Fleet Sizing & Scheduling of Electric Thin-Haul Aircraft

A Column Generation and Large Neighborhood Search Based Method

Bart Debeuckelaere

Fleet Sizing & Scheduling of Electric Thin-Haul Aircraft

Using a Column Generation & Large Neighborhood Search Based Algorithm

by

Bart Debeuckelaere

to obtain the degree of Master of Science
at the Delft University of Technology,
to be defended publicly on June 23rd 2021.

Student number:	4537467
Project duration:	Sep 2020 – Jun 2021
Supervisor:	Dr. Mihaela Mitici
Thesis committee:	Dr. Dick Simons
	Dr. Donatella Zappala

TU Delft

Acknowledgements

I would like to thank my supervisor, Dr. Mihaela Mitici for the guidance throughout the project. Thanks as well to Dr. Theresia Van Essen, Dr. Alessandro Bombelli and Mr. Juseong Lee for the feedback and for helping me overcome difficulties along the way.

To my roomies, Bloempje, Floflo, Mishy, Seba & Skipper, thanks for all the weird and wonderful moments that made the past 5 years enjoyable,

A huge thanks to my family. To my parents for always supporting me in every choice I made. For the love and for being trusted advisors. Thanks as well to my sisters, Elise & Olivia, for low key being there for me.

Delft, July 2020

Contents

List of Figures

List of Tables

List of Abbreviations

AR	Aspect Ratio
ASK	Available Seat Kilometer
B&B	Branch-and-Bound
B&P	Branch-and-Price
CG	Column Generation
E-VRP	Electric-Vehicle Routing Problem
E-VSP	Electric-Vehicle Scheduling Problem
ESK	Empty Seat Kilometer
G-VRP	Green-Vehicle Routing Problem
LNS	Large Neighborhood Search
MILP	Mixed Integer Linear Programming
MTOW	Maximum Take-Off Weight
PSP	Pricing Sub-Problem
RC	Reduced Cost
RMP	Restricted Master Problem
RPK	Revenue-Passenger Kilometer
TAT	Turn-Around-Time
ToD	Time-of-Day
TS	Time-Space
UAM	Urban Air Mobility
VRP	Vehicle Rougint Problem
VSP	Vehicle Scheduling Problem
VTOL	Vertical Take-Off and Landing

1
Introduction

Aviation must become more sustainable, currently emitting 2% of all human-induced CO2 emissions, 12% of all CO2 emissions attributed to transport [2]. Thanks to rapidly advancing technologies, electric aircraft could be part of the solution. While electric VTOL (Vertical Take-Off and Landing) receive a lot of attention from researchers, many hurdles remain in regulation, certification and public acceptance. Meanwhile some fixed-wing electric aircraft with up to 10 seats are already in the flight testing phase. Electric propulsion is already common place for gliders equipped with self-starter motors and Pipistrel and Siemens are already flying electric ultralights. while electric aircraft start entering general aviation, the logical next step is electric aircraft with a seating capacity of 6 to 10 people. MagniX started test flying an electrified Cessna Caravan in June 2020, Eviation also already has a prototype of their electric aircraft and Skylax intends to start flight testing early 2022. These aircraft are entering the ballpark of thin-haul airlines, who provide scheduled services over small distances ranging from just ten to a couple hundreds of kilometers in small aircraft with up to 10 seats. Due to their high costs of €0.40/ASK compared to an airlines €0.11/ASK [1], thin-haul airlines are only found where there are no alternatives for fast transport such as the Hawaiian islands or Alaska. Economic feasibility studies conclude that these electric aircraft can provide huge costs savings for thin-haul airlines [1] [3], [4]. Furthermore, feasibility studies show that lower prices may even unlock a huge amount of latent demand such that electric thin-haul air transport may appear in many more places where it currently does not exist. Yet there is very little research dedicated to their operations.

With these promising applications, the need for operations research in electric-thin haul aviation is clear. With a number of aircraft swiftly advancing in their development, the urgency is also abundantly clear. Prior work explored the feasibility so the logical next step is to investigate how to operate them, which is the gap that this thesis attempts to close. In doing so, determining which service trips to make must be part of the problem and cannot be known in advance. Whereas it may be realistic to assume that the set of service trips is known in a scenario where a traditional fleet is replaced by an electric fleet, this may not be a good solution as it is not tailored to electric aircraft with limited range and long charging times. Furthermore, the service trips cannot be known in advance when the electric aircraft are not replacing an already existing fleet. In the fiercely competitive airline industry, An operator will be interested in designing a schedule that minimises cost per RPK (Revenue-Passenger Kilometer) because this has been the main reason why thin-haul air transport largely missed out on the growth of the airline industry. This metric also allows for a direct comparison of the economic competitiveness versus other modes of transport and gives a lower bound on ticket fares in order to be profitable.

This thesis report is organized as follows: In chapter 2, the scientific paper is presented where we discuss a novel method to solve our problem. We applied the methods to a number of case studies and discuss the results. Chapter 3 contains the Literature Study that preceded and supports this project. We focus on prior work on the E-VRP (Electric-Vehicle Routing Problem), E-VSP (Electric-Vehicle Scheduling Problem) and feasibility studies in electric thin-haul operations. Appendix A presents how the problem was developed. Appendix B discusses a Branch-and-Price. Appendix C presents an alternative LNS (Large Neighborhood Search) algorithm, which was discarded in favor of the LNS algorithm presented in the scientific paper. Lastly, Appendix D summarises the results of all case studies.

2
Scientific Paper

Fleet Sizing and Scheduling of Electric Thin-Haul Aircraft, a Column Generation and Large Neighborhood Search Based Method

Bart Debeuckelaere,*

Delft University of Technology, Delft, The Netherlands

Abstract

This paper proposes an optimisation model to minimize operational cost of a fleet of electric thin-haul aircraft under a minimum RPK (Revenue-Passenger Kilometer) constraint. A solution that minimises cost per RPK can be found by varying the minimum RPK level. The solutions describes how many aircraft are needed, as well as a schedule for each aircraft for a single day. We consider a set of airports where charging infrastructure is assumed to be present. All aircraft must start and end the day from the assigned hub airport. Next, we consider a discretised time space between a start time and end time with constant time steps. The problem is represented on a time-space network where each node uniquely defines a location (airport) and point in time, and arcs connect the nodes. An arc connecting two consecutive nodes at the same airport are ground arcs and represent waiting on the ground. An arc connecting two nodes at different airports are flight arcs. The cost, duration and energy consumption on flight arcs is determined in advance. A schedule is represented as a sequence of arcs. The number of passengers on a flight is limited by the demand. We developed a method to find a schedule that minimises costs while meeting a minimum RPK constraint. This method was then illustrated on a network with 5, 10, 15, 20 and 30 airports. The results show that our method is much faster than a traditional linear programming model. The obtained solution is a local minimum and is very close to the global optimum. We found that cost per RPK shows a sawtooth pattern, gradually decreasing as aircraft utilisation increases but spiking up when an additional aircraft is added to the fleet. Furthermore, the sawtooth pattern shows an increasing trend. The schedule shows a strong preference for connections with a longer distance and with high demand. The algorithm first fills up these connections, often using back-and-forth flights, until they are largely saturated. When there is little demand left on these connections, the algorithm starts adding flights on connections with a shorter distance or less demand.

Keywords: scheduling, electric vehicle, electric aircraft, thin-haul, fleet sizing, linear optimisation, column generation, large neighborhood search

1 Introduction

Aviation must become more sustainable. Aircraft are currently emitting 2% of all human-induced CO2 emissions globally, or 12% of all transport related CO2 emissions [8], and the public opinion is becoming more opposed to flying. Electric aircraft could be part of the solution. Recent developments in battery and charging technology have been very promising. In operations research, a lot of work has been dedicated to UAM (Urban Air Mobility). In this concept, an autonomous VTOL (Vertical Take-Off & Landing) aircraft with a handful of seats available provides on-demand mobility. Though technology is developing at a high pace, there are many hurdles to this concept in regulation, certification and public acceptance. Meanwhile fixed-wing electric aircraft with up to 10 seats are already in the flight testing phase. A number of economic feasibility studies conclude that these electric aircraft provide opportunities for huge costs savings for commuter airlines [3] [27], [31]. Yet there is very little research dedicated to their operations, which is vastly different from fossil fuel powered aircraft.

Some experts still consider electric propulsion infeasible for commercial aviation and advise research to focus on electric regional airlines with a potential market introduction 20 years from now. This is a misconception according to Moore et al. [6], who states the importance of having an agile and adaptable research and market plan, starting with smaller applications. Current commercial aviation is receiving most attention because that is where the bulk of the revenue is but there is a poor fit with current technology levels. Meanwhile, electric propulsion is already common place for gliders equipped with so called self-starter motors and Pipistrel and Siemens are already flying electric ultralights. while electric aircraft start entering general aviation, the logical next step is electric aircraft with a seating capacity of 6 to 10 people. MagniX started test flying an electrified

*Msc Student, Air Transport and Operations, Faculty of Aerospace Engineering, Delft University of Technology

Cessna Caravan in June 2020, Eviation also already has a prototype of their electric aircraft and Skylax intends to start flight testing early 2022.

Clearly, electric aircraft are entering the ballpark of thin-haul commuter airlines, who provide scheduled services over small distances ranging from just ten to a couple hundreds of kilometers in small aircraft with up to 10 seats. Due to their currently high costs of 0.40 €/ASK compared to an airlines 0.11 €/ASK [3], commuter airlines are only found where there are no alternatives for fast transport such as the Hawaiian islands or Alaska. As current developments and research plainly show that electric aircraft can substantially reduce costs, some current operators are already planning to electrify their fleets. Furthermore, feasibility studies show that lower prices may even unlock a huge amount of latent demand such that electric thin-haul air transport may appear in many more places where it currently does not exist. Yet, very little operations research is devoted to operating an electric fleet.

With the very promising applications to replace current thin-haul fleets and even provide new thin-haul services where it currently does not exist, the need for operations research in electric-thin haul aviation is clear. With a number of promising concepts rapidly advancing in their development, the urgency is also abundantly clear. Prior work explored the feasibility of such air transport so the logical next step is to investigate how to operate them. In doing so, determining which service trips to make must be part of the problem and cannot be known in advance. Whereas it may be realistic to assume that the set of service trips is known in a scenario where a traditional fleet is replaced by an electric fleet, this may not be a good solution as it is not tailored to electric aircraft with limited range and long charging times. Furthermore, the service trips cannot be known in advance when the electric aircraft are not replacing an already existing fleet. In the fiercely competitive airline industry, An operator will be interested in designing a schedule that minimises cost per RPK (Revenue-Passenger Kilometer) because this has been the main reason why thin-haul air transport largely missed out on the growth of the airline industry. This metric also allows for a direct comparison of the economic competitiveness versus other modes of transport and gives a lower bound on ticket fares in order to be profitable.

This paper therefore proposes a method to minimise cost under a minimum RPK constraint. A solution that minimizes cost per RPK can be found by varying the minimum RPK level. We are interested in finding i) What fleetsize is required to meet a certain RPK? ii) How should the aircraft be scheduled? What makes up a successful schedule? And iii) Is there a pattern w.r.t. cost per RPK?

The contributions of this paper are as follows. To the best of our knowledge, this is the first paper that considers cost per RPK as an evaluation criterion in an electric vehicle scheduling problem and is the first paper to apply operations research to electric thin-haul aircraft. We propose a new variation of the electric-vehicle scheduling problem where the set of service trips is not known in advance and where a meeting a minimum RPK level must be met. We then propose a method to solve the problem, tailored to the intricacies of the problem. The methods and insights developed in this paper will support future thin-haul airline operators in deciding how many aircraft to acquire and provide a schedule that minimises cost per RPK.

The remainder of this paper is structured as follows. In Sec. 2, the state-of-the-art in electric thin haul aviation and electric-vehicle scheduling and routing is discussed. In Sec. 3, we describe the formulation of our problem. In Sec. 4, we propose a novel method to solve our problem. The case studies to evaluate the methods are described in Sec. 5. Lastly, the results of the case studies are laid out in Sec. 6.

2 Literature Review

This chapter reviews the state-of-the-art in operations research of electric aircraft and closely related fields. The research in the field of electric aircraft scheduling is very limited. Three fields of research are deemed most relevant. Firstly, Sec. 2.1 considers the E-VRP (Electric-Vehicle Routing Problem), which applies the classical VRP (Vehicle Routing Problem) to an electric fleet. A set of customer nodes is represented on a graph. The problem consists in finding an optimal set of routes such that each customer node is visited exactly once. The E-VRP is known to be NP-Hard and most research uses heuristics to efficiently solve the instances. A second relevant field of research considers the E-VSP (Electric-Vehicle Scheduling Problem) and is discussed in Sec. 2.2. Instead of a set of customer nodes as is the case in the VRP, the VSP (Vehicle Scheduling Problem) is concerned with a set of service trips. Each service trip has a start and end location and duration. The problem is to schedule a fleet of electric vehicles such that each service trip is performed under the given constraints. In the single-depot problem, all vehicles start and end their service from the same vehicle depot. A variation to this problem is the multi-depot problem. The single-depot VSP is Polynomial bound ($O(N^3)$) but multi-depot VSP and the E-VSP are NP-Hard [30]. Lastly, a small number of feasibility studies and other considerations towards electric aircraft are presented in Sec. 2.3.

2.1 The Electric-Vehicle Routing Problem (E-VRP)

Erdogan and Miller-Hooks [26] are the first to discuss routing for alternative vehicles. They describe the G-VRP (Green-Vehicle Routing Problem) to help organisations overcome the difficulties of operating vehicles using alternative power sources of limited driving range and limited refuelling infrastructure. The problem is represented on a graph $G(V, E)$. The set of vertices or nodes V consists of the set of customer nodes, refuelling station nodes and one depot node. Each node can only be visited once. To allow for multiple visits to the refuelling stations, the author decided to create the augmented graph $G(V, E)$ which contains a number of copies of the refuelling station nodes. The objective is to minimize the total travelled distance, as shown in Eq. (1).

$$min \sum_{ij} d_{ij} x_{ij} \qquad \forall ij \in V', i \neq j \tag{1}$$

The original G-VRP [26] provided a very good baseline for future work in electric vehicle routing. A number of extensions were proposed to make the model more practically relevant and a number of solution techniques have been developed. These include partial recharge, time-windows and more. Table 1 summarises the current literature concerning the E-VRP. The fourth column denotes whether the servicing is modelled using a constant servicing time (C), a partial recharge (RP) or a full recharge (FR). The problem is NP-Hard and a range of heuristic methods were proposed to solve real life instances. Further work in this field includes studying extensions to the problem that could be relevant in practice while further improving on the methods to efficiently handle larger instances.

Paper	**Objective**	**Time Windows**	**Service**	**#Customers**	**# Vehicles**	**Exact**
[26]	Distance	-	C	500	80	-
[1]	Distance	+	FR	500	80	-
[7]	Operating Cost	-	PR	500	80	-
[11]	Operating Cost	+	PR	50	?	+
[25]	Operating Cost	+	PR	100	15	+
[18]	Distance	+	PR	100	15	-
[20]	Distance	-	C	20	10	+
[5]	Distance	+	C	20	10	+
[2]	Total Cost	+	C/FR	100	15	-
[10]	Charging Time	+	PR	100	15	-

Table 1: Summary of E-VRP research papers

2.2 The Electric-Vehicle Scheduling Problem (E-VSP)

The traditional VSP (Vehicle Scheduling Problem) is a well known problem. The problem consists of designing an optimal schedule for a fleet of vehicles such that a set of pre-determined service trips is performed at a minimal cost. After the publications of Li [13] and Adler [14], addressing the scheduling of electric vehicles, this field of research has been extremely active.

In most cases, the objective is to minimize the operational costs while executing each service trip. The models are usually formulated on a graph. Each arc or node in the graph has certain attributes such as cost and energy consumption. The task is to select a set of arcs to form the schedule of one or more vehicles. A number of variations of the original E-VSP have been proposed in literature as shown in Table 2. In the Service column, C, PR and FR denote constant time, partial recharge and full recharge respectively. Though the problem is NP-Hard, exact methods based on Column-Generation algorithms often do the job. The problem has been applied in a number of case studies of electric bus fleets. One downside to this approach is that the set of service trips must be determined beforehand. Steiner and Inrich [16] present a variation with a dynamic demand pattern as input such that determining which service trips to make becomes part of the problem. Further work is this field will be studying different extensions and variations to the problem in specific case studies.

Paper	Vehicle	Objective	Service	Infr.	#Trips	# Vehicles	Exact
[13]	Bus	Operating Cost	C	-	1,000	60	+
[14]	Bus	Operating Cost	C	-	4,300	500	+
[30]	Bus	Operating Cost	PR	-	500	30	+
[12]	Bus	Annual Cost	C	+	600	30	+
[9]	Bus	Annual Cost	FR	+	300	35	+
[16]	(Diesel) Bus	Profit	N/A	N/A	18	1	+
[29]	Bus	# vehicles	PR	+	700	35	-
[32]	Bus	Operating Cost	C	-	60	13	-
[15]	Bus	Annual Cost	PR	+	?	700	+
[21]	Aircraft	Operating Cost	PR	-	50	25	+

Table 2: Summary of E-VSP research papers

2.3 Electric Flight

Electric technology is developing at an extremely high pace. This has lead to a number of common misconceptions about electric propulsion that are preventing research from reaching its full potential [6]. Therefore, this section will discuss electric aircraft technology and operations in more detail with a focus on thin-haul commuter aircraft.

Thin-haul or Commuter refers to one step smaller than short-haul. Flight legs vary in distance from ten to a couple hundred kilometers using small aircraft with no more than 10 seats. Due to their high cost, they are only found where there is no high speed alternative. Harish et al. [3] assessed the economic feasibility of electric thin-haul operations. The demand on each individual route in a thin-haul network may be small but the cumulative demand over a network provides opportunities. Consider as well that latent demand may emerge if the service becomes successful and prices drop. Although the demand is arguably there, thin-haul airlines have not experienced the growth rates of commercial transport before the COVID crisis. That is mainly caused by the high operating costs. Whereas traditional airlines such as have an operating cost of €0.11 per ASK (Available Seat Kilometer) (Delta Air Lines) compared to €0.40 for commuter airlines (Cape Air) [3]. The high operating costs and alternative fast transport make it very hard for commuter airlines to be successful. From a direct comparison, using electric aircraft results in a staggering 20 to 30% reduction in operating costs, mainly driven by a reduction in energy and maintenance costs and partly offset by an increased acquisition cost of the aircraft and batteries.

To assess the potential demand, Sun et al. [31] estimated door-to-door travel time of existing modes of transport and air taxis, which essentially are commuter aircraft. They found that commuter aircraft are most competitive for distances between 100 and 300km. A similar study by Wei et al. [28] concludes that a suburban commuter airline should target ground commutes exceeding 45 minutes. With a cruise speed of 160 knots and take-off distance of 500 feet (300kph, 150m), travel times can be reduced by 45%. According to Courtin et al. [22], these requirements are realistic.

3 Problem Description

This section describes how the problem is formulated and introduces mathematical notations. We consider a homogeneous set of aircraft K. The exact fleet size is not known in advance but rather a decision to make in our problem. Let a_1, a_2, ... be airports that the aircraft can operate at. Let one of these airports be the hub a_h, where all aircraft are stationed overnight. We consider a discretised time space, starting and ending at t_{min} and t_{max} respectively with a time step t_{step}. Let $i(a,t)$ be a node defining airport a at time t where $t \in [t_{min}, t_{max}]$ and $t = t_{min} + nt_{step}$ ($n \in \mathbb{N}$). Let N denote the set of all nodes. Denote the first and last node at the hub as the origin node $N_o(a_h, t_{min})$ and destination node $N_d(a_h, t_{max})$. Let N' be the set of all nodes except for the origin and destination node such that $N = N_o \cup N_d \cup N'$. Let two consecutive nodes at the same airport $i(a_1, t_1)$ and $j(a_1, t_1 + t_{step})$ be connected by a ground-arc (i,j). Next, consider two airports a_1, a_2 where the distance $d_{(a_1,a_2)}$ between them lies within an interval $[d_{min}, d_{max}]$. Suppose the flight-time $t_{(a_1,a_2)}$ to fly from airport a_1 to airport a_2 is known. Let nodes $i(a_1,t_1)$ and $j(a_2, t_1 + t_{(a_1,a_2)})$ be connected by a flight-arc (i,j). Let A_g and A_f be the set of ground- and flight-arcs respectively. Let $A = A_g \cup A_f$ be the set of all arcs and let δ_i^+ and δ_i^- be the set of incoming and outgoing arcs at node i respectively. Knowing the nodes i and j is sufficient to uniquely define the arc (i,j) connecting them. Each arc $(i,j) \in A$ has a cost $c_{(i,j)}$ indicating how much it costs to use this arc, and an energy $e_{(i,j)}$, indicating how the energy level in the aircraft's battery changes over that arc. On a flight-arc, $e_{(i,j)}$ has a negative value as energy is consumed during the flight. On ground-arcs, the aircraft can recharge it's battery. By assuming a linear charging process, the energy in ground arcs is set equal to the charging power times the duration of the arc $e_{(i,j)} = P_{charge}t_{step}$. Let $G(N, A)$ be a TS (Time-Space) graph, comprising of the set of nodes N and set of arcs A.

Demand is defined for each airport pair in time windows with length t_D. Let $D_{((a_1,a_2),[t_{min}+nt_D,t_{min}+(n+1)t_D))}$ be the maximum number of passengers that can be flown from airport a_1 to airport a_2 on a flight departing in the interval $[t_{min}+nt_D, t_{min}+(n+1)t_D)$ where $n \in [0, 1, \cdots, (t_{max}-t_{min})/t_D]$ to create non-overlapping intervals, spanning the time between t_{min} and t_{max}. In the remainder of this paper, we will refer to a combination of an airport pair and time interval as a demand interval. Let D be the set of all possible demand intervals. The number of passengers on each flight-arc is limited by either the number of seats s available on the aircraft, or by demand. On each flight, the number of passengers may not exceed the number of seats s on the aircraft. Next, let $A_{((a_1,a_2),[t_1,t_2))}$ be the set of all flight-arcs serving demand from airport a_1 to a_2 departing in the time interval $[t_1, t_2)$. The total number of passengers on all flights in this demand interval $((i, j) \in A_{((a_1,a_2),[t_1,t_2))})$ cannot exceed the demand in that interval $D_{((a_1,a_2),[t_1,t_2))}$.

Determining the exact number of passengers on each flight does leave symmetries in the problem, different solutions with the same outcome. Suppose there are 2 flights serving the same demand interval such that capacity exceeds demand $(2s > D_{((a_1,a_2),[t_1,t_2))})$. Note that costs are not affected by the number of empty seats on each flight. The solution where passengers are distributed $(s,\ D_{((a_1,a_2),[t_1,t_2))} - s)$ (the first flight taking s passengers while the second takes the remaining passengers) has the same cost as the solution where passengers are distributed the other way around $(D_{((a_1,a_2),[t_1,t_2))} - s,\ s)$. However, the exact number of passengers on each arc does not have to be known. It suffices to know the number of empty seats ('loss') $l_{((a_1,a_2),[t_1,t_2))}$ in a demand interval. Suppose that a number of f flights are scheduled in the demand interval $((a_1, a_2), [t_1, t_2))$. The value of the related loss variable can be calculated using Eq. (2) without specifying on which flights they occur. This simplifies the model by eliminating these passenger-symmetries. Note that it is also possible to have a higher value for the loss variable. This would mean that seats are left empty while there actually is demand for them. In reality, the operator will always attempt to sell as many seats as possible.

$$l_{(a_1,a_2),[t_1,t_2)} = \begin{cases} 0 & \textbf{if } f \cdot s \leq D_{((a_1,a_2),[t_1,t_2))} \\ f \cdot s - D_{((a_1,a_2),[t_1,t_2))} & \textbf{else} \end{cases} \tag{2}$$

A feasible schedule can be constructed on graph G by selecting a subset of arcs, adhering to the following rules. An example of how a schedule may look on a graph is visualised in Fig. 1

- The aircraft must start at the origin-node N_o and end at the destination-node N_d. This represents starting and ending the day at the assigned Hub airport.
- Between any two consecutive flights, the aircraft must stay on the ground for at least the TAT (Turn-Around-Time) t_{tat} before departing for another flight.
- The aircraft starts at N_o with a fully charged battery. Over each arc, the energy changes according to the energy of the arc. If the energy exceeds the battery capacity e_{max}, it is set equal to e_{max}. Suppose the energy in the battery at node i is e_i. Then if arc (i, j) is used, the energy in the battery at node j is $e_j = min(e_i + e_{(i,j)}, e_{max})$. The energy in the batteries is not allowed to drop below zero at any point in the schedule.
- The selected arcs must form a continuous path.

By slightly altering the graph G, we can ensure that any continuous path from N_o to N_d automatically satisfies the TAT. We achieve this by including the TAT in the flight-arcs. Let the flight-arc (i, j) connect nodes nodes $i(a_1, t_1)$ and $j(a_2, t_1+t_{(a_1,a_2)}+t_{tat})$. During the TAT, the aircraft is on the ground so has the option to recharge the batteries. If $e_{(a_1,a_2)}$ is the energy consumed during the flight and $e_{tat} = P_{charge}t_{tat}$, then the change in energy over this arc is $e_{(i,j)} = e_{(a_1,a_2)} + e_{tat}$. This also implies that now the energy is not allowed to drop below e_{tat} at any node. If it is lower, that means that the energy dropped below zero at the end of a flight and partly recharged during the TAT.

The cost of this schedule equals the sum of the costs of all arcs that make up the schedule. Let the RPK (Revenue-Passenger Kilometer) of a flight equal the number of revenue-generating passengers (so excluding crew) multiplied by the flight distance. The RPK of a schedule is the summation of the RPKs of all flights in that schedule. Recall that the exact number of passengers wasn't known. Instead, the number of available seats is known as well as the number of empty seats in each demand interval. The total RPK of a schedule can thus be computed as the ASK (Available Seat Kilometer) minus the ESK (Empty Seat Kilometer). The ASK

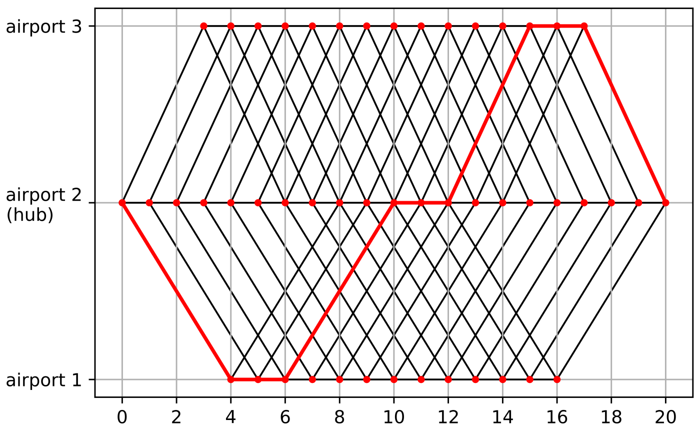

Figure 1: Example of a schedule where the aircraft flies from airport 2, to airport 1, back to airport 2, to airport 3 and back to airport 2. Airport 2 is the hub in this case. In between flights, the aircraft stays on the ground for two time steps.

of a schedule is the summation of ASKs of all flights in the schedule, whereas the ESK is the summation of the ESKs in each demand interval.

The objective is to create a schedule that minimises Cost/RPK, which is a nonlinear objective. Therefore, we change the objective to creating a schedule that minimises cost while meeting a minimum RPK level RPK_{min}. This can be formulated as a MILP (Mixed-Integer Linear Program). By repeatedly solving this problem for various minimum RPK levels, a solution that minimizes Cost/RPK can be found.

Let x_{ijk} be a binary decision variable indicating whether aircraft $k \in K$ is using arc $(i, j) \in A$. Let e_{ik} be a continuous decision variable indicating the energy level of aircraft $k \in K$ at node $i \in N$. Lastly, let $l_{((a_1,a_2),[t_1,t_2))}$ be a continuous 'loss' variable, indicating the number of empty seats on all selected flights in demand interval from airport a_1 to airport a_2 on all flights departing in the time interval $[t_1, t_2)$ where $((a_1, a_2), [t_1, t_2)) \in D$. Because the schedule is constructed by selecting which arcs to include, we call this the 'Arc-Based Formulation'.

(ARC BASED FORMULATION)

Decision Variables

$$x_{ijk} = \begin{cases} 1 & \textbf{if } \text{aircraft } k \text{ is using arc } (i,j) \\ 0 & \textbf{else} \end{cases}$$

$e_{ik} \in [e_{tat}, e_{max}]$ the energy level in the battery of aircraft k at node i

$l_{((a_1,a_2),[t_1,t_2))} \geq 0$ the total number of empty seats on flights from airport a_1 to airport a_2 departing in the time interval $[t_1, t_2)$

MILP formulation

$$min \sum_{(i,j)\in A} \sum_{k\in K} x_{ijk} c_{ij} \tag{3}$$

s.t.

$$\sum_{(i,j)\in\delta_i^+} x_{ijk} = \sum_{(i,j)\in\delta_i^-} x_{ijk} \qquad \forall i \in N', k \in K \tag{4}$$

$$\sum_{(i,j)\in\delta_o^-} x_{ijk} \leq 1 \qquad \forall k \in K \qquad (5)$$

$$\sum_{k\in K} x_{ijk} \leq 1 \qquad \forall (i,j) \in A_f \qquad (6)$$

$$\sum_{(i,j)\in A_f}\sum_{k\in K} x_{ijk}\cdot s\cdot d_{ij} - \sum_{((a_1,a_2),[t_1,t_2])\in D} l_{((a_1,a_2),[t_1,t_2])} d_{(a_1,a_2)} \geq RPK_{min} \qquad (7)$$

$$\sum_{(i,j)\in A_{((a_1,a_2),[t_1,t_2])}}\sum_{k\in K} x_{ijk}\cdot s - l_{((a_1,a_2),[t_1,t_2])} \leq D_{((a_1,a_2),[t_1,t_2])} \qquad \forall((a_1,a_2),[t_1,t_2]) \in D \qquad (8)$$

$$e_{jk} \leq e_{ik} + x_{ijk}e_{ij} + (1-x_{ijk})e_{max} \qquad k\in K, (i,j)\in A \qquad (9)$$

$$e_{tat} \leq e_{ik} \leq e_{max} \qquad \forall k\in K, i\in N \qquad (10)$$

$$x_{ijk} = \{0,1\} \qquad \forall (i,j)\in E, k\in K \qquad (11)$$

$$l_{((a_1,a_2),[t_1,t_2])} \geq 0 \qquad \forall((a_1,a_2),[t_1,t_2]) \in D \qquad (12)$$

The Objective (3) is to minimise total operating costs. Constraint (4) ensures flow conservation for each aircraft on each node, except for the origin and destination nodes. Constraint (5) ensures that each aircraft can be deployed at most once. Constraint (6) certifies that a flight arc can be used by at most one aircraft. Constraint (7) guarantees that a sufficient RPK is served. Constraint (8) verifies the passenger demand limits in each demand interval, where several flight arcs may serve the same demand interval. Constraint (9) regulates the energy level. The last component is needed such that non-used arcs do not impose a restriction on the energy level. Constraints (10), (11) and (12) regulate which values the decision variables may assume. x variables are binary, e are continuous within the limits of the battery and the l variables should be positive integers but are relaxed to positive continuous decision variables. Suppose some l in the solution has a non-integer value. When rounded down to the nearest integer, Constraint (8) will still be satisfied as all other terms in this constraint are integer and constraint (7) will become less tight. This solution is therefore still feasible and thus l can be relaxed without loss of generality.

This problem is extremely difficult to solve. That is because it is a resource-constrained vehicle-scheduling problem, which is NP-Hard [13]. While this formulation is relatively simple and can easily be plugged into a commercial solver like Gurobi, its NP-Hardness and the fact that the graph G can easily contain hundreds or thousands nodes and arcs make it impractical in most realistic cases. Moreover, the arc-based formulation has a lot of seemingly useless of decision variables and symmetries. For each aircraft added to the fleet, there is a new set of decision variables on each arc and on each node, even those that are not used by the aircraft. Symmetries remain in two forms. i) Suppose aircraft #1 uses a path (feasible sequence of arcs) and aircraft #2 takes a different path. Assigning the same path to the other aircraft results in the same solution. When the fleet size increases, the number of permutations for any of set of paths increases exponentially. ii) Suppose a path contains a flight arc $(i(a_1,t_1), j(a_2,t_2))$. Replacing this flight arc by the same flight at the next time step $(i'(a_1,t_1+t_{step}), j'(a_2,t_2+t_{step}))$ often results in the same objective value. Section 4 therefore discusses a novel method to solve the problem.

4 Methodology

Because the arc based formulation in Sec. 3 is only able to solve very small instances to optimality, this section presents a new optimisation algorithm to solve this problem. Firstly, Sec. 4.1 proposes a label correcting algorithm to heuristically construct a good initial solution. Starting from this initial solution, a CG (Column Generation) algorithm explained in Sec. 4.2 computes a strong lower bound by repeatedly solving a smaller version of the problem called the RMP (Restricted Master Problem) until the optimality criterion is met. CG solves the relaxed RMP to find a lower bound. Next, the integer version of the final RMP is used to improve on the initial solution. To obtain a global optimum, CG must be implemented in a B&P (Branch-and-Price) framework as explained in Appendix B. Because B&P proved unsuccessful to find a global minimum, an LNS (Large Neighborhood Search) algorithm was developed in Sec. 4.3 to find a local minimum.

