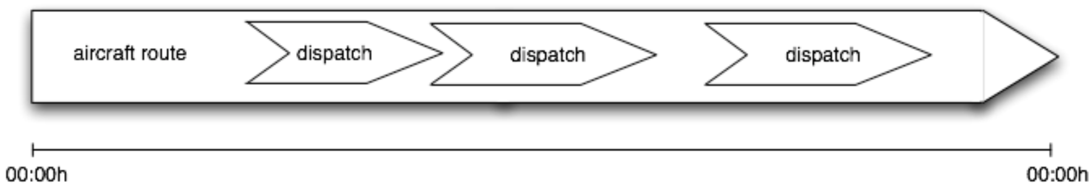

Figure 14 - example of a 24-hour flight route

If the arrival airport of a dispatch does not match the departure airport of the subsequent dispatch in the same route, then the aircraft needs to be relocated. Relocation flights are *non-revenue* trips or *empty legs*.

Figure 15 is a sequence diagram that depicts an example flight route. The solid lines depict a revenue flight; the dashed arrow from C to B depicts a relocation flight. The figure shows that the number of empty legs and the length of the empty legs are determined by the flight schedule and thus by the dispatcher. A route is optimal if the number of relocation flights is minimized.

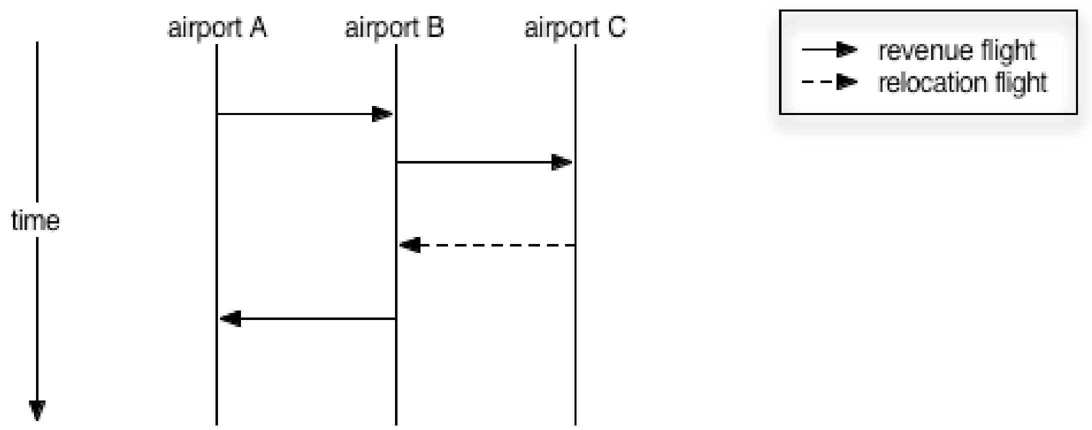

Figure 15 - a relocation flight in an aircraft route

4.2.3 Crews and home bases

During active operation, each aircraft has a single crew aboard to perform flight operations. Multiple crews can serve on the same aircraft on the same day. A crew can be replaced only if the destination of one of its dispatches is the home base of that crew.

The crew regulations for air taxis are the same as those for airliners (see paragraph 3.3.2). An air taxi crew starts its duty rotation (see Figure 10) at its home base. A duty rotation is made up of one or more duty periods. If a duty period ends at a base other than the crew's home base, the crew requires hotel accommodations (hotac). At the end of its rotation, the crew always returns to its home base. The flight that returns the crew to its home base is called the *return-to-base*.

4.2.4 Aircraft maintenance

The Very Light Jets are sent in for periodic maintenance after 50, 300 and 500 active flight hours. Before any aircraft is sent in for maintenance, its crew is first returned to its home base. After maintenance is performed, it is assumed that the aircraft is returned to that home base.

Air taxis have to include maintenance into the flight schedule. Each maintenance shift can be scheduled into the route of an aircraft as a dispatch with a fixed origin and destination. The duration of such a dispatch can be set at any required length, from a couple of hours to several weeks. Contrary to airliners, air taxi maintenance cannot be scheduled far ahead of time. Although the number of flight hours and take-off landing cycles for an air taxis are probably quite constant, the air taxi flight routes are not. Like urban taxis, an *end-of-day procedure* is required to determine when aircraft will be put in for maintenance. The concept of an end-of-day procedure is described in paragraph 4.5.

4.3 Dispatching process

The previous paragraphs showed that the process that the scheduling of a dispatch is initiated by a customer request, maintenance requirements or the end-of-day procedure. The objective of the dispatching process is two-fold. Firstly, to minimize the denials to customers. The second object is the minimization of cost.

While the airline scheduling process has three steps (see Figure 8), the dispatching process has only two. The routing and rostering step are integrated into a single step. The first step in dispatching is to determine if a feasible dispatch can be scheduled to service the request. A potential dispatch is feasible when:

- It can be sequenced with other dispatches of the same aircraft without overlap;
- the aircraft can perform the dispatch without breaking restrictions;
- the crew can perform the dispatch without breaking duty regulations.

If no feasible dispatches are available, the customer requests will be denied service. The percentage of requests that is actually accepted and scheduled as a dispatch is called the *acceptance rate*.

If and only if there is at least one feasible dispatch is the second step of the dispatches process initiated. The second step is to select the best available dispatch and include it into the flight schedule. Usually the selection of a dispatch is based on cost, but it can also be based on quality of service. The operator will then update the flight schedules to include the selected dispatch. The dispatching process is depicted in Figure 16.

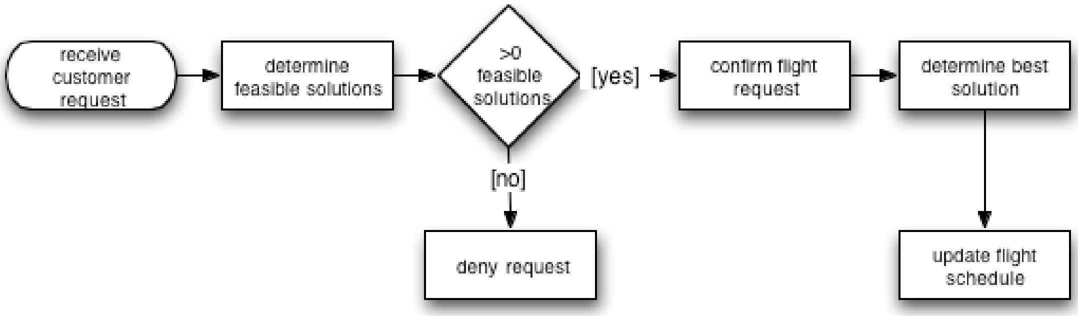

Figure 16 – the dispatch process

4.4 Disruption handling

The previous chapter has shown that aircraft operations suffer from multiple disruptions. Dispatches that cannot be performed as scheduled due to a disruption can be rescheduled to a later time, scheduled for another aircraft route or be cancelled (see paragraph 3.3.6). Air taxi operators want to minimize the number and size of delays that customers suffer and minimize the number of flights that is cancelled due to a disruption. When handling disruptions, Etirc uses this prioritization:

1. Minimize the number of cancelled flights;
2. minimize the delay noticed by the customer;
3. minimize the costs of handling the disruption.

From a scheduling perspective, a disruption is related to a particular aircraft for a particular period (e.g. aircraft breakdown) or to a particular dispatch (e.g. due to a short delay). A breakdown can be included in the schedule by modelling the disruption as a special dispatch with a fixed origin and destination and can be set for a fixed time.

A time window has a special added value if a dispatch needs to be rescheduled due to a disruption. A bigger time window increases the chance that a disrupted dispatch can be rescheduled. Within the Etirc case study, the time window of disrupted dispatches is increased to include a maximum delay of two hours.

Paragraph 3.3.5 discussed measures to prevent disruptions. Air taxi operators can prevent schedule disruptions by including basic maintenance requirements, weather conditions and airport congestion into their flight schedules. Operators can include these considerations into their schedules by including them in an end-of-day procedure.

4.5 End-of-day procedure

The air taxi schedule problem is almost fully static and deterministic during evening hours. The demand during the evening hours is low (see Figure 7) and demand for the next day is largely known. Due to the static and deterministic character of the dispatch problem, the end of the day is the best time to consider rest and maintenance considerations and optimize them accordingly.
The inclusion of crew rest into the flight schedule is one of the primary roles of the end-of-day procedure. The end-of-day procedure can send crew back to rest at their home base. The decision to send crews back to their home base is made on the basis of the scheduled flight routes and cost parameters from the business model. There are three different situations wherein the end-of-day procedure can decide to send a crew to its home base:

1. **Route passing home base**
   The last dispatch of the day ends at the home base and/or the first dispatch of the next day requires the aircraft at the home base. The crew is sent to rest at home as early as possible.
2. **Affordable return trip**
   If the crew can be send home for a cost below a set cost, the end-of-day procedure schedules an extra dispatch that returns the crew to its home base.
3. **Long idle time**
   If the aircraft is idle for a period that is longer than a predetermined duration, then the aircraft is send to the home base by scheduling an extra dispatch.

4.6 Conclusion

Air taxi operations are dynamic due to their on-demand character. This requires that operations be scheduled on a dispatch basis. Customer request and maintenance periods can be scheduled on a dispatch basis. Disruptions can be entered into existing flight schedules as a special kind of dispatch.

The flight schedule can be changed at any moment in time. Most customers place their flight requests up to hours before actual departure. Due to this *scheduling window*, the dispatch schedule is a sliding window that contains a combination of a fully dynamic schedule and a semi-fixed schedule. This is depicted in Figure 17. As a result of the sliding window, the optimization of operations is a continuous process. There is only one event that is always fixed in time: the end-of-day procedure. The end-of-day procedure uses the knowledge of the semi-fixed schedule together with the more dynamic schedule to optimize long-term consideration on the basis of the fixed and the dynamic schedule.

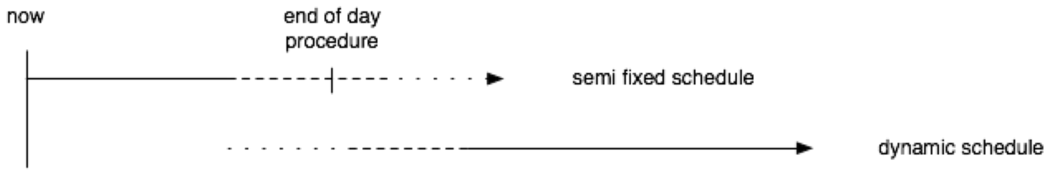

Figure 17 - dynamics in air taxi operation

5 Literature on air taxi operational costs

The previous chapter has defined the operational processes of air taxi operators. Cost optimizing these processes requires a cost model. Up to date there are no detailed cost calculation methods for air taxi services (Fiertz, 1999, Mane and Crossly 2006). Existing cost calculation methods for air taxis merely express the economic value of time saved by air taxi transport. This chapter surveys literature on cost of air operations to be able to build an air taxi cost model in chapter 6.

5.1 Definitions

A lot of cost functions in aviation are expressed using special flight profile terminology. The terminology distinguishes periods of flight and flight preparation.

The total time an aircraft is in the air is called the *flight time*. Before and after each flight, a short handling time is required for fuelling and basic maintenance. The total time of taxi out, take off, flight, landing and taxi is called the *block time*. This block time is equal to the entire duration of a dispatch. These definitions are given in the flight profile of the Association European Airlines (AEA), as visualized in Figure 18.

