The Air Freight Flight Schedule Development Problem

AFFSDP

Anne Bart Beijneveld

TU Delft

The Air Freight Flight Schedule Development Problem

AFFSDP

by

Anne Bart Beijneveld

to obtain the degree of Master of Science
at the Delft University of Technology,
to be defended publicly on 10 December 2021.

Student number:	4390458	
Project duration:	November 2020 - December 2021	
Thesis committee:	Dr. ir. M.D. Pavel	TU Delft, Chair
	Dr. A. Bombelli	TU Delft, Supervisor
	Dr. ir. P.C. Roling	TU Delft, Examiner
	F. Gibellini	External, Supervisor

TU Delft

Acknowledgements

Dear Reader,

The submission of the report that lies in front of you resembles the final step for the completion of the Master Air Transport Operations at the faculty of Aerospace Engineering of the TU Delft. For this last project, the goal was to optimize the air transportation network of a full cargo airline and the flow of goods through it. Since I am fascinated by the air transportation industry and very enthusiastic about optimization problems, I believe I could not have gotten a more exciting thesis assignment.

Throughout the duration of my thesis I have been supported and guided by Alessandro. We have had numerous interesting discussions about networks and the mathematical meaning of certain constraints of my model. Although diving into mathematical expressions is not always easy when youre not in the same room, we managed to come to agreements on these topics. Finally, I appreciate that Alessandro was always open to explore the use of other ideas and concepts. I believe that this has strongly contributed to my joy in performing this research. Therefore: thank you very much Alessandro! Next, I would also like to thank Paul for joining all the milestone meetings and providing new insights that I was able to use in my research. Also, a special thanks to Marilena for chairing my thesis committee.

Furthermore, I got the opportunity to do this work in cooperation with a large full cargo airline. Also from this airline I have been provided with full support to perform this research. Francesco was my daily supervisor from the airline. He was always ready to help me out and thinking along with my problems at any time. But I am also deeply thankful for the many reminder of him that I should not over complicate my live. Thank you for your continuous positive support throughout this thesis Francesco. You have greatly helped me to further improve this research!

Additionally, I would like to thank Koen and Bert for answering the countless questions I had about the operational side of the air freight industry. If they knew the answer, I would be provided with a detailed explanation of it. Otherwise, they would bring me in contact with someone who would. These meetings allowed me to gain very valuable insights for my research. Moreover, they formed a source of inspiration for potential solutions.

Finally, I would very much like to thank my family, friends and boyfriend for their endless support. Explaining what I was working on to people that were not directly involved, helped me to think of different ways to approach the problems at hand. And apart from supporting my thesis, all of them provided many amusing and entertaining distractions that were very welcome in a period where staying inside at home was mostly the norm.

Again, thank you all very much!

Anne Bart
Brussel, November 2021

Contents

List of Tables

Introduction

In the past few decades, technological developments in the air transportation industry have made air transport cheaper, faster and safer. Thereby, it has become indispensable for the global economy. Commercial air transport has been severely impacted by the outbreak of the COVID-19 pandemic. The counterpart of the passenger transportation industry, namely the air freight industry, was affected less dramatically. It even played a vital role by accommodating the transport of medicines and medical supplies in a timely fashion. Moreover, as all physical shops were closed, online shopping became even more popular. This further boosted the size of trade flows around the world.

Designing a schedule that is both feasible and profitable is a very complex task. Especially, the high unpredictability of air freight makes this difficult. Although the average demand for air freight is expected to increase, this does not mean that it will not fluctuate throughout the year. As a result cargo airlines continuously evaluate their network. Will it be able to accommodate all demand or will many aircraft fly partially if nothing is changed?

How a schedule should be designed has been a frequently recurring topic of research. On the contrary, how such a schedule should be adjusted, based on a change in predicted demand, has been given little attention in the literature. This research aims to develop a decision-making tool that can aid full-cargo airlines in deciding on how to expand or decrease their air transportation network most effectively. Expanding or decreasing the size of the network in this sense means that existing flight rotations could be altered to include or exclude certain stopovers such that the network can transport all forecast demand most cost-effectively. The research was performed in co-operation with a large cargo airline. The airline did not only provide data but also gave many valuable insights into the world of air freight. The research had promising results that showed that by giving the designed model a certain level of flexibility it was able to transport more demand for the same level of costs. It is expected that the outcomes can be very useful for airlines like the one used for this research.

This research was performed as a graduation thesis for the master track Air Transport Operations of the Aerospace Engineering faculty of the TU Delft. This master focuses on optimizing the aviation industry, which has as main target points to improve the efficiency, safety, cost and environmental impact caused by this industry. The authors believe to have contributed to this goal by designing a tool that can help cargo airlines improve the efficiency of their networks. As a result, this minimizes the impact of the air transport industry.