4.1 Initial Solution Heuristic Using Label Correction

One may realise that the solution with the minimum cost per RPK will likely have a high aircraft utilisation, which means that the aircraft are performing as many flights as possible. In that case, the fixed, daily ownership cost of the aircraft is shared over a high RPK and contributes less to the overall Cost/RPK. Intuitively, a high aircraft utilisation makes sense. This section proposes an algorithm to heuristically construct an initial feasible solution with a high aircraft utilisation, which can also be interpreted as heuristically minimising the fleet size. The heuristic works by repeatedly finding the path from N_o to N_d with the highest RPK, given the previously selected paths, and adding this to the initial solution. This process repeats until the total RPK of all paths in the initial solution combined exceeds the minimum RPK as shown in Fig. 2.

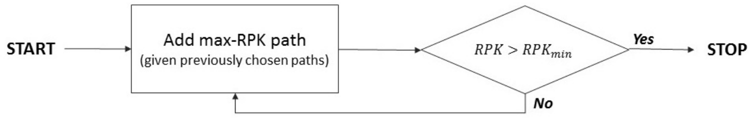

Figure 2: Process of the heuristic to construct an initial feasible solution

The problem of finding the maximum RPK path can be formulated as the following MILP program.

(MAXIMUM RPK PATH)

Decision Variables

$$x_{ij} = \begin{cases} 1 & \textbf{if } \text{arc } (i,j) \text{ is part of the maximum RPK path} \\ 0 & \textbf{else} \end{cases}$$

$e_i \in [e_{tat}, e_{max}]$ the energy level in the battery of the aircraft at node i

$l_{(a_1,a_2),[t_1,t_2)} \geq 0$ the total number of empty seats on flights from airport a_1 to airport a_2 departing in the time interval $[t_1, t_2)$

MILP formulation

$$max \sum_{(i,j)\in A_f} x_{ij} d_{ij} - \sum_{((a_1,a_2),[t_1,t_2))\in D} l_{((a_1,a_2),[t_1,t_2))} d_{(a_1,a_2)} \tag{13}$$

s.t.

$$\sum_{(i,j)\in\delta_i^+} x_{ij} = \sum_{(i,j)\in\delta_i^-} x_{ij} \qquad \forall i \in N' \tag{14}$$

$$\sum_{(i,j)\in A_{((a_1,a_2),[t_1,t_2))}} x_{ij} \cdot s - l_{((a_1,a_2),[t_1,t_2))} d_{(a_1,a_2)} \qquad \forall((a_1,a_2),[t_1,t_2)) \in D \tag{15}$$

$$e_j \leq e_i + x_{ij} e_{ij} + (1 - x_{ij}) e_{max} \qquad \forall (i,j) \in A \tag{16}$$

$$e_{tat} \leq e_i \leq e_{max} \qquad \forall i \in N \tag{17}$$

$$x_{ij} = \{0,1\} \qquad \forall (i,j) \in A \tag{18}$$

$$l_{((a_1,a_2),[t_1,t_2))} \geq 0 \qquad \forall((a_1,a_2),[t_1,t_2)) \in D \tag{19}$$

The Objective (13) is to maximise the RPK over a single path from N_o to N_d. Constraint (14) ensures flow conservation at all nodes, except for the origin and destination. Constraint (15) makes sure that demand is not exceeded. Note that when a path is added to the initial solution, the demand that this path serves must be removed from the demand before searching for the next path. Constraint (16) ensures that the energy consumption is satisfied. Constraints (17), (18) and (19) regulate which values the decision variables may assume. Besides updating the demand before finding the next path, the used flight arcs must be removed from the graph to avoid having a flight arc used by multiple paths.

This formulation is similar to the arc-based formulation presented in Sec. 3. This is a variation of the resource-constrained shortest-path problem and is thus also NP-Hard [13]. Label correcting is commonly used to efficiently solve this problem. Consider a state at a certain node, which consists of an energy- and RPK-level. A label (e, rpk) represents this state and is feasible at node i if there is a feasible path that leads to this state at i. At each node, there is a finite set of feasible states. At the origin there is only one $(e_{max}, 0)$. This label corresponds to a fully charged battery and zero RPK. The label correcting algorithm starts at the origin and chronologically visits all nodes. At each node, it updates the label set of all nodes that are directly reachable from that node. Suppose there is a non-empty set of labels at node i and an arc (i, j). Then for each label (e, rpk) at i, a label $(e + e_{(i,j)}, rpk + d_{(i,j)} * min(s, D_{((a_i,a_j),[t_1,t_2])}))$ is added to the set of labels at j. However, if $e + e_{(i,j)} < e_{tat}$, then this state is infeasible because it does not satisfy the energy limits constraint (17) and is therefore not added to the set of labels at j.

Additionally, we use dominance rules to ignore the labels that cannot lead to the maximum RPK path. A label (e_1, rpk_1) at node i is dominated if there is another label (e_2, rpk_2) at node i such that $e_1 < e_2$ and $rpk_1 < rpk_2$. If a label is dominated, it is removed and no longer considered. Because the second label has more energy, any remaining path from i to N_d that is feasible starting from state (e_1, rpk_1) is also feasible for state (e_2, rpk_2). It is thus impossible for the first state to make up for its lower RPK.

The labels are programmed as keys in a dictionary. Each label (key) in the dictionary points to a list containing the sequence of arcs that led to that state. This enables easy extraction of the path. When the algorithm reaches the destination node N_d, the path corresponding to the label with the highest RPK is extracted and added to the initial solution. This is guaranteed to be the feasible path with the maximum RPK because the algorithm exhaustively explores all possible states and only ignores a state if i) the state is infeasible; or ii) the state cannot lead to the maximum RPK path.

The pseudo-code of the initial solution construction heuristic is shown in Alg. (1)

Algorithm 1 Maximum RPK label correcting

Input TS graph G, demand, RPK_{min}
$RPK = 0$, solution = [ ], labels = [{}, {}, ...] ($|N|$ dictionaries)
while $RPK \leq RPK_{min}$ **do**
    # Label Correcting
    for i in N **do**
        for (i, j) in δ_i^- **do**
            for label in labels(i) **do**
                Update label $((e + e_{(i,j)}, rpk + d_{(i,j)} * min(s, D_{((a_i,a_j),[t_1,t_2])})))$
                Check feasibility, check dominance & add updated label to labels(j)
            end for
        end for
    end for
    # Add best path to the solution
    Get path with highest RPK in labels(d)
    Add path to the initial solution
    Remove served demand & used flight arcs
end while
return solution
Output Initial feasible solution

4.2 Column Generation

The algorithm described in Sec. 4.1 provides a good initial feasible solution. We now want to i) Evaluate how good this solution is by computing a lower bound; and ii) Further improve on this initial solution. The CG (Column Generation) algorithm presented in this section attempts to do both.

In the arc based formulation, the problem was solved by selecting an optimal set of arcs that constitute a set of feasible paths. Let R be the set of all feasible paths from N_o to N_d. An equivalent way of looking at the problem is to select the optimal set of paths from R. Let c_r equal to the total operating cost and ask_r the ASK (Available Seat Kilometers) of path $r \in R$. Let $\delta^r_{(i,j)}$ equal 1 if arc (i,j) is part of path r, and 0 otherwise. Let $p^r_{((a_1,a_2),[t_1,t_2))}$ equal the number of seats available in demand interval $(a_1,a_2),[t_1,t_2))$ in path r. Let y_r be a binary decision variable indicating if path r is included in the solution. This leads to the path based formulation of the problem.

(PATH BASED FORMULATION)

Decision Variables

$$y_r = \begin{cases} 1 & \textbf{if } \text{path } r \text{ is used} \\ 0 & \textbf{else} \end{cases}$$

MILP formulation

$$min \sum_{r \in R} y_r c_r \tag{20}$$

s.t.

$$\sum_{r \in R} y_r \le K \tag{21}$$

$$\sum_{r \in R} y_r \cdot \delta^r_{(i,j)} \le 1 \qquad \forall (i,j) \in A_f \tag{22}$$

$$\sum_{r \in R} y_r \cdot p^r_{((a_1,a_2),[t_1,t_2))} - l_{((a_1,a_2),[t_1,t_2))} \le D_{((a_1,a_2),[t_1,t_2))} \qquad \forall ((a_1,a_2),[t_1,t_2)) \in D \tag{23}$$

$$\sum_{r \in R} y_r \cdot ask_r - \sum_{((a_1,a_2),[t_1,t_2)) \in D} d_{(a_1,a_2)} l_{((a_1,a_2),[t_1,t_2))} \ge rpk_{min} \tag{24}$$

$$y_r = \{0,1\} \qquad \forall r \in R \tag{25}$$

As in the arc based formulation, the objective (20) is to minimize costs. Constraint (21) limits the fleet size. Constraint (22) prevents multiple paths from occupying the same flight arcs. Constraint (23) keeps the solution from exceeding demand. Constraint (24) ensures that the minimum RPK is met. Lastly, constraint (25) indicates that y_r are binary decision variables.

As the path-based formulation is equivalent to the arc based formulation, it is also NP-Hard. Moreover, the number of feasible paths $|R|$ quickly explodes to an impractical size with a growing number of arcs in G. Most variables y_r however will equal 0 in the solution, so instead we only consider a subset of columns (paths) $\overline{R} \subset R$ and iteratively solve i) an *RMP (Restricted Master Problem)* using this subset of columns, starting from the initial solution. ii) a *PSP (Pricing Sub-Problem)* to extract columns with a negative RC (Reduced Cost) to add to $\overline{R}$. The RC of a path r indicates how much the objective function would change per unit increase y_r. When a path with a negative RC is found, it may be beneficial to add this to the subset of paths $\overline{R}$ in the RMP. When no paths exists with a negative RC, that means that the objective value cannot be decreased and thus that the optimal solution is reached. The process of CG is visualised in Fig. 3

Next, let us discuss the PSP. Let α_1 be the dual variable to constraint (21), let $\alpha_2^{(i,j)}$ the dual to constraint (22) for $(i,j) \in A_f$, let $\alpha_3^{((a_1,a_2),[t_1,t_2))}$ correspond to constraint (23) for $((a_1,a_2),[t_1,t_2)) \in D$ and let α_4 the dual variable of constraint (24). The reduced cost c^*_r of any path can then be computed as in Eq. (27).

$$c^*_r = -\alpha_1 + \sum_{(i,j) \in r} c^*_{(i,j)} \tag{26}$$

$$c^*_{(i,j)} = c_{(i,j)} - \alpha_2^{(i,j)} - \alpha_3^{((a_1,a_2),[t_1,t_2))} \cdot s - \alpha_4 \cdot s \cdot d_{(i,j)} \tag{27}$$

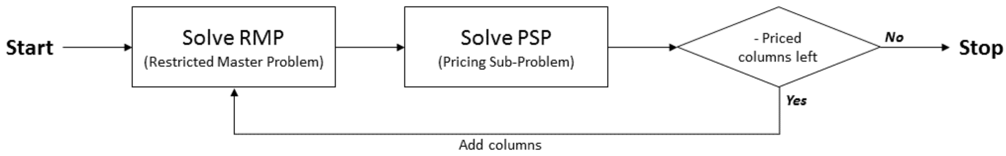

Figure 3: Process of Column Generation to solve the linear relaxation of the path based formulation

The PSP consists in finding a path from N_o to N_d that minimises the reduced cost and can be formulated as follows.

(PRICING SUB-PROBLEM)

Decision Variables

$$x_{ij} = \begin{cases} 1 & \textbf{if } \text{arc } (i,j) \text{ is included in the minimum RC path} \\ 0 & \textbf{else} \end{cases}$$

$e_i \in [e_{tat}, e_{max}]$ the energy level in the battery of the aircraft at node i

MILP formulation

$$min - \alpha_1 + \sum_{(i,j)\in A} x_{ij} c^*_{(i,j)} \tag{28}$$

s.t.

$$\sum_{(i,j)\in\delta_i^+} x_{ij} = \sum_{(i,j)\in\delta_i^-} x_{ij} \qquad \forall i \in N' \tag{29}$$

$$e_j \leq e_i + x_{ij}e_{ij} + (1 - x_{ij})e_{max} \qquad \forall (i,j) \in A \tag{30}$$

$$e_{tat} \leq e_i \leq e_{max} \qquad \forall i \in N \tag{31}$$

$$x_{ij} = \{0, 1\} \qquad \forall (i,j) \in A \tag{32}$$

The objective (28) is to minimise the reduced cost. Constraint (29) ensures flow conservation at all nodes except N_o and N_d. Constraint (30) regulates energy consumption and constraints (31) and (32) indicate which values the decision variables are allowed to take.

This formulation is surprisingly similar to the one to find the maximum RPK path in Sec. 4.1. With the exception of the objective function and absence of the demand constraints, the formulation is identical. This too is a variation on the resource-constrained shortest-path problem and is thus NP-Hard [13]. The label correcting algorithm for Sec. 4.1 is therefore slightly adjusted to be used in solving the pricing subproblem. Consider a state at a certain node, which in this case consists of an energy and RC. A label (e, rc) represents this state. A state is feasible at node i if there is a feasible path that leads to this state at i. At each node, there is a finite set of feasible states. At the origin there is only one $(e_{max}, -\alpha_1)$. The label correcting algorithm starts at the origin and chronologically visits all nodes. At each node, it updates the label set of all nodes that are directly reachable from that node. Suppose there is a non-empty set of labels at node i and an arc (i, j). Then for each label (e, rc) at i, a label $(e + e_{(i,j)}, rc + c^*_{(i,j)})$ is added to the set of labels at j. However, if $e + e_{(i,j)} < e_{tat}$, then this state is infeasible because it does not satisfy the energy limits constraint (31) and is therefore not added to the set of labels at j.

Additionally, the dominance rules to ignore the labels that cannot lead to the minimum RC path are as follows. A label (e_1, rc_1) at node i is dominated if there is another label (e_2, rc_2) at node i such that $e_1 < e_2$ and $rc_1 > rc_2$. If a label is dominated, it is removed and no longer considered. In that case, any remaining path from node i to the destination node N_d that is feasible starting from state (e_1, rc_1) is also feasible from state (e_2, rc_2) such that it is impossible for the first state to make up for its higher RC.

The labels are programmed as keys in a dictionary. Each label (key) in the dictionary points to a list containing the sequence of arcs that led to that state. This enables easy extraction of the paths. Note that there is not one but rather a set of labels at N_d after running label correcting. Instead of only looking for the path that minimises RC, all path with a negative RC at the destination are added to $\overline{R}$. This label correcting algorithm is guaranteed to find the path that minimises RC because it exhaustively explores all possible states and only ignores a state if i) the state is infeasible; or ii) the state cannot lead to the minimum RC path. If the algorithm cannot find a label with a negative RC at N_d, that means that there are none and thus that we obtained the optimal solution of the linear relaxation of the path based formulation. This is a much stronger lower bound than the solution to the linear relaxation of the arc based formulation because there, the energy constraints are also relaxed and the solution can use infeasible paths to obtain a lower solution. The paths in the linear relaxation of the path based formulation do still have to be feasible, leading to a stronger lower bound.

There is no guarantee that this solution will be integral. To achieve integrality, CG must be implemented in a B&B (Branch-and-Bound) framework where CG is executed at each node in the branching tree, resulting in a B&P (Brand-and-Price) algorithm. The B&P algorithm and reasons why it proved unsuccessful to solve the problem to optimality are discussed in Appendix B. Instead, we attempt to improve on the initial solution by solving the integer version of the final RMP, using the final subset of columns $\overline{R}$. This solution will be at least as good as the initial solution because the paths that make up the initial solution are included in $\overline{R}$. There is however no guarantee that the integer solution of the final RMP is the optimal integer solution of our problem. There actually is a major flaw in this approach, namely that $\overline{R}$ only consists of high utilisation paths. Though the optimal solution likely consists primarily of high utilisation paths for the reasons laid out in Sec. 4.1, the optimal solution may require one or more lower utilisation paths. Consider the following example, with $RPK_{min} = 1$. The minimum cost solution is found when flying one aircraft on a single out-and-return flight to the nearest airport, which is a low utilisation path. Any additional flights just increase the cost while the minimum RPK is already met. Suppose for simplicity that each flight arc has a reduced cost equal to -1. Then the RC of the optimal path is $c_r^* = -2$. Now suppose that there is a feasible path that performs the same out-and-return flight twice. This higher utilisation path has a reduced cost of -4 and will dominate the lower utilisation path, which will therefore not appear in $\overline{R}$.

While CG is useful to compute a strong lower bound and improve on the integer solution by solving the integer version of the final RMP, we would still like to keep improving. As B&P proved unsuccessful to find the global minimum, the next section proposes an LNS (Large Neighborhood Search) algorithm to improve the best-known solution to a local minimum.

4.3 Large Neighborhood Search

This section proposes an LNS (Large Neighborhood Search) algorithm to improve the best known solution until a local minimum is found. In general, any LNS starts from a known solution, iteratively explores its neighborhood and moves to the best neighbor. When there is no better solution in the neighborhood, the solution is locally optimal. To fully define the LNS, we need to define i) the neighborhood; and ii) a mechanism to explore the neighborhood.

Let the neighborhood be the set of all feasible solutions where all but one path in the solution are equal. If in a solution, one path is replaced by a low-utilisation path such that the overall solution is still feasible, this is a valid neighbor. The best known solution up until now was achieved by solving the integer version of the final RMP, as explained in Sec. 4.2. This had the major drawback that the subset of paths $\overline{R}$ in the RMP only contained high utilisation paths. This LNS algorithm directly addresses this concern by allowing high-utilisation paths to be replaced by more suitable paths. On each iteration of the LNS, the optimal replacement path is computed for each path in the solution. The algorithm then moves to the best neighbor, with the lowest total cost as visualised in Fig. 4.

The mechanism to explore the neighborhood and find the best neighbor is a variation of the arc based MILP formulation from Sec. 3. While the arc based formulation is impractical to solve large instances, it does work to solve small cases with only 1 aircraft. Consider a solution consisting of a number of paths. Suppose that one path is removed from this solution such that the remaining solution does not meet the minimum RPK. Let $RPK_{remaining}$ be the total RPK of the remaining solution. We will now compute the optimal replacement path such that the solution is feasible again. Let $RPK_{gap} = RPK_{min} - RPK_{remaining}$ be the RPK that is required from this replacement path, as visualised in Fig. 5. The following MILP formulation is used to find the optimal

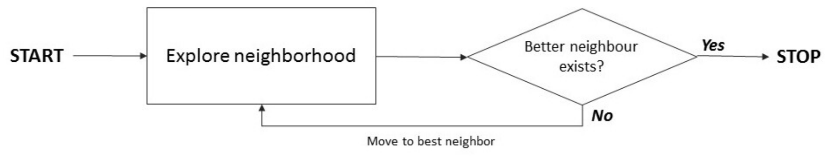

Figure 4: Process of the Large Neighborhood Search to improve the best known solution to a local minimum

replacement path.

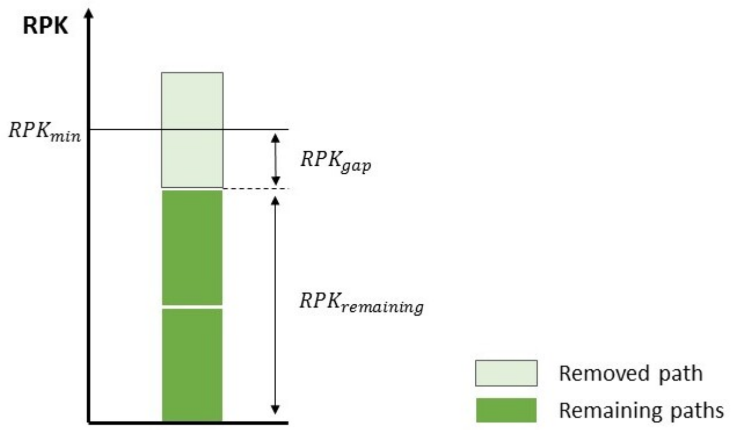

Figure 5: Visual representation of the relation between RPK_{min}, $RPKgap$ and $RPK_{remaining}$

(ARC BASED FORMULATION - REPLACEMENT PATH)

Decision Variables

$$x_{ij} = \begin{cases} 1 & \textbf{if } \text{the aircraft is using arc } (i,j) \\ 0 & \textbf{else} \end{cases}$$

$e_i \in [e_{tat}, e_{max}]$ the energy level in the battery of the aircraft at node i

$l_{(a_1,a_2),[t_1,t_2)} \geq 0$ the total number of empty seats on flights from airport a_1 to airport a_2 departing in the time interval $[t_1, t_2)$

MILP formulation

$$min \sum_{(i,j)\in A} c_{ij}x_{ij} \tag{33}$$

s.t.

$$\sum_{(i,j)\in\delta_i^+} x_{ij} = \sum_{(i,j)\in\delta_i^-} x_{ij} \qquad \forall i \in N' \tag{34}$$

$$\sum_{(i,j)\in\delta_o^-} x_{ij} \leq 1 \tag{35}$$

$$\sum_{(i,j)\in A_f} x_{ij} \cdot s \cdot d_{ij} - \sum_{((a_1,a_2),[t_1,t_2))\in D} l_{((a_1,a_2),[t_1,t_2))} d_{(a_1,a_2)} \geq RPK_{gap} \tag{36}$$

$$\sum_{(i,j)\in A_{((a_1,a_2),[t_1,t_2])}} x_{ij}\cdot s - l_{((a_1,a_2),[t_1,t_2])} \leq D_{((a_1,a_2),[t_1,t_2])} \qquad \forall((a_1,a_2),[t_1,t_2]) \in D \tag{37}$$

$$e_j \leq e_i + x_{ij}e_{ij} + (1-x_{ij})e_{max} \qquad \forall(i,j)\in A \tag{38}$$

$$e_{tat} \leq e_i \leq e_{max} \qquad \forall i \in N \tag{39}$$

$$x_{ij} = \{0,1\} \qquad \forall(i,j)\in A \tag{40}$$

$$l_{((a_1,a_2),[t_1,t_2])} \geq 0 \qquad \forall((a_1,a_2),[t_1,t_2]) \in D \tag{41}$$

The objective (33) is to minimise operating costs. Constraint (34) ensures flow conservation on each node, except for the origin and destination nodes. Constraint (35) ensures that the aircraft can be deployed at most once. Constraint (36) guarantees that a sufficient RPK is served. Constraint (37) verifies the passenger demand limits in each demand bucket. Here the demand served by the remaining paths must be removed from the demand available for the replacement path. Constraint (38) regulates the energy level. Constraints (39), (40) and (41) regulate which values the decision variables may assume.

This formulation is very similar to the arc based formulation in Sec. 3. The program presented here differs in that it only considers one aircraft, instead of a fleet of $|K|$ aircraft, and in that the constraint that prohibits multiple paths to use the same flight arc is absent here. Instead, the flight arcs that are already used by the remaining paths are temporarily removed from the graph. While this formulation is still NP-Hard, it is computationally tractable when applied for just a single aircraft.

4.4 Overall Methodology

The full algorithm consists of the steps described in the previous sections. These steps and their purposes are summarized in Fig. 6. Firstly, an initial feasible solution is constructed by heuristically maximising aircraft utilisation. This is achieved by iteratively finding the maximum-RPK path using a label correcting algorithm and adding it to the solution. This initial solution serves as a starting point for Column Generation, which solves the linear relaxation of the path based formulation and results in a strong lower bound, much stronger than the one from the linear relaxation of the arc based formulation presented in Sec. 3. This is particularly useful to evaluate the best-known integer solution and deciding how many resources should be deployed to try and improve it. Column Generation iteratively solves an RMP (Restricted Master Problem), the path-based formulation using a subset of columns (paths), until no column can be added that would result in an improvement. When Column Generation has run to optimality, the final RMP is solved with integer restrictions to improve on the initial solution. Lastly, an LNS is started from this solution. The LNS iteratively explores the neighborhood where all but one paths in the solution are equal. On each iteration, the algorithm moves to the best neighbor, which is found by using a variation of arc based formulation to find the optimal replacement path for all paths in the solution. When no better neighbor can be found, a local optimum is reached.

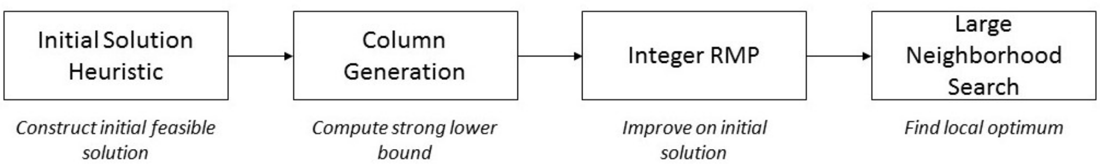

Figure 6: Schematic of the full algorithm proposed in this thesis

5 Description of the Case Studies

This Section describes the case studies on which the methods have been tested. Sec. 5.1 lays out the aircraft model, which is based on the all-electric Skylax E10. Sec. 5.2 introduces the cost model to calculate the cost of each arc. The demand model is presented in Sec. 5.3. Sec. 5.4 presents the Time-Space network and lastly, Sec. 5.5 presents the instances on which the algorithms will be tested.

5.1 Aircraft Model

The aircraft considered in this thesis is the Skylax E10, a very promising concept currently in development. The performance estimates in Tab. 3 are publicly available. This section builds an aircraft model to validate that these claims are realistic and to estimate power consumption for a flight.

Parameter	Value
# seats s	9 + 1 pilot
MTOW	3,500 kg
Empty weight	2,700 kg
Battery weight	1,000 kg
Engine power	260 kW (x2)
Cruise speed	300 kph
Range at MTOW	300 km
Energy costs	47 €/h
Variable costs	109 €/h

Table 3: Aircraft parameters claimed for the Skylax E10

The difference between the MTOW and empty weight leaves 800kg for payload, equivalent to 10 people of 80kg. We assume a parabolic drag polar (Eq. (42)).

$$C_D = C_{D_{min}} + k(C_L - C_{L_{minD}})^2 \tag{42}$$

Where $C_{L_{minD}}$ is the lift coefficient corresponding to the minimum drag coefficient. The lift induced drag constant is calculated as $k = 1/(\pi \cdot AR \cdot e)$ where AR is the wing aspect ratio and e the Oswald efficiency factor. We assume $AR = 10$ and $e = 0.9$ such that $k = 0.035$. Next, we assume $C_{L_{minD}} = 0.1$. Lastly, two methods exist to estimate $C_{D_{min}}$. The first one is by comparing to similar aircraft. This is a quick estimation, useful in the early stages when the aircraft layout is not yet available. According to [23], this results in a $C_{D_{min}}$ of 0.0225 for a light, twin engine aircraft. Alternatively, Eq. (43) provides an estimation based on a build up of drag of the various aircraft components.

$$C_{D0} = \frac{1}{S_{ref}} \sum C_{f_C} \cdot FF_C \cdot IF_C \cdot S_{wet} \tag{43}$$

Where C_{f_C} is the flat plate skin friction coefficient, FF_C is a form factor to account for pressure drag,IF_C is an interference factor to account for the drag increase due to proximity of multiple objects. Based on the assumptions provided in Table 4, a C_{D0} of 0.02055 is obtained. This is very close to the first estimation. We therefore continue with the value in between $C_{D0} = 0.021$. These result in a L/D in cruise of 20.5.

Component	C_f	FF	IF	S_{wet}
Wing	0.0015	1.13	1	32
Tail	0.0015	1.13	1.04	8
Fuselage	0.00027	1.49	1	63
Nacelles	0.00027	1.19	1.5	0.9

Table 4: Assumed values for Eq. (43) [23]

The lift coefficient in each flight phase is calculated from the lift equation $L = \frac{1}{2}\rho V^2 S C_L$, where we assume that the wing surface $S = 16m^2$. Subsequently, the drag coefficient is calculated by inserting the lift coefficient in Eq. (42). Total drag then equals $D = \frac{1}{2}\rho V^2 S C_D$. The power required from the propellors then becomes $P = DV$. We assume an airspeed of 300 km/h in cruise, climb and descend. For the final approach, an airspeed of 120km/h was used. In climb and descend, a vertical speed of $V_z = 4m/s$ adds $V_z mg$ to the power required, where m is the aircraft mass and g the gravitational acceleration. In an electric aircraft m remains constant during a flight as no fuel is burned. The jet velocity V_j is calculated from $P_r = \dot{m}(V_j - V)$, with mass flow $\dot{m} = \rho A_p V$ where A_p is the area of the propellers. Additionally, the propulsive efficiency was estimated by assuming a propeller radius of $1m$ and using the formula $\eta_j = 2/(1 + V_j/V)$. This efficiency converts power required to shaft power $P_r = P_s \cdot \eta_j$. An efficiency of 78% is applied to convert battery power to shaft power $P_s = P_b \cdot 0.78$ [17]. We assume that 40 kW of battery power is used during taxi. During take-off, we assume full that full power is used, which is around 500kW for the Skylax E10. During landing, power is brought to idle and no power is consumed. Lastly, a 45 minute reserve at minimum power consumption is assumed. This leads to the estimated power in each flight phase in Figure 7.

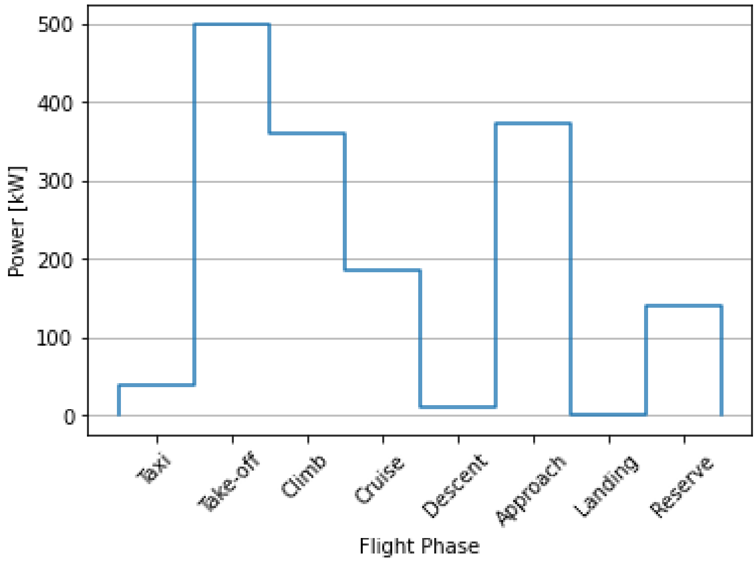

Figure 7: Power estimations per flight phase for the Skylax E10

We assume that all flight phases except for the cruise are fixed in duration as shown in Tab. 5. All fixed flight phases require a combined energy of $e_f = 108$kWh. We also assume that the distance of $d_f = 92$km covered during climb and descent also contributes to the range, meaning that 42 minutes of cruise are required to reach 300km. The total power consumption of this trip is 238 kWh. A battery reserve for 45 minutes is also reported. Minimum power consumption occurs around 200 kph at sea level. Including efficiency losses, 140kW is required. For a 45 minutes reserve, 105 kWh must be available. This brings to total battery capacity to 343kWh. The specific energy of 343W/kg is in line with the expected battery advancements by 2025.

Flight phase	Power [kW]	Duration [min]	Energy consumption [kWh]
Taxi	40	10	6.7
Take-off	500	1	8.3
Climb	360	10	56.7
Cruise	186	42	119
Descent	11	10	3.3
Approach	373	5	23.3
Landing	0	1	0
Reserve	140	45	105
Total			343

Table 5: Estimated power consumption and duration of the flight phases

Finally, the energy required to perform a certain flight can be computed using Eq. (44), in kWh. In reality, aircraft cannot simply fly the shortest possible route. Therefore, the distance is multiplied by a routing factor r. A routing factor of 1.08 is used in this thesis. As all required values have been estimates, the energy required to perform any flight can now be estimated.

$$e = e_f + (r \cdot d - d_f) \cdot P_{cruise}/V_{cruise} \tag{44}$$

5.2 Cost Model

This section documents the estimation of direct operating costs in function of the schedule. This consist of the four parts listed below. Indirect costs also exist and include passenger services, ground equipment, administrative expenses, etc. The first three elements are calculated per flight, whereas ownership costs are a fixed daily cost.

- Energy
- Maintenance
- Battery depreciation
- Labor

- Ownership

Energy costs are difficult to estimate because they vary between cities and different providers. Additionally, prices depend on the peak power and total volume of electricity used. The energy bill can be split in two parts. The supply is concerned with production of electricity and charges a cost per kWh. The delivery part is concerned with transporting the electricity to the user [4]. Delivery charges a price per kWh, a premium for peak power and a fixed charge for the services and being connected to the grid. To reliably determine the electricity rates applicable, one would need to know the total energy consumption and peak power at each individual airport. Electricity costs are included in the direct operating costs. Ownership of charging infrastructure is not. In practice, charging infrastructure may be owned and operated by a third party such as the airport itself, in which case the aircraft operator only incurs the cost of electricity. A cost of €0.20/kWh was chosen. Knowing the energy consumption for a flight, the energy cost can also be determined.

Maintenance costs of a disruptive technology are hard to assess. Two approaches exist. The first one being a bottom up estimation of the maintenance related to all components or assemblies. This approach requires extensive data about the individual components and their failure rates. As this information is not available, a second approach is used. This approach uses top-down semi-empirical relations like the one found in [24]. It estimates 6 sources of maintenance costs. To take the new technologies into account, adjustment factors found in literature were used. According to Boeing, a composite airframe reduces airframe labor hours by 35% due to the lack of corrosion. Information from the automotive industry shows that electric engine material and labor costs are reduced by 20% because of less moving parts and consumables. Finally, the engine overhaul reserves are neglected as electric engines' lifetime exceeds that of the aircraft. They will therefore not be overhauled during the aircraft's lifetime. As **labor cost** is an hourly cost too, the two are combined. A single pilot suffices to fly the Skylax E10 and a rate of 25€/h is assumed. Whereas Skylax reports a variable cost of €109/h, including maintenance, inspections, labor, parts and miscellaneous costs. This is 22% lower then the estimate of €140/h for maintenance and labor. The difference could be in the assumed labor hour rate and other cost factors but seems optimistic. This thesis therefore proceeds with the estimate of €140/h.