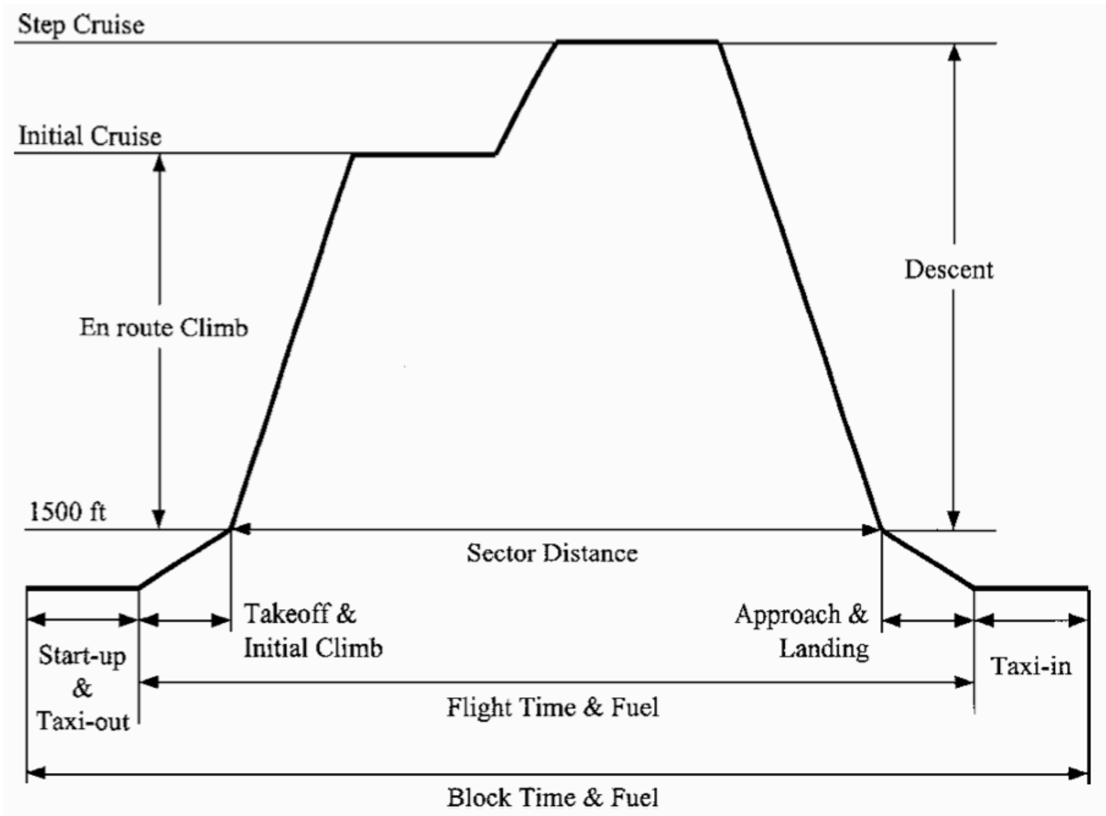

Figure 18 - flight profile as defined by the EAE

Source: Isikveren, 2002

5.2 Direct and indirect operating costs

The operational costs of air operations can be split into direct costs and indirect costs. The direct costs are related directly to an aircraft and its flight operation and are usually expressed as cost per mile or cost per hour. The most common method for direct cost calculation is the "*Standard method of estimating comparative direct operating costs of turbine powered transport airplanes*" of the Air Transport Association of America (ATA). The model expresses costs based on aircraft characteristics like cruise speed, weight and crew size. The rest of this paragraph discusses these cost related parameters.

5.2.1 Costs of a flight operation

The cost of a flight operation is based on crew costs, aircraft costs, fuel costs, airport fees and any additional costs. The cost of the crew can be considered fixed, but can also be considered variable on the basis of flight hours (Pant et al., 1995) or can be spread over the cost of other utilized resources (Isikveren, 2002).

5.2.2 Fuel cost

Fuel cost is one of the biggest costs of a flight operation. The ATA standard has a complete cost model for fuel costs. The costs are a function of the fuel per flight block, the cost per fuel unit, the cost of oil for the engines, the number of engines and the block distance. In extensive empirical research into optimal flight techniques, Isikveren (2002) has created a calculation model for minimum fuel costs. Attributes that increase fuel costs are cruising speed and climb/decent speed. Fuel cost is decreased by a greater flight height. However, forty years of statistics shows that the cost function can be simplified to a simple calculation of the required fuel weight in kilograms for a block distance and the price per kilogram of fuel (Pant et. al, 1995).

5.2.3 Aircraft maintenance cost

Aircraft maintenance cost can be split into time-dependent and cyclic components (Isikveren, 2002). A lot of maintenance costs, like air conditioning, instruments and windows, are dominantly dependent on time and usage. Other costs, like maintenance of landing gear and wings, are dominantly time-cyclic and require maintenance on a time-cyclic base. However, these costs can be correlated to the average sector distance flight time (ATA 1969, Isikveren 2002). As a result maintenance costs can be defined as a function with an almost linear dependency on flight time.

5.2.4 Aircraft cost

Two constructs are considered for financing an aircraft: acquisition and leasing. The cost of acquisition is related to repayment of an initial capital outlay, plus any interest costs. Most literature on aircraft depreciation assumes that an aircraft is depreciated to a residual value of 0% over a depreciation period (Pant et al., 1995). The aircraft interest costs are based on the aircraft's acquisition price times an interest rate. Fixed depreciation and interest cost can be used (Isikveren, 2002). It may also be economical to lease air taxis. The cost of leasing is of specific interest if the number of required aircraft is flexible (Isikveren, 2002).

5.2.5 Insurance cost

The ATA standard considers aircraft insurance (or hull insurance) costs a function of the aircraft acquisition cost times a insurance rate, split over the number of utilized flight hours. The passenger insurance can be determined as either the number of flight hours (Pant et al., 1995) or the function of the number of passengers and the distance covered by an aircraft (Isikveren, 2002).

5.3 Indirect operating cost

Indirect operating costs (IOC) are those airline costs not directly connected with the actual flight of the aircraft. An exhaustive list of indirect costs is supplied by Kroo (2006):

- Aircraft Ground Handling
- Landing Fees
- Aircraft Service
- Cabin Attendants
- Food and Beverage
- Passenger Handling
- Reservations and Sales
- Baggage/Cargo handling
- Passenger Commissions
- Passenger Advertising
- Cargo Commission
- General and Administration

The size of these indirect costs can only be estimated from statistics (Kroo, 2006). On average, the total of indirect costs for an airliner shows a relationship with the direct operational costs. As a result, most indirect cost can be associated with direct costs, which simplifies calculation models.

5.4 Conclusion

The cost structure of air operations is visualized in Figure 19. The direct variable cost can all be aggregated to form a fixed cost per flight hour. The indirect variable costs are related to the scheduled flight routes and crew rotations. Based on this knowledge, the next chapter builds a cost model for air taxis.

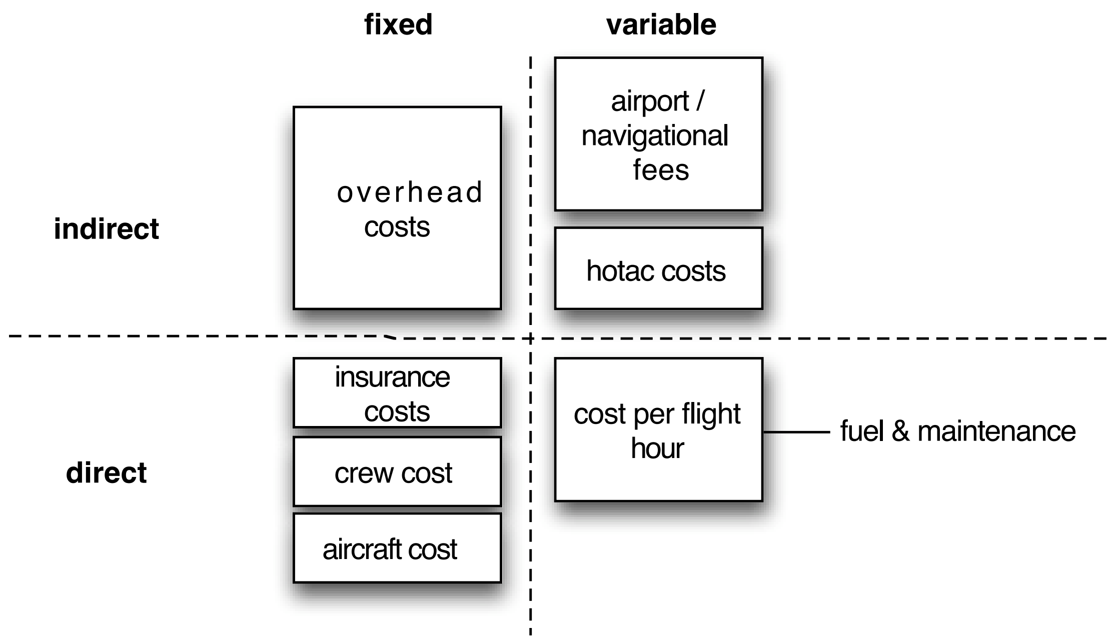

Figure 19 - cost structure of air operations

6 Cost modelling for air taxi operations

The previous chapter surveyed cost aspects of air operations. Most cost can be simplified to hourly flight cost and the cost of fees. This chapter builds a cost model for the variable cost of air taxi operations. This cost model will serve as the basis of cost optimization of air taxi operations.

6.1 Cost types in air taxi operations

The cost modelling of air taxis operations is simplified by assuming a fixed cost price per flight hour. This is a valid assumption, as cost per flight hour is very constant, even at different utilization rates. For the same reason, Etirc adopted an hourly rate for its services. The cost per flight hour includes the following costs:

- Fuel consumption;
- crew cost;
- third party fees (Etirc services);
- progressive and periodic maintenance.

The cost per trip is dependent on the cost per flight hour and the length of the trip, but also the airport fees, navigational fees and the required crew hours. With a fixed cost per flight hour, a fixed flight length and fixed airport and navigational fees, the cost of every trip can be pre-determined. Costs that are not included in the cost per trip are overhead costs and aircraft acquisition or leasing. Figure 20 depicts the cost decomposition of air taxi operations. The aircraft and aircraft insurance costs are summarized as 'fleet costs'.

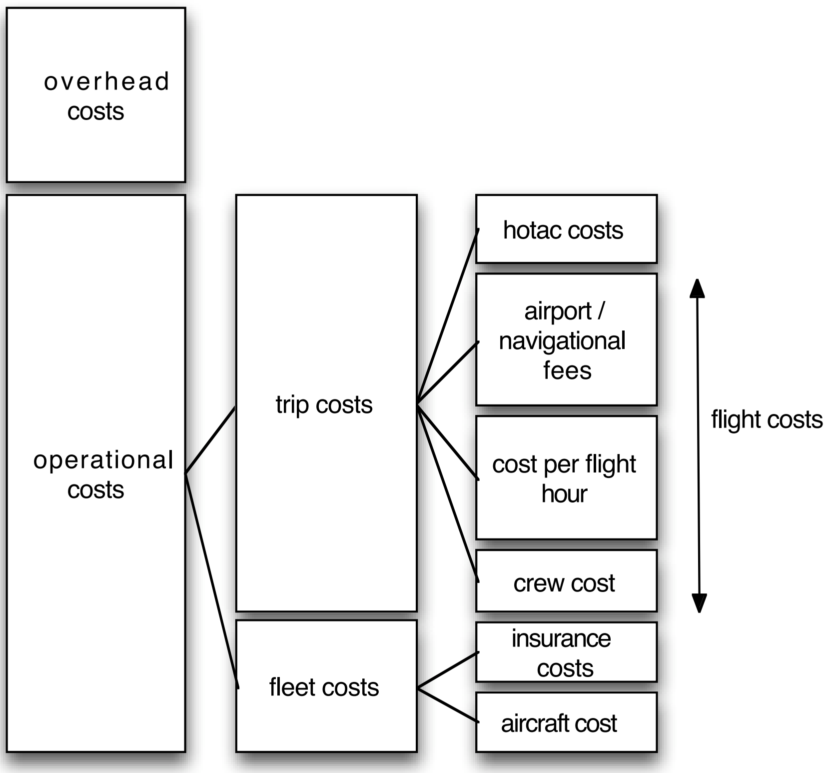

Figure 20 – operational cost decomposition by cost type

6.2 Relating air taxi operations and their variable costs

The previous paragraph defined the types of variable costs that air taxi operators incur. The variable operational costs are mainly formed by the trip costs. Chapter 4 has shown that all trip-related operations of air taxi operators are dependent on the efficiency of the dispatch process. It can be concluded that variable cost can be optimized by an efficient dispatch mechanism. The rest of this paragraph describes the relation between the dispatch process of air taxi operators and the operational costs.

6.2.1 Cost of flight hours

The cost of flight hours is directly dependent on the sum of flight times from all trips. The total number of flight hours is a combination of the revenue and non-revenue trips. Revenue flights have been agreed with the customer and are therefore a fixed cost factor. The non-revenue flights however are not fixed. Their cost can be optimized. There are three types of non-revenue flights (or empty legs):

1. A relocation that is required to reach a customers origin;
2. a relocation to bring the aircraft crew to its home base;
3. a relocation to bring the aircraft to a specific base (e.g. for maintenance).

Each of these three types of empty legs requires different optimization considerations. Only the maintenance related empty legs are ignored in this study.

6.2.2 Airport and navigational fees

Each aircraft takeoff and landing is associated with a fee. The size of this fee is different for every airfield. As a result, each paid fee is directly related to a specific flight, whether that is a revenue flight or not. As said before, revenue flights are a fixed cost factor. Their fees can also be considered fixed. The part of airport and navigational fees that can be optimized is limited to the fee costs of empty legs and the fee cost of dispatches that are scheduled by the end-of-day procedure.

6.2.3 Hotel accommodation costs

The third and last type of variable operational cost is hotel accommodations (Hotac). The costs of hotel accommodation are related to crew activities. An air taxi crew has a duty rotation of maximally five days and requires hotel accommodations when it is at rest at a base other than its home base. The cost of hotel accommodations is assumed to be comparable for all airports in a single operator's network.