Electric aircraft do have one additional cost component, namely that of **battery depreciation** as they must be replaced at the end of its useful life. The battery lifetime depends on a range of factors such as the DoD (Depth-of-Discharge), charging power, operating temperature, etc. For this thesis, we assume a battery lifetime of 1,000 cycles, after which a new battery back must be acquired. Current market prices are around €200/kWh but these are expected to drop to 100-150 by 2025. We therefore assume a price of €125/kWh, resulting in a battery price of €43,000. By assuming linear depreciation over 1,000 cycles, We obtain a cost of €0.125/kWh of consumed energy.

Lastly, the **ownership costs** refer to the costs for an operator to own an asset over a period of time. This includes financing the acquisition, depreciation and insurance. As suggested in [3], we linearly depreciate the aircraft over a period of 15 years and apply an insurance of 2% of the market value. Additionally, we assume that the operator sells the aircraft after 9 years to trade it for a newer model. With an original price of €2m, the aircraft is still worth 800k after 9 years. Over those 9 years, the insurance was on average 28k/year. That leaves the operator with a total ownership cost of 1.452m over the course of 9 years or €440 per day per aircraft.

The elements of the direct operating costs are summarised in Tab. 6. As the energy consumption and flight time for each flight can be determined in advance, these costs are pre-computed and attached to the flight arcs in the TS graph.

Cost	**Value**
Energy	0.2 €/kWh
Battery depreciation	0.125 €/kWh
Labor & maintenance	140 €/h
Aircraft ownership	440 €/day

Table 6: Cost model for the Skylax E10

5.3 Demand Model

In this section we discuss a simple demand model. While in reality demand will depend on many factors such as service level, travel time, availability of alternatives, public acceptance, fare prices, marketing, etc. a thorough demand analysis is far beyond the scope of this thesis. This sections presents a simple demand model used in the case studies but the methods developed in this thesis are flexible to use any demand model as input.

The demand model proposed in this thesis considers three factors: i) population of the origin and destination city, ii) the distance and iii) ToD (Time of Day), Firstly, a variation of the gravity model is used. Let P_{a_1} and P_{a_2} be the population of the cities where airport a_1 and a_2 are located respectively. These are multiplied and instead of dividing by a power of distance as is the case in a classical gravity model, this is multiplied by a Gaussian distribution based on distance $N(d_{(a_1,a_2)}|\mu = 200, \sigma = 100)$ because this reflects the competitiveness of thin-haul air transport in Europe [31]. Next, demand is scaled by three Gaussian distributions, depending on the ToD t: $N(t|\mu = 8, \sigma = 2)$, $N(t|\mu = 16, \sigma = 2)$ and $N(t|\mu = 12, \sigma = 6)$ to model daily fluctuations in demand [19]. Lastly, the demand is scaled by a factor f to get within reasonable bounds. Because no thin-haul airline service is currently available in Europe, one option is to calibrate f by benchmarking against existing modes of transport. However, as these other modes transport huge amounts of people every day, small calibration errors could have a very large influence. Therefore, f chosen manually and fixed at a value of 0.01 to achieve a reasonable demand. This leads to Eq. (45) to calculate the demand between cities a_1 and a_2 at time t. Naturally, demand is rounded to the nearest integer.

$$D(a_1, a_2, t) = P_{a_1} \cdot P_{a_2} \cdot N(d_{(a_1,a_2)}|200, 100) \cdot N(t|8, 2) \cdot N(t|16, 2) \cdot N(t|12, 6) \cdot f \tag{45}$$

Demand is divided in time windows. For this thesis, a time window of 1 hour was used and the center of each time interval was used to compute demand in that interval. Fig. 8(a) shows the effect of the variation on the gravity model. Observe how the connections around London for example have a higher demand because of the large population. Fig. 8(b) shows the variation of demand on some connections around Brussels throughout the day. A clear morning peak at 8am and afternoon peak at 16pm can be observed.

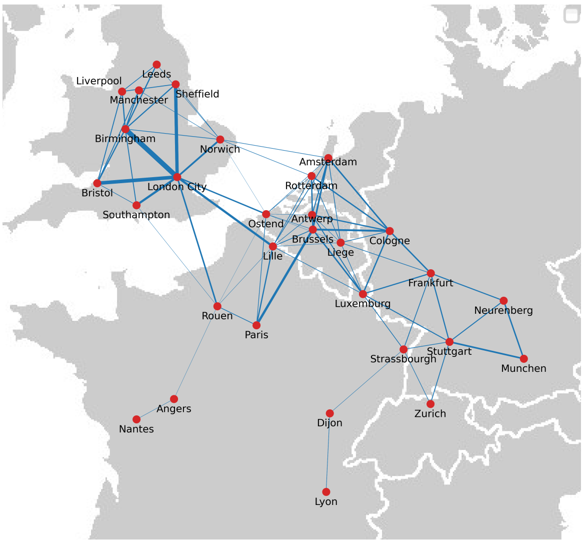

(a) Network with 20 airports and demand weighted links

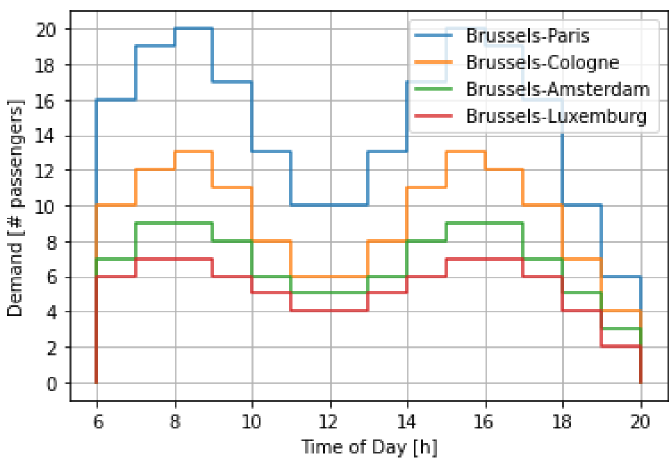

(b) Variation of demand on a few connection depending on ToD

Figure 8: Visualisation of the demand model between some airports

5.4 Time-Space Network

This section lays out how the TS (Time-Space) network G is created. The parameters to create the network are summarised in Tab. 7.

Parameter	Value
$[t_{min}, t_{max}]$	[6.00 AM, 8.00 PM]
t_{step}	15 minutes
t_{tat}	30 minutes
P_{charge}	250 kW
a_h	Brussels
$[d_{min}, d_{max}]$	[100 km, 300 km]
r	1.08

Table 7: Parameters to create a Time-Space network G

The operating hours are from 6AM until 8 PM and nodes are created at 15 minute time steps. For any airport a, there are nodes $(a, 06.00)$, $(a, 06.15)$, $(a, 06.30)$, ..., $(a, 20.00)$. The TAT (Turn-Around-Time) is set at 30 minutes. This is in between the average TAT for Cape Air (52 min) and Mokulele Airlines (19 min) [27]. Airports are chosen near large cities in the area of Belgium and The Netherlands with an extension in the North of France, South-West of Germany and South of the UK because these regions likely have good demand for such an air transport service [31]. Table 8 presents the list of airports, divided in several networks. The airports in the first column are used to construct a TS network of 5 airports. The ones in the next column are added to create a TS network of 10 airports, etc.

30 airport network					
20 airport network					
15 airport network					
10 airport network					
5 airport network					
Brussels (hub)	Frankfurt	Antwerp	Liege	Ostend	Nantes
Amsterdam	Stuttgart	Rotterdam	Munchen	Birmingham	Angers
Paris	Strassbourgh	Zurich	Neurenberg	Bristol	Leeds
Luxemburg	Dijon	Lyon	Rouen	Manchester	Sheffield
Cologne	Lille	London	Southampton	Norwich	Liverpool

Table 8: List of airports in the networks

Brussels was assigned as the hub airport a_h because of its central location on the map. The distance between any airport pair is calculated as the great circle distance, which is the shortest distance on a globe. Distance is subsequently multiplied by a routing factor of 1.08 to account for the fact that the aircraft may not be able to fly the shortest possible route. Flights are considered between any airport pair where the resulting distance lies in the range of 100km to 300km. Below 100km, cars and trains become more attractive modes of transport and above, traditional airlines become more competitive [31]. Besides, the range of the Skylax E10 is currently limited to 300km. The network of airports is visualised in on a map in Fig. 9

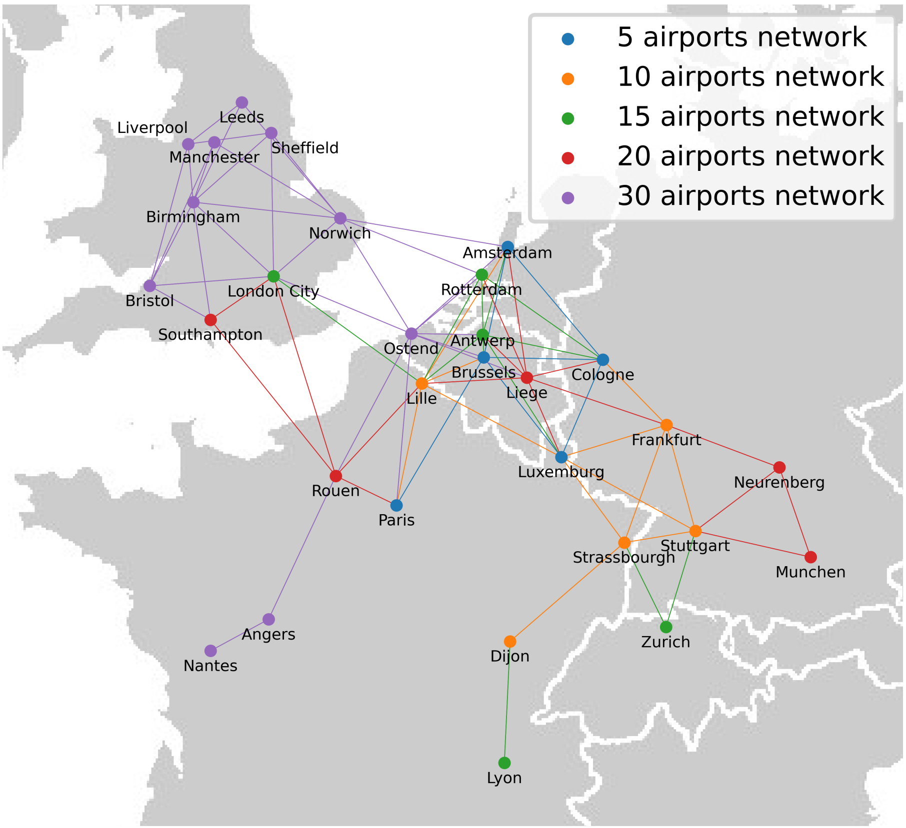

Figure 9: Map of the networks of airports

In our problem, all aircraft must start and end the day at the origin- and destination nodes N_o and N_d. As a result, a number of nodes and arcs in the TS network G cannot be reached. Consider the node $i(a_1, t_{min})$ where $a_1 \neq a_h$. This node is obsolete because there is no way for an aircraft to ever reach it. Any arc starting from this node $(i,j) \in \delta_i^+$ is therefore also obsolete. After the TS networks are created, a post-processing step removes these obsolete nodes and arcs. Firstly, a forward travelling recursive algorithm is used, starting from N_o, to label all nodes that are reachable from N_o without considering energy consumption. Next, a backward travelling recursive algorithm starting from N_d is used to mark all nodes that can reach N_d. If a node is not marked as reachable in both cases, the node and all connected arcs are removed from G. The TS network G for 5 airports is shown if Fig. 10.

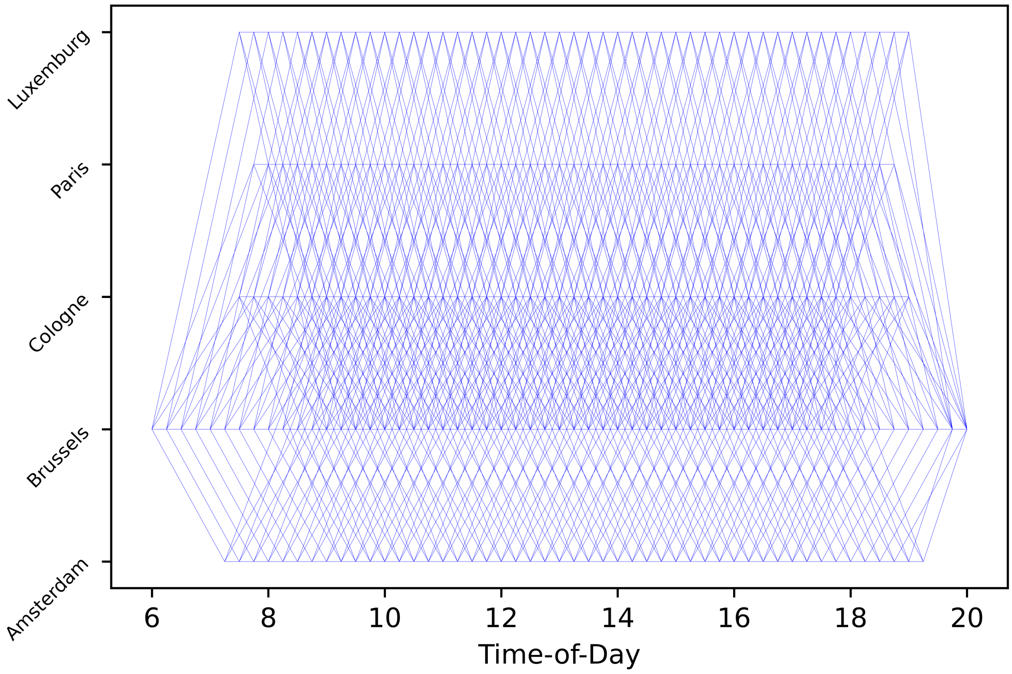

Figure 10: Time-Space network with 5 airports

5.5 Instances

The algorithm will be tested on different networks with a different number of airports. Also within the same network, the algorithm will be tested on a varying minimum RPK. Because the network of Cape Air, the largest thin-haul airline in the world, consists of 20 airports [27], the algorithm will be tested up until 20 airports in steps of 5. To see how far we can push the algorithm, it will then also be tested on a 30 airport network.

Let the market size RPK_{max} be the sum of all RPK that can be served, which is the demand multiplied by the distance in each demand bucket. ($RPK_{max} = \sum_{((a_1,a_2),[t_1,t_2))\in D} D_{((a_1,a_2),[t_1,t_2))} d_{(a_1,a_2)}$). The minimum RPK can then be chosen as a percentage market share. E.g. if a 30% market share is desired, then $RPK_{min} = 0.30 RPK_{max}$. The network with 5 airports was evaluated in steps of 1% market share (1, 2, 3, ...) whereas all other networks were evaluated in steps of 10% market share (10, 20, 30, ...) as indicated in Tab. 9

# Airports	Market Share
5	1, 2, 3, ...
10	10, 20, 30, ...
15	10, 20, 30, ...
20	10, 20, 30, ...
30	10, 20, 30, ...

Table 9: Instances on which the algorithm will be tested

A market share of 100% is not feasible however. This has three reasons: i) The number of flights serving the same piece of demand is limited. With a time step of 15 minutes, there is a maximum of 4 flights per hour from any airport a_1 to airport a_2. Demand may exceed the number of seats available, even if all 4 flights are scheduled. There may also be fewer than 4 flights available. Suppose it takes 1 hour to fly from the hub to airport a_1. Then an aircraft can arrive at a_1 no sooner than 7AM. Using a TAT of 30 minutes, the earliest time at which an aircraft can serve demand starting from a_1 is 7.30AM. This means that for the demand buckets $b = ((a_1, a_2), [7AM, 8AM))$, there are only 2 flight arcs present in G, one departing at 7.30 AM and one departing at 7.45 AM. ii) The energy consumption may make it impossible to use a certain flight arc. Expanding on the previous example. Suppose that upon arrival at airport a_1, the aircraft must first charge its batteries for at least 1 hour before it has enough energy to fly to airport a_2. In that case the flight arcs from a_1 to a_2 departing at 7.30 AM and 7.45 AM are infeasible. This effect becomes more pronounced in larger networks with more airports. The further away an airport is from the hub, the more likely it is that the aircraft has to make intermediary stops to recharge the batteries before continuing. One could filter the flight arcs that are infeasible due to the energy consumption. However, that would make the graph G charger-specific. This means that if one would change the charging power of all chargers, or even just at one airport, then some arcs may become feasible again or others may also become infeasible. Therefore, they were not removed from the graph G. Lastly, iii) There are groups of flight arcs were it is impossible to serve all of them. Suppose there is only one flight arc (N_o, i) leading to node i. If there are two different flight arcs (i, j_1) and (i, j_2), then the aircraft can only choose to operate at most one of these two. If the aircraft is schedule to fly (i, j_2), it is impossible to schedule another aircraft for (i, j_2) because the only arc to reach i is by using the flight arc (N_o, i), which is already used by the first aircraft. We therefore run the case studies up to the maximum attainable market share where a feasible initial solution could be constructed.

6 Case Study Results

The case studies consist of two parts. Firstly, the performance of the algorithms is evaluated in Sec. 6.1. There the improvements and running time of each step in the algorithm presented in Sec. 4 are compared to that of the arc based formulation presented in Sec. 3. Secondly, Sec.6.2 analyses a solution in detail. We investigate why the given schedule looks the way it does. Lastly, in Sec. 6.3 we look for a pattern between different solutions and what makes up a successful schedule. The algorithms were implemented in Python 3.6 and the Gurobi 9.1 default solver was used to solve the MILP models. The experiments were run on a Toshiba laptop with an Intel(R) Core(TM) i7-6500U CPU at 2.50GHz and 8 GB of RAM.

6.1 Algorithmic Performance

Let us analyse how the objective value evolves with a growing market share. The full results of all instances can be found in Appendix D. Here, let us consider the relation between cost and market share on the network with 5 airports, as visualised in Fig. 11. Costs were minimised for each percentage market share (1%, 2%, 3%, ...). This is a discrete graph where the consecutive data points are connected.

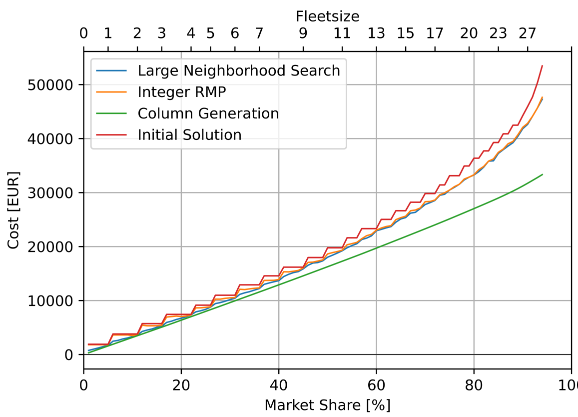

(a) Cost for all market shares

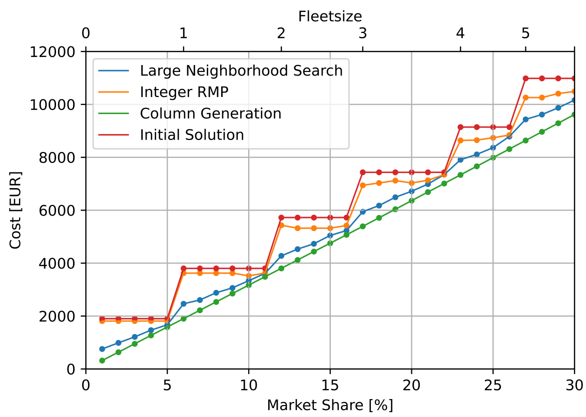

(b) Detailed view of 0-30% market share

Figure 11: Cost-market share relation in the network with 5 airports

Several interesting observations can be made. Firstly, the construction heuristic used to create the initial solution results in clear step increases in cost. These step increases occur each time an aircraft is added to the fleet. A single aircraft is sufficient up to 5% market share but a second aircraft is needed to serve 6% market share.

As costs consist of i) a daily ownership cost of €440 and ii) a cost per flight, the step increase will be at least €440. We observe however that the step increases are roughly 5 times higher than that. That is a result of the way the initial solution is constructed. Each time an aircraft is added to the fleet, this aircraft is assigned a high utilisation schedule that maximises RPK for this aircraft. As a result, its schedule includes multiple flights and thus additional costs, raising the step increases. These high utilisation schedules often overshoot the required RPK. Consider the initial solution for 6% market share. This initial solution has enough capacity to serve up to 11% market share and thus remains unchanged until it can no longer meet the required market share. At 12% market share, a third aircraft is added to the fleet. The fleetsize in the initial solution can thus be found by counting the number of steps. To do that, the discretisation of market shares has to be sufficiently small to see these steps. Here, market share was discretised in steps of 1%. For the larger networks with more than 5 airports, market share was discretised in steps of 10%. This discretisation is not fine enough and no steps can be observed. Additionally, the market size of a larger network is also larger, A 5% market share in the 5 airport network equates to $RPK_{min} = 12,851$ and can be served by 1 aircraft, whereas a 5% market share in the 10 airport network equal $RPK_{min} = 17,626$, which needs two aircraft. Also notice how the step increases become shorter with an increasing market share. That happens because each time an aircraft is added to the fleet, its schedule maximizes RPK. As this aircraft imposes new restrictions on the schedule of subsequent aircraft, the next aircraft will be given a schedule with an RPK that is equal or lower than the previous. At high market shares, the RPK of additional aircraft is limited by the demand already satisfied by other aircraft. Additional aircraft may have to perform flights with only a small number of passengers on board.

These steps are still visible in the integer RMP solution, but less pronounced. As noted in Sec. 4.2, the columns $\overline{R}$ in the RMP only contain high utilisation paths. As a result, when an aircraft is added to the fleet, this aircraft can only be assigned a high utilisation schedule, causing these step increases. The reason that these steps are less pronounced compared to the step increases in the initial solution is that after running column generation, $\overline{R}$ contains tens or hundreds of paths such that a more optimal schedule can almost always be found among those. At a market share of around 50%, the steps in the integer RMP solution start to smooth out. This happens because i) $\overline{R}$ contains more paths with increasing market share. At market shares 10%, 50% and 90%, the number of paths $|\overline{R}| = 43$, 292 and 688 respectively; and ii) The fleetsize is larger so more paths must be selected. When more paths must be selected, it is possible to select a set of paths that does not overshoot the RPK_{min} by a large amount. The integer RMP solution is therefore more suitable for a large fleet, which occurs at large market shares.

This issue of overshooting the RPK_{min} with high utilisation paths was the motivation to develop the Large Neighborhood Search. Figure 11(b) shows that the LNS is very effective at reducing the size of the steps. There is still a small inevitable step increase due to the fixed cost of €440 each time an aircraft is added to the fleet. After this smaller increase follows a roughly linear increase in cost until a subsequent aircraft is needed. This linear increase is the result of the aircraft performing more and/or longer flights to meet the RPK_{min}. However, at market shares above 50%, the improvements of LNS become less noticeable. As mentioned, the larger set of paths $\overline{R}$ and larger fleetsize result in an integer RMP solution that is already very good and leaves little room for the LNS to improve further. The motivation behind LNS becomes less of a problem at these higher market shares and naturally, LNS leads to smaller improvements.

The lower bound provided by Column Generation proves to be a strong bound at low market shares. Near the end of the step increases in costs, the objective value of the integer solution, integer RMP solution and LNS solution get very close to the lower bound, often within a few percentages. To better understand how good these solutions are, the Cost/RPK is plotted in Fig. 12(a). The steps in cost translate to a sawtooth pattern in Cost/RPK here. The solutions that are very close to the lower bound are minimise Cost/RPK for that fleetsize. This confirms the hypothesis upon which the initial solution construction heuristic was built. Namely that the schedule that minimises Cost/RPK is a high utilisation schedule. We also observe that the best known integer solution (LNS solution) diverges from the lower bound with increasing market share. That is because at high market shares, when most of the remaining flights are left with low demand. This causes the integer solution to become more expensive quickly while the impact on the relaxed solution (lower bound) is less severe. Suppose all flights in path r have demand for just a single passenger. The linear relaxation can give a value of $1/s$ to the related decision variable y_r, and thereby receive the benefit of flying this passenger for only $1/s$'th of the cost. The lower bound is thus not severely affected by low demand while the integer solution is.

To verify this hypothesis, we attempted to solve the arc based formulation for the network with 5 airports using the Gurobi 9.1 default solver on the faculty's server with 64 cores, 128 threads and 256 GB of RAM. When the server was unable to solve the arc based formulation to optimality within 1 hour, we saved the best known solution as well as the lower bound provided by Gurobi. From Fig. 12(b) we see that at 10% market share, the LNS actually found the global minimum. In the other cases, Gurobi was able to find a solution that was a few percentages better than the LNS solution but was unable to solve these cases to optimality. Still, the lower bound provided by Gurobi was stronger than the one provided by Column Generation and confirms our hypothesis that when the CG lower bound and LNS solution start diverging, the global optimum is closer to the LNS solution than the CG lower bound. The LNS solution is thus a very good one.

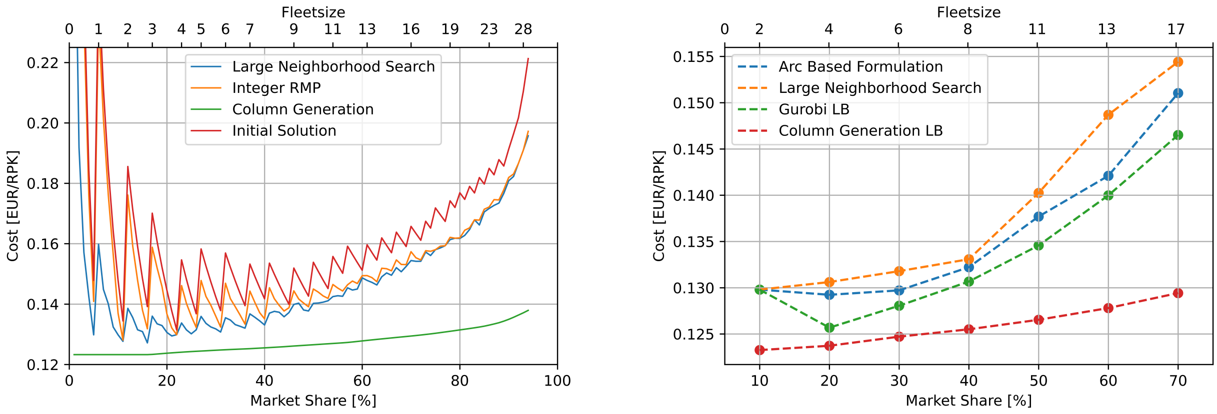

(a) Cost/RPK for our own methods

(b) The objective value from the arc based formulation is closer to the local minimum from LNS than the CG lower bound

Figure 12: Cost/RPK for discrete market shares in a network with 5 airports

Besides the objective value, we are also interested in the running time. The running times in the networks with 5 airports and 20 airports is plotted in Fig. 13. For visual purposes, the number of datapoints in Fig. 13(a) was reduced from one every 1% to one every 6%.

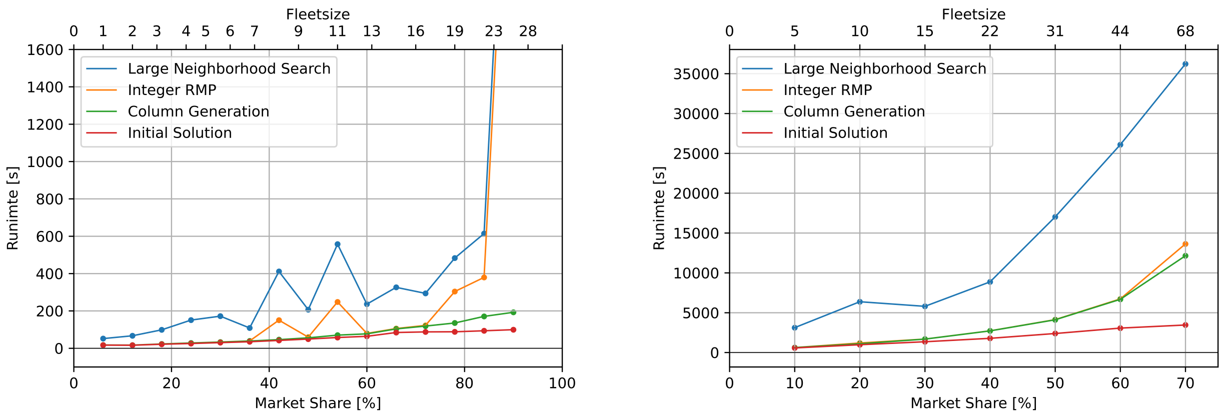

(a) Runtime in the 5 airport network

(b) Runtime in the 20 airport network

Figure 13: Running time for the different algorithms in the 5 and 20 airport networks

As the heuristic to construct the initial solution performs the same label correcting algorithm for each aircraft in the fleet, the runtime of this heuristic increases linearly with fleetsize. Note that fleetsize initially increases linearly with market share but at a certain moment, there is simply not much demand left such that the fleetsize start growing slighly faster. As discussed in Sec. 4.1, this algorithm is NP hard and thus, runtime increases exponentially w.r.t. the graph size of G, in terms of number of nodes and arcs $|N|, |A|$. Luckily, this only takes a small amount of the total runtime and remains computationally tractable, even in the 30 airports network.

The runtime of the Column Generation increases exponentially w.r.t. market share, caused by a rapidly growing number of iterations required to to run CG to optimality, where each iteration takes roughly the same time. Runtime also increases expontentially w.r.t. the size of the graph G because of the pricing subproblem, which is NP-Hard and is solved using a label correcting algorithm. Next, the integer version of the final RMP is solved using the Gurobi 9.1 default solver. Usually, this takes only a fraction of the total runtime but on some rare occasions the optimiser needs much more time, resulting in sudden spikes in the runtime for some market shares.

Lastly, the runtime of the LNS increases linearly with increasing fleetsize because on each iteration of the LNS, the same procedure of finding the optimal replacement path, which takes a comparable time for each replacement path, is performed for each aircraft in the fleet. The number of LNS iterations varies only a little bit. Three or four iterations suffice in most cases. LNS runtime does increase exponentially with increasing size of the graph G, because finding the replacement path happens by solving an adjusted version of the arc based formulation, which is known to be NP-Hard.

Only CG scales exponentially with increasing market share. All methods however scale exponentially with increasing graph size. Whereas in most cases in the 5 airport network could be solved in a matter of a couple of minutes, this took multiple hours for the 20 airport network. Interestingly, at high market shares, the LNS leads to very little improvement in the objective value, often less than 1% while it uses the largest part of the running time. Depending on the time and computational power available, it may not be worth it to execute the LNS when it is observed that it led to little improvements for market shares just below.

6.2 Detailed Analysis of a Solution

When a potential operator is investigating the options and costs related to offering air transport services using electric thin-haul aircraft, they can use the methods presented in this paper to construct a graph such as the one in Fig. 12(a). For each fleetsize, there is a solution with the lowest Cost/RPK. Which fleetsize to choose depends on the objectives and resources of the operator. E.g., the budget may only allow for the acquisition of a set number of aircraft, the operator may have a target market share in mind, the Cost/RPK may need to be below a certain threshold to ensure profitability, etc. Let us consider the case where an operator envisions a target market share of 30% and analyse the solution in detail. As seen in Fig. 14, the fleetsize for 30% and 31% market share is equal but the Cost/RPK is lower at 31%. Therefore, the solution at 31% which is indicated by a star in the figure is the chosen schedule.

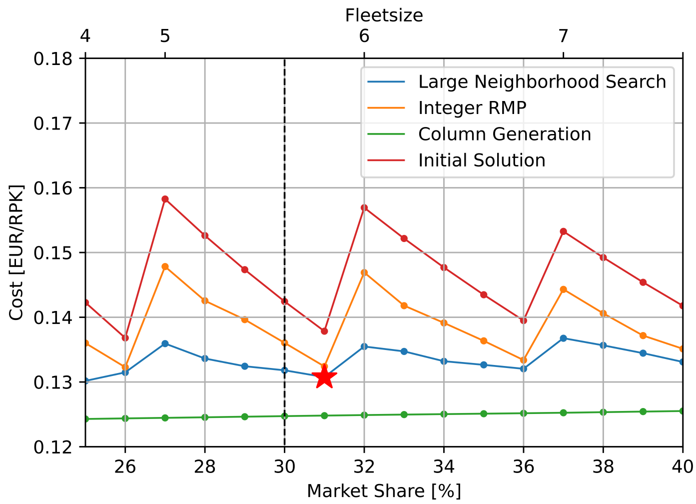

Figure 14: The solution indicated by the star is the chosen schedule

The schedule has a Cost of 13.08 cents per RPK. The lower bound lies 4.81% lower at 12.48 cents per RPK. The 13.08 cents can be interpreted as a lower bound on the ticket prices for the operator to be profitable. Suppose a passenger is flown from Brussels to Paris, which is a 272km flight, this costs the airline €36.6 in direct operating costs. The ticket price must be higher than €36.6 to cover this. Besides direct operating costs, other expenses such as overhead or airport fees also have to be added so in reality, the fares will be higher. This schedule uses a fleet of 6 aircraft. At an estimated sales price of €2M per aircraft, the initial investment in the fleet will be €12M.

The schedule is given in Tab. 10. Let us analyse this schedule in detail.

We observe that the schedule has a strong preference to fly back-and-forth between Brussels and Paris. Table 11 shows that only on three occasions, there is a flight to and from Cologne and only one to and from Amsterdam. Secondly, we see that all flights are out-and-return flights to the hub at Brussels Airport. Let us analyse why these patterns exist.

Aircraft	Origin	Destination	Departure	Arrival	Distance	Passengers
1	Brussels	Paris	6:45	8:00	272 km	9
1	Paris	Brussels	8:45	10:00	272 km	9
1	Brussels	Cologne	10:45	11:45	201 km	8
1	Cologne	Brussels	14:00	15:00	201 km	9
1	Brussels	Paris	15:45	17:00	272 km	9
1	Paris	Brussels	17:45	19:00	272 km	9
2	Brussels	Paris	6:00	7:15	272 km	7
2	Paris	Brussels	8:00	9:15	272 km	9
2	Brussels	Paris	10:15	11:30	272 km	9
2	Paris	Brussels	13:15	14:30	272 km	9
2	Brussels	Paris	15:15	16:30	272 km	9
2	Paris	Brussels	17:30	18:45	272 km	7
3	Brussels	Paris	7:15	8:30	272 km	9
3	Paris	Brussels	9:15	10:30	272 km	9
3	Brussels	Paris	11:30	12:45	272 km	9
3	Paris	Brussels	14:00	15:15	272 km	9
3	Brussels	Paris	16:15	17:30	272 km	9
3	Paris	Brussels	18:45	19:30	272 km	9
4	Brussels	Amsterdam	7:00	7:45	171 km	9
4	Amsterdam	Brussels	8:15	9:00	171 km	9
4	Brussels	Paris	9:45	11:00	272 km	9
4	Paris	Brussels	12:00	13:15	272 km	9
4	Brussels	Paris	14:30	15:45	272 km	9
4	Paris	Brussels	16:45	18:00	272 km	9
5	Brussels	Paris	7:00	8:15	272 km	9
5	Paris	Brussels	10:15	11:30	272 km	9
5	Brussels	Paris	12:15	13:30	272 km	9
5	Paris	Brussels	14:30	15:45	272 km	8
5	Brussels	Cologne	16:45	17:45	201 km	9
5	Cologne	Brussels	18:45	19:15	201 km	7
6	Brussels	Cologne	6:00	7:00	201 km	9
6	Cologne	Brussels	7:45	8:45	201 km	9
6	Brussels	Paris	9:15	10:30	272 km	8
6	Paris	Brussels	11:45	13:00	272 km	9
6	Brussels	Paris	14:15	15:30	272 km	8
6	Paris	Brussels	16:30	17:45	272 km	9

Table 10: Complete Schedule for 31% market share on the 5 airport network

From \ To	Brussels	Paris	Cologne	Amsterdam	Luxemburg
Brussels	-	14	3	1	0
Paris	14	-	-	-	-
Cologne	3	-	-	0	0
Amsterdam	1	-	0	-	-
Luxemburg	0	-	0	-	-

Table 11: Distribution of flights for 31% market share on the 5 airport network

	Paris	Brussels	Cologne	Amsterdam	Luxemburg
Population	2.148 million	1.209 million	1.061 million	0.822 million	0.614 million

Table 12: The population of the cities in the 5 airport network causes large differences in demand

Firstly, the link between Brussels and Paris has the largest demand and thus, the aircraft can fly at full capacity more often whereas on other links, there isn't always enough demand to fill up all seats. The differences in demand are shown in Fig. 15 and are mainly caused by the population of the cities as shown in Tab. 12. Throughout the entire day, demand between Brussels and Paris exceeds the number of seats on the aircraft, which means that one flight per hour can be performed in both directions (Bru-Par and Par-Bru) without leaving a single seat empty. In the time intervals [07:00 - 09:00) and [15:00-17:00) there is even enough demand for a second flight at full capacity. The connection with the second highest demand is Brussels-Cologne, but already here, there is a gap from 10:00 to 14:00 and after 18:00 where there is not enough demand for even just one flight per hour without leaving any seats empty. Of the 6 flights between Brussels and Cologne, 2 are scheduled within this time period and must therefore leave 1 and 2 seats empty respectively. Next is the Brussels-Amsterdam connection. Only between [07:00, 09:00) and [15:00, 17:00) can an aircraft fly at full capacity between these two cities. The only two flights between Brussels and Amsterdam are therefore scheduled within these times and don't leave any seat empty.

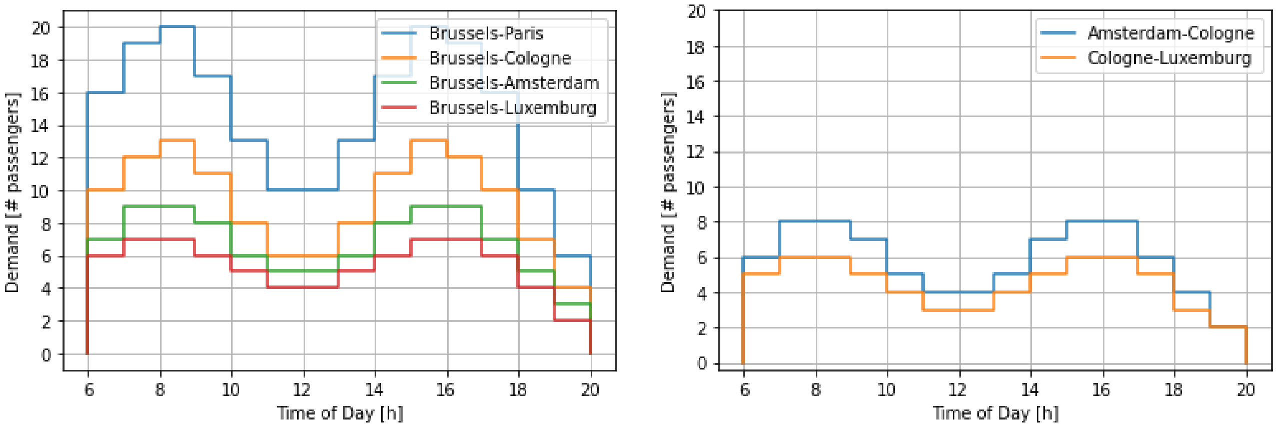

(a) (b)

Figure 15: Visualisation of the demand on all connections in the 5 airport network

Secondly, the distance between Brussels and Paris is almost at the maximum flight range of the aircraft. Recall from Sec. 5.1 that energy consumption of a flight was composed of a fixed amount of energy and a part varying linearly with distance. As a result of the fixed part, the energy per kilometer decreases with flight distance as shown in Fig. 16(a). As the cost of a flight are partly based on a per kWh basis, also cost per kilometer decrease with increasing flight distance as visualised in Fig. 16(b). The solution therefore prefers to schedule flights with a longer distance. In some cases it may still be wise to include shorter distance flights as well. Longer flights also have a longer flight time and total energy consumption, so they also take more time to recharge. If more flights can be scheduled for an aircraft by using shorter flights, the fixed daily ownership cost of €440/aircraft can be distributed over a larger RPK, resulting in a lower overall cost per RPK. Furthermore, a lack of demand may also prevent the solution from scheduling longer flights.

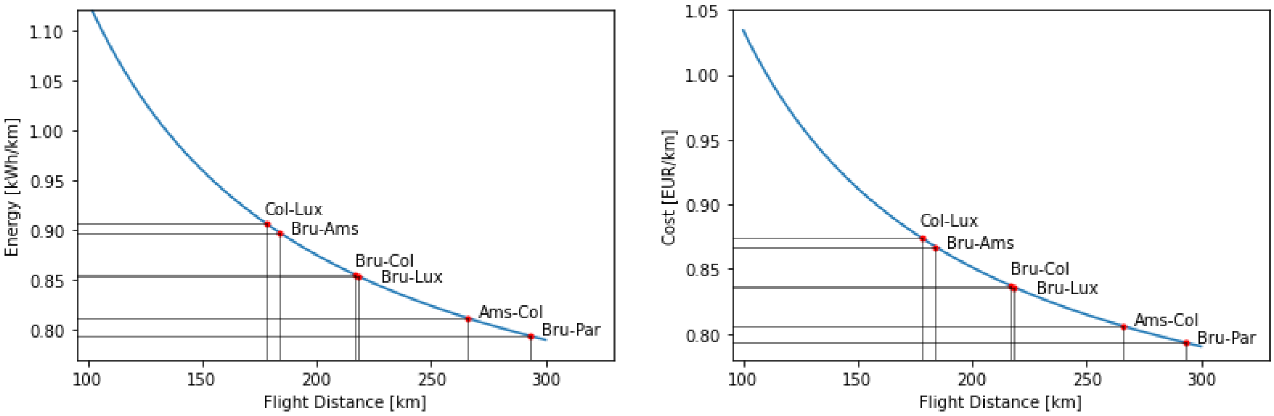

(a) Energy consumption per km (b) Cost per km

Figure 16: Energy and Cost per RPK decrease with flight distance, assuming a constant number of s passengers

Lastly, the hub at Brussels Airport lies at a convenient central location in the network and has access to the 3 connections with the highest demand (Bru-Par, Bru-Col, Bru-Ams) out of 5 connections in total. These connections are the first to be exploited when a market share of only 31% is required. Note that the links are non-directional, meaning that demand, cost and energy consumption are identical for a flight from a_1 to a_2 and from a_2 to a_1. If the schedule includes a flight from Brussels to Paris for example, then the return flight is also worth considering. However, when a larger market share is required, at some point the demand to and from Brussels will be largely exploited and it may be necessary to include other flights such as Paris-Luxemburg, breaking the out-and-return pattern.

Each aircraft is assigned six flights to get the most out of each one. As a result, 6 peaks can be seen in the number of simultaneous flights in Fig. 17. Note that not all six aircraft start flying immediately. As aircraft #1 and #2 occupy the demand from Brussels to Paris in the first hour $[06.00, 07.00)$, aircraft #3 and #5 stay on the ground until 07.15 and 07.00 respectively to serve demand on the same route but in the next hour. By the time all aircraft took-off for their first flight of the day at 07.00, two aircraft already landed (#2 at 07.15, #6 at 07.00). Hence, the first peak in simultaneous flights does not exceed 4 at this stage. Recall that demand varied throughout the day but even at the lowest point, between 11.00 and 13.00, there is enough demand to schedule 4 flights without empty seats, both ways between Brussels and Paris during each hour. The schedule in Tab. 10 shows that indeed 4 flights are scheduled as such. It makes more sense to schedule as many flights per aircraft as possible instead of waiting on the ground to avoid the off-peak hours. As a result, the variation in demand throughout the day is not really reflected in number of simultaneous flights throughout the day.

Also shown in Fig. 17 is the number of chargers in use at Brussels. Throughout the day, at most two chargers are in use simultaneously. Near the end of the day the number of chargers peaks to 5. However, as the aircraft have all night to recharge, two chargers would actually suffice. Next to the two chargers at Brussels, there must also be two charger at Paris, and one at Amsterdam and Cologne. This makes 6 chargers in total to keep this fleet of 6 aircraft airborne.

6.3 Comparison of Multiple Solutions

Whereas the previous section analysed the solution for 31% market share in detail, this sections compares the solutions for the market shares indicated in Fig. 18 and looks for a pattern between different solutions.

In Sec. 6.2, we learned that at 31% market share, flights between Paris and Brussels make up the bulk of the schedule because of the long distance between them and high demand. The distribution of flights when the market share is increased to 49% is shown in Tab. 13. The additional market share is mainly achieved by adding flights between Brussels-Cologne and Brussels-Amsterdam, the two pairs with the highest demand after Brussels-Paris. There is also a flight from Amsterdam to Cologne, which has the second largest flight distance after Brussels-Paris. Lasty, one out-and-return flight between Brussels and Luxemburg is scheduled.

In Tab. 14 and Tab. 15, the market share is increased to 63% and 81% respectively. More flights are added on the Brussels-Amsterdam and Brussels-Cologne connections, which become largely saturated. At 81% market

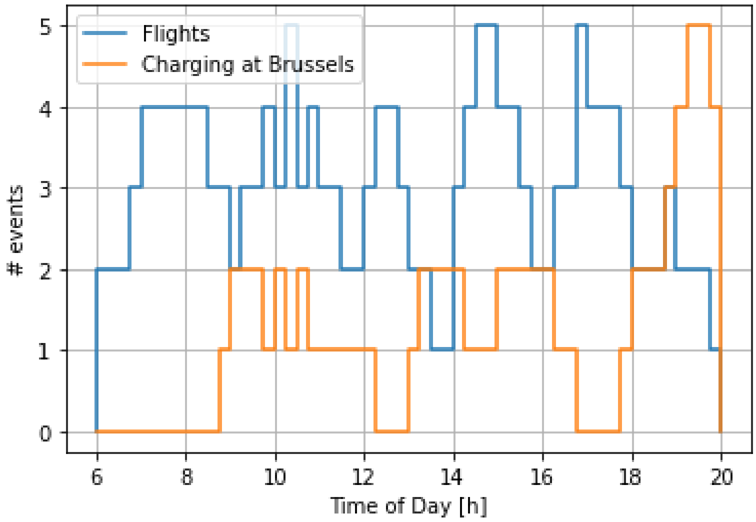

Figure 17: Number of flights occurring simultaneously and number of chargers in use at Brussels throughout the day

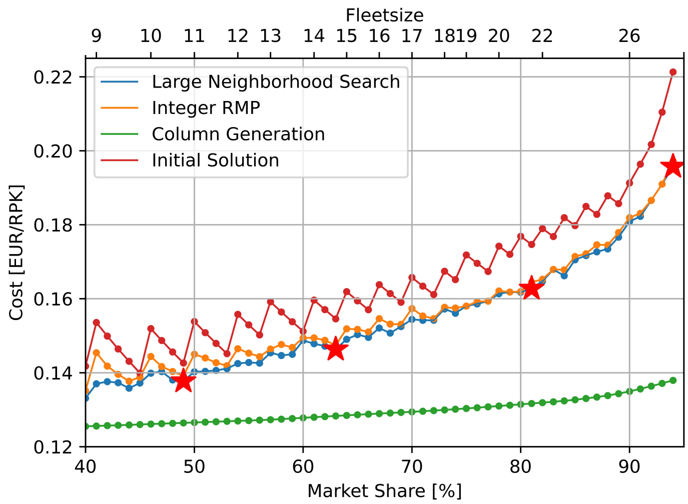

Figure 18: The solutions of 49%, 63%, 81% and 94% market share are compared in this Section

share, flights are also scheduled between Cologne and Luxemburg. Even though the short flight distance and low demand between these two cities makes it unattractive to schedule flights between them, they are necessary to achieve the desired market share and can no longer be avoided. The higher cost per km related to the short flight distance is amplified by the lack of demand. The cost per km between Luxemburg and Cologne is €0.87/km as seen in Fig. 16(b). If the aircraft could fly at full capacity, with 9 passengers on board, this would result in a Cost/RPK of 9.7 cents on this flight. However, demand here peaks at 6 as shown in Fig. 15(b). As the cost per km can only be distributed over 6 passengers, the cost per RPK increases by 49% to 14.5 cents.

When the market share is increased, the schedule is forced to include more and more flight with a lower demand. This can be seen in the average load factor of the schedule. The load factor of a flight is the percentage of seats that is occupied. The load factor of a schedule is the average load factor over all flights in the schedule. The inclusion of of low demand flights results in a decreasing average load factor, as shown in Fig. 19.

Lastly, Tab. 16 shows the distribution of flights at the maximum feasible market share of 94%. Predictably, the largest number of flights is found where most demand is, between Brussels and Paris. The lowest number of flights is found between Cologne-Luxemburg and Cologne-Amsterdam. An aircraft first needs to fly from the hub (Brussels) to Amsterdam, Cologne or Luxemburg before a flight between them can be scheduled. The window to schedule these flights is therefore shorter and that is reflected in the distribution of flights.

From \ To	Brussels	Paris	Cologne	Amsterdam	Luxemburg
Brussels	-	16	8	7	1
Paris	16	-	-	-	-
Cologne	9	-	-	0	0
Amsterdam	6	-	1	-	-
Luxemburg	1	-	0	-	-

Table 13: Distribution of flights for 49% market share on the 5 airport network

From \ To	Brussels	Paris	Cologne	Amsterdam	Luxemburg
Brussels	-	18	11	9	3
Paris	18	-	-	-	-
Cologne	12	-	-	3	0
Amsterdam	8	-	4	-	-
Luxemburg	3	-	0	-	-

Table 14: Distribution of flights for 63% market share on the 5 airport network

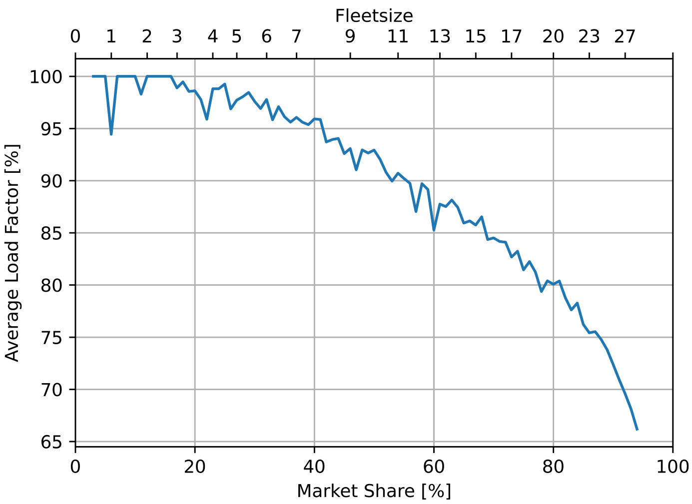

Figure 19: The average load factor of the schedule decreases with increasing market share

From \ To	Brussels	Paris	Cologne	Amsterdam	Luxemburg
Brussels	-	21	18	17	16
Paris	21	-	-	-	-
Cologne	19	-	-	10	10
Amsterdam	16	-	11	-	-
Luxemburg	16	-	10	-	-

Table 16: Distribution of flights for 94% market share on the 5 airport network

From \ To	Brussels	Paris	Cologne	Amsterdam	Luxemburg
Brussels	-	19	13	11	9
Paris	19	-	-	-	-
Cologne	13	-	-	8	4
Amsterdam	10	-	9	-	-
Luxemburg	10	-	3	-	-

Table 15: Distribution of flights for 81% market share on the 5 airport network

To achieve the minimum RPK, the algorithm tends to schedule back-and-forth flights between airports with a long distance and high demand. These flights are the first to be scheduled until they are largely saturated. When this happens and the market share is increased further, the schedule will start planning flights with a shorter distance and/or lower demand.

7 Conclusions and Recommendations

In this paper, we addressed a scheduling problem, aimed at electric thin-haul aircraft. The problem was to determine a schedule that minimises cost under a minimum RPK (Revenue-Passenger Kilometer) constraint, which can also be interpreted as a minimum market share. By repeatedly solving this problem for various RPKs, a solution that minimises cost per RPK can be found. Because the problem is NP-hard and could not be solved to optimality using the MILP (Mixed Integer Linear Programming) arc based formulation, we developed our own algorithm. The algorithm was implemented and tested on a range of networks, varying from 5 airports to 30 airports for a scheduling horizon of one day.

The algorithm started by heuristically minimising the fleetsize to satisfy the minimum RPK constraint. A lower bound to the objective value was obtained by solving the linear relaxation of the path based formulation through CG (Column Generation). Though the integer version of the path based formulation is equivalent to that of the arc based formulation, its linear relaxation provides a much stronger lower bound. Next, we improve on the initial solution by solving the integer version of the final RMP (Restricted Master Problem). Lastly, as Branch-and-Price proved unsuccessful to find the global optimum, a LNS (Large Neighborhood Search) was designed to improve on the best-known integer solution until a local optimum.

The algorithm proved to be successful in finding a very good local optimum. The case studies revealed that the Cost per RPK is not a smooth function, rather showing a sawtooth pattern. Each time an aircraft was added to the fleet to meet the desired market share, cost per RPK shoot up and gradually decrease when the minimum market share is increased further until another additional aircraft is required. For a given fleetsize, the lowest Cost/RPK is obtained when the fleet is operating near full capacity.
The schedule has a strong preference for flights with i) high demand; and ii) long distance. When demand is large enough, no seats must be left empty and a larger RPK is achieved for the same cost. Long distance flights are preferred because the energy and cost per kilometer is lower for long distance flights. To achieve the minimum RPK, the algorithm tends to schedule back-and-forth flights between airports with a long distance and high demand. These flights are the first to be scheduled until they are largely saturated. When this happens and the market share is increased further, the schedule will start planning flights with a shorter distance and/or lower demand. The general trend in the sawtooth pattern is therefore an increasing Cost per RPK with increasing market share.

The contributions of this paper are as follows. To the best of our knowledge, this is the first paper that considers cost per RPK as an evaluation criterion in an E-VSP (Electric-Vehicle Scheduling Problem) and is the first paper to apply operations research to electric thin-haul aircraft. We proposed a new variation of the E-VSP where the set of service trips is not known in advance but rather had a minimum RPK constraint. We then propose a method to solve the problem, which proved successful in finding a good local minimum. The methods and insights developed in this paper will support future thin-haul airline operators in deciding how many aircraft to acquire and provide a schedule that minimises cost per RPK.

Further research could focus on the following extensions to the problem. i) The multi-depot variation. In this paper, there was only 1 hub. It may however be interesting to have multiple hubs where the aircraft can stay overnight. This variation would provide access to more flights and the flights around the hub would be less congested at high market shares; ii) Adding a charger constraint. In this paper, the number of chargers required at any airport is assumed to be present. In reality, the number of chargers will be limited. This could be achieved by adding variables to track whether an aircraft is charging on a ground arc or not, and adding a constraint to limit the number of aircraft charging simultaneously on the same ground arc.; iii) A battery swap variation. Instead of using chargers, another variation may use battery swapping or a combination of charging and swapping.; and lastly iv) Include variable electricity prices. The cost of electricity may vary drastically by geographic location, time of day, peak charging power and total charging volume. A variation of the problem where electricity costs are incurred on the ground arcs instead of on flight arcs may allow to include variations based on location and time of day. An additional variable will be required to track whether an aircraft on the ground is in fact charging, or just waiting. This thesis took the first steps in the field of electric thin-haul airline operations. With the current plethora of opportunities for further development, it promises to remain a fascinating field for the upcoming years.

References

[1] Schneider M. Stenger A. and Goeke D. The electric vehicle-routing problem with time windows and recharging stations. *Transportation Science*, 2020.

[2] Verma A. Electric vehicle routing problem with time windows recharging stations and battery swapping stations. *European Journal of Tranportation and Logistics*, 2018.

[3] Harish A. Perron C. Bavaro D. Ahuja J. Ozcan M. Justin C.Y. Briceno S.I. German B.J. and Mavris D. Economics of advanced thin-haul concepts and operations. *Journal of Aerospace Information Systems*, 2016.

[4] Justin C.Y. Payan A.P. Briceno S.I. German B.J. and Mavris D.N. Power optimized battery swap and recharge strategies for electric aircraft operations. *Transportation Research Part C: Emerging Technologies*, 2020.

[5] Koc C. and Karaoglan I. The green vehicle routing problem: A heuristic based exact solution approach. *Applied Soft Computing*, 2016.

[6] Moore M. D. et al. Misconceptions of electric aircraft and their emerging aviation markets. *Aerospace Sciences Meeting*, 2014.

[7] Felipe A. Ortuno M.T. Righini G. and Tirado G. A heuristic approach for the green vehicle routing problem with multiple technologies and partial recharges. *Transportation Research Part E: Logistics and Transportation Review*, 2014.

[8] Air Transport Action Group. Facts & figures. `https://www.atag.org/facts-figures`, 2021.

[9] Li. L. Lo H.K. and Xiao F. Mixed bus fleet scheduling under range and refueling constraints. *Transportation Research Part C: Emerging Technologies*, 2019.

[10] Cortes-Murcia D.L. Prodhon C. Asfar H.M. The electric vehicle routing problem with time windows, partial recharges and satellite customers. *Transportation Research Part E: Logistics and Transportation Review*, 2019.

[11] Koyuncu I. and Yavuz M. Duplicating nodes or arcs in green vehicle routing: A computational comparison of two formulations. *Transportation Research Part E: Logistics and Transportation Review*, 2019.

[12] Wang Y. Huang Y. Xu J. and Barclay N. Optimal recharging scheduling for urban electric buses: A case study in davis. *Transportation Research Part E: Logistics and Transportation Review*, 2017.

[13] Li J.-Q. Transit bus scheduling with limited energy. *Transportation Science*, 2014.

[14] Adler J.D. Routing and scheduling of electric and alternative-fuel vehicles. *Ph.D. thesis Arizona State University*, 2014.

[15] An K. Battery electric bus infrastructure planning under demand uncertainty. *Transportation Research Part C: Emerging Technologies*, 2020.

[16] Steiner K. and Irnich S. Schedule-based integrated intercity bus line planning via branch-and-cut. *Transportation Science*, 2018.

[17] Silva H. L. and Guimaraes T. A. Conceptual design of a thin-haul aircraft by energy sizing optimization including aero-propulsive interactions. *Journal of Aerospace Information Systems*, 2020.

[18] Keskin M. and Catay B. Partial recharge strategies for the electric vehicle routing problem with time window. *Transportation Research Part C: Emerging Technologies*, 2016.

[19] Kleinbekman I.C. Mitici M. and Wei P. Rolling-horizon electric vertical takeoff and landing arrival scheduling for on-demand urban air mobility. *Aerospace Information Systems*, 2019.

[20] Leggieri V. Haouari M. A practical solution approach for the green vehicle routing problem. *Transportation Research Part E: Logistics and Transportation Review*, 2017.

[21] Pereira M. Short-range route scheduling for electric aircraft with battery-charging and battery-swapping constraints. *MSc thesis TU Delft*, 2019.

[22] Courtin C. Burton M. Butler P. Ye A. Vascik P. and Hansman J. Feasibility study of short takeoff and landing urban air mobility vehicles using geometric programming. *Aviation technology, integration and operations conference*, 2018.

[23] D.P. Raymer. *Aircraft design: A conceptual approach.* American Institute of Aeronautics and Astronautics, Washington D.C., 1989.

[24] J. Roskam. *Airplane Design.* DARcorporation, 1997.

[25] Dessaulniers G. Errico F. Irnich S. and Schneider M. About sections exact algorithms for electric vehicle-routing problems with time windows. *Operations Research*, 2016.

[26] Erdogan S. and Miller-Hooks E. A green vehicle routing problem. *Transportation Research Part E: Logistics and Transportation Review*, 2012.

[27] Justin C.Y. Payan A.P. Briceno S.I. and Marvis D.N. Operational and economic feasibility of electric thin haul transportation. *Journal of Aerospace Information Systems*, 2017.

[28] Wei L. Justin C.Y. Briceno S.I. and Mavris D.N. Door-to-door travel time comparative assessment for conventional transportation methods and short takeoff and landing on demand mobility concepts. *Transportation Research Part E: Logistics and Transportation Review*, 2018.

[29] Liu T. and Ceder A. Battery-electric transit vehicle scheduling with optimal number of stationary chargers. *Transportation Research Part C: Emerging Technologies*, 2020.

[30] van Kooten Niekerk M.E. van den Akker J.M. and Hoogeveen J.A. Scheduling electric vehicle. *Public Transport*, 2017.

[31] Sun X. Wandelt W. and Stumpf E. Competitiveness of on-demand air taxis regarding door-to-door travel time: A race through europe. *Transportation Research Part E: Logistics and Transportation Review*, 2018.

[32] Tang X. Lin X. and He F. Robust scheduling strategies of electric buses under stochastic traffic conditions. *Transportation Research Part C: Emerging Technologies*, 2019.

3
Literature Study
previously graded under AE4020

Fleet Sizing & Scheduling of Electric Aircraft in a Thin-Haul Network

Literature Review

Bart Debeuckelaere

TUDelft

Contents

List of Figures

List of Tables

Abbreviations and Nomenclature

AC	Alternating Current
ASK	Available Seat Kilometer
BMS	Battery Management system
DC	Direct Current
DEP	Distributed Electric Propulsion
DF	Deficit Function
DoD	Depth-of-Discharge
EPRI	Electric Power Research Institute
ETS	Emission Trading Scheme
EU	European Union
EV	Electric Vehicle
E-VRP	Electric-Vehicle Routing Problem
E-VSP	Electric-Vehicle Scheduling Problem
G-VRP	Green Vehicle Routing Problem
KPI	Key Performance Indicator
Li-air	Lithium-air
Li-ion	Lithium-ion
Li-S	Lithium-Sulfur
LTO	Landing and Take-Off
MILP	Mixed Integer Linear Programming
MTOW	Maximum Take-Off Weight
Ni-Cd	Nickel-Cadmium
Ni-MH	Nickel-Metal Hybrid
RC	Radio-Controlled
RMP	Restricted Master Problem
RPK	Revenue Passenger Kilometer
SAE	Society of Automotive Engineers
SF	Surplus Function
SoC	State-of-Charge
STOL	Short Take-Off and Landing
TAT	Turn Around Time
TS	Time-Space
TSE	Time-Space-Energy
UAM	Urban Air Mobility
VRP	Vehicle Routing Problem
VSP	Vehicle Scheduling Problem
VTOL	Vertical Take-Off and Landing

1 Introduction

Commercial aviation must become more sustainable. Aircraft are currently emitting 4% of all greenhouse gasses and the public opinion is becoming more opposed to flying. Electric aircraft could be part of the solution. Recent developments in battery and charging technology have been very promising. In operations research, a lot of work has been dedicated to Urban Air Mobility (UAM). In this concept, an autonomous Vertical Take-Off and Landing (VTOL) aircraft provides on-demand mobility. Though technology is developing at a high pace, there are many hurdles to this concept in regulation, certification and public acceptance. Meanwhile electric aircraft with up to 10 seats are already in the flight testing phase. A number of economic feasibility studies conclude that electric aircraft provide opportunities for huge costs savings for commuter airlines [8] [53], cite46. Yet there is very little research dedicated to their operations, which is vastly different from fossil fuel powered aircraft.

Some experts still consider electric propulsion infeasible for commercial aviation and advise research to focus on electric regional airlines with a potential market introduction 20 years from now. [17] consider this as a misconception and states the importance of having an agile and adaptable research and market plan, starting with smaller applications. Current commercial aviation is receiving most attention because that is where the bulk of the revenue is but there is a poor fit with current technology levels. Meanwhile, electric propulsion is already common place for gliders equipped with so called self-starter motors and Pipistrel and Siemens are already flying electric ultralights. Electric aircraft are slowly entering general aviation. The logical next step is electric aircraft with a seating capacity of 6 to 10 people. MagniX started test flying an electrified Cessna Caravan in June 2020, Eviation also already has a prototype of their electric aircraft and Skylax intends to have one next year. This type of aircraft is entering the ballpark of Thin-Haul commuter airlines. This type of airline provides scheduled services over small distances ranging from just ten to a couple hundreds of kilometers in small aircraft with up to 10 seats. Due to their high costs of 0.40 €/Revenue Passenger Kilometer (RPK) compared to an airline's 0.11 €/RPK, commuter airlines are only found where there are no alternatives for fast transport such as the Hawaiian islands or Alaska. Some of these operators are already planning to electrify their fleets and yet, very little research is devoted to operating an electric fleet. As some research has shown that electric aircraft are able to substantially reduce costs, more opportunities for commuter airline services may appear. The urgency of operations research on electric thin-haul commuter aircraft is clear.

This literature study reviews the state-of-the-art in operations research of electric aircraft and closely related fields. Firstly, this study will look which optimization models have been developed in the past. Though there is very little research on electric commuter aircraft, there is a lot of research on two very closely related topics. i) Electric-Vehicle Routing Problem (E-VRP) research applies the famous Vehicle Routing Problem (VRP) to electric vehicles, often representing cars or busses. A lot of progress has been made on E-VRP research, making the material useful to apply in more specific topics; ii) Electric-Vehicle Scheduling Problem (E-VSP), mostly applied to electric bus scheduling as electric busses are becoming a viable alternative to fossil fuel powered busses, This area of research attempts to develop efficient and useful optimisation models to answer questions around fleet sizing, charging infrastructure and scheduling.

Each of the optimization models need some input to start from. Firstly, they require an estimation of demand. Previous research in E-VRP and E-VSP usually assume a predefined set of service trips the fleet must conduct. That is realistic where an operator intends to transition an existing service to electric vehicles. A number of variation on this problem exist. One such example is the use of time-windows for the service trips.

Equally important are the aircraft model parameters. These can be obtained from a comparison and critical assessment of the claims of the companies currently working on an electric aircraft. A number of companies are showing very promising results. Nevertheless, their claims are likely to change when technology matures, either because they overestimated themselves or as technology advances. Therefore an extensive sensitivity analysis is required to draw valid and general conclusions. The battery and charging model are also paramount. Past research has often assumed the battery State-of-Charge to decrease linearly in operations and increase linearly up until a certain threshold while charging. More recent papers sometimes treat the charging process as a non-linear process or bi-linear process to make the model more realistic. A small number of feasibility studies of electric thin-haul commuter airlines provide a baseline to estimate operational costs. Lastly, emissions are important to consider because these also give electric aircraft an advantage over gas powered aircraft. Though electric aircraft produce no local emissions, the production of electricity does. Taking this into account, the effective CO_2 emissions of electric aircraft are over an order of magnitude smaller than for fuel powered aircraft and in the same range as gasoline powered cars.

In the remainder of this literature study, the Electric-Vehicle Routing Problem (E-VRP) and Electric-Vehicle Scheduling Problem (E-VSP) will be discussed in Section 2.1 and Section 2.2 respectively. The developments in electric aircraft and batteries will be discussed in Section 3.1 and Section 3.2 respectively. Section 3.3 presents the current literature on electric thin-haul aviation. Lastly, considerations towards emissions are presented in Section 3.4.