Crews are assumed to travel along with a single aircraft until their rotation returns them to their home base. This ends their duty rotation. The end-of-day procedure determines when crew are sent home. As a result, the cost of hotel accommodations is directly related to the route that is scheduled for an aircraft and its crew and the efficiency of the end-of-day procedure.

6.3 Conclusion

The previous paragraphs have shown that the all optimizable costs are related to the scheduled routes of aircraft and the specific dispatches in those routes. These relations between operations and their optimizable costs are depicted in Figure 21.

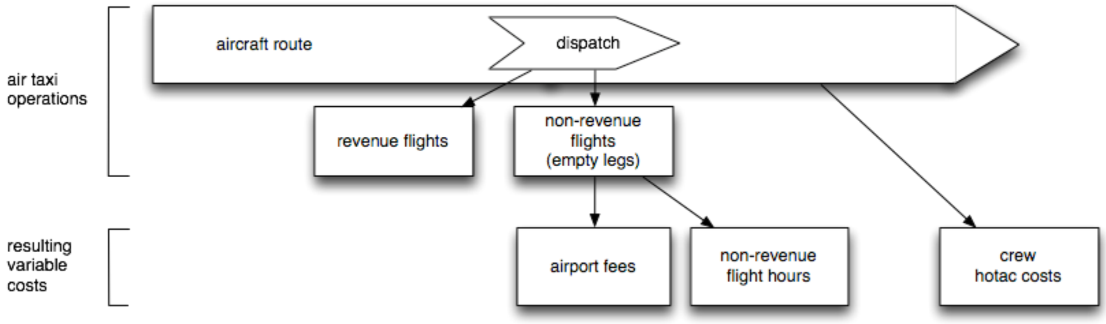

Figure 21 - cost decomposition for optimizable operational costs

Minimizing operational costs for air taxis is directly linked to optimizing the aircraft routes. To optimize costs, a dispatch model is required that models flight routes, individual dispatches, crew rotations, and their costs. An optimization algorithm is required to minimize the operational costs while taking all resource constraints, maintenance requirements and crew duty restrictions into account.

The next chapter contains a literature survey for transport optimization models and algorithms. Chapter 8 will build on this survey to select suitable techniques to optimize costs of air taxi operations.

7 Transportation optimization models and algorithms

The previous chapters have characterized the optimization problem for air taxi operations as a dispatch problem and have modelled the operations-cost relation of the dispatch problem. This chapter classifies this dispatch problem in an optimization perspective and describes optimization solutions that are provided by literature. This knowledge will be used in chapter 8 to select a suitable solution for the air taxi optimization problem.

7.1 Definitions

An algorithm is a finite list of well-defined instructions for accomplishing some task (Wikipedia, June 2007). A conceptual algorithm can be calibrated for a specific problem. A calibrated algorithm can be implemented into software. Algorithm calibrations and implementations that do not change the classification of its conceptual algorithm are considered the same algorithm and not a new algorithm.

7.2 Problem classification

The air taxi dispatch problem is classified by literature as a Dial-A-Ride Problem (DARP). This is a special form of the Vehicle Routing Problem (see Figure 22). In the standard DARP, transport is provided by a homogeneous fleet with a fixed size.

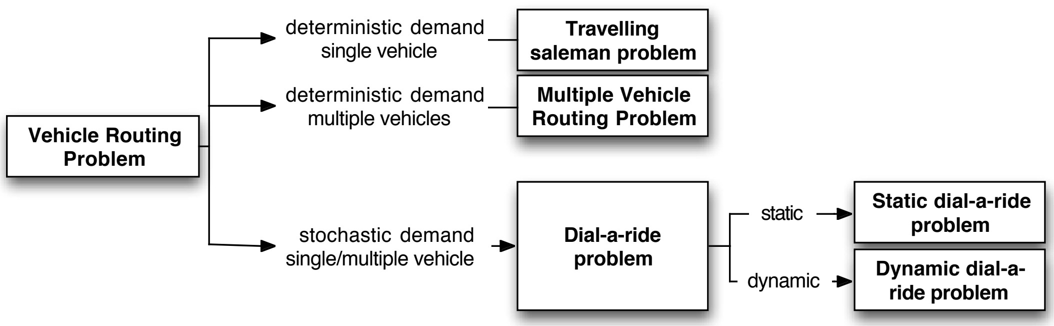

Figure 22 - classification of the dial-a-ride problem

The air taxi dispatch problem is a dynamic stochastic problem. Stochastic dynamic problems are the most complex problems. In a stochastic dynamic routing problem, there are events that may lead to route modification (Ghiani, 2003). These events are the arrival of new customer requests, disruptions and maintenance events. How these events are handled is based on the policy set by the operator.

There is no clear distinction between single and multiple vehicle DARPs. The distinction between static and dynamic DARPs is also often blurred in practice since requests are often cancelled and, as a result, planners may allow the introduction of new requests in a solution designed for a static problem (Cordeau, Laporte 2002). For that reason, both static and dynamic models and algorithms will be researched in this chapter.

7.3 Dial-a-ride optimization algorithms

There are three 'schools' of dial-a-ride algorithms: exact algorithms, heuristics and meta-heuristics.

7.3.1 Exact algorithms

Exact algorithms were the earliest transportation optimization algorithms. Exact algorithms are able to guarantee a global optimal solution, but require huge computational efforts for large optimization problems. Due to more computational power and more efficient algorithms, exact algorithms have recently regained interest. Exact methods for the Dial-A-Ride Problem can be classified into three approaches (Hjorring, 1995): direct tree search, dynamic programming and linear programming.

Direct tree methods consist of incrementally building vehicle routes by means of a branch and bound tree. This method is inefficient and as a result unsuitable for all but small problems.
A dynamic program can solve problems that exhibit properties of overlapping sub problems and an optimal substructure (Wikipedia, April 2007). Dynamic programming algorithms start at an initial state and create a solution through incremental steps. The method is much faster than naïve methods, while it can still deliver an optimal solution (Trick, 1998).
Most research in transport optimization has been focused on the third approach: linear programming. A linear optimization model can include complex controller functions, useful for an end-of-day procedure. This control functions functionality is limited to linear relations however (Boyd, 1991). There are several formulations for a linear program: set partitioning, vehicle flow and commodity flow. The set partitioning approach generates a constraint matrix in which columns represent feasible routes. The rows in the matrix correspond to customers (Hoffman 1993, Barnhart et al., 1998). This matrix can be solved using mathematical techniques.
A vehicle flow is a optimization model that uses *decision variables* to determine what vehicle will service which paths. The formulation is also called "three-index formulation" because the decision variables have three indices: one index to indicate what vehicle is used and two indices to indicate the origin and destination of a single trip. In the air taxi problem, there is a homogeneous fleet. As a result, the vehicle index can be discarded. A commodity flow is an extension on the vehicle flow to model the quantity of goods that are transported. This extension is not required for the air taxi problem.

7.3.2 Heuristics

A heuristic is a method of directing problem solving using knowledge based on experience. Heuristics, while often not able to find the optimal solution, are able to find a suitable solution quickly. Often, the first result of a heuristic is 10-15% from optimal. Four heuristic approaches can be classified: construction methods, mathematical programming, iterative improvement and interactive methods. All these heuristics stop at a local optimum (Hjorring, 1995).

Construction methods are generally based on the savings method of Clarke and Wright (1964). The method computes the savings achieved by combining the end of a route with the start of another route for all feasible route combinations. The second approach, mathematical programming, assigns customers to vehicles based on approximate delivery costs. The third approach, iterative improvement, is based on an initial solution. The most used interactive improvements are the 2-optimal and 3-optimal method.

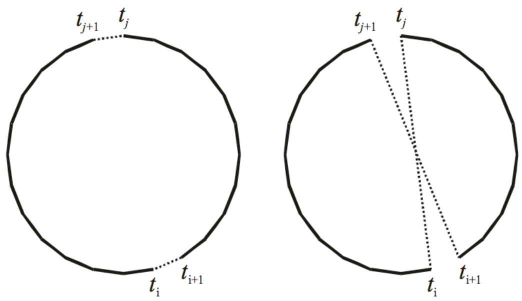

The 2-optimal method breaks two edges forming two disconnected chains. It then connects disconnected chains with different edges, until the optimal solution is found. An example is shown in

Figure 23. The 3-optimal method works on the same principle, but is able to find better solutions.

Figure 23 - A 2-opt optimization change

(original on the left, 2-opt improvement on the right. Source: Lucic, 2002)

The fourth heuristic approach, interactive methods, offer a combination of human experience and algorithm logic. They do not replace human dispatchers, but assist them by allowing the user to define routes and examine route possibilities.

7.3.3 Meta-heuristics

Meta heuristics are heuristic method for solving a very general class of computational problems. Meta-heuristics are generally applied to problems for which there is no satisfactory problem-specific algorithm or heuristic (Wikipedia, April 2007). Meta-heuristics that are used for transportation problems are evolutionary algorithms like Ant System and Genetic Algorithms, along with Tabu Search and Simulated Annealing. These are all stochastic search techniques. The stochastic approach of meta-heuristics enables these algorithms to escape from local optima and potentially find the global optimum. However, meta-heuristics can never guarantee a global maximum solution.

Simulated Annealing (SA) is an optimization approach that is drawn from physics. SA can solve difficult optimization problem by a combination of a randomized method and a deterministic descent process. Its inner working is depicted in Figure 24. The SA algorithm runs several iterations. In each iteration, the existing solution is slightly changed. If an improvement is obtained, the current best solution is replaced by the generated one. The algorithm stops when no further improvements can be reached in the neighbourhood of the current solution. The focus on a *neighbourhood search* can prevent SA from finding the global optimum.

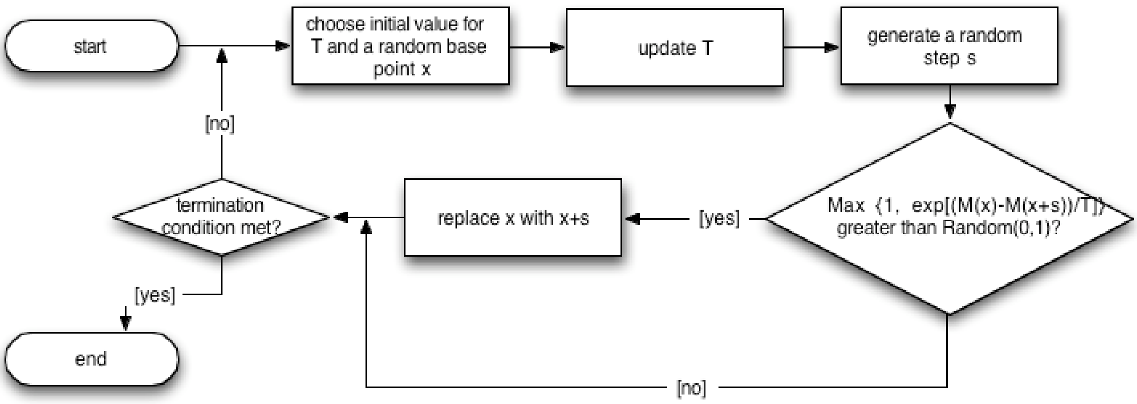

Figure 24 - schematic representation of Simulated Annealing Algorithm

Adapted from Forbes, Jones, 1993

An Ant System is a relatively modern approach in the field of transport optimization. It uses multi-agent technology. In the Ant System (AS), the routing problem is modelled as several nodes in a search space (Dorigo, Gambarella, 1996). Every node is an airport that has to be visited. A set of artificial ants is used to search the problem space. Artificial ants have some memory; they can perceive local information and live in a discrete world. By giving the ants some basic behaviour, their actions and their interactions can be steered. For example: each individual ant's behaviour can be set to minimize the cost of its own movements. Artificial ants are able to define part of a solution at every iteration. At the end of every iteration, an ant can use a pheromone trail to exchange basic information with other ants. After a certain number of cycles, the process is ended (Lucic, 2002). However, only individual ant behaviour can be adapted. As a consequence, the actions of the system as a whole cannot be steered directly.

Genetic Algorithms (GAs) are inspired by biological processes of populations of organism to adapt to their surrounding environment. A GA has three steps that build solutions. First is the *selection* of suitable solutions. Second is *crossover*, where two suitable solutions are combined to create two new solutions. The third step is *mutation*. The mutation step introduces noise into the population to prevent the algorithm from sticking to a local optimum (Lucic, 2002). Each Genetic Algorithm can differ in mainly three aspects. Firstly, in the ways to *encode* (or model) the individuals; Secondly, in the method used to test solution fitness; and thirdly in the strategies used for mutation and crossing over (Li et al., 2005, Goncalves, 2004).