2 Optimisation Models

The research in the field of electric aircraft scheduling is very limited. There exist extensive amounts of research in traditional airline operations, which is very different to operating small electric aircraft. Recently, UAM has become a very active field of research. These assume on-demand mobility using VTOL aircraft and is thus also very different from operating a scheduled electric aircraft fleet. Two fields of operations research are deemed most relevant. Firstly, the E-VRP considers the classical VRP for an electric vehicle. A set of customer nodes is represented on a graph. The problem consists in finding an optimal set of routes such that each customer node is visited exactly once. This field of research is discussed in Section 2.1. Note that the E-VRP is known to be NP-Hard. Therefore, most research uses heuristics to efficiently solve the instances. This literature study will focus on the Mixed Integer Linear Programming (MILP) formulations and considerations given to operations of a fleet of electric vehicles. This study will not go into the heuristic methods that have been developed to solve real life instances of the E-VRP in detail.

A second relevant field of research considers the E-VSP and is discussed in Section 2.2. Instead of a set of customer nodes as is the case in the VRP, the Vehicle Scheduling Problem (VSP) is concerned with a set of service trips. Each service trip has a start and end location and duration. The problem is to schedule a fleet of electric vehicles such that each service trip is performed under the given constraints. In the single-depot problem, all vehicles start and end their service from the same vehicle depot. A variation to this problem is the multi-depot problem. The single-depot VSP is Polynomial bound ($O(N^3)$) but multi-depot VSP and the E-VSP are NP-Hard [60]. Some case studies apply the E-VSP to an intra- or inter-city bus network. The problem is usually represented on a graph. Obviously a bus is very different from an aircraft but many analogies can be drawn, making the operations research very relevant for the scheduling of an electric aircraft. Firstly, the charging infrastructure and technology is practically identical. The same issues arise for electric aircraft as for electric busses such as long charging time, expensive charging infrastructure, etc. Both are carrying a relatively small amount of passengers on scheduled trips where demand varies drastically throughout the day. Both are more expensive in acquisition and have a lower performance than their fossil fuel powered equivalents, which is offset by lower operational and maintenance costs.

2.1. The Electric-Vehicle Routing Problem (E-VRP)

The first paper discussing alternative vehicles for vehicle routing was [52]. The author describes the Green Vehicle Routing Problem (G-VRP) to help organisations overcome the difficulties of operating vehicles using alternative power sources of limited driving range and limited refuelling infrastructure. The problem is represented on a graph $G(V,E)$. The set of vertices or nodes V consists of the set of customer nodes, refuelling station nodes and one depot node. Each node can only be visited once. To allow for multiple visits to the refuelling stations, the author decided to create the augmented graph $G'(V',E')$ which contains a number of copies of the refuelling station nodes. The objective is to minimize the total travelled distance, as shown in Equation 2.1.

$$min \sum_{i,j} d_{ij} x_{ij} \qquad \forall i,j \in V', i \neq j \tag{2.1}$$

Note that this paper is not considering electric vehicles. However, as the service time required at each refuelling station is assumed to be 15 minutes, this can be interpreted as fully changing, which takes 15 minutes to charge. That may be a very optimistic charging assumption but the important thing is that this is a variable that can be changed in the model. The distance travelled is minimized under a maximum time constraint using a MILP optimisation model. As the VRP is NP-Hard and the VRP is a special case of the G-VRP, the G-VRP is also NP-Hard, meaning that the exact optimum cannot be found in polynomial time. As the instance size increases, the computation time will increase exponentially, making it impossible to efficiently solve large instances exactly. Therefore, the author solves the optimisation using a combination of heuristic methods. Numerical experiments have shown that the heuristics perform well compared to exact methods for small instances and that they can be used to solve large instances as well. This paper provides a solid foundation for further research. Recommendations include expanding the model with additional constraints to make it more relevant in practice and developing better heuristic approaches.

2.1.1. Time-Windows, Limited Capacity & Partial Recharge

[2] introduces the electric vehicle routing problem with time windows and recharging stations, which builds on the work of [52] by including customer time-windows and limited capacity. Customer time windows and capacity are important constraint for delivery services in practice. The authors present a hybrid heuristic method to solve the problem using a combination of a variable neighbourhood search algorithm and tabu search. Constraint 2.2 is added to ensure the correct changing time until a full charge upon visit to a charging station, where τ_i is the time of arrival at node i, g the recharge rate, Q the battery capacity and y_i the remaining battery capacity at node i. This constraint ensures that the battery is fully charged. Constraint 2.3 ensures that each customer is visited within its time window. Lastly, Constraints 2.4 and 2.5 ensure that the capacity at the start of a tour is within the limits and that it never falls below zero during the tour, where u_i is the remaining cargo at node i and C the total capacity of the vehicle.

$$\tau_i + (t_{ij}x_{ij}) + g(Q - y_i) - (l_0 + gQ)(1 - x_{ij}) \leq \tau_i \qquad \forall i \in F', j \in V'_{N+1}, i \neq j \tag{2.2}$$

$$e_j \leq \tau_i \leq l_j \qquad \forall j \in V' \tag{2.3}$$

$$0 \leq u_j \leq u_i - q_i x_{ij} + C(1 - x_{ij}) \qquad \forall i \in V'_0, j \in V'_{N+1}, i \neq j \tag{2.4}$$

$$0 \leq u_0 \leq C \tag{2.5}$$

[20] also expands on the work of [52] but in a different way by introducing the Green Vehicle Routing Problem with Multiple Technologies and Partial Recharges. The contributions of this paper are the inclusion of partial recharges and minimisation of costs. Partial charging allows for a more optimal solution by expanding the feasible space. This paper minimises costs instead of distance travelled. In reality, the costs will be a decisive factor for companies to decide on the electrification of their vehicle fleet. In doing so, the authors made a very valuable contribution towards operationalizing electric fleets. The paper also considered multiple charging technologies. The simplest technologies are the cheapest and take more time to charge. They discovered that no technology outperformed the others in all cases. Firstly, the objective function is changed to Equation 2.6 to minimize costs instead of distance. Because costs depend on multiple components, the objective function becomes more complex. The first term represents the cost of overnight charging at the depot, the second term the cost of charging en-route and the third term the fixed to travel a link. The cost of electricity is represented by γ, g_v is the amount charged at the depot and Z_{jt} is the amount charged at charging stations using technology t. Because this is a decision variable, partial charges are now possible

using multiple technologies. Constraint 2.7 ensures that a partial recharge is allowed and that multiple charging technologies can be considered. Constraint 2.2 to ensure the correct charging time also has to adapt accordingly. A numerical study revealed that of the heuristic methods developed, the local exhaustive search heuristic performed best for instances with less than 200 customers while a simulated annealing algorithm worked best for larger instances.

$$min \sum \gamma g_v + \sum \sum \gamma_t z_{jt} + \theta \sum xij \tag{2.6}$$

$$y_j^L = y_j^A + \sum_{t \in T_j} z_{it} \qquad \forall j \in F_0 \tag{2.7}$$

[24] presents a similar model with time windows and partial charging using multiple technologies. Their contribution was the development of a math-heuristic method to find the exact solution. In a first stage, an Adaptive Large Neighbourhood Search is used to find a good solution. Starting from this good solution, a commercial solver can be used to get to the exact solution.

In [50], the authors. present 4 variations of the model presented in [52]; i) A single, full recharge per route; ii) Multiple full charges per route; iii) A single partial recharge per route and iv) Multiple partial recharges per route. Additionally, instead of heuristic methods, the paper presents exact branch-price-and-cut algorithms to get to an exact solution. The study reveals that allowing multiple, partial recharges allows to reduce costs up to 5% at the cost of increased computational complexity. Essentially, the variations of the model that this paper discusses don't reveal anything new. The main contribution is in branch-price-and-cut methods to solve the problem exactly.

[38] proposes a variation on [2], extended to allow for partial recharge. The authors propose an adaptive large neighbourhood search algorithm to solve the instances. The average improvement by allowing partial recharge was roughly 2.9% over a full recharge. This improvement comes at the cost of an increase in computation time by 10% to 25% depending on the instance.

In all of the previous papers, the model is based on an augmented graph which includes a number of copies of recharging nodes to allow for multiple visits. This introduces additional binary variables to an already large problem. To avoid this, [11] proposes a different method. The authors do this by adding 3-point binary variables $z_{i,r,j}$ to indicate when a vehicle travels from depot or customer node i to depot or customer node j via recharging node r. These variables replace the variables $x_{i,r}$ and $x_{r,i}$. As a result, the time and energy level variables only exist for the depot and customer nodes and not for the recharging nodes anymore as the value of the time and energy level after passing by a recharge station can immediately be calculated for at the next customer node. In this model, charging nodes are no longer required. This method is known as the arc-duplication method. There are two arcs between each pair of nodes, one direct arc and one recharging arc. The recharging arc passes by the closest recharging station. [24] compares the computational effort of the node duplicating and arc duplicating method and found that the less common arc-duplication method outperforms node-duplication.

[40] further improves on this formulation. The author achieves this by developing a nonlinear formulation that is linearized using the Reformulation-Linearization Technique. The paper subsequently proves that their newly proposed formulation can be solved using commercially available solvers and consistently outperforms the state-of-the-art branch-and-cut methods as previously proposed in [50] and [11]. [46] proposes a three-stage heuristic method but the method does not appear to be very competitive.

2.1.2. Battery Swap

[3] is the first paper to explicitly model battery swap instead of battery charging, although the original G-VRP from [52] used a fixed service time, which realistically models a battery swap. When

a vehicle reaches a service station, it can choose to charge the batteries or to perform a battery swap, in which case the vehicle receives a new, fully charged battery. A battery swap takes a constant amount of time and assumes a fixed cost. This process is much quicker than fully charging an empty battery and as a result, the number of vehicles can be reduced. Electric vehicles are very expensive so a reduced fleet size is a huge advantage. However, additional batteries must be acquired and batteries happen to be the most expensive part of an electric vehicle. The mathematical model is almost identical to the model that only allows for a full recharge. The time to charge is replaced by a constant swap time and a fixed cost per battery swap is assumed if the model is optimizing for costs.

2.1.3. Other Variations

[23] proposes a variant of the E-VRPTW by allowing satellite customers. The idea is that while a vehicle is charging, once customer can be visited using an alternative mode of transport. The driver could for example walk a few hundred meter to deliver a package. To model the alternative mode of transport, a certain distance may be covered during the charging period at a given speed ratio.This allows to reduce the distance travelled by the vehicles and makes use of the idle time while charging.

2.1.4. Final Remarks

The original G-VRP [52] provided a very good baseline for future work in Electric Vehicle Routing. A number of extensions were proposed to make the model more practically relevant and a number of solution techniques have been developed. Table 2.1 summarises the current literature concerning the E-VRP. The fourth column denotes whether the servicing is modelled using a constant servicing time (C), a partial recharge (RP) or a full recharge (FR). The problem is NP-Hard and a range of heuristic methods were proposed to solve real life instances. Further work in this field is may be to study different extensions to the problem that could be relevant in practice while further improving on the methods to solve the problem to efficiently handle larger instances.

Paper	Objective	Time Windows	Service	#Customers	# Vehicles	Exact
[52]	Distance	-	C	500	80	-
[2]	Distance	+	FR	500	80	-
[20]	Operational Cost	-	PR	500	80	-
[24]	Operational Costs	+	PR	50	?	+
[50]	Operational Cost	+	PR	100	15	+
[38]	Distance	+	PR	100	15	-
[40]	Distance	-	C	20	10	+
[11]	Distance	+	C	20	10	+
[3]	Total Cost	+	C/FR	100	15	-
[23]	Charging Time	+	PR	100	15	-

Table 2.1: Summary of E-VRP research papers

2.2. The Electric-Vehicle Scheduling Problem (E-VSP)

The traditional Vehicle Scheduling Problem (VSP) is a well known problem. The problem consists of designing an optimal schedule for a fleet of vehicles such that a set of pre-determined service trips is performed at a minimal cost. The first papers found in literature that address the scheduling of electric vehicles are [28] and [29]. After that, this field of research has been extremely active.

[28] and [29] propose an Electric Vehicle (EV) scheduling model using either battery charging or battery swap and a vehicle scheduling model with a maximum distance constraint. Both models

are NP-hard. A column-generation algorithm was developed to solve the instances. The first model is referred to as EV scheduling with battery renewal and is represented on a graph as shown in Figure 2.1. This may imply charging or battery swap.

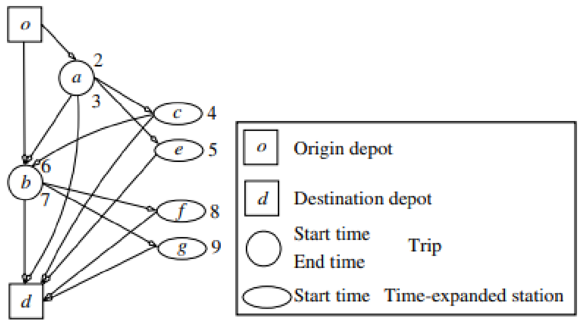

Figure 2.1: Graph representation of the model for Electric Vehicle scheduling with battery renewal [28]

The graph consists of an origin and destination depot, the set of service trip nodes S and set of time-expanded station nodes T. Each service trip is characterised by a fixed start and end time. Each station is represented by a number of nodes called time-expanded station nodes. Because vehicles may arrive at the station nodes at any time, they are time-discretised and a node exists for each time step, The set of arcs A define the possible routes the busses can take. If a route is infeasible, for example, trip a ends at t=5 and trip b starts at t=4, no arc exists connecting the two. A maximum waiting time is used to limit the number of arcs to the station nodes.

The objective is to minimize the operational costs while executing each service trip. The costs are a combination of distance related costs, cost for waiting time, service costs at battery stations and a daily maintenance cost collected upon arrival at the destination depot. All costs can be assigned to the arcs and are calculated beforehand. For convenience, the distance of a service trip is included in the distance of the arc going to that service trip. For example, if the distance from node A to node B is 3 and the service trip at B has a distance of 1, then the arc from A to B has a distance of 4. This has no effect on the outcome of the minimisation because each service trip must be performed exactly once. The objective function is given in Equation 2.8. The binary decision variable x_{ij} is 1 if a bus is assigned to node j directly after visiting node i.

$$min \sum_{(i,j)\in A} c_{i,j} x_{i,j} \tag{2.8}$$

Constraints 2.9 computes the distance that a vehicle has already accumulated at node j. This is the combination of the distance accumulated at note i and the distance from i to j. Constraint 2.10 resets the accumulated distance to zero at the start depot or at a service station. Lastly, Constraint 2.11 ensures that the accumulated distance plus the distance to a service station does not exceed the vehicle's range.

$$g_j = \sum_{j\in A} (g_i + d_{ij}) x_{ij} \qquad \forall j \in S \tag{2.9}$$

$$g_i = 0 \qquad \forall i \in T \cup o \tag{2.10}$$

$$(g_i + d_{it}) x_{it} \leq D \qquad \forall t \in T \cup d, (i,t) \in A \tag{2.11}$$

Using a fixed service time at the service stations is a good representation of battery swap or refueling. Battery charging on the other hand can be variable in duration based on how much energy is left in the batteries and how much the vehicle should be recharged. A conservative estimation would be to use the time to fully charge a completely depleted battery. This yields a sub-optimal solution because too much time would be spent at the service stations. The second model discussed in [28] is the same but without the service nodes. That means that there is no possibility of recharging or swapping a battery along the way.

The optimisation is solved using a Branch-and-Price for small scale instances and a column-generation based heuristic algorithm for large instances. Branch-and-Price works by applying Branch-and-Bound to the original problem and solving the problem at each node in the branching tree using Column-Generation. Branch-and-Bound starts by fixing one binary variable. This adds two nodes in the branching tree, where the value of the variable is either 1 or 0. For each node in the branching tree, the linear relaxation of the path-based formulation of the problem is solved using Column-Generation. If the binary variables in the relaxed solution are still binary, the solution is feasible. If this is not the case, the relaxed solution gives a lower bound. If the lower bound of a branch is higher then the best known feasible solution, that branch is fathomed and no longer considered. In the path based formulation, a route is defined as a sequence of service trips and charging events. To problem consists in selecting an optimal set of routes such that each service trip is served at minimal cost. Because the number of possible routes grows exponentially, the Restricted Master Problem (RMP) starts with a small subset of all possible routes. Once an optimal set is selected, the pricing problem is solved to identify arcs with a negative reduced cost. When the reduced costs of the arcs is known, a dynamic programming algorithm is able to extract the routes with the minimal reduced cost. These are then added to the RMP and the problem is solved again until no columns (routes) with a negative reduced cost remain. The optimal solution of the relaxed problem is then found and the branching and bounding resumes. The Branch-and-Price method works in case of a fixed service time or full recharge. If partial recharge is allowed, the number of routes is infinite and column generation is no longer possible. One way to get around this is to use discretised partial charging but also that drastically increases the number of possible routes for column generation.

The author of [60] defines 2 different models. Though using a slightly different formulation, the first model is equivalent to the ones from [28] and [29]. The cost minimizing objective function consists of a fixed cost per vehicle and a cost per arc. The fixed cost per vehicle was previously described as a daily maintenance cost. This model also assumes linear charging. The difference in formulation is that no additional nodes exist for the chargers or battery swap stations. Instead, the waiting time, service time and change in battery level are included in the arc. This makes the model simpler but does not allow to constrain capacity of the service stations. The second model takes into account nonlinear charging and a variable electricity price throughout the day. To include the variable price, an arc variable is added that says how much energy was charged on that arc. The electricity costs are no longer based on the electricity consumed per arc but on the electricity charged per arc. This way, the cost can vary throughout the day. To model the nonlinear charging behaviour, a set of nodes is created for each trip that represent the combination of a trip and SoC level as depicted in Figure 2.2. A benefit of this method is that it also allows for partial charging but only in discrete steps. The discretisation causes the model size to grow significantly, especially as the discretisation step decreases. Only one of the trip nodes must be visited so the constraint forcing each trip to be executed must change accordingly. Two column-generation algorithms were proposed to solve real life instances.

Interestingly, both models produced very comparable results. Therefore, linear charging can be the model of choice if uncertainties are large. In that case, it might not make sense to try to present a really detailed model because other uncertainties are larger while linear charging is able to capture the general behaviour.

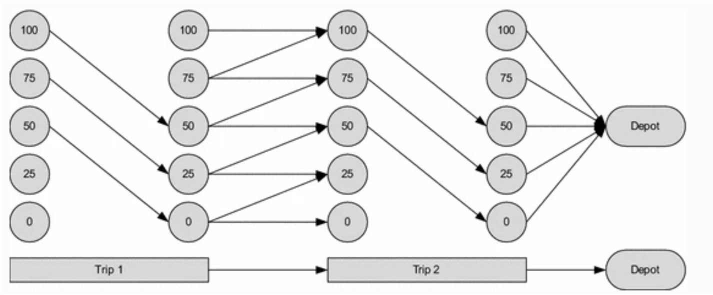

Figure 2.2: Node set for a discretized SoC per trip

[27] attempts to integrate the planning of charger infrastructure and scheduling of electric busses to achieve a global optimum. This is particularly useful for operators where no charging infrastructure is present and must therefore also be provided by the operator. The model is applied to an existing bus network in Davis, California. To incorporate the infrastructure planning, a number of candidate stations were identified beforehand. Each candidate station has a finite number of candidate chargers. The binary variables Z_{nk} assumes a value of 1 if the k'th charger at station n is used and 0 if it is not used. The binary variable Z'_n indicates if a station is used and equals 1 if any of its chargers are used. Constraint 2.12 ensures the correct value for each candidate station n, where M is a sufficiently large positive number. If there is any nonzero Y_{it}^{nk}, a bus coming from trip i assigned to charger k of station n, than Z_{nk} must be one. Similarly, Constraint 2.13 ensures that Z'_n assumes a value of 1 if any of its chargers are used.

$$\sum Y_{it}^{nk} \leq MZ_{nk} \qquad \forall k \in K, n \in N \tag{2.12}$$

$$\sum_{k}^{K_n} Z_{nk} \leq MZ'_n \qquad \forall n \in N \tag{2.13}$$

The objective function becomes the minimization of the annual electric bus recharging system operating costs. This now also includes installation and maintenance costs per charging station and per charger on top of the vehicle and electricity costs as discussed in previous papers. [58] also considers the total cost of ownership to combine scheduling and infrastructure planning. They take it one step further by considering two distinct bus types.

Intercity bus transport is different from urban transit bus transport. In urban transit, passengers usually pick a line and busses are scheduled with a high frequency. For intercity busses, passengers choose a specific scheduled service beforehand based on a number of factors including time-of-day, trip duration and service level. This way, intercity bus operations closely resemble scheduled commuter airline operations. [34] recognizes the interaction between the schedule and demand. His model determines which service trips to make, starting from a dynamic demand pattern on a single bus corridor for a non-electric bus. A corridor is a sequence of stations which a bus may choose to serve. D_{ijkl} represents the demand between stations i and j for a bus leaving in time-window k with trip duration l. The profit is maximised by choosing which stations to include in the bus line. The profits consist of revenues minus fixed and variable costs. The novelty of this paper is using dynamic demand as an input. Whereas previous scheduling papers used a predefined set of service trips, determining which trips to make is decided based on the input. This makes the model very useful in a context where EVs are offering a new service instead of replacing an existing service. In case of a new service, determining which service trips to make should be part of the problem. The instances are solved using a Branch-and-Cut algorithm.

[55] is the first and so far only research paper to apply Deficit Function (DF) theory to solve the E-VSP for battery electric transit vehicles. The DF and related Surplus Function (SF) are determined based on a given set of service trips. It is assumed that all service trips start and end at a terminal and chargers can be installed at the terminals. From DF theory, a lower bound on the required number of vehicles and an upper bound on the number of chargers at each terminal can be determined. With some additional considerations, a Pareto-front is constructed relating the required number of vehicles with the accompanying required number of chargers.

In real life applications, it is possible that multiple vehicle types are used. [56] Attempts to solve the E-VSP for multiple vehicle types. The objective is to minimize the total annual scheduling costs. To solve the Multi-Deport VSP with different vehicle types, [22] models the vehicle flow on a Time-Space-Energy (TSE) network and passenger flow on a Time-Space (TS) network.

The only paper that considers the E-VSP for aircraft is [41]. She optimises operational costs, which are a combination of aircraft and battery acquisition costs and charging costs. The aircraft under consideration is a conceptual electric regional aircraft. The parameters are based on a Class I estimation using the Embraer E175 as a starting point. Both battery swap and fast charging are allowed. In a first stage, a fixed cost per battery swap is assumed to optimise the scheduling. In a second stage, the number of batteries for battery swapping and scheduling the charging of swappable batteries is optimised. The author found that during off-peak hours, fast charging is the preferred option because this is cheaper than battery swaps. During peak hours, battery swaps are preferred because they prevent the fleet size to grow to satisfy the peak demand, resulting in a lower overall cost.

2.2.1. Robust Scheduling

Whereas the previously discussed papers assume a deterministic network, in reality trip times and energy consumption are affected by road and traffic conditions. To increase the robustness of the schedule, [62] proposes two strategies. Both models optimize the fleet size and schedule. The first one is a static model that includes a buffer distance such that the busses are unlikely to end up stranded with a depleted battery. This is fairly straight forward. The range Constraint 2.11 is adjusted with parameter λ to keep keep a buffer in the driving range at all times, as seen in Equation 2.14.

$$(g_i + d_{it})x_{it} \le D/\lambda \qquad \forall t \in T \cup d, (i,t) \in A \tag{2.14}$$

The static model is capable of increasing robustness of the schedule but is unable to make use of any real-time information that might be available. The second method is a dynamic method that allows to periodically reschedule based on real-time the traffic information. The day is divided in a number of time periods. At the start of each period, the fleet is rescheduled. The situation and schedule for the rest of the day are known at that moment. The schedule is represented by arc set A_1. Two alternative arc sets A_2 and A_3 are possible but not executed. Based on the current traffic information, the likelihood of en-route breakdowns of the schedule are estimated. A penalty factor is applied to en-route breakdowns in the objective function. The problem is now to select the optimal arc set for the remainder of the day to minimize the costs, thus preferring routes with a low likelihood of breakdowns.

[33] considers a bus operators who wants to transition to an electric fleet while preserving the schedule. The objective is to minimize the total costs, consisting of charging station costs, bus costs, average costs for the relocation of buses from a terminal to a charging station, charging costs and a penalty incurred in case of a charging demand overflow. The model is formulated as an integer stochastic program and is solved using a Lagrangian relaxation and a branch-and-bound algo-

rithm,

Note that these are applied to a bus schedule. Also in air traffic, trip times are not deterministic. Wind may affect the estimated flight time. Delays due to weather or ground handling activities may propagate through the schedule. Robust scheduling is a very active field in airline operations research. [16] proposed a stochastic optimisation model. [35] attempts to minimize the probability of a passenger missing a connection in a hub-and-spoke network. [51] applies fuzzy logic while [21] creates a scenario tree to have multiple plans ready in case of deviations to the original plan. [1] integrate aircraft routing and scheduling with an eye on robustness. These papers are a small selection of the wide pool of robust airline scheduling research to illustrate that this is an active field of research and that there are many ways to approach robustness.

2.2.2. Fleet Transition

It may happen that an existing operator intends to partially or fully electrify its fleet. A target is then set at some point in the future. The author of [30] introduced the Electric Bus Fleet Transition Problem. The objective is to minimize the costs during the planning horizon such that the electrification target is met in time. Costs include the necessary acquisition of electric busses and charging infrastructure.

2.2.3. Infrastructure Planning

A number of papers include infrastructure costs in their objective function. By doing so, they are able to make a decision on which candidate charging stations to include. While this is an elegant modelling framework, this may not always be realistic because this method does not take interactions with the power grid into account. [6] developed a method to determine an infrastructure planning taking both the electrical grid and transportation system into account. To account for a gradual deployment of electrical busses, the planning is multi-staged and was found to be robust to changes. Though it is good to be aware of this, this is beyond the scope of this literature study.

2.2.4. Final Remarks

The E-VSP has been discussed in a number of different variations as shown in Table 2.2. Again C, PR and FR denote constant time, partial recharge and full recharge respectively. Though the problem is NP-Hard, exact methods based on Column-Generation algorithms often seem to do the job. The problem has been applied in a number of case studies of electric bus fleets. One downside to this approach is that the set of service trips must be determined beforehand. [34] Presents a variation with a dynamic demand pattern as input such that determining which service trips to make becomes part of the problem. Further work is this field will be studying different extensions and variations to the problem that are useful for specific case studies.

Paper	Vehicle	Objective	Service	Infr.	#Trips	# Vehicles	Exact
[28]	Bus	Operational Costs	C	-	1,000	60	+
[29]	Bus	Operational Costs	C	-	4,300	500	+
[60]	Bus	Operational Costs	PR	-	500	30	+
[27]	Bus	Total Annual Costs	C	+	600	30	+
[22]	Bus	Total Annual Costs	FR	+	300	35	+
[34]	(Diesel) Bus	Profit	N/A	N/A	18	1	+
[55]	Bus	# vehicles & batteries	PR	+	700	35	-
[62]	Bus	Operational Costs	C	-	60	13	-
[33]	Bus	Total Annual Costs	PR	+	?	700	+
[41]	Regional Aircraft	Operational Costs	PR	-	50	25	+

Table 2.2: Summary of E-VSP research papers

3 Electric Flight

Electric technology is developing at an extremely high pace. This has lead to a number of common misconceptions about electric propulsion that are preventing research from reaching its full potential [17]. Therefore, this section will discuss electric aircraft technology and operations in more detail with a focus on thin-haul commuter aircraft. The terms commuter and thin-haul are often used interchangeably in terms of air transport. Thin-haul refers to one step smaller than short-haul. It involves from the General Aviation market with seating for 6 to 10 passengers. Flight legs vary in distance from ten to a couple hundred kilometers. Section 3.1 will discuss the electric aircraft and power consumption model, Section 3.2 discusses battery technologies, Section 3.3 discusses considerations about thin-haul operations and related costs and lastly, Section 3.4 briefly touches on the emissions attached to energy consumption.

3.1. Aircraft Model

Electric aircraft have been around for a while. Already way back in 1940, Fred Militky attempted to build an electric model aircraft. Only with the introduction of Nickel-Cadmium (Ni-Cd) batteries and more efficient motors, he was able to commercially produce the first Radio-Controlled (RC) aircraft in 1972. One year later, the first manned electric aircraft took to the skies. The MB-E1 was able to fly for 15 minutes, more than enough for a proof of concept [37]. Especially considering that the Wright brothers only flew for 12 seconds on their first flight and are still remembered for it.

Developments in battery and electric motor technology enable significantly improved flight performance. Modern Lithium-ion (Li-ion) batteries achieve a battery density of up to 400Wh/kg, over a 10-fold improvement compared to the Ni-Cd batteries from the 70s with 30Wh/kg but still roughly 60 times lower than Kerosene. The author of [37] compares the efficiency through the entire energy chain, from energy source to propulsion. They find that the overall efficiency of current electric propulsion is 73%. That's almost double the efficiency of a Kerosene powered propeller-driven aircraft with 39%. From this direct comparison between electric propulsion and fossil fuel powered aircraft and considering the requirements of regional aircraft, [37] concludes that battery energy density must increase to at least 1000kWh/kg but preferably even to 2,000kWh/kg to be commercially interesting.

This is a misconception according to [17]. There are two flaws in this idea. Firstly, the battery density requirement follows from a direct comparison between electric and fossil fuel based propulsion while electric aircraft are fundamentally different. Because electric motors are scalable, that is a 10W motor has approximately the same efficiency and power-to-weight as a 1,000W motor, there is a lot more design freedom that allows for aerodynamic coupling between the propulsion, control surfaces and wing. One very promising concept is Distributed Electric Propulsion (DEP) with a large number of small electric motors over the entire span of the wing. This would allow to eliminate the need for high-lift devices because of artificially created airflow over the wing and allow for a smaller wing, designed for cruise. This concept is being explored by NASA with the X-57 prototype. The second flaw is that research and development should focus on regional and large airlines. This ignores the development path that may lead to this. Their assumption is that as technologies develop, the research and market plan should be adaptable. While it is tempting to focus on airlines because that's where the bulk of revenues are, there is a poor fit with current technology and the evolution that it will take cannot be ignored. The first applications will be found in General Aviation. Some motorgliders and a number of ultralights have already adopted electric propulsion. One step larger and the aircraft become attractive for commuter airlines. Only after that can one realistically start planning the introduction of an electric regional airliner.

There are three ways to obtain input parameters to an aircraft model for an optimisation problem. Firstly, a number of studies have performed a conceptual design and sizing of electric aircraft. For example, the author of [36] performed a conceptual design of a 4-seater electric-combustion hybrid. No research was found considering an all-electric aircraft for 6 to 10 passengers, the range for a commuter aircraft. Therefore, a second option is to perform the conceptual sizing as part of the thesis. This method is far beyond the scope of the thesis. The last method is to critically evaluate the performance figures from real aircraft and/or claims from aircraft in development. The next section will therefore summarise the current state in the development of electric aircraft.

3.1.1. Current State of Aircraft Development

An overview of the current state of relevant aircraft in development is provided in Table 3.1. This overview presents the electric aircraft that appear to have a high potential for thin-haul operations or are technologically relevant.

Aircraft	Seats	Range [km]	Cruise Speed[kph]	VTOL
MangiX eCaravan	10	240	200	No
Skylax E6	6	300	300	No
Skylax E10	10	300	300	No
Eviation Alice	10	1,000	440	No
X-57 Maxwell (Mod III)	2	160	277	No
Lilium Jet	5	300	300	Yes
Jaunt Journey	4	250	280	Yes
Joby S4	4	280	370	Yes

Table 3.1: Summary of promising aircraft for thin-haul operations

Electrical equipment company MangiX transformed an existing Cessna Caravan to a fully electric aircraft. Their expertise is in electrical equipment, not aircraft design. Therefore they opted to start from an existing airframe. The Cessna Caravan has been around for almost 40 years and with over 2,600 units produced, it is one of the most used airframes in the world. The seating capacity for 10 people makes it ideal for thin-haul operations. MangiX started flight testing in June 2020 and plans to put the aircraft in service as early as 2021. The fact that they are this far in development makes the performance claims thrust worthy. Starting from an existing airframe is the cheapest and fastest way to put an electric aircraft in operations. As [17] discussed, they are not able to exploit the additional design freedom of electric technology.

The Skylax E6, Skylax E10 an Eviation Alice are starting from scratch. As a result, their designs are more suited for electric propulsion and have an improved performance. Skylax is planning an initial version with a 300km range and intends to double that range by 2030. It must be noted that Skylax does not have a flying prototype yet so their claims must be interpreted with a grain of salt. However, their performance seems realistic compared to the eCaravan which has an older airframe that was not designed for electric propulsion. While electric technology gives more design freedom for exotic concepts such as DEP, Skylax still uses the concept of podded engines as traditional aircraft do. In doing so, they also do not fully exploit the benefits of electric technology yet. The reason is likely to be a rapid market introduction, currently planned in 2025. This follows the logic laid out by [17] that the market plan must be agile. Eviation Alice uses two wingtip propellers that fold back during cruise. Cruise power is provided by a single pusher propeller in the back. Though full scale models exist, no flight testing has been performed so far. Eviation Alice has a big lead in performance over Skylax but that comes at a cost. The Alice is estimated at €4 million compared to €1.9 million for the Skylax E10. In fact, the Skylax' list price comes very close the the fuel powered Cessna Caravan at €1.8 million.

Meanwhile NASA is developing the next X-plane, the X-57 Maxwell. The aircraft has a much lower performance than the previously discussed ones. That is because this aircraft has a different purpose. It is a research platform and flight demonstrator for new technologies and has no commercial intentions. A series of modifications or mods is planned for flight testing. In Mod II, a regular Tecnam P2006T is electrified, resulting in a configuration that is similar to the Skylax aircraft. Mod III and IV will test a wing 60% smaller than the original and DEP.

The last three aircraft in the list are VTOLs and have a lower seating capacity. That is because these aircraft are intended for UAM. A number of hurdles such as regulations about noise, safety and privacy must be overcome before UAM can become a reality. While this research will focus on electric thin-haul operations because that might very well become a reality in a few years, it is interesting to consider the performance achieved by these aircraft. Of these, only Lilium is already in the flight testing phase. An array of electric fans tilt downwards for VTOL. In cruise mode, the electric fans provide DEP on top of the wing.

One may observe that both the range and cruise speed of these aircraft are spread around 300km and 300kph. The Eviation Alice claims a much higher performance but is the only one to do so. Meanwhile, the low range of the X-57 can be attributed to the non-commercial purpose of this aircraft. A realistic aircraft model for research that intends to be practically relevant could therefore be assumed to have a range of 300km, a cruise speed of 300kph and 10 seats.

There are two very important side notes. Firstly, to critically assess whether these claims are realistic, the numbers must be validated. This does not have to be an extensive conceptual design but the parameters should make sense when subjected to basic aircraft performance calculations. Secondly, these parameters are by no means the only realistic parameters. Once the aircraft presented in Table 3.1 will be operational, their actual performance will be different from the current estimations. Over time, technology will improve and performance will increase. To be practically relevant, the thesis must address the uncertainty with extensive sensitivity analysis.

[41] proposes Equation 3.1, an adaptation for electric aircraft of the range equation proposed by [37]. The difference from the original range equation is that the mass of an electric aircraft stays constant throughout the flight.

$$R = E^* \cdot \eta_{total} \cdot \frac{1}{g} \cdot \frac{L}{D} \cdot \frac{m_{battery}}{m} \tag{3.1}$$

Where E^* is the energy density of the battery, η_{total} is the total system efficiency, g the gravitational acceleration, $\frac{L}{D}$ the lift-drag ratio in cruise and $\frac{m_{battery}}{m}$ the ratio of battery mass over total mass. From this equation, it is clear that range depends on total mass and thus, occupation of the aircraft. A static range is therefore not realistic when the number of passengers is varying. The ranges discussed previously refer to the range at Maximum Take-Off Weight (MTOW).

3.1.2. Phases of Flight

How much power is required to fly a given flight leg depends on the distance. Because the aircraft goes through a number of flight phases which each have different power requirements, a simple linear relation between range and energy consumed might be oversimplified. Figure 3.1 shows the different flight phases and related power usage for a conceptual electrified Tecnam P2012T Traveller as investigated by [53].

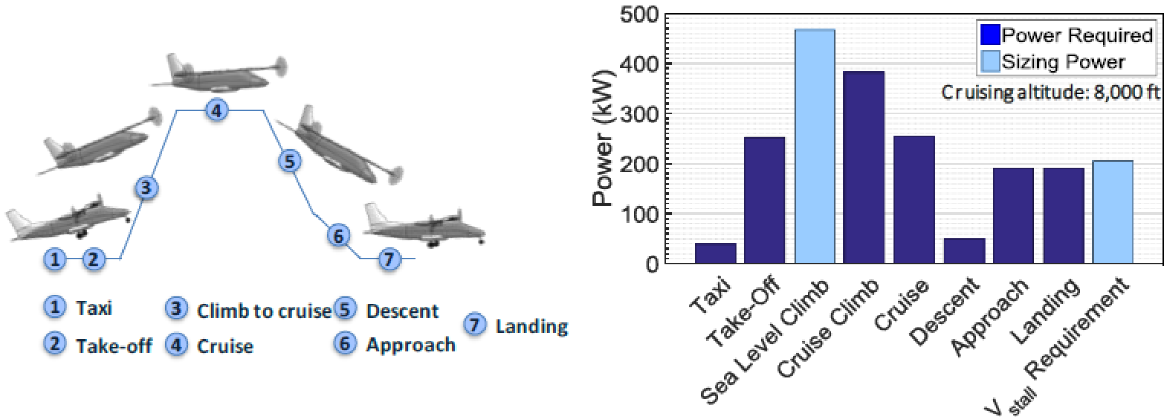

Figure 3.1: Power Consumption for the different flight phases of a conceptual electric Tecnam P2012T Traveller [53]

If one assumes that for all flight legs, all phases are identical in duration and power except for cruise, then the power required for a flight leg can be calculated using Equation 3.2. Where c_0 is the energy required for all fixed flight phases, c_1 the energy required for a unit distance in cruise, d_0 the distance already covered in climb and descent and r a routing factor to take into account that the aircraft might not be able to fly the shortest route. [44] uses a routing factor of 0.42 while [53] uses a 0.06 in the network of Cape Air and 0.28 for Mokulele Airlines, two commuter airlines. This model is simple and flexible. If for example it is known that at a given airport, the aircraft will have to taxi a bit longer, than c_0 can be adjusted accordingly. If the aircraft has to make a large detour to fly from A to B, then r can be increased for that particular flight leg.

$$E = c_0 + c_1 \cdot r \cdot (d - d_0) \tag{3.2}$$

3.2. Battery Model

The battery is an essential part of the electric aircraft. It must be able to store enough energy and deliver it at high power to the motors under all conditions. Battery performance is subject to degradation and changes in temperature. Therefore, batteries are designed to operate under worst case conditions. The following sections will discuss the battery types available on the market and in development, the definition of SoC, the battery charging process and considerations for battery swap.

3.2.1. Battery Types

A vast array of battery types is available nowadays and choosing the right one is crucial. The most important factors for an electric aircraft are i) Energy density in Wh/kg; ii) Price in €/kWh and iii) Other characteristics such as degradation, etc.

Promising advancements have been made. Lithium-air batteries reach a theoretical energy density close to that of gasoline but issues with stability and air purity render this battery type impractical for real applications. Lithium-ion batteries are currently the most widely used type. Lithium-Sulfur is close to commercialisation but suffers from high degradation and thus a short battery life. Table 3.2 provides an overview of battery types. [57]

The most widely used battery today is the Li-ion battery. This battery is completely maintenance free and has a very low self-discharge rate. Its high performance make this the battery of choice in

Type	Specific energy [Wh/kg]	Price [€/kWh]	Cycle-life
Li-ion	250	150	1200
Lead-acid	40	50	350
Li-metal	300	-	2,500
Li-Sulfur	600	125	50
Ni-Cd	80	100	1,000
Ni-MH	120	100	500
Li-air	13,000	-	50

Table 3.2: Summary of current and promising battery types

cellphones, where size is an important factor and in electric vehicles, where weight is important. Research is showing very promising results for Li-metal batteries. These have a higher specific energy and longer cycle life than Li-ion batteries. However, uncontrolled Lithium depositions can penetrate the separator and cause a short circuit. Safety is therefore the main hurdle to be overcome before Li-metal batteries can be widely adopted. Lithium-Sulfur (Li-S) batteries are also showing promising results. To be commercially interesting, the cycle-life has to be improved. Lastly, the Lithium-air (Li-air) battery has a mind-boggling theoretical specific energy, similar to that of gasoline. However, it is currently not possible to create a Li-air battery outside the lab. One issue is air purity. Because the air in cities and airports is not clean, filters and compressors would be required, consuming a large portion of the energy stored in the battery. Even if the battery had an efficiency of just 25%, it would still be competitive with gasoline. A combustion engine has a thermal efficiency of roughly 25% while electric motors are over 90% efficient. Another issue is the sudden-death syndrome. The reaction between Lithium and Oxygen creates a Lithium peroxide film, blocking electron movements and killing the battery. It is currently not possible to build a practically usable Li-air battery. [57]

The Lead-acid battery is the battery of choice for many industrial applications where weight is not a factor because of its low cost. They are not sensitive to abuse in the form of deep cycle discharging and overcharging. Because they have a shorter cycle-life, the price per cycle is similar to that of Li-ion batteries. Still, the up-front investment is significantly lower and if a consumer has to replace the lead-acid battery, he or she can benefit from technology improvements and/or reduced battery prices over time. This battery requires some maintenance. [57]

Lastly, the Ni-Cd and Nickel-Metal Hybrid (Ni-MH) batteries are well understood and relatively low in maintenance. The low performance make this battery useful for applications where weight is not important but a relatively high cycle-life is preferred. [57]

3.2.2. State-of-Charge

Batteries degrade over time. One factor that influences degradation is the Depth-of-Discharge (DoD). A Battery Management system (BMS) is required to keep the battery in the right operating window to minimize degradation [41]. The BMS measures the SoC, the capacity remaining in the battery q over the nominal capacity Q_n as seen if Equation 3.3. The DoD is the complement of the SoC.

$$SoC = 1 - DoD = \frac{Q_n - Q_b}{Q_n} = \frac{q}{Q_n} \tag{3.3}$$

[41] found that there is no common consensus in literature over a definition for the SoC. That is because degradation causes the available capacity to reduce and because the BMS limits the SoC range, reducing the available capacity further. [43] uses three formulations, the SoC, relative SoC and practical SoC, The motivation behind these formulations is that batteries degrade over time. [7]

showed a reduction in capacity of 10% to 40% over 450 charge cycles of Li-ion batteries, the most widely used battery type in the world. Over time the measured capacity Q_m reduces. Degradation depends on many variables. Most notably is the SoC during each cycle. To limit battery degradation, batteries are not charged and discharged to their limits, resulting in the practical capacity Q_p and practical SoC_p.

$$SoC_r = \frac{Q_m - Q_b}{Q_m} \tag{3.4}$$

$$SoC_p = \frac{Q_p - Q_b}{Q_p} \tag{3.5}$$

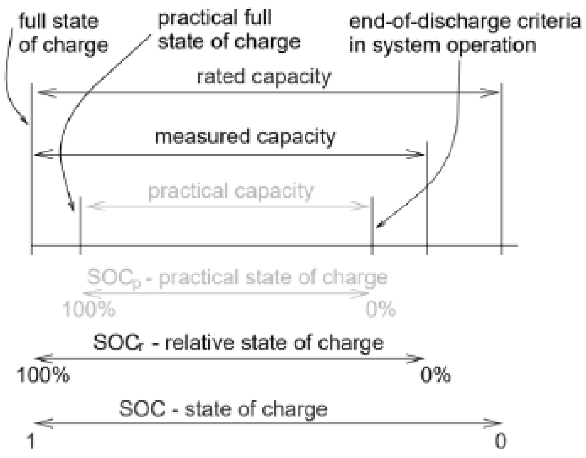

Figure 3.2: Assessment of geographical regions to enter the MRO market

The State-of-Health SoH is the ration between Q_m and Q_n and gives an indication of the degradation of the battery. [48] proposed a method to determine the SoH under temperature and cycle aging but this was too complex for use in practise.

The previous equations relate to a single battery cell. For applications requiring a larger capacity, the battery is composed of a group of individual cells placed either in series or parallel. A battery constructed of cells in parallel can be considered as a single large cell and Equations 3.3 through 3.5 apply. For placement in series, Equation 3.6 is to be used [26] where N is the total number of cells.

$$SoC = \frac{\sum_i^N SoC_i Q_{n,i}}{\sum_i^N Q_{n,i}} \tag{3.6}$$

If all individual cells are identical, this equation simplifies to Equation 3.7

$$SoC = \frac{\sum_i^N SoC_i}{N} \tag{3.7}$$

Due to the chemical nature of batteries, an accurate determination of the SoC is difficult and remains a highly theoretical exercise. A number of methods are proposed in literature. [42] summarised the most common methods and discuss which use cases they are most suitable for. The most reliable test is the discharge test. The major drawback is that it interrupts the battery usage and is time consuming. Ampere-Hour counting is the most common technique. Starting from a known SoC, the current from the battery is measured and integrated over time to determine the SoC at that moment. This method has two issues. Firstly, accurate current measuring equipment is expensive. Secondly, losses in the system could account for substantial errors. Re-calibration is required to keep the errors from growing. The most simple SoC determination method is by applying a constant charge/discharge factor. Sometimes, knowing the chemical properties of the reactants in the battery suffice to determine the SoC. For example, the acid density in a lead-acid battery is directly related to the SoC. This is not possible for all types of batteries. The open-circuit voltage is also related to the SoC. However, long waiting times are required to reach a steady-state, making the open-circuit voltage an impractical parameter to use.

3.2.3. Battery Charging & Battery Swapping

At a certain moment, the battery reaches a point where the energy level is insufficient to complete the next flight leg. At that moment, the battery must either be charged to a sufficient level or swapped with a fully charged battery.

Chargers can be classified by the time it takes to fully charge. The Electric Power Research Institute (EPRI) and the Society of Automotive Engineers (SAE) define three levels of charging [41].

- **Alternating Current (AC) Level 1** is characterised by a maximum power of up to 1.9kW in the US and 7.4 kW in Europe. This type of charging is convenient for overnight charging when the battery is not used.
- **AC Level 2** is characterised by a maximum power of 19.2kW in the US and 43 kW in Europe. A level 2 charger has a price tag between €2,000 and €15,000.
- **Direct Current (DC) Level 3**, also known as Fast-Charging, is the most expensive type with prices ranging from €50,000 to €150,000. Level 3 includes all charging levels above level 2.

[57] calls these levels slow, rapid and fast charging respectively and adds the level of Ultra-Fast Charging. An Ultra-Fast charger is able to charge a battery up to 70% in 10 to 60 minutes. Note that the battery should be specifically designed for this and will feature a reduced specific energy. Even if the battery is designed for Ultra-Fast Charging, this will expedite degradation. Fast chargers are naturally more expensive. Besides more expensive infrastructure, the peak power drawn from the grid is higher and this results in increased electricity prices. A famous example of a fast charger is the Tesla Supercharger. The V1 and V2 Superchargers have a power output of 150Kw. For the V3, this is increased to 250kW. Tesla is even planning the introduction of the Megacharger, intended for their electric truck, the Semi. The Megacharger is reported to have an output level of over 1 MW. It must be noted that charging time cannot simply be decreased with more power as battery and charger design go hand-in-hand. The increased charger power is serving larger batteries. With the current technology, Tesla is able to recharge batteries to 80% in 30 minutes and 100% in 75 minutes.

To model the battery process, most papers in operations research assume linear charging and discharging behaviour as discussed in Section 2. That is, the SoC varies linearly over time. A number of papers use a nonlinear relation or approximate the nonlinear behaviour by assuming a multi-linear graph. The added accuracy comes at the cost of computational complexity.

If a linear charging process is assumed, the charging time can be determined from Equation 3.8 [14]. The main benefit of this equation is its simplicity. According to [60], though it is not a perfect representation of reality, linear charging is able to capture the general behaviour of the system. Still, a better approximation is highly desired in a scheduling problem because of the large difference in charging time.

$$t_{charge} = \frac{Q_n \Delta SoC}{P_{charger}} \tag{3.8}$$

Linear charging is an approximation of the actual charging process as can be seen in Figure 3.3a. A common argument for accepting linear behaviour is that the range of batteries is never fully utilized because the rate of degradation increases when the DoD is pushed to the limits. That line of thought is not entirely correct. Batteries are never fully discharged. Fully charging is common practice and the difference in charge rate is large when approaching a full charge. For Tesla batteries, the last 20% charges approximately 6 times slower than the first 80%. Below a SoC of 80%, linear charging is a very good approximation. With increasing SoC, the charge acceptance decreases and the process starts slowing down. [41] used a bi-linear approximation of the charging process as shown in Figure 3.3b to take this reduction in speed into account. This bi-linear process is an elegant solution because it is already much better approximation of reality without making the model overly complex.

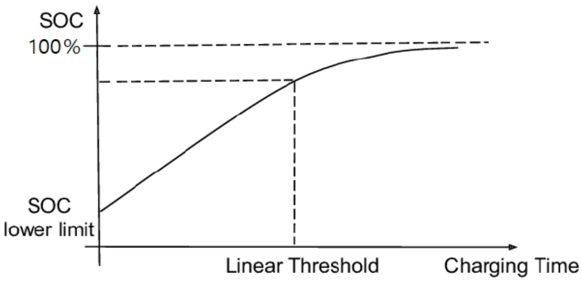

(a) The actual charging process is nonlinear once the SoC surpasses 80%

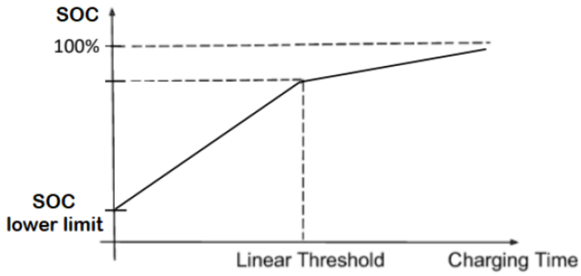

(b) Bi-linear charging process approximation

To model the bi-linear charging, the author of [41] introduces additional variables a, b and c as defined in Equations 3.9 through 3.11. In these equations, q_R is the energy level required to make the next flight and SF the safety factor to have reserve capacity in the battery. The energy level present in the batteries before starting the charging process is q_E. The charging time is then computed using Equation 3.12.

$$a = \begin{cases} 1, \text{ if } q_R + SF \geq 0.9Q \\ 0, \text{ otherwise} \end{cases} \tag{3.9}$$

$$b = \begin{cases} 1, \text{ if } a = 1 \text{ and } q_E \geq 0.9Q \\ 0, \text{ otherwise} \end{cases} \tag{3.10}$$

$$c = \begin{cases} 1, \text{ if } a = 1 \text{ and } b = 0 \\ 0, \text{ otherwise} \end{cases} \tag{3.11}$$

$$t_c = \begin{cases} \frac{q_c}{p_c} \text{ if } a = 0 \\ \frac{q_c}{p_c/10} \text{ if } b = 1 \\ \frac{0.9Q - q_E}{p_c} + \frac{q_R + SF - 0.9Q}{p_c/10} \text{ if } c = 1 \end{cases} \tag{3.12}$$

The operational costs of charging depend on the costs of the charger infrastructure and cost of electricity. [53] breaks the electricity price down into supply, transmission and delivery. Though not visible on your electricity bills, this is commonly performed by different companies. Usually, there is a fixed monthly charge for being connected to the grid besides the charge for the amount of electricity consumed. The electricity price itself depends on the location, time of day and peak power draw. In [53], the author discusses an optimisation method for charging batteries in a battery swap station. This method takes into account that lower charging power and charging in off-peak hours results in lower electricity prices. [31] only looked at price variations throughout the day. From a cost-benefit analysis, they found that expanding their fleet from the minimum of 59 vehicles to 62 vehicles results in a lower total cost because charging can be scheduled in the off-peak hours. In April 2020, the average Dutch household paid €0.1707 cents per kWh of electricity. A Tesla Supercharger charges €0.28 per kWh. The increased price covers the infrastructure costs as well as increased electricity price due to the high power levels.

Now matter how advanced chargers become, they will never be as quick as refueling. One alternative method is to swap the depleted battery entirely with a fully charged one. This may lead to a reduced Turn Around Time (TAT). According to [41], the battery swap on the Siemens eFusion takes 5 minutes. Typical battery swaps for EVs take between 5 and 15 minutes according to [63] while Tesla's battery swap stations needed just 1.5 minute [45] before the project was scrapped.

Charging takes a lot longer but the actual time saved due to battery swap is smaller. That is because scheduled aircraft have a minimum TAT to perform a sequence of ground handling tasks. These tasks may include unloading and loading passengers, cargo and crew, refueling, providing catering, cleaning the cabin and performing line maintenance. Legacy airlines usually have a TAT between 45 and 60 minutes. Low cost airlines like Ryanair achieve TATs of down to 25 minutes. They achieve this by cleaning less and speeding up the boarding of passengers. The TAT for commuter aircraft can be brought down further because the cabin does not have to be cleaned between every short flight and because boarding and unboarding of a small number of passengers goes a lot quicker. Two famous commuter airlines are Mokulele Airlines and Cape Air. They have an median TAT of 19 minutes and 35 minutes respectively. During those 19-35 minutes, the battery can already charge. Only when the charging time exceeds the minimum TAT is battery swap actually faster.

Despite lowering the electricity prices thanks to off-peak charging at lower power levels, the logistics and upfront investment form a large obstacle. Additional batteries must be acquired as well as the battery swap station itself. Battery swap is therefore more expensive than charging. [41] used both charging and battery swap in her optimisation model and found that battery swaps during peak hours reduced the total number of aircraft requires while charging was used during off-peak hours. The battery swap stations bring an additional scheduling problem. Namely how many batteries to store and when to charge them. [53] proposed a method to hierarchically minimize the number of batteries and peak-power demand during charging. With this method, they are able to reduce peak power by over 50% and reduce electricity prices by 20% compared to the heuristic benchmark method which minimizes the number of batteries and charges immediately. [5] studied the number of batteries required for different charging speeds at Munich's airport. In their example, 290 batteries are required if fully charging takes 1 hour. This increases to 499, 1230 and 2318 batteries required for 2, 6 and 12-hour charging times respectively. Clearly, the number of batteries grows rapidly as charging power decreases.

To use battery swap in a scheduling optimization model, one may assume a fixed cost per swap. From the determined schedule, the cost of battery swaps can be validated by minimizing the number of batteries and charging power required to comply with the schedule. If the assumed cost per swap is not close enough, the assumed cost must be changed and the aircraft scheduling optimisation must be re-done. This method does not guarantee to find the overall optimal solution. No literature currently exists that combines the aircraft scheduling and battery swap optimisation to find the global optimum.

A number of hurdles remain for battery swap stations. First is the upfront investment. An operator intending to launch electric aircraft services may not have the capital to build a number of battery swap stations. Alternatively, airports may decide to build those themselves but that arises questions with ownership. Who owns the batteries? What if a brand-new battery is swapped for a heavily degraded one? Homogeneity is preferred as a swap station can only store and charge a finite number of batteries but currently, no standards exist. Logistical issues add to the already complex operations. An operator entering the market with a new fleet of electric aircraft may want to limit complexity because this has never been done before. Lastly, flexibility is limited. Whereas an electric aircraft would be able to fly to any airport that has a charger, this might not be possible for battery swaps if the operator does not have spare batteries stored at an airport. The author of this literature study is of the opinion that battery swapping is only interesting at the home base from where the airline operates. At all other airports, only charging should be considered. Only after the successful introduction of an electric fleet may the operator identify improvement potentials with additional swap stations.

3.3. Thin-haul Operations & Costs

The terms commuter and thin-haul are often used interchangeably for air transport. Thin-haul refers to one step smaller than short-haul. It involves from the General Aviation market with space for up to 10 passengers. Flight legs vary in distance from ten to a couple hundred kilometers. Due to their high cost, they are only found where there is no high speed alternative. Mokelulu Airlines for example provides fast transport between the Hawaiian Islands where the only alternative is to travel by boat.

The author of [59] applies the VRP to an on-demand air taxi network. This is different from scheduled operations because the passengers determine the service trips. The optimisation assigns a routing to each aircraft based on the current state of the fleet and a requested trip. If the problem turns out to be infeasible, the newly requested trip is rejected. The model presented in this paper lies at the basis of Fly Aeolus, a commercial air taxi service operating a fleet of 15 Cirrus SR22 aircraft. Besides 1 pilot, there is space for 3 passengers.

[49] differentiates UAM and thin-haul as summarised in Table 3.3. Note that this is an interpretation of the author. A thin-haul commuter may be powered by fossil fuels and the mission range may be smaller than 150 miles. Nevertheless a number of clear differences can be identified. Thin-haul commuter aircraft offer more seats than UAM and offers mobility over longer distances to travel between cities. An additional difference is that current UAM concepts are on-demand whereas thin-haul operations are scheduled services. A lot of VTOL aircraft are currently in development and those are aiming at UAM.

Thin-Haul Commuter	**Intra-City UAM**
4-9 passengers	1-4 passengers
Fixed wing	tiltrotor or rotorcraft configuration
Standard runway or Short Take-Off and Landing (STOL)	VTOL
(Distributed) electric or hybrid/electric propulsion	Quiet electric or hybrid/electric propulsion
200-300 knot cruise	mission range 5-50 miles
mission range 150-300 miles	enhanced autonomy

Table 3.3: Characteristics of a thin-haul commuter aircraft and intra-city UAM aircraft [49]

The author of [8] build on this envisioned thin-haul commuter concept to assess the economic feasibility of electric thin-haul operations. According to [8] only 10% of all available airports are used by commercial airlines. The demand on each individual route in a thin-haul network may be small but the cumulative demand over the entire network provides opportunities. Consider as well that latent demand may emerge if the service becomes successful and prices can be reduced. Although the demand is arguably there and a large network of airports is available, thin-haul airlines have not experienced the growth rates of commercial transport before the COVID crisis. That is mainly caused by the high operating costs. Whereas traditional airlines such as have an operating cost of $0.13 per Available Seat Kilometer (ASK) (Delta Air Lines) compared to $0.47 for commuter airlines (Cape Air) [8]. The high operating costs and alternative fast transport make it very hard for commuter airlines to be successful. To address the economic feasibility, the operational costs of a Cessna 402 used by Cape Air and a concept electric aircraft are compared. From a direct comparison, using electric aircraft results in a staggering 20% reduction in operating costs, mainly driven by a reduction in energy and maintenance costs and partly offset by an increased acquisition cost of the aircraft and batteries. The benefits of electric propulsion are clear.

In [53] the authors build on the work of [8] by developing a "power-optimized" strategy to minimize the peak power used for charging in a battery swap station. Firstly, they performed the conceptual sizing on an electric version of the Tecnam P2012T Traveller and considered two philosophies. In the first one, the battery is sized to perform a full day of operations and only charging during the night. The other concept is to use swappable batteries. Due to the huge battery required to perform uninterrupted operations, they went on with the latter. The strategy was simulated on the network of Cape Air and Mokulele Airlines, two successful commuter airlines. Table 3.5 provides an overview of key information of these two airlines. The "power-optimized" strategy is able to reduce energy costs by 20% compared to the benchmark "Power-as-needed" charging strategy and by 70% compared to fossil fuel powered aircraft. However, lowering the peak-power and number of chargers may increase the number of batteries required. Therefore, [9] develops the "power-investment-optimized" strategy. This strategy starts from the "power-optimized" solution as an upper bound on the number of batteries and lower bound on the number of chargers. A pareto-front is constructed by incrementally increasing the number of chargers to find the solution with the minimum overall cost. Interestingly, fast charging was not considered as it was deemed impractical with the short TATs usually operated by commuter airlines. However, modern fast charging technology as described in Section 3.2.3 has numerous advantages over battery swap and cannot be ignored. The TAT window of 20 to 35 minutes provides an opportunity to at least partly recharge the batteries. Table 3.4 shows the cost assumptions made in [8], [53] and [9], based on an entry into service in 2030. These assumptions appear to be very realistic. Skylax claims that its aircraft will have price of €1.9m. That claim is likely to increase when the aircraft actually enters production. Eviation claims a price of €4m for an electric aircraft with a better performance. As Tesla Supercharger already achieve a power output of 250W, the assumed charger power seems conservative. As battery swap is used in these papers, that is fine. If fast-charging is used instead, the charger power will likely be higher.

Aircraft Acquisition Cost	€2.5m [8]
Charger Power	125kW [53] [9]
Charger Efficiency	90% [53] [9]
Charger Useful Like	7 years [9]
Battery Specific Energy	350 Wh/kg [53]
Battery Useful like	1,000 cycles [53] [9]
Charger Cost	$100,000 [53] [9]
Battery Specific Cost	125 $/kWh [53], 100$/kWh [9]
Discount Factor	8.1% [53] [9]

Table 3.4: Cost assumptions in literature

	Cape Air	**Mokulele Airlines**
Fleet	48 Cessna 402	8 Cessna 208
Weekly flights	1839	732
Cities Served	19	9
Median TAT	35 min	19 min

Table 3.5: Key information of existing commuter airlines Cape Air and Mokulele Airlines

On average, each aircraft in the Cape Air fleet conducts 5 flights per day while Molukele aircraft conduct 10 flights per day on average. The difference is partly in the lower TAT but mostly in the shorter distances flown by Mokulele. From Figure 3.4 it can be seen that 12% of the Cape Air flights exceeds 300km, which was identified as a realistic range for electric aircraft in the near future.

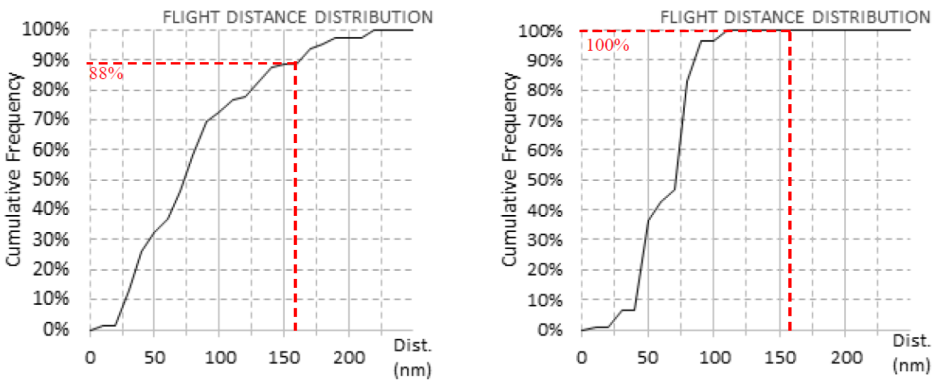

Figure 3.4: Flight distance distribution of Cape Air (left) and Mokulele Airlines (right)

To assess the potential demand, [61] estimated door-to-door travel time of existing modes of transport and air taxis, which essentially are commuter aircraft. They found that commuter aircraft are most competitive for distances between 100 and 300km. In that range, railways is the biggest competing mode of transport. Potential demand therefore depends on how well the railway network is developed. Western-Europe is very well connected by railways. Figure 3.5 shows the top 1000 links based on travel demand and travel time competitiveness of a commuter aircraft. Even though Western-Europe is well connected by railway, the high travel demand is appealing. An extension to the UK also seems promising, especially up north where railway connections are more scarce. Other potential regions include Madrid and Barcelona and large cities in Eastern-Europe. This paper does not consider the average income and living standards. With this in mind, the region of Western-Europe and the UK appears to be most interesting to start an initial service.

A similar study was performed on suburban areas in the United States in [54]. They also vary take-off distance in their sensitivity analysis because this affects the number of available airports. They conclude that a suburban commuter airline should target ground commutes exceeding 45 minutes. With a cruise speed of 160 knots and take-off distance of 500 feet (300kph, 150m), travel times can be reduced by 45%. If the take-off distance doubles to 1,000 feet, travel time is reduced by 40%. According to [47], these requirements are realistic. A 300 feet take-off distance is readily feasible and can be reduced to 100 feet with near-future technology and a reduction in cruise speed.

Also for UAM, many demand estimations have been made. Though demand for this concept of transport is different than for thin-haul aviation, there are many similarities. [18] expects such a demand that small airports which are almost not utilized today will need upgrading and additional short runways will be required. In [10], an agent based model called BaySim is used. This model generates passengers with trip requests. The numbers of passengers that are generated is based on

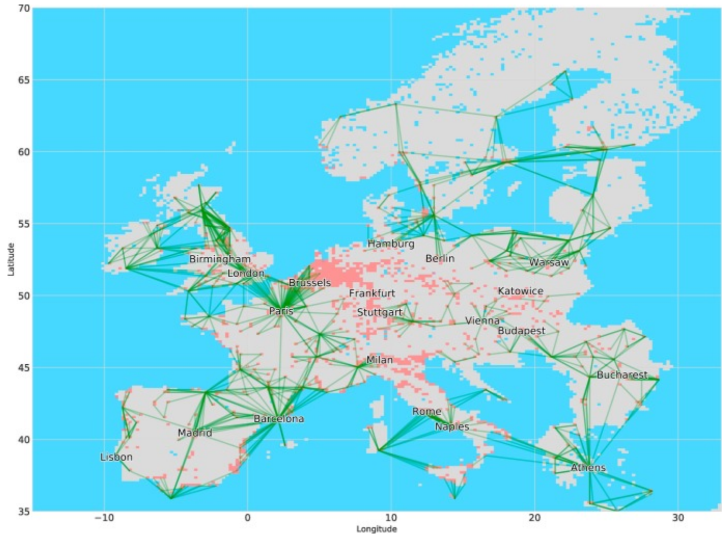

Figure 3.5: Top 1000 links based on travel time competitiveness and travel demand

demographic studies of the Bay Area, California. The author of [39] uses the combination of three normal distributions, N(4,2), N(16,2) and N(12,6), to model the variation of demand throughout the day. [44] models demand by creating individual people with trip requests. The trips can be either commuter, churn or airport transfer. A commuter travels in the morning and back in the afternoon, a churn trip is a one-way trip at random times throughout the day. The outcome is similar to the simple model used by [39]. This supports the use of a simple demand model based on normal distributions in a generalised context.

[41] proposed an aircraft scheduling optimisation for an electric version of a regional airliner with both battery swapping and battery charging. Since no such aircraft exists to date, she based her aircraft model on a Class I estimation starting from the Embreaer 175. For fast charging, she assumes a charger power of 3MW. This by far exceeds any existing charging technology. Tesla Superchargers have up to 250kW and the currently in development Tesla Megacharger has 1MW. A 3MW charger is not realistic in the foreseeable future. Then again, this paper considers an electric regional airlines so it is acceptable to assume future advancements in technology and look further than the near future. That is in contrasted with the assumed battery price of over 1,402 €/kWh, equal to the battery price of the Pipistrel Taurus Electro G2. This aircraft has only been build on a small scale and with already relatively old technology. In contrast, other papers assume a battery price down to 100 €/kWh. Continuing along the lines of assuming future technologies, this would be a more realistic assumption.