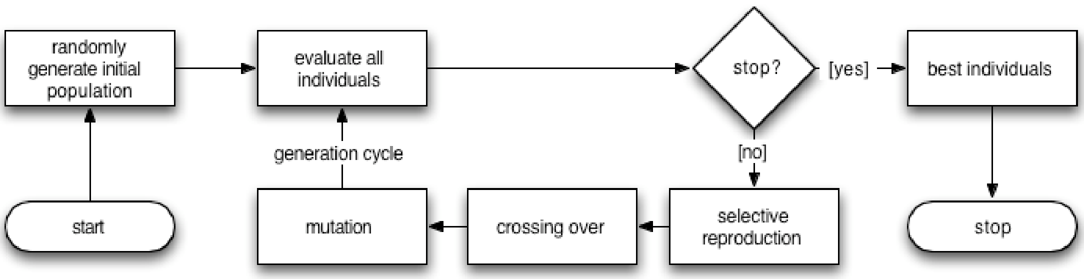

Figure 25 - schematic display of a Genetic Algorithm

Adapted from Fontain, 1992

Tabu search (Cavique, 1999, Lucic, 2002) is a general heuristic procedure that is suitable for guiding a search to obtain a good solution in a complex solution space. The technique helps to search beyond local extreme points by using memory structures. Tabu search uses a fixed number of iterations. Its goal is not to find an optimal solution but a solution that is good enough. From each iteration to another, two solutions exist: the current and the best found. Every new solution is found as a move to a neighbouring solution. Tabu search is limited to a local optimization. New solutions are always limited to the neighbourhood of the current solution and best-found solution.

7.4 Conclusion

The main difference between several Dial-a-ride algorithms is the optimacy of the results and the speed of their solution. Another difference is the method that is used to build solutions. The solutions are constructed, adapted, replaced and/or improved. As a result, some algorithms provide a reasonably ‘stable’ solution, while others can completely alter its solution in every step of the problem-solving process.

The next chapter will weight the qualities of each of the described algorithm schools against the requirements of operational air taxi optimization.

8 An optimization model for air taxis

The previous chapter has surveyed a broad range of optimization algorithms that can solve Dial-A-Ride Problems. The air taxi problem is a special version of the Dial-a-ride problem, with some additional solution requirements. This chapter sets these requirements before selecting a suitable optimization solution for the air taxi problem.

8.1 Requirements

The additional requirements for the air taxi dispatch problem are closely related to the business model. On top of the standard speed and flexibility requirements for an on-demand solution, the air taxi business model has specific requirements for the optimization objectives, the flexibility of the fleet size, the inclusion of an end-of-day procedure and the 'stability' of the results.

The optimization objective is three-fold: firstly to minimize customer rejects, secondly to minimize cost and thirdly to maximize flexibility to handle disruptions. The optimization solution is required to provide an optimal solution. An optimal solution is defined as the combined best performance on all three of these objectives.

Additionally, the optimization solution is required to enable the exchange of capacity. Aircraft resources can either be added to the model or removed from the model for any period of time.

Thirdly, the optimization solution is required to be able to implement an end-of-day procedure. The end-of-day procedure can perform a control function on the basis of all available demand data and operations planning.

Lastly, the optimization solution is required to minimize the number and size of changes to the schedule in order to create a 'stable' schedule. In the real world, every change has to be communicated to crews and airports. Changes in the existing schedule are only acceptable if these changes offer a cost benefit.

8.2 Assumptions

The above requirements are directly related to the air taxi business model. It is assumed that the dispatch problem can be simplified to a scale that can be solved within 30 seconds. Further, it is assumed that the level of uncertainty is low. Departure times are fixed when scheduled and passengers always show up at the agreed time.

The number of available crews is not in the optimization model as a constraint. For simplification, it is assumed that a crew is always available for dispatch. This assumption is reasonable for the business case. Aircraft are the most costly resources and operators wish to prevent aircraft being grounded due to crew-shortages. Although crew are not a constraint in the optimization, the cost of crew is included into the optimization model, enabling basic optimization of crew cost.

8.3 Algorithm selection

Paragraph 8.1 showed that the most important requirements for the selection of optimizing algorithm are the ability to minimize rejects, optimize cost, the support for an end-of-day procedure and the stability of the solution.

Heuristic algorithms do not meet these requirements. Heuristics can be caught in locally optimal solutions. Exact algorithms are preferred over the traditional heuristics because they promise better cost optimization.
Ant and Bee systems cannot implement an end-of-day procedure. To do so, the Ant or Bee behaviour is required to implement special end-of-day behaviour. Ants càn be made aware of the progress of time. However, the agents are unaware of their future movements (which model future dispatches). Their behaviour is limited to ad-hoc reactions and incremental solution building. As a result, the end-of-day procedure cannot be implemented in an Ant or Bee System.
Simulated Annealing, Genetic algorithm and Tabu Search are unable to meet the requirement for a stable solution due to their “improvement function”. The function stochastically changes the solutions in every iteration. This makes it hard to guarantee the stability of solutions.
Likewise, a linear programming approach cannot guarantee stable solutions. At every change in the problem, the linear model completely rebuilds the previous solution, which can result in a completely rescheduled schedule.

A dynamic programming solution is the most suitable. It can implement an end-of-day procedure as a separate procedure that considers the solution of the current schedule as a sub problem. The exchange of capacity can be modelled by adding resources to or removing resources from the model. A dynamic programming approach can also guarantee solution stability.

8.4 Conclusion

A dynamic programming approach is probably the best solution to optimize air taxi operations. The dynamic programming approach stands out specifically due to its flexibility in handling dynamic optimization problems and the stability of its results.

Etirc has developed a Dynamic Programming algorithm. Its logic promises to offer a relatively fast and efficient capacity allocation. Its basis is simple, robust and extendable. It should be able to support both simple and complex extensions (shift time limits, time windows, etc). Further efforts will focus on implementing a Dynamic Programming approach in order to prove or disprove it to be an effective solution.

9 Dispatch algorithm simulation

The previous chapter has concluded that a dynamic programming algorithm provides a fitting solution for the dispatch problem. This conclusion needs to be proven in practice. It is unsure if a dynamic programming algorithm is able to supply an optimal solution. To decrease the complexity of the problem, the Etirc algorithm limits the searched solution space. This may prevent the algorithm from finding an optimal solution. To test the fitness of the dynamic programming approach, a prototype has been build and used in simulations of a specific real life case study.

9.1 Building the algorithm

A prototype of a dynamic programming algorithm has been created. This algorithm will serve as the basis for proving the effectiveness of a dynamic programming solution. This prototype is based on the working paper "*An optimal dispatch algorithm for n aircraft and m airfields*", by T. Mentzel of Etirc aviation.

9.2 Simulation setup

9.2.1 Case study

All tests are set up within a simulation environment with a fixed configuration. This simulation is based on a real life case study of Etirc. It is assumed this case study is a typical representation of air taxis in the network of Etirc.

The basic simulation is built on these numbers:
- Ten airports;
- thirty-eight routes between these airports that are regularly flown;
- a weekly demand of 270 flight requests.

All simulations for this case study are based on ten test runs. The average of these test runs indicates the algorithm performance.

9.2.2 Benchmark

Although dynamic programming algorithms are often exact, the tested algorithm is a non-exact dynamic programme. The outcomes of non-exact algorithm are not fully predictable. To prove a non-exact algorithm, it needs to be proven that the algorithm is better than existing algorithms for a standard problem. In the air taxi domain, a standard problem can be defined by a typical demand scenario.

The insertion algorithm as described in "Efficient Insertion Heuristics for Vehicle Routing and Scheduling Problems" by A. Campbell and M. Savelsbergh, 2004 is utilized as a benchmark. This benchmark algorithm fits the identified problem area: dial-a-ride problems (DARPs). The algorithm of Campbell and Savelsbergh is oriented on the same business model as air taxis: cost optimizing through cost minimization in routing and scheduling in an on-demand environment.

The insertion heuristic is leading in its sort. Both Campbell and Savelsbergh are at the front-line of current logistic research. Their algorithm supports all DARP problems that are researched in current literature: time windows, shift time limits (crew regulations), variable delivery volumes, variable delivery times and multiple routes per vehicle. The algorithm is also successfully used in practice. It has been utilized for home grocery delivery and shuttle bus services as recently as 2007.

Comparative research of non-exact (approximating) algorithms by Golden et al. (1980) show that the basic form of the benchmark algorithm can be classified as a ‘cheapest insertion’ algorithm. This type of algorithm will produce results that are on average 10-15% from optimal for a 100-node problem.

The results of the benchmark are directly compared to those of the test algorithm. The benchmark does not implement a time window. The results of tests with time window are benchmarked against the same test algorithm without time window.

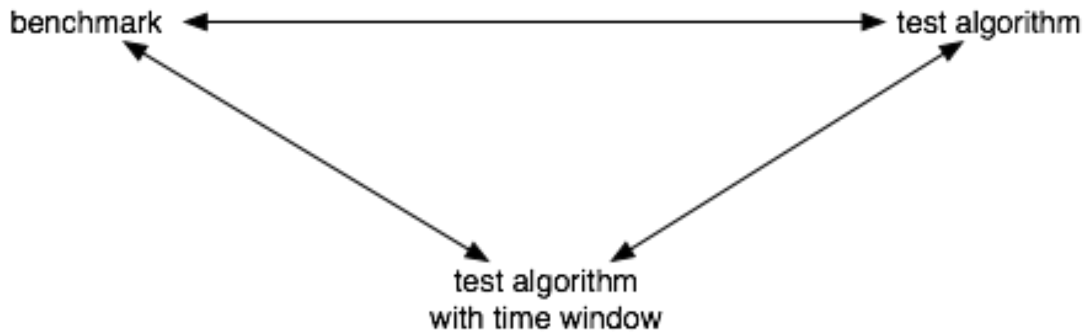

Figure 26 - benchmarking with and without time windows

9.2.3 Logical tests

The output of the test algorithm and the benchmark is a flight schedule. Every schedule is tested by a few basic logic tests (see Table 2). These tests are required to prove the validity of the results. These tests do not test the performance or optimacy of the algorithms.

Check	Logical test
Sequence check	Check per aircraft that all sequenced trips can actually be sequenced.
Overlap check	Check that there is no overlap between any fixed or provisional dispatches of the same aircraft in the final schedule.
Flight duration check	Check that enough time is scheduled to perform every flight, together with the required handling times
Demand association	Check that all accepted requests are associated to exactly one dispatch in the dispatch-list. Check that all denied requests are indeed not in the dispatch-list.

Table 2 - logic test on the algorithm results

9.2.4 Simulation machine configuration

The prototype is built on the basis of Apache, PHP and MySQL technology. The prototype runs on a MacBook Pro, 2.16 GHz Intel Core Duo, 2 GB 667 MHz DDR2 SDRAM using Apache 1.3.33, PHP v5.2.0, MySQL 4.1.22

9.3 Algorithm calibration

The basic algorithm prototype was improved and calibrated through an iterative process of algorithm extensions and explorative testing. These calibrations where tested for a small test run (<40 customer requests). This set of requests was specifically designed to test each algorithm extension's ability to find improved solutions.

The algorithm was improved in three iterations. In each iteration, the search space of the algorithm was increased, at the cost of a complexity increase in the algorithm. The first three iterations offered an improvement of the algorithm performance. Further iterations offered no further performance improvements. The calibrated algorithm will function as test algorithm further tests.

9.4 Simulation environment and variables

The simulation environment is formed by:

- The selected simulation technique;
- the selected fleet size;
- the utilized scenarios;
- the utilized time windows;
- a benchmark for comparison of the algorithm performance.

The only variable through tests is the fleet size. Each variable is described in the next paragraphs.

9.4.1 Simulation technique

The simulation uses discrete event simulation (see Paragraph 3.1). Discrete event simulations are useful for stochastic problems. A discrete model is chosen because a continuous model is more complex while that complexity is not required.

9.4.2 Fleet size

Figures from early calibration indicate a fleet size of fourteen will service the required 95% of customers. To be able to determine the relation between fleet size and each metric, five independent measurements have been made. The fleet size is varied from twelve up to and including sixteen aircraft.

9.4.3 Scenarios

The customer requests have been created by a randomizer that simulates a spread in the requested departure times that is equal to the demand spread in the daily demand curve for airlines. The demand volume and trip distribution are based on demand estimations of the case study.

9.4.4 Time window size

The benchmark uses no time window. The test algorithm implementation is tested in two versions. The tested algorithm uses a two-hour time window. In a two-hour time window, a customer request can be scheduled to depart one hour earlier or one hour later than the requested departure time.

9.5 Test setup

The performance testing of the test algorithm testing is split into three separate tests. The first test indicates the general performance of the algorithm under normal operating conditions. The second test indicates the flexibility of the algorithm to handle disruptions. The third test indicates the ability of the algorithm to be extended for with an end-of-day procedure. The algorithm performance is set out against a benchmark. The tests are described in the next paragraphs.

9.5.1 Test 1 : Performance under normal operations

The first test is oriented at proving the strong performance under basic operational circumstances. The performance of the algorithm is tested on the basis of cost of its resulting schedule and its flexibility to service a maximum number of requests. The results of the test are visualized by a bar graph. The relative performance for different fleet sizes is tested through four metrics (see Table 3).

Metric	Calculation
Acceptance rate	Number accepted requests/number of placed requests
Flight hours / aircraft year	Total number of flight hours*52 / fleet size
Average empty leg cost per dispatch	Total empty leg cost/number of accepted dispatches
Empty leg percentage	Number of empty legs / (number of accepted requests+ number of empty legs)*100%

Table 3 - metrics for testing algorithm performance

The average empty leg cost per dispatch is assumed to be related to the fleet size and the algorithm performance. To visualize the algorithm performance, the average empty leg cost per dispatch is displayed in a linear graph for different fleet sizes. The results of the algorithm and the benchmark are combined in a single line graph.

The empty leg percentage is assumed to be related to the specifics of a demand scenario and to the fleet size. The empty leg percentage is averaged over all scenarios to minimize the influence of specific demand scenarios. The relation of empty leg percentage to fleet size is visualized in a bar graph.

The hours that aircraft are required to be in maintenance hours are not included in the metric of flight hours per aircraft year. The actual number of active flight hours will therefore be lower than the metric shows. The metric is still relevant from a relative point of view and can be adjusted to include maintenance relatively easy.

As mentioned in paragraph 8.2, the optimizations are not restricted by the number of crews. However, the total number of flight hours can be utilized to give an indication of the required number of crew. This is outside the scope of this research.

9.5.2 Test 2.1 : Flexibility for handling disruptions

The performance of the algorithm in handling disruptions is tested on three performance indicators. Firstly, the number of cancelled dispatches has to be minimized. Secondly, the delay has to be minimized. Thirdly, the cost of handling the disruption has to be minimized. The flexibility to handle disruptions is tested through three metrics. These are explained in Table 4. The results are set out against the benchmark.

Metric	Calculation
Reschedule rate	Number rescheduled requests/number of requests that has to be rescheduled due to a disruption
Average empty leg cost per rescheduled dispatch	Total empty leg cost of rescheduled flights/number of rescheduled dispatches
Delay	Average delay in minutes compared to the agreed departure time

Table 4 - metrics for testing algorithm flexibility

Simulated short delays are between 20 and 45 minutes, with 5-minute intervals. Each delay is associated with an aircraft and a time. The delay is linked to the first dispatch during or after that time. Simulated disturbances last between 1.5 hour and 4.5 hours, with half hour intervals. Each disruption is associated with an aircraft and a time. The simulation links the disturbance to the dispatch that is performed during that time. The disturbance is handled from the moment the aircraft can first land. Table 5 summarizes the types, durations and count of disruption tests.

Type	Test cases	Duration	Effects
Short delay	96	20 up to 45 minutes with 5 minute intervals	Delays subsequent dispatch
Disturbance	96	1.5 hour up to 4.5 hours with half hour intervals	Cancels all dispatches for a given period

Table 5 - disruption types, duration and counts for simulation

The delays and disturbance are tested against in all 10 test runs, resulting in 960 tested delays and 960 tested disturbances. The effect of delays and disruptions are measured by the resulting delay on the subsequent dispatch(es) and the percentage of rescheduled flights. The flexibility of the test algorithm is made visible by setting out the reschedule rate and average delay against the benchmark.

9.5.3 Test 2.2 : Flexibility for implementation of an end-of-day procedure

The end-of-day procedure that is tested implements a procedure to send crew back to their home base. This procedure is run at 5 PM. At this time of the day, the second demand peak of the day is almost past, offering a near-static problem with enough data to consider future operations. Initial test results indicated that the cost of such an end-of-day procedure is lowest when all aircraft are stationed at the airfield that is visited most. Therefore, this test uses a single home base for all crews.

The implemented end-of-day procedure is a rule-based control system with relatively few rules. Its effect on the speed performance of the test algorithm is marginal.

Within the one-week period of each scenario, the end-of-day procedure can be run six times: once between each consecutive two days. To simulate the passing of time, the end-of-day procedure is run after every fortieth request.

The frequency of rest at home and the cost of return flights are measured through two metrics. The first measures the number of times a crew ends its duty period at its home base. The second measures the extra cost of the return-to-base flights on top of normal operating cost.

Metric	**Calculation**
Number of visits per type	Summarized per type: - route passing home base - long idle period - low cost return flight
Cost of return-to-base	Cost of additional return-to-base dispatches

Table 6 - metrics for testing effective end-of-day procedures

10 Results

This chapter summarizes the results from the three tests that have been defined in the previous chapter. The results of the test show the performance of the test algorithm under normal conditions, its flexibility to handle unforeseen disruptions and its extendibility to include an end-of-day procedure.

10.1 Cost optimization under normal conditions

10.1.1 Results on the customer acceptance rate

The business model requires that at least 95% of customer requests is accepted and serviced. Simulations show that the test algorithm can service 95% of customer requests with a fleet of fourteen aircraft. The test results of customer acceptance rate are shown in Figure 27 and Figure 28. The first figure shows that the benchmark and test algorithm are able to service the same number of customer requests.

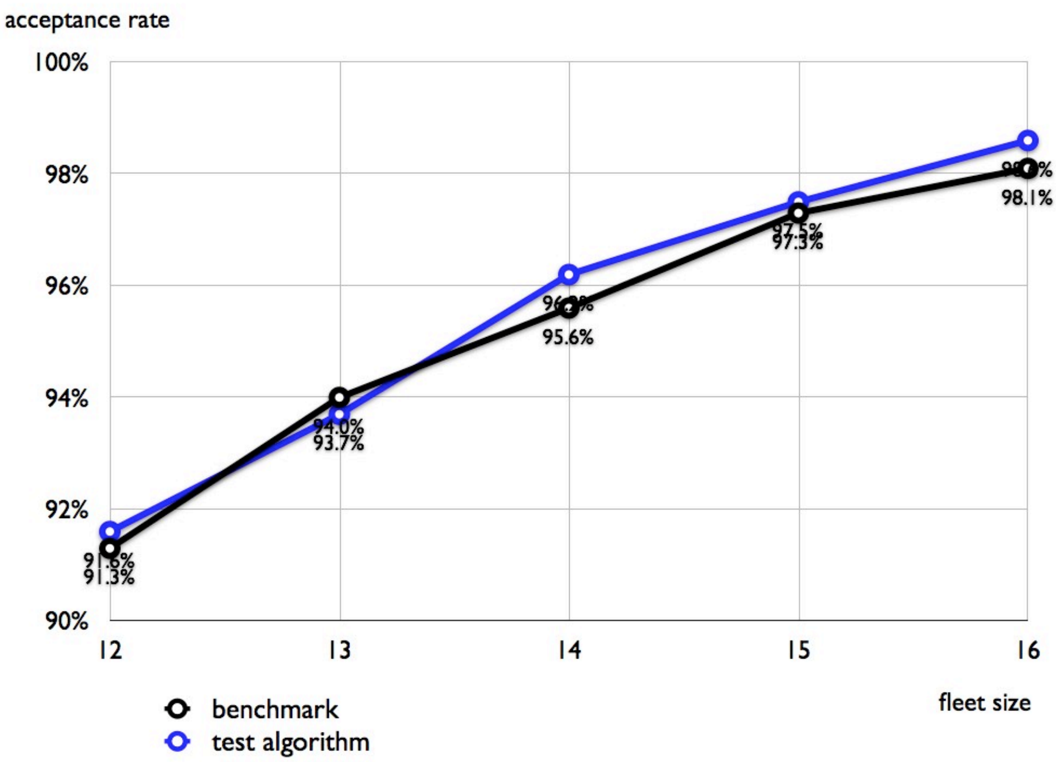

Figure 27 - acceptance rate without time windows

In a second test step, a 2-hour time window was added. The test algorithm with time window has a customer acceptance rate that is 1-2% higher than the test algorithm without time window. This is visualized in Figure 28.

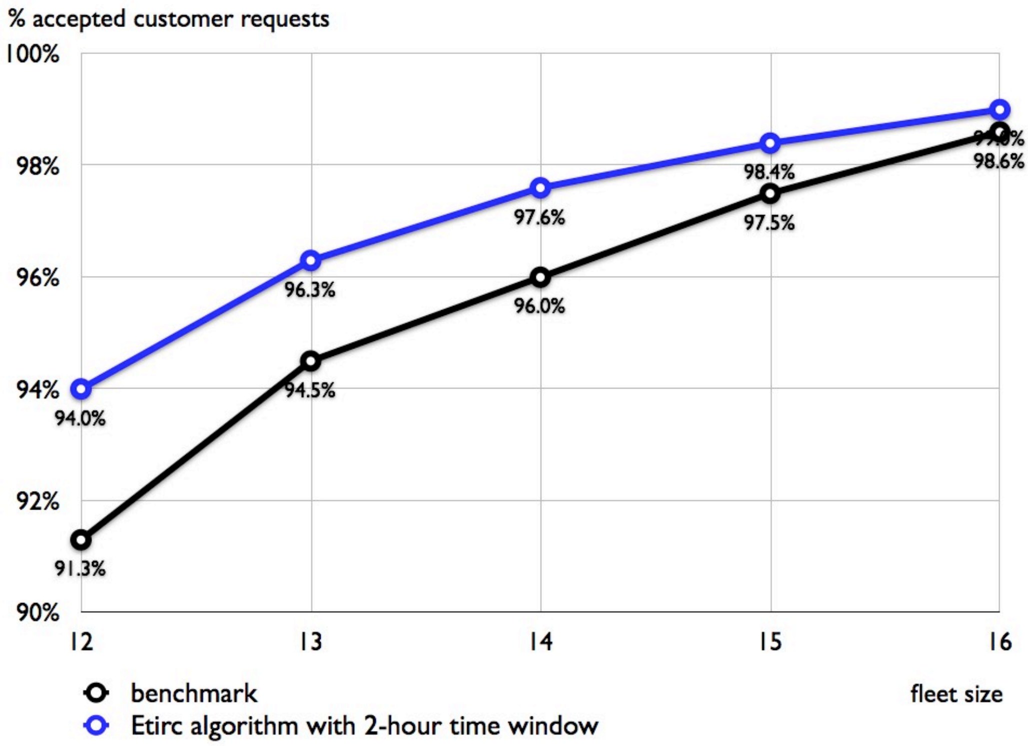

Figure 28 - acceptance rate with and without time window

With the capacity advantage of a time window, a smaller fleet can be utilized to service the same amount of customers.

10.1.2 Results on minimization of empty leg cost

Empty legs are the biggest form of variable cost. The algorithm performance to minimize the variable cost is measured through the average cost of an empty leg per revenue flight and the total number of empty legs. This paragraph focuses on the cost of empty legs. The next paragraph will focus on the number of empty legs.

Simulations show that with the test algorithm and a fleet of fourteen aircraft, an average empty leg costs 514 for the benchmark. At the same fleet size and with the test algorithm, the average empty leg cost is 482. The average empty leg cost of the test algorithm is around 5-6% below the benchmark.

Figure 29 displays the results for the benchmark, and the test algorithm with time window. On average, the addition of a time window introduces no cost advantage or disadvantage over the test algorithm without time window.

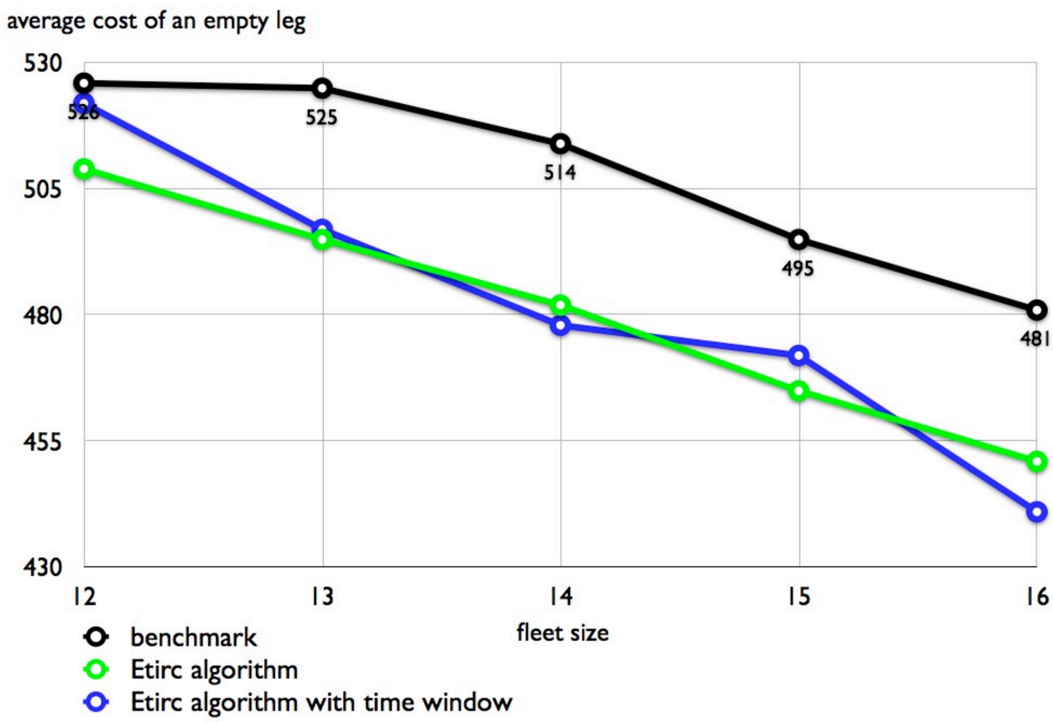

Figure 29 - average empty leg cost

Firstly, it can be concluded that the test algorithm is able to be more cost effective than the benchmark. Further, given the higher acceptance rate of the test algorithm with the addition of a time window and the near equal costs, it can be concluded that the addition of a time window is beneficial.