3.4. Energy and Emissions

Electric aircraft have no tailpipe emissions, meaning they do not emit locally. Obviously, the energy consumed is produced somewhere and that can involve emission of harmful gasses. To understand how environmentally friendly electric aircraft actually are, the production of electric energy must be investigated. Furthermore, The European Union (EU) has implemented the Emission Trading Scheme (ETS). This already covers all intra-European flights. Policies like these provide additional opportunities for operations involving lower emissions so it is important to be aware of them.

Figure 3.6 provides an overview of the CO_2 emissions related to electricity production in the EU [4]. There is a clear downwards trend as renewable energy sources are becoming more prevalent but the situation varies drastically between different countries. In 2016, The Netherlands was still above 500 g/kWh while Sweden went as low as 13g/kWh. The EU wide average is around 300g/kWh.

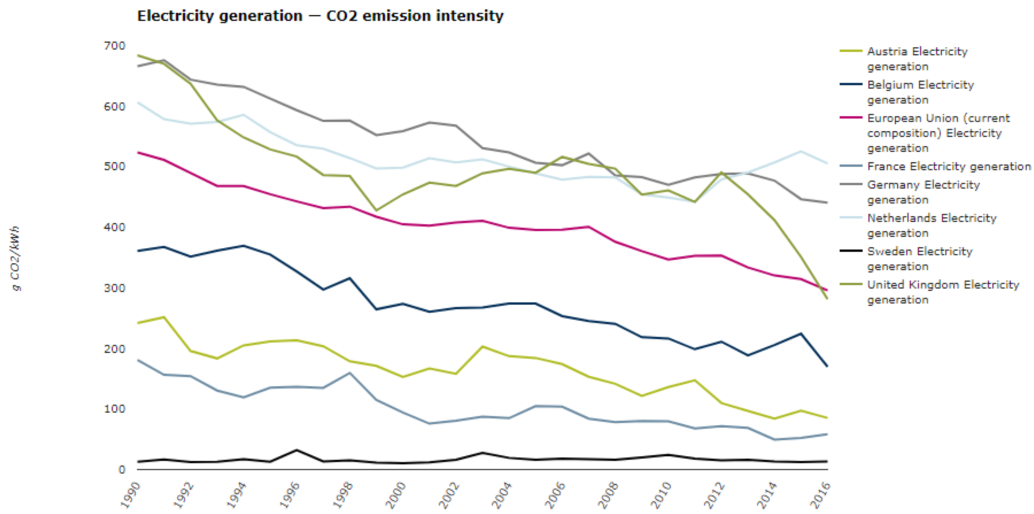

Figure 3.6: CO_2 emissions related to the production of electricity

To put this into perspective, current airlines are estimated to produce 115g of CO_2/RPK [25]. The developers of the Skylax E10 report that the aircraft will have a battery weight of 1,000 kg and range of 300 km. Assuming an energy density of 250 Wh/kg results in a 250kWh battery. At EU average, a full charge is therefore equivalent to 75kg of CO_2. This 75kg is able to transport 10 passengers over a distance of 300km, resulting in 25g/RPK, just 22% of the estimated emissions for airlines. This is already a substantial improvement over traditional airlines but renewable energy sources can further reduce this number. If the aircraft was charged in Sweden, the effective emissions would only be 1.1g/RPK. Compared to cars, a consumption of 5 liter of gasoline per 100 km results in roughly 120g/km [15]. If the car contains 5 people, that means 24 g/RPK. Note that these are simplified estimations but it shows that electric aircraft are competitive with gasoline powered cars in terms of CO_2 emissions. Trains still do better, at just 10g/RPK [32].

3.4.1. NO_x Emissions Charge

Five countries in Europe are charging NO_x emission charges for aircraft (Switzerland, Sweden, United Kingdom, Denmark and Germany). The purpose of these is to reduce local emission of nitrogen oxides. Switzerland and Sweden were the first to introduce these charges in 1997 and 1998 based on the aircraft emission class. Switzerland used five classes and Sweden seven. To harmonise the policies in Europe, the European Civil Aviation Conference published guidelines for emission classifications to use for emission charges. Following these guidelines, London Heathrow introduced their own charges in 2004 and Gatwick in 2005. Germany followed in 2008. [19]

The charges are based on emission value units. The emission value units are determined for each aircraft type during the Landing and Take-Off (LTO) cycle and are based on the number of engines, fuel flow during the LTO cycle and an index indicating the NO_x emissions per weight unit of fuel. The amount that is charged per emission unit varies per country. In Germany, this is set at €3/unit. Sweden charges €5.5, Switzerland charges between €1 and €3 and the UK charges €1.6. As a result, aircraft such as the A320 may be charged up to €27 and the Boeing 747 up to €385 per landing. For aircraft lighter than 5.7 tonnes, a fixed amount of €3 is charged per landing. [19]

Germany intends the charges to be revenue neutral, meaning that this is not an additional cost for airlines. Instead, the landing fees based on the aircraft weight are lowered by an equivalent amount. This benefits the airlines with lower emissions and makes environmentally friendly aircraft and engines a more attractive option. As these charges are aimed at reducing local emissions, electric aircraft should be entirely exempt.

3.4.2. CO_2 Emission Policies

Aviation accounts for roughly 3.5% of CO_2 emissions in the EU and became part of the ETS in 2008 [19]. In total, roughly 45% of emissions within Europe fall under this scheme. The ETS is referred to as a 'cap and trade' system. There is a limit to how much each sector is allowed to emit. If a sector surpasses this cap, they can trade the additional emission rights with another sector who is under its emissions cap. Set-up in 2005, the ETS became the first carbon trade market in the world. Emission rights are traded in auctions so there is no fixed price. At the start of 2020, the average price was €25 per tonne of CO_2. This dropped by 40% to €15 per tonne following the outbreak of COVID-19 [12]. Given that airlines are estimated to emit 115g of CO_2 per RPK [25], this translates into a cost of €0.0029/RPK. Airline costs are estimated at €0.13/RPK, meaning that the acquisition of emission rights increases costs by 2.2%. Electric aircraft may allow the aviation industry to reduce its emissions below the cap to avoid this extra cost.

The ETS is in its third phase, which lasts until 2021. This phase is mainly characterised by the inclusion of additional gasses and sectors in the trading scheme. The fourth phase will last from 2021 until 2030 and will see additional policies to reduce emissions. This will be part of the European Green Deal which is currently being developed and negotiated by the European Commission.

3.4.3. Noise Charges

Noise is also seen as pollution. Zurich was the first airport to introduce noise charges in 1980. Today, over 100 airports in Europe apply noise charges based on the time-of-day, the weight category and noise class of the aircraft. There are 9 weight categories based on the maximum take-off weight. Light aircraft with a MTOW between 6 and 15 tonnes fall under category 1, the heaviest aircraft such as the A380 fall under category 9. General Aviation is typically lighter than category 1 and is therefore exempt from the noise charges and so are commuter aircraft. For aircraft that are too heavy to be exempt, there are 4 noise classes where the quietest aircraft fall under class 4. Modern aircraft are usually classified in class 4. This means that even if electric aircraft are quieter, this will not give them an economic advantage over aircraft from the same weight class. Besides, electric aircraft are not completely silent as much of the noise is generated by the propeller, airframe, landing gear, etc. [13].

In an optimisation problem, noise and emission charges may be represented by a fixed cost per flight leg. For electric commuter aircraft, no such costs actually apply. Only once electric aircraft become larger to carry more passengers or cargo, the noise charges will apply. With current policies, no emission charges apply to electric aircraft, even though they are not entirely zero-emission because of electricity production.

4 Research Outline

The feasibility and economic considerations in electric thin-haul aviation have been investigated in previous research. However, no one has proposed a fleet sizing and scheduling optimisation model using fast charging technology. [41] comes close by doing this for a conceptual electric regional airline and is based on a given set of service trips and given charger location. While a large electric aircraft seems very far off, a number of companies are developing 10-seater electric aircraft with commercial potential in thin-haul operations. It would be possible to adjust this model to apply it at a thin-haul commuter network but this requires the timetable of scheduled service trips to be known in advance, Because the operating costs of electric aircraft are over 20% lower and because they may provide an environmentally friendly alternative in more and more congested ground transport, there may be opportunities for electric air transport services where currently no network exists. In that case, determining which service trips to make must be part of the question. This has been recognized in research and some papers have estimated demand and competitiveness.

4.1. Research Gap

Electric thin-haul aviation lies in the overlap between airline research and electric vehicle or bus scheduling research. The reason being that electric thin-haul aviation has many similarities with electric busses, probably more than with airliners. The electric bus scheduling problem can be extended to also include where to install charging infrastructure. Airline research on the other hand is in a very advanced stage and focuses on robustness of the schedules, which is not applicable for thin-haul yet as this field is not mature enough. The papers discussing the E-VSP usually take a given set of service trips as input. This is not applicable to electric aircraft for two reasons; i) Electric aircraft have very different performance figures. The service trips that an electric aircraft should make could be different to what a fossil fuel powered aircraft should take; ii) Hardly any scheduled thin-haul networks currently exist. An operator intending to deploy a fleet of electric aircraft will have to **determine which service trips to make.**

Whereas E-VRP and E-VSP studies minimize either travelled distance or costs, this approach will not work for this thesis because determining the service trips is part of the problem. Minimizing costs would simply return an empty set of service trips unless if a minimum RPK constraint is added. In the end, every company wants to maximise profits but also this is not possible because prices and thus revenues cannot be estimated reliably. The author of this literature study is of the opinion that the most commercially relevant objective is to **minimize the cost per RPK.** That will allow a direct comparison with other modes of transport and give an impression of the minimum ticket price an operator should apply to be profitable.

4.2. Research Questions

These considerations lead to the following **Research Question** and Sub-questions.

Given a daily time discretised dynamic demand pattern, how should a fleet of electric aircraft be sized and scheduled given charging infrastructure and current airport infrastructure to minimise costs/RPK?

- How do we model an electric aircraft, capturing it's behavior in a network, in the simplest way?
- What infrastructure do electric aircraft require?
- How do we estimate the costs of operating a fleet of electric commuter aircraft and charging

infrastructure?

- How can we model demand in a thin-haul network, such that is captures dynamic behavior while remaining flexible to new input?
- How can we model a network of electric commuter aircraft using current airport infrastructure?
- How can a linear optimisation problem be formulated to minimize costs of a thin-haul network based on dynamic demand with a minimum RPK constraint?
- What is the computational complexity of such optimisation problem?
- Can an algorithm based on Branch-and-Bound and Column Generation solve the optimisation problem exactly in reasonable time on a network with up to 20 airports and 50 aircraft?
- What is the relation between computational time and instance size?
- What are the characteristics of an optimal schedule in terms of ASK, RPK, costs per ASK and RPK, load factors and aircraft productivity?
- Which routes are most prominent in an optimal schedule?
- What is the sensitivity of the Key Performance Indicator (KPI)s to variations in the model parameters?
- How do costs- and emissions per RPK compare with other modes of transport?
- What are the remaining obstacles towards commercialization of an electric thin-haul network?
- Which topics are most relevant for further research?

4.3. Method

The main contribution of this thesis will be proposing a new model that is able to use a complex demand pattern to determine an optimal set of service trips, fleet sizing and scheduling to have the largest chance to succeed commercially. Because determining service trips, fleet size and scheduling are interrelated, this thesis will attempt to combine these in a single mathematical model. The objective will be to minimize the cost per RPK. Both costs and RPK are variable so this implies a non-linear objective function. Therefore, the model will incorporate a minimum RPK constraint and the objective will be to minimize the costs. The problem will be solved iteratively for different minimum RPK levels to construct a graph indicating the cost-RPK relation. This graph will start at a minimum cost of acquiring a single aircraft and operating it on the shortest route as this is required to achieve a non-zero RPK. The maximum RPK is determined by the demand model. When taking a low RPK, acquisition costs will be high relative to operational costs. When taking the maximum RPK value, a less optimal schedule is expected to ensure inclusion of all passengers, even those on expensive routes requiring deadheading. The optimal cost/RPK is expected to be somewhere in between. The planning for the research project can be found in Appendix A.

5 Conclusion

Electric aircraft are key to sustainable aviation in the future. Technology is developing at a very high pace and even has commercial potential in the near future. A number of companies are developing electric commuter aircraft and intend to start production in the coming years. These aircraft will be very different to operate than fossil fuel powered aircraft. Operations research is required to ensure optimal fleet sizing and scheduling. This literature review summarized the state-of-the-art in relevant research fields.

Although a number of hurdles currently prevent commercial operations of UAM or large electric aircraft, that is where operations research in electric aviation has focused on. This ignores the need for an agile market and research plan. Electric commuter aircraft will provide a step up to larger electric aircraft but receive little attention in research. Economic and operational feasibility studies have indicated that electric aircraft reduce costs in thin-haul commuter operations over 20%.

Electric technology has been studied extensively for the VRP and VSP. Heuristic methods are usually used to solve the E-VRP whereas exact algorithms based on Branch-and-Bound and Column-Generation are more common for the E-VSP. The E-VSP has been applied to a number of case studies to investigate scheduling of electric bus fleets. Scheduled electric bus operations have many similarities with scheduled electric thin-haul aviation. The models developed in these papers provide a starting point to develop a model for the fleet sizing and scheduling of electric aircraft.

The feasibility studies in electric thin-haul aviation assume battery swap. As fast-charging technology is already very advanced, the batteries can already be charged substantially during the TAT. Battery swapping induces additional logistic challenges in the scheduling of spare batteries and ownership issues in case the swap stations are operated by a third party provider. It also requires a larger upfront cost because of the additional batteries and the swap station. Fast-charging appears to be the most practical way to replenish the energy level.

Commuter airlines missed out on the industry growth due to high operating cost. Electric aircraft will allow costs to decrease and optimisation algorithms must ensure optimal use of the aircraft in a network. The reduced cost may allow new entrants to enter the market and offer a scheduled thin-haul airline service. A decisive KPI that determines success is the cost/RPK. Therefore, the research project following this literature study will attempt the following:

- Develop a fleet sizing and scheduling optimisation for electric commuter aircraft using fast-charging technology based on dynamic demand, charging infrastructure and current airport infrastructure. Past research assumes a known set of service trips. No past research has investigated the fleet sizing and scheduling of electric aircraft based on dynamic demand and no past research has minimized cost/RPK.
- Develop an algorithm to solve the optimisation exactly for a network with up to 20 nodes and 50 aircraft. This will ensure that the model is useful in practice.
- Conduct sensitivity analysis in aircraft and network parameters to draw conclusions on the general behaviour of the optimisation and address uncertainty in the assumptions.

A Gantt Chart

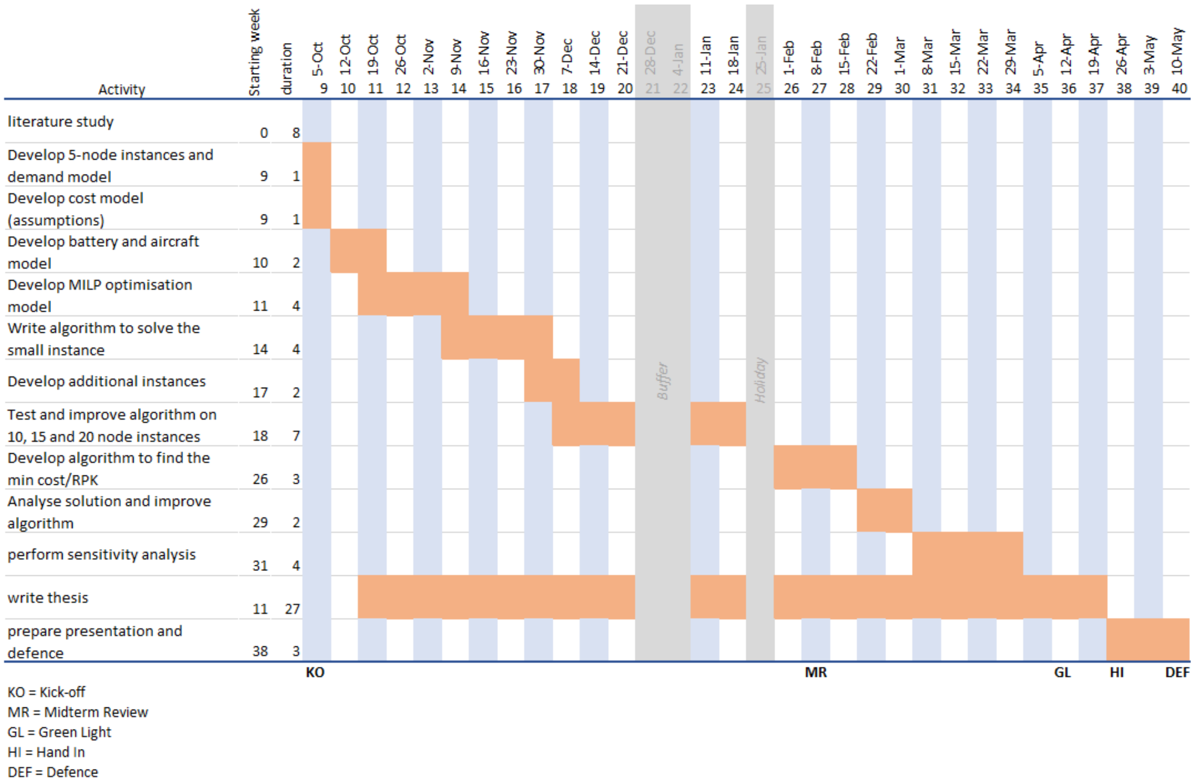

Figure A.1: Planning for the research project

Bibliography

[1] Jamili A. A robust mathematical model and heuristic algorithms for integrated aircraft routing and scheduling, with consideration of fleet assignment problem. *Journal of Air Transport Management*, 2016.

[2] Schneider M. Stenger A. and Goeke D. The electric vehicle-routing problem with time windows and recharging stations. *Transportation Science*, 2020.

[3] Verma A. Electric vehicle routing problem with time windows recharging stations and battery swapping stations. *European Journal of Tranportation and Logistics*, 2018.

[4] European Environment Agency. Overview of electricity production and use in europe. `https://www.eea.europa.eu/data-and-maps/indicators/overview-of-the-electricity-production-2/assessment-4`, 2019. Accessed: 25 Sep 2020.

[5] Pl¨otner K.O. Vratny P. C. Schmidt M. Isikveren A.T. and M. Hornung. Impact of electrically powered transport aircraft on energy and battery demand for germany. *Proceedings of the 62nd German Aerospace Congress (DLRK 2013)*, 2013.

[6] Lin Y. Zhang K. Shen Z.-J. M. Ye B. and Miao L. Multistage large-scale charging station planning for electric buses considering transportation network and power grid. *Transportation Research Part C: Emerging Technologies*, 2019.

[7] Johnson B.A. and White R.E. Characterization of commercially available lithium-ion batteries. *Journal of power sources*, 1998.

[8] Harish A. Perron C. Bavaro D. Ahuja J. Ozcan M. Justin C.Y. Briceno S.I. German B.J. and Mavris D. Economics of advanced thin-haul concepts and operations. *Journal of Aerospace Information Systems*, 2016.

[9] Justin C.Y. Payan A.P. Briceno S.I. German B.J. and Mavris D.N. Power optimized battery swap and recharge strategies for electric aircraft operations. *Transportation Research Part C: Emerging Technologies*, 2020.

[10] Alonso J.J. Arneson H.M. Melton J.E. Vegh M. Walker C. and Young L.A. System-of-systems considerations in the notional development of a metropolitan aerial transportation system. *NASA*, 2014.

[11] Koc C. and Karaoglan I. The green vehicle routing problem: A heuristic based exact solution approach. *Applied Soft Computing*, 2016.

[12] Van den plas S. When covid-19 met the eu ets. `https://carbonmarketwatch.org/2020/03/26/when-covid-19-met-the-eu-ets/`, 2020. Accessed: 22 Sep 2020.

[13] Simons D.G. Introduction to aircraft noise. Technical report, Delft University of Technology, 2019.

[14] J. Artal-Sevil J. Dom´ınguez-Navarro J. Dufo-L´opez, R. Yusta-Loyo and Bernal-Agust´ın J. Design of an electric vehicle fast-charging station with integration of renewable energy and storage systems. *International Journal of Electrical Power Energy Systems*, 2019.

[15] Ecoscore. Hoe bereken je de co2-uitstoot uit het brandstofverbruik? `https://ecoscore.be/nl/info/ecoscore/co2`, 2020. Accessed: 22 Sep 2020.

[16] List et al. Robust optimization for fleet planning under uncertainty. *Transportation Research Part E: Logistics and Transportation Review*, 2003.

[17] Moore M. D. et al. Misconceptions of electric aircraft and their emerging aviation markets. *Aerospace Sciences Meeting*, 2014.

[18] Moore M.D. et al. Projected demand and potential impacts to the national airspace system of autonomous, electric, on-demand small aircraft. *12th AIAA Aviation Technology, Integration, and Operations (ATIO) Conference*, 2012.

[19] Bigoni F. Impact of hybrid-electric aircraft operations. *MSc thesis Politecnico di Milano*, 2017.

[20] Felipe A. Ortuno M.T. Righini G. and Tirado G. A heuristic approach for the green vehicle routing problem with multiple technologies and partial recharges. *Transportation Research Part E: Logistics and Transportation Review*, 2014.

[21] Repko M. G.J. and Santos B. F. Scenario tree airline fleet planning for demand uncertainty. *Journal of Air Transport Management*, 2017.

[22] Li. L. Lo H.K. and Xiao F. Mixed bus fleet scheduling under range and refueling constraints. *Transportation Research Part C: Emerging Technologies*, 2019.

[23] Cortes-Murcia D.L. Prodhon C. Asfar H.M. The electric vehicle routing problem with time windows, partial recharges and satellite customers. *Transportation Research Part E: Logistics and Transportation Review*, 2019.

[24] Koyuncu I. and Yavuz M. Duplicating nodes or arcs in green vehicle routing: A computational comparison of two formulations. *Transportation Research Part E: Logistics and Transportation Review*, 2019.

[25] Carbon Independent. Aviation co2. `https://www.carbonindependent.org/22.html`, 2019. Accessed: 17 Sep 2020.

[26] Lu L. Han X. Li J. Hua J. and Ouyang M. A review on the key issues for lithium-ion battery management in electric vehicles. *Journal of power sources*, 2013.

[27] Wang Y. Huang Y. Xu J. and Barclay N. Optimal recharging scheduling for urban electric buses: A case study in davis. *Transportation Research Part E: Logistics and Transportation Review*, 2017.

[28] Li J.-Q. Transit bus scheduling with limited energy. *Transportation Science*, 2014.

[29] Adler J.D. Routing and scheduling of electric and alternative-fuel vehicles. *Ph.D. thesis Arizona State University*, 2014.

[30] Pelletier S. Jabali O. Mendoza J.E. and Laporte G. The electric bus fleet transition problem. *Transportation Research Part C: Emerging Technologies*, 2019.

[31] G. Jiang, M. Zhang Y. Zhang C. Zhang K. Zhang and Zhao Z. Operation and scheduling of pure electric buses under regular charging mode. *2018 21st International Conference on Intelligent Transportation Systems (ITSC)*, 2018.

[32] Amit K. Is holidaying by train really that much better for the environment? `https://www.wired.co.uk/article/trains-planes-emissions-co2-comparison#:~:text=According%20to%20the%20European%20Environment,miles%20from%20journeys%20in%20cars`, 2019. Accessed: 22 Sep 2020.

[33] An K. Battery electric bus infrastructure planning under demand uncertainty. *Transportation Research Part C: Emerging Technologies*, 2020.

[34] Steiner K. and Irnich S. Schedule-based integrated intercity bus line planning via branch-and-cut. *Transportation Science*, 2018.

[35] Cadarso L. and Marín Á. Robust passenger oriented timetable and fleet assignment integration in airline planning. *Journal of Air Transport Management*, 2012.

[36] Silva H. L. and Guimaraes T. A. Conceptual design of a thin-haul aircraft by energy sizing optimization including aero-propulsive interactions. *Journal of Aerospace Information Systems*, 2020.

[37] Hepperle M. Electric flight - potential and limitations. *DLR*, 2012.

[38] Keskin M. and Catay B. Partial recharge strategies for the electric vehicle routing problem with time window. *Transportation Research Part C: Emerging Technologies*, 2016.

[39] Kleinbekman I.C. Mitici M. and Wei P. Rolling-horizon electric vertical takeoff and landing arrival scheduling for on-demand urban air mobility. *Aerospace Information Systems*, 2019.

[40] Leggieri V. Haouari M. A practical solution approach for the green vehicle routing problem. *Transportation Research Part E: Logistics and Transportation Review*, 2017.

[41] Pereira M. Short-range route scheduling for electric aircraft with battery-charging and battery-swapping constraints. *MSc thesis TU Delft*, 2019.

[42] Piller S. Perrin M. and Jossen A. Methods for state-of-charge determination and their applications. *Journal of power sources*, 2001.

[43] Sauer D.U. Bopp G. Jossen A. Garche J. Rothert M. and M. Wollny. State of charge–what do we really speak about. *The 21st international telecommunications energy conference*, 1999.

[44] Kohlan L. Patterson M.D. and Raabe B.E. Urban air mobility network and vehicle type - modeling and assessment. *NASA*, 2019.

[45] Shareef H. Islam M.M. and Mohamed A. A review of the stage-of-the-art charging technologies, placement methodologies, and impacts of electric vehicles. *A review of the stage-of-the-art charging technologies, placement methodologies, and impacts of electric vehicles.*, 2016.

[46] Bruglieri M. Mancini S. Pezzella F. Pisacane O. and S. Suraci. A three-phase matheuristic for the time-effective electric vehicle routing problem with partial recharges. *Electronic Notes in Discrete Mathematics*, 2017.

[47] Courtin C. Burton M. Butler P. Ye A. Vascik P. and Hansman J. Feasibility study of short takeoff and landing urban air mobility vehicles using geometric programming. *Aviation technology, integration and operations conference*, 2018.

[48] Rong P. and Pedram M. An analytical model for predicting the remaining battery capacity of lithium-ion batteries. *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, 2006.

[49] Hansman R.J. and Vascik P.D. Operational aspects of aircraft-based on-demand mobility, 2016.

[50] Dessaulniers G. Errico F. Irnich S. and Schneider M. About sections exact algorithms for electric vehicle-routing problems with time windows. *Operations Research*, 2016.

[51] Dožić S. and Kalić M. Three-stage airline fleet planning model. *Journal of Air Transport Management*, 2015.

[52] Erdogan S. and Miller-Hooks E. A green vehicle routing problem. *Transportation Research Part E: Logistics and Transportation Review*, 2012.

[53] Justin C.Y. Payan A.P. Briceno S.I. and Marvis D.N. Operational and economic feasibility of electric thin haul transportation. *Journal of Aerospace Information Systems*, 2017.

[54] Wei L. Justin C.Y. Briceno S.I. and Mavris D.N. Door-to-door travel time comparative assessment for conventional transportation methods and short takeoff and landing on demand mobility concepts. *Transportation Research Part E: Logistics and Transportation Review*, 2018.

[55] Liu T. and Ceder A. Battery-electric transit vehicle scheduling with optimal number of stationary chargers. *Transportation Research Part C: Emerging Technologies*, 2020.

[56] Yao E. Lio T. Lu T. and Yang Y. Optimization of electric vehicle scheduling with multiple vehicle types in public transport. *Sustainable Cities and Society*, 2020.

[57] Battery University. Battery university. `https://batteryuniversity.com/`, 2020. Accessed: 20 Sep 2020.

[58] Rogge M. van der Hurk E. Larsen A. and Uwe Sauer D. Electric bus fleet size and mix problem with optimization of charging infrastructure. *Applied Energy*, 2018.

[59] van der Zwan F.M. Wils K. and Ghijs S.S.A. Development of an aircraft routing system for an air taxi operator. *Aeronautics and Astronautics*, 2011.

[60] van Kooten Niekerk M.E. van den Akker J.M. and Hoogeveen J.A. Scheduling electric vehicle. *Public Transport*, 2017.

[61] Sun X. Wandelt W. and Stumpf E. Competitiveness of on-demand air taxis regarding door-to-door travel time: A race through europe. *Transportation Research Part E: Logistics and Transportation Review*, 2018.

[62] Tang X. Lin X. and He F. Robust scheduling strategies of electric buses under stochastic traffic conditions. *Transportation Research Part C: Emerging Technologies*, 2019.

[63] Zhang X. and Rao R. A benefit analysis of electric vehicle battery swapping and leasing modes in china. *Emerging Markets Finance and Trade*, 2016.

4 Conclusions

In this MSc Thesis, we formulated an optimisation model to schedule a fleet of electric thin-haul aircraft. The problem was to determine a schedule that minimises cost under a minimum RPK (Revenue-Passenger Kilometer) constraint. By varying the minimum RPK, a solution that minimises cost per RPK can be found. Because the problem is NP-hard and could not be solved to optimality using the pure MILP (Mixed Integer Linear Programming) formulation, we developed a novel method. The algorithm was implemented and tested on various of networks, ranging from 5 to 30 airports for a single day of scheduling.

The algorithm proved to be successful in finding a very good local optimum. The case studies revealed that the Cost per RPK is not a smooth function, rather showing a sawtooth pattern. For a given fleetsize, the minimum Cost/RPK is found when the aircraft are flying at high-utilisation, meaning that they are performing as many flights as possible. The schedule shows a strong preference for flights with i) high demand; and ii) long flight distance. When demand is large enough, few seats must be left empty and a larger RPK is achieved for the same cost. Long distance flights are preferred because they have a lower energy and cost per flown kilometer. To achieve the minimum RPK, the algorithm tends to schedule back-and-forth flights between airports with a long distance and high demand. These flights are the first to be scheduled until they are largely saturated and there is little demand on them left. When this happens and the minimum RPK is increased further, the schedule will start planning flights with a shorter distance and/or lower demand. This causes the sawtooth pattern for Cost/RPK to have an increasing trend.

The contributions of this MSc Thesis are as follows. To the best of our knowledge, this is the first paper that considers cost per RPK as an evaluation criterion in an E-VSP (Electric-Vehicle Scheduling Problem) and is the first paper to apply operations research to electric thin-haul aircraft. We proposed a new scheduling problem where the set of service trips is not known in advance but rather had a minimum RPK constraint. We then proposed a method to solve the problem, which proved successful in finding a good local minimum. The methods and insights developed for this MSc Thesis will support future thin-haul airline operators in decisions on fleetsize and in scheduling to minimise the cost per RPK.

Further research could focus on the a number of extensions to the problem formulated in this Thesis. i) A multi-depot variation. This Thesis considers only 1 hub where all aircraft start and end their day. It may however be interesting to have multiple hubs to provide access to more flights and prevent the flights around the hub to be congested; ii) Adding a charger constraint. In this Thesis, the number of chargers required at any airport was assumed to be present. In reality, the number of chargers will be limited; iii) A battery swap variation. Instead of using fast-charging, swapping depleted batteries by fully charged ones also has some benefits; and lastly iv) Variable electricity prices. The electricity prices can vary drastically by geographic location, time of day, peak power and total energy consumption. A variation of the problem where electricity costs are incurred on the ground arcs instead of on flight arcs may allow to include variations based on location and time of day. An additional variable will be required to track whether an aircraft on the ground is charging or not. This Thesis took the first steps in the field of electric thin-haul airline operations, which promises to remain a fascinating field for the upcoming years with many opportunities for further research.

A

Improvements on the arc based formulation

The first step in the optimisation was laying out the mathematical formulation of the problem. This is the arc based formulation and is explained in detail in Sec. 3 of the scientific paper. This section discusses how the problem was originally formulated and what improvements were made to finally get to the arc based formulation.

A.1. Enforcing the Turn-Around-Time

In the original network, the TAT (Turn-Around-Time) was not included in the flight arcs as plotted in Fig. A.1(a). Constraint A.1 is installed to enforce the TAT. We realise that when an aircraft arrives at node $i(a,t)$, then the earliest possible node where the aircraft may depart for a subsequent flight is node $j(a,t+t_{tat})$. We therefore augment the TS (Time-Space) network by making the following adjustments. For each node $i(a,t)$ in the network, we create a copy $i^*(a,t)$. Let the original node i be a ground node and the new node be an arrival node. Each flight arc (i,j) is replaced by flight arc (i,j^*) which, apart from arriving at the arrival node, has exactly the same properties as the original arc. Lastly, let the arrival nodes be connected to the ground nodes using TAT arcs $(i^*(a,t), j(a,t+t_{tat}))$. This is visualised in Fig. A.1(b) Upon arrival at an arrival node after a flight, the aircraft must first use a TAT arc before it has access to new flight arcs. The TAT is therefore automatically respected and constraint A.1 is no longer necessary.

$$x_{i''(a,t-2t_{step})i'(a,t-t_{step})} + x_{i'(a,t-t_{step})i(a,t)} \geq 2 \sum_{(i,j)\in(\delta_i^- \cup A_f)} x_{ij} \qquad \forall i(a,t) \in N - (i(a_h,t_{min}) \cup i(a_h, t_{min}+t_{step})) \tag{A.1}$$

Augmenting the network comes at the cost. The number of nodes is doubled and a new set of arcs is introduced. As each aircraft has a decision variable for each node and for each arc, this introduces a lot of additional decision variables. Furthermore, additional constraints are needed to ensure flow conservation on the arrival nodes and to regulate the change in energy over the TAT arcs. Augmenting the network does more harm than good. Luckily, we can simplify the network. After using a flight arc, there is only one arc available, the TAT arc. The flight- and TAT arc can be merged without loss of generality. By doing so, the arrival nodes and separate TAT arcs are no longer needed while retaining the advantage that no TAT constraint is needed. This results in the simplified augmented network and is visualised in Fig. A.1(c). This network even has a few nodes and arcs less than the original network. Suppose that that TAT is 30 minutes and the time step in the TS network is 15 minutes. Suppose it takes a flight time $t(a_h,a_1)$ to fly from the hub a_h to airport a_1. In the original network, the first node at a_1 would be $i_1(a_1, t_{min}+t(a_h,a_1))$. This node is followed by $i_2(a_1, t_{min}+t(a_h,a_1)+t_{step})$ and $i_3(a_1, t_{min}+t(a_h,a_1)+2t_{step})$. These are connected by the ground arcs (i_1,i_2) and (i_2,i_3). The first node in the simplified augmented network is $i_3(a_1, t_{min}+t(a_h,a_1)+t_{tat})$ because $t_{tat} = 2t_{step}$. The nodes i_1 and i_2 and the ground arcs (i_1,i_2) and (i_2,i_3) are not present anymore. On top of eliminating the TAT constraints, the number nodes and arcs is reduced slightly, which reduces the number of decision variables and number of other constraints.