Because of its higher customer acceptance rate, all further test results focus on the test algorithm with the added 2-hour time window.

10.1.3 Results on the percentage of empty legs

Empty legs form the biggest part of optimizable costs. Simulations show that for at a fleet size of fourteen, 33.8% of flights are empty legs. This percentage decreases when the fleet size is increased.

Figure 30 shows the percentage of empty legs for the use case from the benchmark and the test algorithm (with time window).

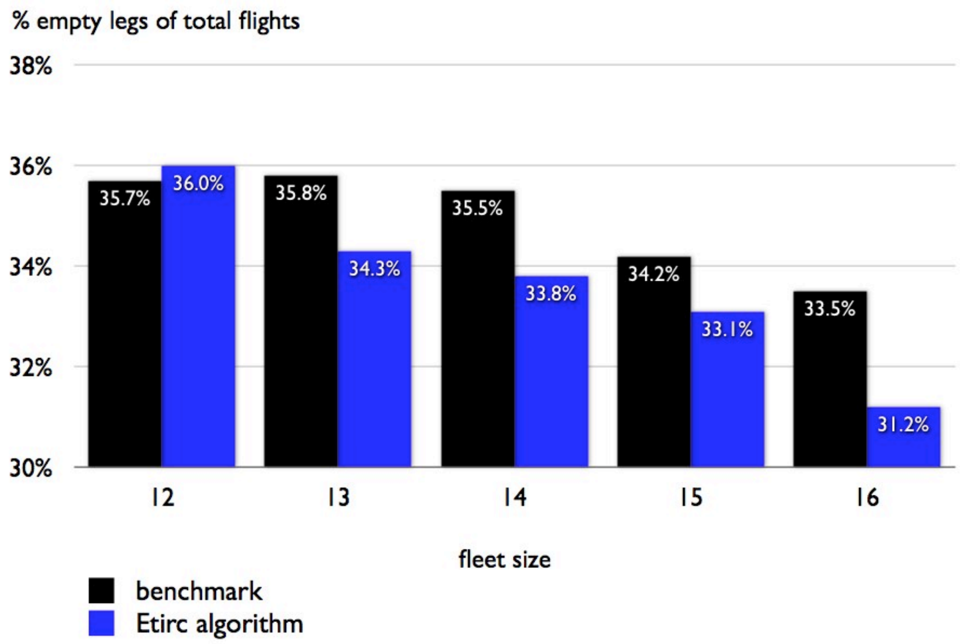

Figure 30 - percentage empty legs of the test algorithm

It is visible that the test algorithm is able to minimize the number of empty legs. Additionally, with a bigger fleet size, the number of empty legs is smaller. The previous test has shown that the cost of these empty legs also decreases. It can be concluded that a bigger fleet has less empty legs costs.

10.1.4 Results on the average flight hours per aircraft year

For a fleet size of fourteen, an average aircraft spends 2040 hours a year in flight. Further, each aircraft spends 1420 hours on "handling time", which includes taxiing to and from the airport.

Aircraft maintenance is not included in these numbers. However, the time required for aircraft maintenance is relatively easy to calculate from the flight hours per aircraft, the maintenance requirement and the average maintenance duration.

In total, an aircraft is occupied 3460 hours a year. This means an aircraft is busy 40% of the time, while 60% of the time it is 'idle'. The idle time is a form of 'time-flexibility' that can be put to efficient use in time of disruptions.

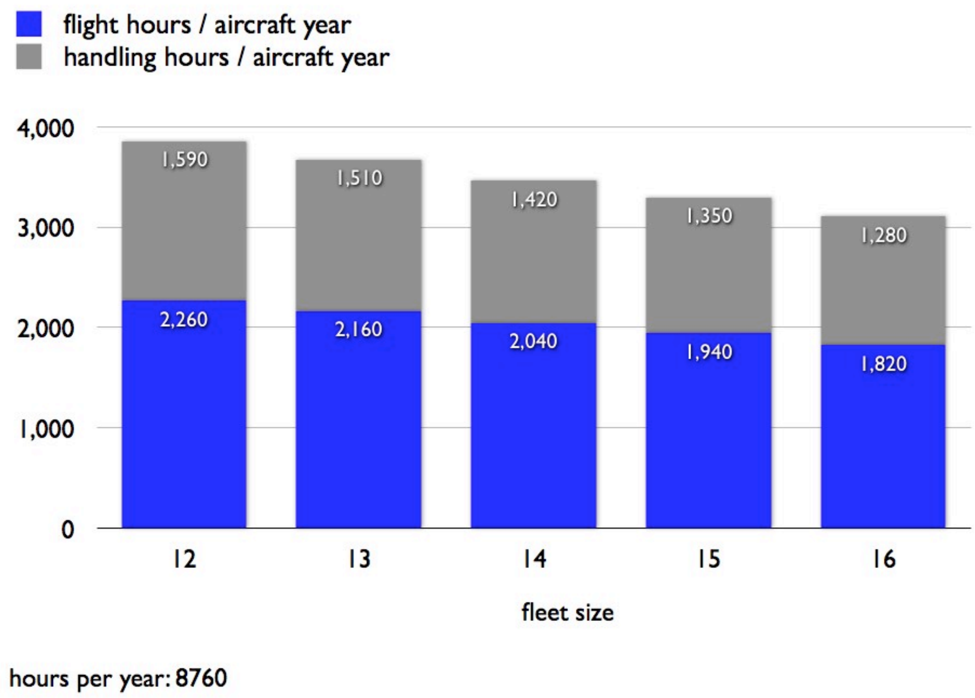

Figure 31 - hours of active aircraft usage per year

It can be concluded that a smaller fleet, while being more cost effective, is less flexible in situations of disruptions.

The 2000 flight hours a year equal the maximum recommended flight hours per aircraft for Very Light Jets. It might be advisable to increase the fleet size to limit the strain of flight hours on aircraft.

10.2 Cost optimization in time of disruptions

10.2.1 Results on minimization of delay

A delayed dispatch may result a subsequent dispatch to be delayed further. The simulations have tested the effects of a prior delay on the flight schedule. Simulations with the test algorithm show that a prior 45 minutes delay will delay the subsequent dispatch with 3 minutes (on average). The benchmark algorithm suffers a slightly smaller delay.

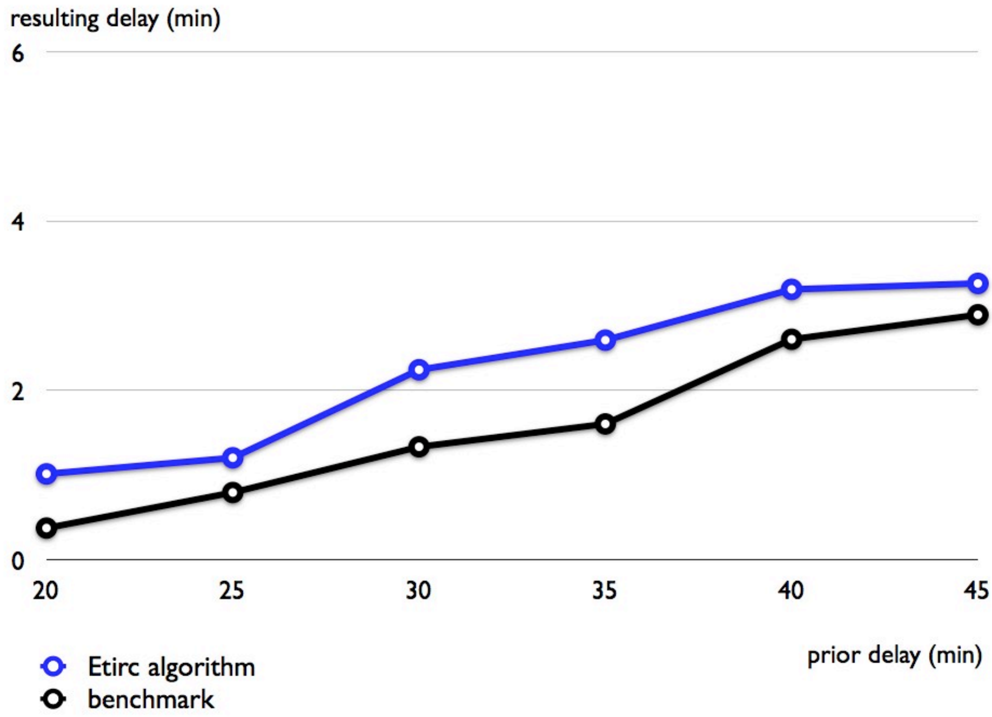

Figure 32 - duration of a delay due to a prior delay

The first conclusion that can be drawn is that delays have a small effect on the schedule. In airline industries, a 45-minute delay often leads to a 2-hour delay. This is not so for air taxis. Delays seem to be buffered by the unutilized idle time between subsequent dispatches.

In general, the test algorithm shows a slightly higher resulting delay. The longer delay of the test algorithm can be explained by its time window. The time windows create a more dense schedule that enables a higher acceptance rate. However, in a more dense schedule, there is less time between dispatches to buffer the effect of delays.

Even thought the test algorithm suffers from bigger resulting delays, it can be concluded that the delays rarely result in undesired delays. Overall, the resulting delays are well below the accepted 15-minute delay.

10.2.2 Results on rescheduling flexibility

Disruptions can cause scheduled flights to be cancelled. Most disruptions have a duration of 1.5 up to 4.5 hours. Simulations show that the test algorithm is able to reschedule around 60% of all flights that are affected by a disruption. This is well above the 25% rescheduling of the benchmark. The test results are depicted in Figure 33.

Prior tests have revealed that the visible dip and peak in the graph can be attributed to randomness of the input.

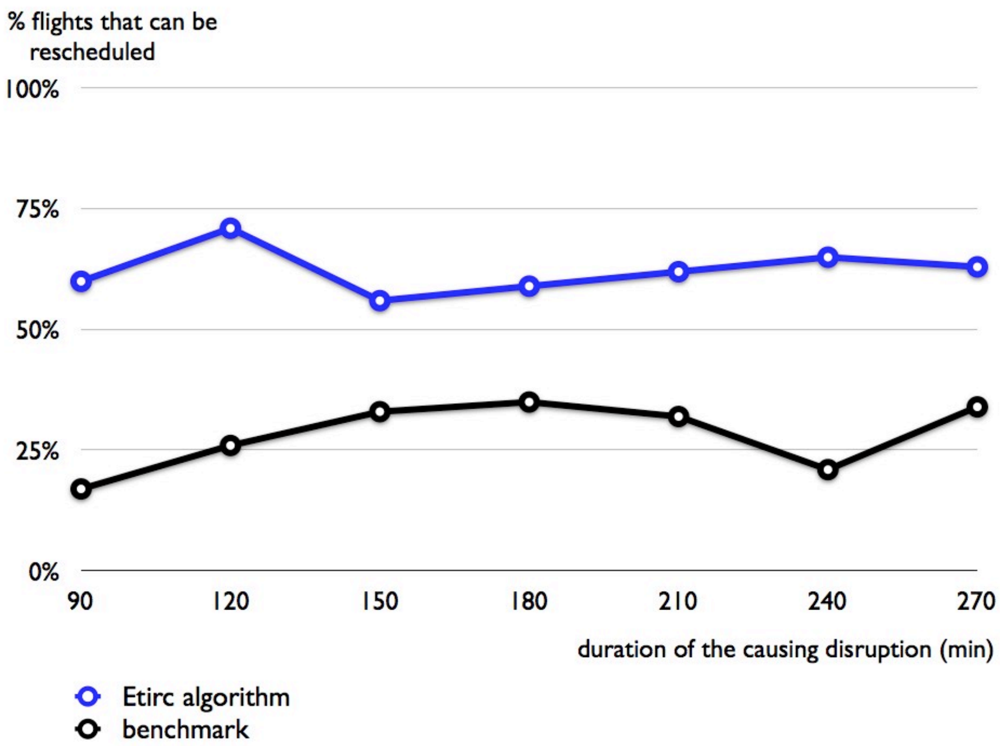

Figure 33 - ability to reschedule disrupted flights

The reschedule rate (20 to 65%) is well below the 95+% acceptance rate for regular customer requests. The lower rate can be explained by the short time period between the disruption being known and the execution of potential alternative solutions. Often, this time period is too short to relocate aircraft to reschedule cancelled flights. However, the 60% reschedule rate of the test algorithm is well above the 25% reschedule rate of the benchmark. The much higher rescheduling rate can be explained by the test algorithm's increased ability to search the solution space.