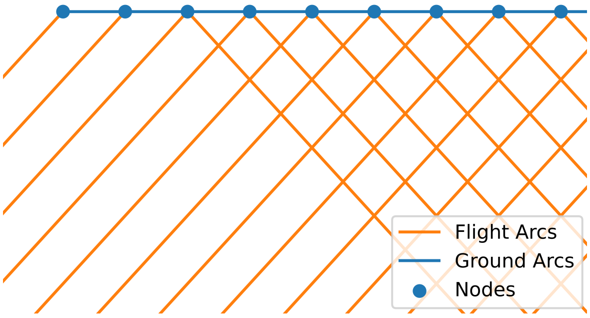

(a) Original network

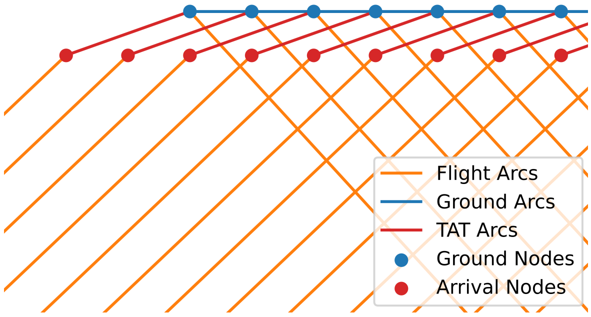

(b) Augmented network

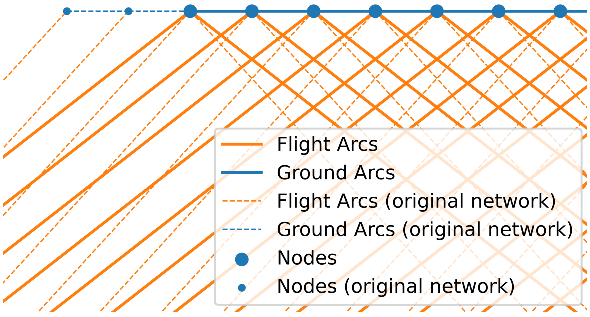

(c) Simplified augmented network. The underlying original network shows the elimination of the first few nodes and ground arcs.

Figure A.1: Comparison of the different versions of the TS network

A.2. Determining the Number of Passengers

The number of passengers has to be determined to properly compute the RPK and to prevent exceeding the demand. In the original formulation, the number of passengers was set as an attribute of a flight arc. To allow for a $0, 1, \cdots, s$ passengers on a flight, there were $s + 1$ copies of each flight arc. As there is a binary decision variable on each arc for each aircraft, there are $(s+1)|K|$ decision variables for each pair of nodes with flight arcs.

A first improvement was to eliminate these copies of the flight arcs. Instead, the integer decision variable p_{ij} was introduced to choose to the number of passengers on the flight arc (i, j). Constraint (A.2) is added to ensure that the number of passengers is zero when the arc is not used. Note that the p decision variable is independent of which aircraft it applies to. This leaves $|K|$ binary decision variables and one integer decision variable for a pair of nodes with a flight arc.

$$p_{ij} \leq \sum_{k \in K} s \cdot x_{ijk} \qquad \forall (i, j) \in A_f \tag{A.2}$$

Integer decision variables remain difficult to handle and determining the exact number of passengers on each flight leaves symmetries in the problem, different solutions with the same outcome. Suppose there are 2 flights serving the same demand interval such that capacity exceeds demand ($2s > D_{((a_1,a_2),[t_1,t_2))}$). Note that costs are not affected by the number of empty seats on each flight. The solution where passengers are distributed $(s,\ D_{((a_1,a_2),[t_1,t_2))} - s)$ (the first flight taking s passengers while the second takes the remaining passengers) has the same cost as the solution where passengers are distributed the other way around

$(D_{((a_1,a_2),[t_1,t_2))} - s,\ s)$. However, the exact number of passengers on each arc does not have to be known. It suffices to know the number of empty seats ('loss') in a demand interval. Let $l_{((a_1,a_2),[t_1,t_2))}$ an integer decision variable, representing the number of empty seats in the demand interval. Instead of an integer decision variable on each flight arc, there now is only one integer decision variable for all flights in the same demand interval.

As a final improvement, the integer loss variable can be relaxed without loss of generality. As explained in Sec. 3 of the scientific paper, when a loss variable assumes a non-integer value, it can always be rounded down to the nearest integer without breaking any of the other constraints nor affecting the objective value.

B

Branch-and-Price

The method laid out in the scientific paper is able to find a local optimum to the problem as well as a lower bound. The lower bound indicates how far the local optimum is from the global optimum at most, but there is no way to prove exactly where this global optimum is located without computing it. As the arc based formulation was unsuccessful in solving the problem to optimality, this section describes a B&P (Branch-and-Price) algorithm to find the global optimum. We also explain why also this was unsuccessful.

The path based formulation laid out in Sec. 4 of the scientific paper is equivalent to the arc based formulation from Sec. 3. CG was used to solve the linear relaxation of the path based formulation to optimality. If the solution obtained by CG happens to be integral, then this is also the optimal solution to the integer version of the problem and thus the global optimum we are seeking. There is however no guarantee that this will be the case. While solving the integer version of the final RMP usually results in an improved integer solution, no claims can be made about local nor global optimality of this solution. The computational results show that LNS often results in an improvement over the integer RMP solution, proving that it is not necessarily locally/globally optimal. To find a global optimum, CG must be implemented in a B&B (Branch-and-Bound) framework where CG is solved at each node in the search tree, resulting in a B&P algorithm.

The underlying idea behind B&B is to divide and conquer. Because the original problem is too difficult to be solved directly, it is divided in smaller subproblems that can be conquered individually. The dividing is done by partitioning the feasible space in smaller and smaller subsets called branches. Conquering is done by bounding how good the solution of a branch can be. If the bound of a branch is worse than the incumbent solution (the currently best known integer solution), the branch is fathomed (discarded from further consideration). The three steps of branching, bounding and fathoming make up the B&B framework.

We start by solving the path based formulation without any binary restrictions using CG. Next, a binary decision variable is selected to branch on. This gives rise to two subproblems, one where the variable is 0, one where it is 1. This is enforced by setting Constraint (B.1) to either = 1 or = 0.

$$\sum_{r\in R} y_r \cdot \delta^r_{(i,j)} \leq 1 \qquad\qquad \forall (i,j) \in A_f \tag{B.1}$$

Each of the subproblems can be solved using CG. Consider one branch. Even if the solution obtained by CG is not integral, it provides a bound to how good any integer solution in this branch can be. A branch is fathomed when one of the following situations occurs:

- The bound is worse than the incumbent solution
- The branch has no feasible solution
- The solution of the branch is integer

After a branch is solved and the fathoming criteria are checked, the next variable to branch on is selected. This continuous branching results in a search tree as depicted in Fig. B.1.

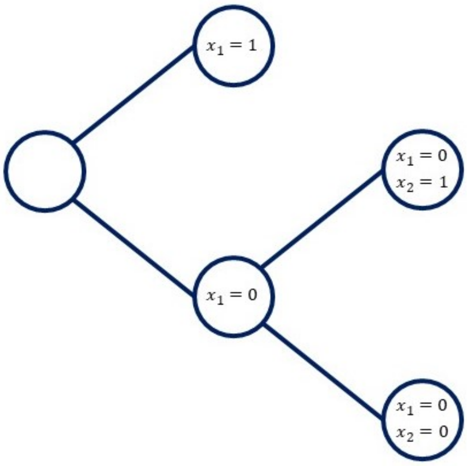

Figure B.1: Example of a branch-and-Bound search tree where x_1 is branched first, and x_2 is branched next on the $x_1 = 0$ side.

Before selecting which variable to branch one, we must first select which node in the search tree will be branched. In this paper, we use best-first search, meaning that we always branch on the node with the best bound (lowest) because this is the fastest way to get to the global optimum. If this global optimum is buried in this branch, the fastest way to it is by branching there. If it is not, then it must be in one of the other branches with a higher bound. So when the global optimum where to be found, we would still need to branch on the nodes with a lower bound to know for sure that we have indeed found the global optimum. Either way, the node with the lowest bound must be solved at some point.

Branching is done on the flight arc with a cumulative arc flow closest to 0.5 because this is keeping the solution from being integer. Additionally, branching where the arc flow is either 0 or 1 will result in the same solution and bound in one of the sub-branches.

In a best-case scenario, the incumbent solution is actually the global optimum. In that case, B&P needs to branch on all nodes in the search tree where the bound is lower than the incumbent solution, until these are all fathomed. The B&P was tested on the network with 5 airports at a very low market share. This instance was sufficiently small to find the global optimum using the arc based formulation and the global optimum was set as the incumbent solution. Even in this best-case scenario in a very small instance, B&P was unable to reduce the gap between the incumbent solution and the lowest bound in the search tree. The main reason for this was that setting an arc variable equal to 0 imposes little restrictions on the solution, which almost always finds a way around the prohibited arcs to obtain a solution that remains lower than the incumbent solution. In the branch where the variable was forced to 1, the solution was usually worse than the incumbent solution and the branch was fathomed. This causes the search tree to be very unbalanced and continuously branch on the 0-side without improving the lower bound.

C
Alternative Large Neighborhood Search with Time Windows

An alternative LNS (Large Neighborhood Search) based on time-windows was considered. Let $(t_1, t_2]$ be a time-window. let $A_{(t_1,t_2]}$ be the set of arcs that start in the time-window and let $A^C_{(t_1,t_2]}$ be its complement such that $A = A_{(t_1,t_2]} \cup A^C_{(t_1,t_2]}$. Let two different solutions be neighbors if the decision variables related to all arcs in $A^C_{(t_1,t_2]}$ are equal. So if $(i, j) \in A^C_{(t_1,t_2]}$ and it is included in one solution, then it must also be part in the neighboring solution. If it is not included in one solution, then it is also not present in its neighbor.

To find the best neighbor, a variation of the arc based formulation MILP is solved, using a small TS graph that is extracted from the full TS graph. Because all arcs starting outside the time-window are fixed, the nodes where a path must start and end inside the time-window are also fixed. The smaller TS graph is extracted by labelling all reachable arcs using i) A forward travelling recursive labelling algorithm, starting from each node where a path must start, and ii) A backward travelling recursive labelling algorithm, starting from each node where a path must finish. The result is an extracted graph such as the one depicted in Fig. C.1.

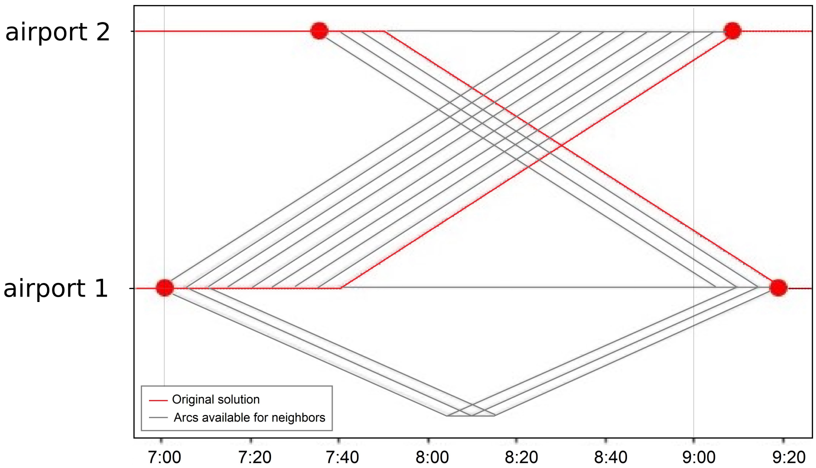

Figure C.1: Example of an extracted graph for the LNS with a time-window (7:00, 9:00]

This variation of an LNS was tested using Time-windows of 2,3 and 4 hours. However, because the solution outside of the time-window is already fixed, short windows of a couple of hours leave little room for improvement. When the time-windows are extended too much, the graph becomes too large and starts resembling the full-sized problem, which cannot be solved exactly in most cases. This variation of the LNS was therefore discarded and is not presented in the scientific paper.

D

Complete Computational Results

The algorithm was tested on various instances as explained in Sec. 5 of the scientific paper. The computational results of all instances are summarised in Tab. D.1. The first two columns indicate the number of airports in the network and the market share. Next, the runtime and objective value of the initial construction heuristic is given. The runtime of CG is shown in the next column, as well as the obtained lower bound, number of iterations performed and number of columns. The runtime and objective value of solving the integer RMP is displayed next as well as the percentage improvement compared to the initial solution. For LNS, the runtime, objective, percentage improvement compared to the integer RMP solution and number of iterations performed is given. Finally, the number of aircraft is the LNS solution is shown.

Table D.1: Computational Results

		construction		**Column Generation**				**Integer RMP**		**LNS**			
airports	**share**	**runtime**	**obj.**	**runtime**	**LB**	**#it.**	**col.**	**runtime**	**obj.**	**runtime**	**obj.**	**#it.**	**#ac.**
5	1%	0min 10s	1900	0min 1s	316	2	21	0min 0s	1810 (-4.7%)	0min 11s	754 (-58.3%)	2	1
5	2%	0min 8s	1900	0min 0s	633	2	21	0min 0s	1810 (-4.7%)	0min 16s	986 (-45.5%)	2	1
5	3%	0min 9s	1900	0min 0s	950	2	21	0min 0s	1810 (-4.7%)	0min 14s	1211 (-33.1%)	2	1
5	4%	0min 8s	1900	0min 0s	1267	2	21	0min 0s	1810 (-4.7%)	0min 28s	1465 (-19.0%)	2	1
5	5%	0min 6s	1900	0min 0s	1584	2	21	0min 0s	1810 (-4.7%)	0min 7s	1668 (-7.9%)	2	1
5	6%	0min 16s	3800	0min 0s	1900	2	22	0min 0s	3621 (-4.7%)	0min 34s	2464 (-32.0%)	3	2
5	7%	0min 11s	3800	0min 0s	2217	2	22	0min 0s	3621 (-4.7%)	0min 17s	2606 (-28.0%)	2	2
5	8%	0min 11s	3800	0min 0s	2534	2	22	0min 0s	3621 (-4.7%)	0min 42s	2879 (-20.5%)	2	2
5	9%	0min 10s	3800	0min 0s	2851	2	22	0min 0s	3621 (-4.7%)	0min 17s	3063 (-15.4%)	2	2
5	10%	0min 11s	3800	0min 0s	3168	3	43	0min 0s	3520 (-7.4%)	0min 21s	3336 (-5.2%)	2	2
5	11%	0min 11s	3800	0min 1s	3485	4	69	0min 0s	3621 (-4.7%)	1min 0s	3609 (-0.3%)	2	2
5	12%	0min 16s	5722	0min 0s	3801	2	26	0min 0s	5431 (-5.1%)	0min 50s	4274 (-21.3%)	3	3
5	13%	0min 15s	5722	0min 0s	4118	4	50	0min 0s	5320 (-7.0%)	0min 53s	4529 (-14.9%)	3	3
5	14%	0min 16s	5722	0min 1s	4435	4	49	0min 0s	5320 (-7.0%)	1min 7s	4731 (-11.1%)	3	3
5	15%	0min 16s	5722	0min 1s	4752	4	53	0min 0s	5320 (-7.0%)	1min 24s	5046 (-5.2%)	3	3
5	16%	0min 24s	5722	0min 1s	5069	4	67	0min 0s	5420 (-5.3%)	0min 30s	5229 (-3.5%)	2	3
5	17%	0min 27s	7432	0min 1s	5389	5	58	0min 0s	6939 (-6.6%)	2min 7s	5942 (-14.4%)	4	4
5	18%	0min 21s	7432	0min 1s	5712	5	70	0min 0s	7029 (-5.4%)	1min 15s	6174 (-12.2%)	3	4
5	19%	0min 20s	7432	0min 2s	6035	8	73	0min 0s	7119 (-4.2%)	1min 4s	6489 (-8.8%)	2	4
5	20%	0min 21s	7432	0min 2s	6360	8	92	0min 0s	7029 (-5.4%)	2min 36s	6714 (-4.5%)	4	4
5	21%	0min 21s	7432	0min 1s	6685	6	122	0min 0s	7130 (-4.1%)	1min 18s	6987 (-2.0%)	2	4
5	22%	0min 20s	7432	0min 1s	7009	6	114	0min 0s	7342 (-1.2%)	0min 31s	7342 (-0.0%)	1	4
5	23%	0min 25s	9141	0min 2s	7334	8	129	0min 1s	8638 (-5.5%)	3min 10s	7906 (-8.5%)	4	5
5	24%	0min 25s	9141	0min 2s	7660	9	129	0min 0s	8650 (-5.4%)	2min 2s	8109 (-6.3%)	4	5
5	25%	0min 26s	9141	0min 2s	7985	10	133	0min 0s	8738 (-4.4%)	2min 7s	8363 (-4.3%)	4	5
5	26%	0min 25s	9141	0min 2s	8311	9	157	0min 0s	8839 (-3.3%)	1min 43s	8786 (-0.6%)	2	5
5	27%	0min 31s	10982	0min 3s	8636	11	171	15min 20s	10258 (-6.6%)	2min 34s	9432 (-8.0%)	3	6
5	28%	0min 31s	10982	0min 3s	8962	12	202	1min 21s	10258 (-6.6%)	1min 49s	9616 (-6.3%)	3	6
5	29%	0min 32s	10982	0min 3s	9289	10	180	0min 1s	10407 (-5.2%)	4min 42s	9870 (-5.2%)	6	6
5	30%	0min 30s	10982	0min 2s	9616	9	171	0min 0s	10489 (-4.5%)	2min 18s	10162 (-3.1%)	3	6
5	31%	0min 31s	10982	0min 4s	9943	15	229	0min 0s	10549 (-3.9%)	1min 0s	10417 (-1.2%)	2	6
5	32%	0min 35s	12905	0min 3s	10271	12	171	3min 44s	12082 (-6.4%)	2min 47s	11142 (-7.8%)	4	7
5	33%	0min 35s	12905	0min 3s	10598	13	212	0min 51s	12026 (-6.8%)	1min 45s	11426 (-5.0%)	2	7
5	34%	0min 34s	12905	0min 3s	10925	12	197	0min 2s	12157 (-5.8%)	2min 30s	11640 (-4.3%)	4	7
5	35%	0min 37s	12905	0min 3s	11253	12	225	0min 0s	12265 (-5.0%)	4min 39s	11932 (-2.7%)	6	7
5	36%	0min 34s	12905	0min 3s	11581	13	218	0min 0s	12341 (-4.4%)	1min 9s	12216 (-1.0%)	2	7
5	37%	0min 38s	14574	0min 4s	11910	16	230	41min 2s	13723 (-5.8%)	1min 56s	13005 (-5.2%)	3	8
5	38%	0min 29s	14574	0min 2s	12240	10	178	0min 17s	13731 (-5.8%)	2min 38s	13248 (-3.5%)	3	8
5	39%	0min 38s	14574	0min 5s	12571	18	282	0min 7s	13749 (-5.7%)	2min 12s	13476 (-2.0%)	3	8
5	40%	0min 38s	14574	0min 4s	12903	14	250	0min 0s	13891 (-4.7%)	1min 21s	13682 (-1.5%)	2	8
5	41%	0min 44s	16182	0min 4s	13235	15	275	60min 0s	15323 (-5.3%)	1min 34s	14437 (-5.8%)	2	9
5	42%	0min 42s	16182	0min 4s	13568	13	206	1min 44s	15305 (-5.4%)	4min 20s	14853 (-3.0%)	5	9

Continued on next page

Table D.1 – *Continued from previous page*

		construction		Column Generation				Integer RMP		LNS			
airports	**share**	**runtime**	**obj.**	**runtime**	**LB**	**#it.**	**col.**	**runtime**	**obj.**	**runtime**	**obj.**	**#it.**	**#ac**
5	43%	0min 41s	16182	0min 5s	13902	20	258	0min 9s	15424 (-4.7%)	1min 59s	15174 (-1.6%)	3	9
5	44%	0min 47s	16182	0min 7s	14236	23	225	0min 1s	15573 (-3.8%)	2min 31s	15358 (-1.4%)	4	9
5	45%	0min 46s	16182	0min 7s	14573	23	298	0min 1s	16050 (-0.8%)	1min 33s	15867 (-1.1%)	2	9
5	46%	0min 59s	17962	0min 7s	14910	21	299	60min 0s	17070 (-5.0%)	3min 4s	16538 (-3.1%)	4	10
5	47%	0min 49s	17962	0min 7s	15247	22	298	0min 21s	17114 (-4.7%)	3min 26s	16949 (-1.0%)	4	10
5	48%	0min 49s	17962	0min 6s	15585	21	258	0min 3s	17313 (-3.6%)	2min 27s	17030 (-1.6%)	3	10
5	49%	0min 49s	17962	0min 7s	15923	22	305	0min 0s	17520 (-2.5%)	1min 21s	17345 (-1.0%)	2	10
5	50%	0min 54s	19771	0min 9s	16261	30	292	12min 52s	18633 (-5.8%)	3min 51s	18023 (-3.3%)	4	11
5	51%	0min 53s	19771	0min 9s	16600	29	285	0min 53s	18872 (-4.5%)	4min 25s	18397 (-2.5%)	5	11
5	52%	0min 55s	19771	0min 13s	16940	36	377	0min 10s	19072 (-3.5%)	2min 46s	18795 (-1.4%)	3	11
5	53%	0min 53s	19771	0min 10s	17280	30	357	0min 1s	19328 (-2.2%)	1min 33s	19220 (-0.6%)	2	11
5	54%	0min 57s	21619	0min 12s	17620	37	332	2min 57s	20331 (-6.0%)	5min 9s	19775 (-2.7%)	5	12
5	55%	0min 58s	21619	0min 12s	17961	34	386	0min 11s	20540 (-5.0%)	3min 38s	20181 (-1.7%)	4	12
5	56%	1min 7s	21619	0min 10s	18304	32	328	0min 1s	20774 (-3.9%)	2min 34s	20523 (-1.2%)	3	12
5	57%	1min 3s	23318	0min 13s	18648	37	374	0min 7s	21446 (-8.0%)	3min 0s	21300 (-0.7%)	3	12
5	58%	1min 4s	23318	0min 13s	19000	40	313	0min 15s	22007 (-5.6%)	4min 27s	21563 (-2.0%)	4	13
5	59%	1min 5s	23318	0min 11s	19353	34	322	0min 2s	22263 (-4.5%)	2min 54s	21985 (-1.2%)	3	13
5	60%	1min 3s	23318	0min 13s	19709	38	333	0min 2s	23039 (-1.2%)	2min 36s	22931 (-0.5%)	2	13
5	61%	1min 47s	25028	0min 19s	20064	42	334	1min 12s	23429 (-6.4%)	4min 7s	23179 (-1.1%)	3	14
5	62%	1min 21s	25028	0min 16s	20420	46	327	0min 3s	23706 (-5.3%)	3min 38s	23454 (-1.1%)	3	14
5	63%	1min 11s	25028	0min 15s	20776	47	362	0min 3s	23869 (-4.6%)	1min 50s	23695 (-0.7%)	2	14
5	64%	1min 14s	26634	0min 21s	21132	56	424	1min 19s	24982 (-6.2%)	3min 53s	24515 (-1.9%)	4	15
5	65%	1min 15s	26634	0min 18s	21489	59	299	0min 17s	25345 (-4.8%)	2min 22s	25112 (-0.9%)	2	15
5	66%	1min 24s	26634	0min 19s	21847	62	297	0min 2s	25619 (-3.8%)	3min 40s	25369 (-1.0%)	3	15
5	67%	1min 16s	28209	0min 21s	22206	68	317	8min 34s	26628 (-5.6%)	4min 44s	26185 (-1.7%)	4	16
5	68%	1min 26s	28209	0min 22s	22565	65	344	0min 18s	26763 (-5.1%)	4min 3s	26340 (-1.6%)	4	16
5	69%	1min 17s	28209	0min 26s	22924	69	400	0min 7s	27147 (-3.8%)	2min 34s	27039 (-0.4%)	2	16
5	70%	1min 22s	29818	0min 25s	23285	70	372	7min 35s	28310 (-5.1%)	3min 31s	27782 (-1.9%)	3	17
5	71%	1min 21s	29818	0min 31s	23647	83	387	0min 33s	28348 (-4.9%)	4min 12s	28133 (-0.8%)	4	17
5	72%	1min 27s	29818	0min 31s	24012	88	382	0min 3s	28611 (-4.0%)	2min 51s	28531 (-0.3%)	2	17
5	73%	1min 21s	31413	0min 34s	24380	95	394	1min 23s	29584 (-5.8%)	2min 58s	29501 (-0.3%)	2	18
5	74%	1min 21s	31413	0min 36s	24750	95	441	0min 26s	29952 (-4.7%)	3min 16s	29686 (-0.9%)	3	18
5	75%	1min 33s	33125	0min 43s	25123	110	413	0min 34s	30461 (-8.0%)	1min 8s	30461 (-0.0%)	1	18
5	76%	1min 28s	33125	0min 45s	25501	110	481	2min 27s	31093 (-6.1%)	2min 43s	30980 (-0.4%)	2	19
5	77%	1min 25s	33125	0min 48s	25880	109	427	0min 21s	31533 (-4.8%)	1min 4s	31533 (-0.0%)	1	19
5	78%	1min 28s	34922	0min 47s	26261	118	453	2min 48s	32497 (-6.9%)	2min 59s	32345 (-0.5%)	2	19
5	79%	1min 38s	34922	0min 50s	26643	123	476	0min 36s	32855 (-5.9%)	2min 15s	32854 (-0.0%)	2	20
5	80%	1min 33s	36365	1min 0s	27028	142	498	0min 36s	33294 (-8.4%)	2min 27s	33255 (-0.1%)	2	20
5	81%	1min 31s	36365	1min 24s	27413	153	522	4min 32s	34238 (-5.8%)	3min 52s	33882 (-1.0%)	3	21
5	82%	1min 31s	37719	1min 0s	27800	143	449	1min 16s	34828 (-7.7%)	3min 38s	34713 (-0.3%)	3	21
5	83%	1min 42s	37719	1min 5s	28190	151	499	10min 26s	35820 (-5.0%)	1min 6s	35820 (-0.0%)	1	22
5	84%	1min 33s	39271	1min 17s	28585	155	538	3min 28s	36226 (-7.8%)	3min 54s	35882 (-0.9%)	3	22
5	85%	1min 32s	39271	1min 11s	28987	168	520	5min 34s	37448 (-4.6%)	5min 5s	37252 (-0.5%)	3	23
5	86%	1min 36s	40879	1min 38s	29406	193	577	6min 46s	38052 (-6.9%)	2min 37s	37944 (-0.3%)	2	23
5	87%	1min 39s	40879	1min 21s	29832	179	558	8min 40s	39037 (-4.5%)	4min 48s	38612 (-1.1%)	3	24
5	88%	1min 35s	42483	1min 37s	30268	186	659	6min 27s	39466 (-7.1%)	4min 43s	39230 (-0.6%)	3	24
5	89%	1min 36s	42483	1min 34s	30732	190	654	60min 0s	40688 (-4.2%)	7min 49s	40408 (-0.7%)	5	25
5	90%	1min 39s	44243	1min 32s	31216	185	688	53min 38s	42081 (-4.9%)	4min 15s	41828 (-0.6%)	3	26
5	91%	1min 45s	45927	1min 48s	31721	195	800	36min 7s	42819 (-6.8%)	3min 16s	42629 (-0.4%)	2	26
5	92%	1min 43s	47692	1min 59s	32245	195	845	58min 53s	44121 (-7.5%)	1min 16s	44121 (-0.0%)	1	27
5	93%	4min 10s	50299	2min 43s	32773	177	824	60min 1s	45647 (-9.2%)	4min 19s	45647 (-0.0%)	1	28
5	94%	2min 47s	53463	2min 33s	33322	183	999	60min 1s	47645 (-10.9%)	6min 50s	47288 (-0.7%)	2	29
10	10%	2min 15s	5722	0min 3s	4345	4	64	0min 0s	5431 (-5.1%)	5min 4s	4690 (-13.6%)	2	3
10	20%	4min 22s	10982	0min 22s	8741	20	353	3min 18s	10569 (-3.8%)	11min 3s	9534 (-9.8%)	3	6
10	30%	5min 21s	16182	0min 50s	13235	35	619	31min 4s	15489 (-4.3%)	27min 14s	14485 (-6.5%)	4	9
10	40%	7min 19s	21656	1min 12s	17833	53	662	0min 27s	20732 (-4.3%)	19min 40s	20175 (-2.7%)	3	12
10	50%	8min 41s	28500	2min 13s	22624	101	654	0min 6s	27296 (-4.2%)	26min 6s	26729 (-2.1%)	3	16
10	60%	10min 48s	36986	5min 5s	27647	199	761	0min 10s	36181 (-2.2%)	28min 34s	35635 (-1.5%)	2	21
10	70%	13min 14s	48738	9min 6s	32953	287	1002	0min 12s	48217 (-1.1%)	42min 45s	47548 (-1.4%)	2	28
10	80%	14min 18s	68957	23min 52s	38746	437	1653	19min 41s	66792 (-3.1%)	134min 20s	65659 (-1.7%)	4	41
15	10%	4min 43s	7432	0min 8s	6173	7	126	0min 0s	7169 (-3.5%)	9min 21s	7026 (-2.0%)	2	4
15	20%	8min 35s	14576	1min 9s	12464	30	540	0min 6s	14051 (-3.6%)	25min 22s	13367 (-4.9%)	3	8
15	30%	14min 0s	23459	2min 26s	18940	54	791	26min 18s	22795 (-2.8%)	76min 2s	21416 (-6.0%)	4	13
15	40%	17min 55s	32022	5min 3s	25708	123	815	1min 15s	30759 (-3.9%)	52min 47s	29977 (-2.5%)	2	18
15	50%	20min 44s	42134	10min 11s	32875	234	883	0min 11s	40784 (-3.2%)	131min 16s	40307 (-1.2%)	4	24
15	60%	25min 37s	54274	17min 8s	40425	325	1141	0min 17s	53672 (-1.1%)	76min 49s	53566 (-0.2%)	2	31
15	70%	30min 46s	71964	28min 41s	48402	395	1566	0min 29s	71511 (-0.6%)	162min 59s	70938 (-0.8%)	3	42
15	80%	36min 22s	104973	65min 0s	56921	580	2361	1min 35s	104973 (–0.0%)	186min 22s	104476 (-0.5%)	2	64
20	10%	9min 31s	9349	0min 45s	7481	21	267	0min 2s	9030 (-3.4%)	41min 40s	8007 (-11.3%)	4	5
20	20%	16min 17s	18150	2min 5s	15160	40	682	1min 29s	17396 (-4.2%)	86min 19s	16681 (-4.1%)	4	10
20	30%	22min 30s	27034	5min 30s	23153	100	727	0min 4s	27034 (-0.0%)	68min 30s	26843 (-0.7%)	2	15
20	40%	29min 39s	38855	15min 33s	31673	239	927	0min 7s	38749 (-0.3%)	102min 24s	38598 (-0.4%)	2	22
20	50%	39min 43s	54091	28min 43s	40740	351	1417	0min 18s	53392 (-1.3%)	214min 57s	53069 (-0.6%)	3	31
20	60%	51min 4s	75400	60min 10s	50363	490	2032	0min 50s	74335 (-1.4%)	322min 42s	74029 (-0.4%)	3	44
20	70%	57min 31s	110873	144min 50s	60588	691	3263	24min 48s	110595 (-0.3%)	376min 33s	109879 (-0.6%)	2	68
30	10%	28min 0s	14582	6min 3s	12103	40	758	2min 55s	13989 (-4.1%)	126min 55s	13248 (-5.3%)	3	8
30	20%	51min 21s	28942	41min 11s	24899	187	1271	0min 20s	28735 (-0.7%)	295min 49s	28457 (-1.0%)	3	16
30	30%	80min 23s	48093	101min 47s	38942	404	1602	0min 28s	47680 (-0.9%)	472min 11s	47370 (-0.7%)	3	27
30	40%	109min 59s	72048	379min 42s	54135	742	3154	1min 0s	72048 (-0.0%)	283min 17s	72048 (-0.0%)	1	41
30	50%	108min 52s	106751	788min 36s	70723	1259	5977	3min 51s	105810 (-0.9%)	1744min 11s	105442 (-0.3%)	4	63

Bibliography

[1] Harish A. Perron C. Bavaro D. Ahuja J. Ozcan M. Justin C.Y. Briceno S.I. German B.J. and Mavris D. Economics of advanced thin-haul concepts and operations. *Journal of Aerospace Information Systems*, 2016.

[2] Air Transport Action Group. Facts & figures. https://www.atag.org/facts-figures, 2021.

[3] Justin C.Y. Payan A.P. Briceno S.I. and Marvis D.N. Operational and economic feasibility of electric thin haul transportation. *Journal of Aerospace Information Systems*, 2017.

[4] Sun X. Wandelt W. and Stumpf E. Competitiveness of on-demand air taxis regarding door-to-door travel time: A race through europe. *Transportation Research Part E: Logistics and Transportation Review*, 2018.