10.3 Results of the end-of-day procedure

The implemented end-of-day procedure determines if individual crews are sent home to rest. The decision is made on the basis of their current location and their flight schedule for the next day. All test results for the end-of-day procedure are based on a fleet of fourteen aircraft and an end-of-day procedure that returns aircraft if their schedule is empty for a 24-hour period or if the aircraft can return to its home base at a cost of less than 500. In the case study, a single home base is utilized for all aircraft.

The test results show that, on average, aircraft return to their home base at the end of the day 4.2 times a week. The aircraft are sent home for one of three reasons (see paragraph 4.5). The foremost reason for return to the home base is an empty schedule. Together, all aircraft are sent home 33 times a week because their schedule is empty for at least the 24-hour period. The second most frequent reason for return to the home base is the route of the aircraft passing the home base during the evening hours. This happens 25 times at week. Only a very small number of crews is sent to their home bases because a return flight is below a pre-set cost parameter. These numbers are displayed in Figure 34.

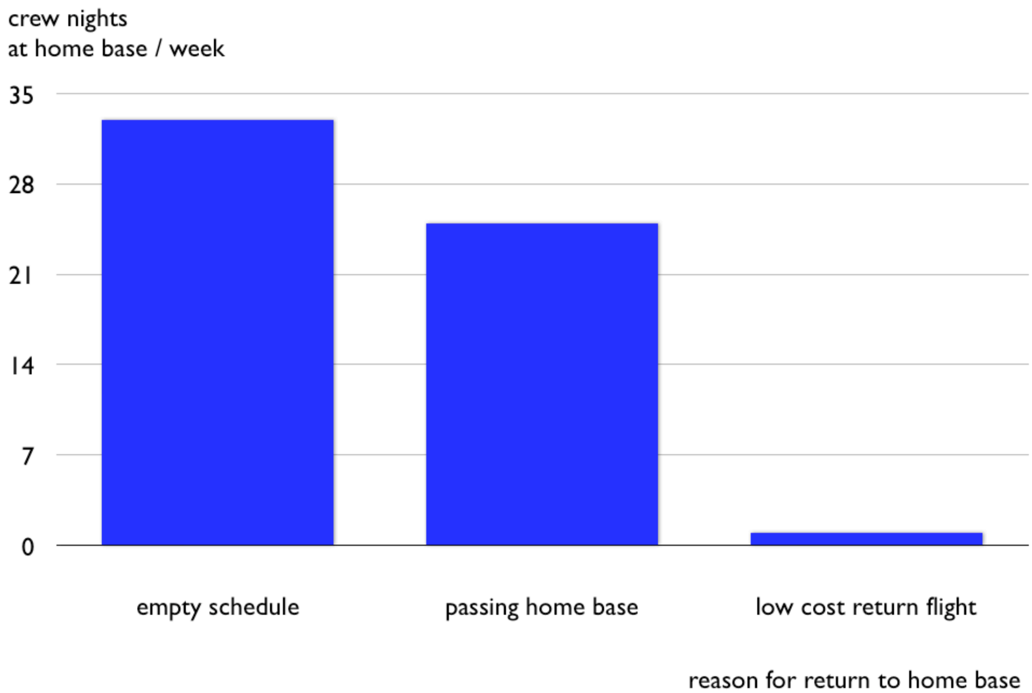

Figure 34 - frequency of crew rest at home

An operator will employ more than one crew per active aircraft. Assuming an operator employs two crews per aircraft, each crew will return to its base around 2 times per week. This is well above the required 1.4 times per week (based on a 5 day rotation with 2 days rest). It can be concluded that the implemented end-of-day procedure is sufficient to include basic duty time regulations.
Further, the cost of return flights is relatively low. The average cost of a return-to-base flight is 407. These cost are below the average empty leg cost of a revenue trip, meaning that the return-to-base flights are scheduled cost-efficiently.

11 Conclusions and recommendations

This chapter will summarize the results of this research to answer the main research question:

> *What are the operational costs of Air taxi operations, which factors determine these costs and*
> *which solutions could contribute to the minimization of these costs through services oriented at air taxi operators?*

11.1 Answers to research questions

1) What are the operational costs of Air taxi operations?
The operational costs that air taxi operators incur can be simplified to three types of costs: cost per flight hour, navigational and other fees and the cost of hotel accommodations.

2) Which factors determine these costs?
The size of the above costs are fully dependent on the schedules of aircraft and their crews. Schedules with efficient flight routes minimize the sum of costs of flight hours, fees and hotel accommodation. Due to the on-demand character of air taxi services, operational schedules cannot be created and optimized for long periods ahead of time. A dispatch mechanism is required that optimizes operations whenever a new event requires a new schedule. The size of an operators operational costs is determined by the cost effectiveness and the flexibility of that dispatch mechanism.

3) Which solutions could contribute to the minimization of these costs through services oriented at air taxi operators?
An effective dispatch mechanism offers cost effectiveness and flexibility. A dynamic programming approach is able to offer both cost effectiveness and flexibility. A good alternative to dynamic programming is a linear programming approach. However, a linear approach completely rebuilds existing schedules after every event. This creates an unwanted amount of uncertainty in the scheduling process. A linear programming approach is also limited in its flexibility to include an end-of-day procedure.

To prove the cost effectiveness and flexibility of a dynamic programming approach, a new algorithm has been designed and implemented. The created algorithm provides a solution that offers both the required cost and flexibility advantages. The created algorithm supports the dispatching and routing of aircraft. This basic version does not include considerations like aircraft maintenance and disruption prediction. Tests with a simple end-of-day procedure show the algorithm can be extended with extra considerations. These advantages are summarized in Figure 35.

Solution characteristic	Results
Effectiveness	5.5% lower average empty leg cost 1.7% fewer empty legs
Flexibility	30-35% higher rescheduling in disruptive situations
Extendibility	Extendable with an end-of-day procedure

Figure 35 - performance of the tested solution

11.2 Recommendations

The dispatching of aircraft can be solved using a dynamic programming approach. A dynamic programming approach can offer the required cost effectiveness and flexibility and provides a stable. The algorithm that was implemented for the purpose of this research is an effective solution.

Literature advises that operational optimizations should never be purely oriented on reaching the theoretical optimum. Reaching optimacy is often impossible due to the on-demand character of demand. Flexibility to handle uncertain demand and schedule disruptions is at least as important to cost effectiveness as pure cost optimization.

The results with the simulations have shown that the following aspects are of importance:

1. **Align the configuration of home bases to network demand**
   The home base of aircraft and crews have a big impact on the return-to-base costs of crews. The home bases of an operator network can be carefully selected by testing different configurations and their cost efficiency.

2. **Select fleet size to match the required acceptance rate**
   The cost effectiveness of the dispatch algorithm or its flexibility is heavily influenced by the fleet size. The first consideration in the determination of fleet size should be the percentage of demand that the operator wishes to service.
   Simulations show that an efficient fleet size also pushes the aircraft usage close to maximum advisable flight hours. Estimations of flight hours per aircraft per year should be a second consideration in selecting the fleet size.

3. **Utilize a time window for scheduling**
   A time window increases the number of customers that can be serviced with a limited fleet size. The smaller the fleet, the more useful a time window. The size of the time window should be aligned with the wishes of the customer.

4. **Calibrate the optimizer for each new operator/network**
   Real demand data may differ from the utilized demand assumptions. If demand data is different, the home base configuration, fleet size and time windows can be reconfigured to optimize them. This is a continuous process.

5. **Test the optimization algorithm with real demand**
   Reality may manifest possibilities for further optimization. Examples of potential future calibration are:
   - A pre-set cost ceiling for dispatch cost;
   - different end-of-day procedures.

6. **Keep procedures simple**
   A basic end-of-day procedure is able to include basic crew regulations. Each increase in complexity hugely increases the number of feasible solutions that has to be considered. Minimizing the feasible solutions increases the speed and ability of finding an optimal solution.

11.3 Research reflection

11.3.1 Scientific value

The tested algorithm was compared to existing Dial-a-Ride algorithms (see chapter 7). The algorithm offers an effectiveness and flexibility improvement while sacrificing little in the speed of the solution. The complexity of the required calculation power is multiplied by a factor smaller than 10. The calculation efficiency of algorithm can be optimized further through improved implementations.

11.3.2 Business value

The primary business values of the selected and tested solution are flexibility, cost effectiveness and the extra capacity offered by the solution. These will strengthen the business case of air taxi operators.

The simulator can be used as prototype for the actual implementation of the algorithm in Etirc's reservation system.

The tool that was created to perform this research is reusable for different (air taxi) operators. The difference between operators is their network, their fleet specification, their cost data and the demand for their services. Each of these environmental and parametric variables (see Figure 4, paragraph 2.3) can be changed without making changes to the algorithm.

11.3.3 Evaluation of research method

Due to the non-existence of air taxi operators in Europe, no empirical research could be performed. A literature study could not provide solutions that are directly applicable to the specific air taxi problem. The chosen research method (simulations) is the only usable alternative. Simulations are a suitable method when the results of representative tests can be aggregated to form reliable averages. The business case was suitable to supply the required volume of tests. The utilized scenarios are directly based on the assumptions of the case study. The validity of research results is based on the validity of these assumptions.

11.3.4 Internal validation

The test results prove that the selected solution is a suitable solution for the utilized business case. The cost and flexibility results are based on extensive simulation runs of respectively 13500 customer requests and 2000 simulated disruptions. Their results show a high reliability.

The results on the end-of-day procedure are based on a much smaller test base (70 measurement points). The end-of-day results are purely indicative.

11.3.5 External validation

The external validity of the selected solution can be measured at two levels:

1. Other business cases (other operators)
2. Other transportation related fields (other that air taxi)

The created algorithm can be used for other business cases by changing the environmental and parametric variables that have been used in the simulations to describe the network, fleet specification, demand and costs. In the same fashion, the algorithm can be adapted to support any transport related field with similar ondemand characteristics.

The cost effectiveness and flexibility benefits of the selected solution may be different for other business cases and transportation fields. A sensitivity analysis of these changes in circumstances could be performed by using the existing simulation tool with new parameters and running the same tests as described by paragraph 9.5.

11.3.6 Further research

Further research could be performed to calibrate the created algorithm further for the given business case and other business cases. This research can be performed through further simulation. These calibrations would serve more use however if they are performed in live usage.

The current implementation of the selected dispatch algorithm supports customer requests, time windows, disruptions and an end-of-day procedure. Other operational aspects like pre-scheduled maintenance and crew duty duration are not included. Further research could focus on the effect of these variables on total costs.

A basic two-hour time window has been utilized throughout all tests. Further research could be performed on the influence of time windows on the cost effectiveness and flexibility of the dispatch algorithm. This research should include research into the wishes of customers considering time flexibility.

List of definitions

air taxi	a personal means of business travel that uses Very Light Jets (VLJs) to provide on-demand flight services
air-taxi operator	a provider of air taxi services. Owns one or more Very Light Jet aircrafts (VLJ) and uses these to offer personalized flights.
air taxi operations	Operations undertaken by air taxi operators for offering air taxi services.
aircraft route	a path that a single aircraft performs to execute any number of scheduled flights
customer request	an individual trip requested by a customer. Each customer request consists of an origin airport, a destination airport and a requested departure time. synonyms: *request for transport*, *flight reservation*
crew rotation	the route that crew is sent on to perform flight operations. A crew rotation starts with a flight starting from the home base of a crew and ends in a rest period at the home base of that crew. synonyms: *crew routing,crew pairing* or *air duty*
decision variable	a variable in a linear optimization that determines whether a vehicle is assigned to a specific dispatch
dispatch	a customer request that is assigned to a particular aircraft as a part of its route.
dispatch strategy	Strategy that determines what vehicle (taxi/air taxi) is sent out to service a customer request. Urban taxi operators often use a nearest-first or longest idle-time-first strategy or a combination of the two. synonym: *dispatch policy*
empty leg	A flight trip that carries no paying customers. synonyms: *one-way availability*, *deadhead*, *ferry leg*, *non-revenue flight*, *repositioning flight*
end-of-day procedure	A procedure that is run at the end of the day to relocate resources and include mid term optimizations

heuristic	a method of directing problem solving using knowledge based on experience
hotac	Hotel accommodations for crew. A crew requires hotel accommodations if it is required to stay the night at a base other than its home base.
leg	A single non-stop flight between two airports. synonyms: *trip*
non-revenue flight	A flight trip that carries no paying customers. Is the opposite of a revenue flight. See "empty leg".
operational costs	The recurring expenses for air taxi operators which are related to the execution of air taxi flights.
revenue flight	A flight trip that carries passengers that pay for that flight. Is the opposite of a non-revenue flight or empty leg
schedule window	The period between a request for transport being entered and its requested departure time. In other terms: the time between booking and travelling.
scheduling	Creating a planning for the use of aircraft and personnel resources, wherein no resources are used double on any moment in time.
shortest path problem	The problem of finding a path between two nodes such that the sum of the weights of its constituent arcs is minimized (adapted from Wikipedia, July 2007)
VLJ	Very Light Jet aircraft. These aircraft have a maximum take-off weight of less than 4540 kilogram (10.000 lb)

References

[ATA67] Air Transport Association of America (ATA), *Standard method of estimating comparative direct operating costs of turbine powered transport airplanes*, 1967, http://adg.stanford.edu/aa241/cost/atamethod.html

[BAI92] Bailey, W.A. Jr., Clark, T. D. Jr, *Taxi Management and Route Control: A systems study and simulation experiment*, proceedings of the 1992 Winter Simulation Conference, 1992

[BAL98] Ball, M., Roberts, A., *A graph partitioning approach to crew scheduling*, Transportation Science, 1985, vol. 19, no2, pp. 107-12

[BAR98] Barnhart, C., Johnson, E.L., Nemhauser, G.L., et al., *Branch-and-price: column generation for solving huge integer programs*, Georgia Institute of technology, Operations Research, Vol. 46, No. 3, may-june 1998

[BAR03] Barnhart, C., *1.206J / 16.77J / ESD.215J Airline Schedule Planning, Course sheets*, Massachusetts Institute of Technology, Spring 2003

[BOY91] Boyd, S., Barratt, C., *Linear Controller Design: Limits of Performance*, Prentice-Hall, 1991

[BOY06] Boyd, J.A., Bass, E.J., et al, *Investigation of Future Air Taxi Services Using Discrete-Event Simulation*, in 2006 IEEE Systems and Information Engineering Design Symposium.

[CAM04] Campbell, A.M., Savelsbergh, M., *Efficient Insertion Heuristics for Vehicle Routing and Scheduling Problems*, TRANSPORTATION SCIENCE, Vol. 38, No. 3, August 2004, pp. 369–378

[CAV99] Cavique L., Rego C., Themido, I., *3 new heuristics for the crew scheduling problem*, meta-heuristics international conference No2 , FRANCE (21/07/1997) 1999, pp. 37-47

[CLA64] Clarke, G., Wright, J.W., *Scheduling of Vehicles from a Central Depot to a Number of Delivery Points,* Operations Research, Vol. 12, No. 4 (Jul. - Aug., 1964), pp. 568-581

[COR02] Cordeau, J.-F., Laporte, G., 2002. *The dial-a-ride problem: Variants, modeling issues and algorithms*, Quarterly Journal of the Belgian, French and Italian Operations Research Societies, forthcoming.

[COR04] Cordeau, J.F., Laporte, G., *Transportation on Demand*, CRT-2004-25, October 14, 2004

[CRA06] CRA International, *The Economic Impact of Very Light Jets, Final Report*, May 10, 2006

[DES83] Desrosiers, J., Soumis, F., Desrochers, M. Sauve, M., *Routing and Scheduling by branch-and-bound on time windows*, , G-83-16, Ecole des Hautes Etudes Commerciales, Universite de Montreal, 1983

[DES92] Desrochers, M., Desrosiers, J., Solomon, M., *A new Optimization Algorithm for the vehicle routing problem with time windows*, Operations Research, Vol. 40, No. 2. (Mar. - Apr., 1992), pp. 342-354.

[DOR96] Dorigo, M., Gambardella, L.M., *Ant Colony System: A Cooperative Learning Approach to the Traveling Salesman Problem*, IEEE Transactions on Evolutionary Computation, Vol.1, No.1, 1997.

[EUP06] European Union, *REGULATION (EC) No 1899/2006 OF THE EUROPEAN PARLIAMENT AND OF THE COUNCIL*, 27.12.2006, Official Journal of the European Union

[FAA07] http://www.fly.faa.gov/flyfaa/usmap.jsp, July 2007

[FEO89] Feo, T.A., Bard, J.F., *Flight Scheduling and Maintenance Base Planning*, Management Science, Vol. 35, No. 12, December 1989

[FIE99] Fiertz, R., Monico, C., *Cost Allocation Whitepaper*, National Civil Aviation Review Commission, september 1999

[FON92] E. Fontain, *Application of Genetic Algorithms in the Field of Constitutional Similarity*, Journal of Chemical Information and Computer Sciences 32, 748-752, 1992

[FRE03] Freling, R., Huisman, D., Wagelmans, A.P.M., *Models and Algorithms for Integration of Vehicle and Crew Scheduling*, Volume 6, Number 1 / January, 2003

[GHI03] Ghiani, G., Guerriero, F., Laporte, G., Musmanno, R., *Real-time vehicle routing: solution concepts, algorithms and parallel computing strategies*, European Journal of Operational Research, Volume 151, Number 1, 16 November 2003, pp. 1-11(11)

[GOL80] Golden, B., Bodin, L., et al, Approximate Travelling Salesman Algorithms, Operations Research Vol. 28, No. 3, Part II, May-June 1980, p. 694 of 694-711

[GON04] J. F. Goncalves, J.F., Mendes, J.J.M., Resence, M.G.C., A Genetic Algorithm for the Resource Constrained Multi-Project Scheduling Problem, AT&T Labs Research Technical Report: TD-668LM4, October 29, 2004

[GOP98] Gopalan, R., Talluri, K.T., *The Aircraft Maintenance Routing Problem*, Operations Research, Vol. 46, No. 2. (Mar. - Apr., 1998), pp. 260-271

[HOF93] Hoffman, K.L., Padberg, M., *Solving airline crew scheduling problems by branch-and-cut*, Management Science, vol. 39, No. 6, June 1993

[HJO95] Hjorring, C., *The Vehicle Routing Problem and Local Search Metaheuristics*, PhD thesis, Department of Engineering Science, The University of Auckland, 1995

[ISI02] Isikveren, A.T., *Identifying Economically optimal Flight Techniques of Transport Aircraft*, Journal of Aircraft Vol. 39, No. 4 July-August 2002

[JUR07] Jurkat, M.P., *Discrete Event Simulation Systems (DESS)*, Course CS452/Mgt532 Simulation for Managerial Decisions, The Robert O. Anderson Schools of Management, University of New Mexico

[KIT00] Kitamura, R., Chen, C., Pendyala, R.M., Narayanan, R., *Micro-simulation of daily activity-travel patterns for travel demand forecasting*, Business and Economics, Volume 27, Number 1 / February, 2000, Springer Netherlands

[KRO06] Kroon, *Aircraft Design: Synthesis and Analysis*, course textbook, 2006

[LEE06] Lee, D.W., Bass, E.J., et al*., A Traffic engineering model for air taxi services*, August 2006

[LEE06-2] Lee, D.W., Bass, E.J., Patek, S.D., *Characterizing the Aggregate Behavior of an Air Taxi Network*, 11th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference
6 - 8 September 2006, Portsmouth, VirginiaAIAA 2006-7023

[LI05] Li, J., Liu, W., *A Hybrid Genetic Algorithm for the Resource Constrained Multi-Project Scheduling Problem*, Conference on Emerging Technologies and Factory Automation, 2005. ETFA 2005. 10th IEEE, Publication Date: 19-22 Sept. 2005, Volume: 2, On page(s): 6 pp.-

[LUC99] Lucic, P., Teodorovic, D., *Simulated annealing for the multi-objective aircrew rostering problem*, Transportation Research Part A33, 1999 19-45

[LUC02] Lucic, P., *Modeling Transportation Problems Using Concepts of Swarm Intelligence and Soft Computing*, Doctorate dissertation, Virginia Polytechnic Institute, Februari 2002

[MAN06] Mane, M., Crossley, W.A., *Preliminary cost feasibility study of air taxi operations*, 6th IAA Aviation Technology, Integration and Operations Conference, 25-27 September 2006

[MER05] Mercier, A., Soumis, F., *An Integrated Aircraft Routing, Crew Scheduling and Flight Retiming Mode*, Technical Report G-2003-48, GERAD.

[NEM04] Nemhauser, G., *Large Scale Neightborhood Serch in Scheduling an Air Taxi service*, Matheuristics 2006

[OMB04] Ombuki, B. Ross, B.J., Hanshar, F., *Multi-objective Genetic Algorithms for Vehicle Routing Problem with Time Windows*, Technical Report # CS-04-02 , January 2004

[PAN95] Pant, R.K., Gogate, S.D., Arora, P., *Economic Parameters in the Conceptual Design of an Air-Taxi Aircraft*, JOURNAL OF AIRCRAFT, Vol. 32, No. 4, July-August 1995

[SAV05] Savelsbergh, M.W.P., *Scheduling and Air Taxi Service*, Abstract 2005-10-17, http://www.ieor.berkeley.edu/Seminars/Abstracts/2005-10-17.htm

[SCH83] Schoeter, J.R., *A model of taxi service under fare structure and fleet size regulation*, The Bell Journal of Economics, Vol. 14, No. 1 (Spring 1983), pp. 81-96

[SHR98] Shrivasta, M., Chande, P.K., Monga, A.S., *Taxi despatch: A fuzzy rule approach*, 1998

[SIN93] Forbes, G. W., Andrew, Jones, A. E. W., *Global Optimization in Lens Design*, Sinclair Optics Design Note, 1993, http://www.sinopt.com/learning1/desnotes/globopt.htm

[TEO88] Teodorovic, D., *Airline Operations Research*, Transportation Studies Volume 10, Gordon and Breach Science Publishers, 1988

[TRI98] Trick, M.A., *A Tutorial on Dynamic Programming*, http://mat.gsia.cmu.edu/classes/dynamic/node4.html

[WON03] Wong, K.I., Wong S.C., Yang, H., Wu, J. H., *Modeling Urban Taxi services with Multiple User Classes and Vehicle Modes*, 2003

Appendix A: European legislation on flight time limitations

Summarized version of European legislation. This new legislation will be enforced as of July 16th 2008.
The minimum rest which must be provided before undertaking a flight duty period starting at home base shall be at least as long as the preceding duty period or 12 hours whichever is the greater;
The minimum rest which must be provided before undertaking a flight duty period starting away from home base shall be at least as long as the preceding duty period or 10 hours whichever is the greater; when on minimum rest away from home base, the operator must allow for an eight-hour sleep opportunity taking due account of travelling and other physiological needs;
An operator shall ensure that the minimum rest provided as outlined above is increased periodically to a weekly rest period, being a 36-hour period including two local nights, such that there shall never be more than 168 hours between the end of one weekly rest period and the start of the next. As an exception to OPS 1.1095(1.9), the Authority may decide that the second of those local nights may start from 20:00 hours if the weekly rest period has a duration of at least 40 hours.

Duty
Standard Flight Duty Period (FDP) (without extensions) is 13h for up to 3 sectors (minus 30 mins per additional sector up to 2hrs). If FDP starts on between 02:00 and 05:59 (WOCL) the flight duty is reduced by amount of time that falls within this period (with a maximum 2hr reduction). If FDP ends between 02:00 and 05:59, FDP is reduced by 50% of the time that falls into this period. Extensions: FDP may be extended one hour by the operator, maximum 2 times a week. Limited to 6 sectors (4 sectors if WOCL encroached less than 2hs, 2 sectors if more). Increased rest: 2 extra rest hours pre and post FDP or 4 post FDP. If FDP starts between 22:00 and 04:59 maximum FDP with extensions is 11:45h

Rest

Daily at home:	as long as preceding FDP or 12hrs (whichever is greater)
Daily away:	as long as preceding FDP or 10hrs (whichever is greater)
Weekly:	36hrs with 2 local nights in 7consecutive days (no more than 7 days between 2 successive weekly rest periods)
Yearly:	96 days
Augmented crew:	13hrs + 3 if "suitable place outside the cockpit" 13hrs+ 8 if "horizontal rest facilities available”

Standby
Airport standby counts in full for cumulative duty
Other types of Standby and standby’s impact on FDP will be defined nationally (no harmonisation)

Source of summary http://www.eurocockpit.be/media/FTL_In_Short_NT-05-1108-F.pdf
Full version:
http://eur-lex.europa.eu/LexUriServ/site/en/oj/2006/l_377/l_37720061227en00010175.pdf