The report that lies in front of you comprises three main parts. In the first part, the scientific article can be found. This article contains all the key findings of this research and an extensive explanation of the steps taken to achieve the final results. In the following section, the literature study, that was conducted to start this research, can be found. This study contains background information and an elaborate discussion of the state of the art research currently available. Moreover, it addresses the key research questions that guided this research. Finally, in Section III, the supplementary work can be found. This section informs the reader on the verification and validation process of this research and gives a more detailed overview of certain results.

I

Scientific Paper

The Air Freight Flight Schedule Development Problem

Anne Bart Beijneveld*

Delft University of Technology, Delft, The Netherlands

Abstract

In this paper, we introduce the Air Freight Flight Schedule Development Problem (AFFSDP). In this problem, a cargo-only carrier has to adjust an existing flight schedule based on changes in the predicted demand. The result of this research is a model that combines airport selection, fleet routing and cargo routing, together with the use of a (random) mandatory flight list and timetable setting for all other optional flights. The model is formulated as a Mixed Integer Linear Program (MILP). To reduce the pre-processing time of the model two meta-heuristics have been implemented that reduced the run time from several hours to minutes. To improve the performance of the model, all fractional numbers have been reshaped to integer numbers. The symmetry of the model was brought to a minimum by implementing a path-preference model. The model was tested with 3 aircraft types, for 3 different demand scenarios and 3 different freedoms. This sums up to 27 unique test cases. Overall, the results show that against a benchmark solution the total serviced demand could be increased by 20-50%.

Keywords: *Flight schedule development, disruption management, air freight transportation, schedule planning, path-flow model, matheuristic.*

1 Introduction

In the past few decades, technological developments in the air transportation industry have made air transport cheaper, faster, and safer. Thereby, it has become indispensable for the global economy. Boeing has predicted an average growth rate of 4.2% for the air freight industry due to the world trade growth [Boeing, 2018]. Although this prediction was made before the outbreak of the COVID-19 pandemic, it is expected that the fast-growing e-commerce market remains a key driver for the growth of the air freight industry [ReportLinker, 2021].

Two types of carriers are mainly responsible for the transport of air freight demand: combination carriers and cargo-only carriers. Combination carriers combine the transport of passengers with that of cargo. Cargo-only carriers focus their attention on cargo only. Thereby, their network is explicitly designed for air freight demand. Designing a profitable flight schedule for cargo-only carriers is a difficult task that requires the input of experienced network planners. In general, trade imbalances make it hard to create an efficient network. Therefore, for every airline, the flight schedule design is the central element of the planning process. To start this process, four main questions have to be answered: which markets should be connected by direct or indirect flights, how frequently should these markets be connected, how much capacity should be offered on these routes and which of the available aircraft will fly these routes [Derigs et al., 2009]?

These planning decisions have to be made several months in advance based on the then-predicted demand. However, accurately predicting this demand is complex. This is the result of a few factors that play an important role in this industry. First of all, the capacity of aircraft is limited either by volume or by weight, whichever becomes limiting first. This is highly dependable on the type of commodities that are being shipped. As a result, the limiting factor is often route-specific. Furthermore, an estimation is made of the weight and volume of the goods to be shipped when the booking is made, but this often differs from the actual size on the day of transport [Huang and Lu, 2015]. In addition to this, cargo bookings are often placed relatively close to the day of the actual flight, making it challenging to ensure that enough capacity is available on each route [Sandhu and Klabjan, 2006]. Finally, it also frequently occurs that bookings do not show up at all [Amaruchkul et al., 2007]. Nonetheless, thanks to sophisticated forecasting models, a reasonable estimate can be made several months in advance.

Usually, the forecast becomes more accurate as the day of operation becomes closer. Therefore, it is constantly verified whether the current schedule would be able to accommodate the expected demand or if too much capacity is offered. If this is not the case, then the flight schedule has to be adjusted accordingly. In many airlines, this process is performed manually by experienced planners. For some smaller networks, it might

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

be possible to find an optimal solution manually. However, for larger networks, the solution space may become rather large. In this case, the chances of manually determining the most cost-effective solution are minimal. Therefore, the aim of this paper is to solve this rescheduling problem. Formally, we refer to this problem as the *Air Freight Flight Schedule Development Problem* (AFFSDP). The term "development" is used, because an existing schedule is developed into an improved schedule. The objective of this research is to design a tool that can find a solution to the AFFSDP. The output of this model should provide insights into potential solutions to accommodate the re-predicted demand most cost-effectively. It was found that the previous studies that focused on flight schedule development frequently used a predefined master flight list. This list consisted of mandatory and optional flights. As a result, the quality of the solution would be highly dependent on the quality of this master list. In this paper, this dependency is brought to a minimum as the designed model only requires the mandatory flight list as an input. All other flights are optional and the model will determine the most optimal set of flights. To the best of our knowledge this is the first paper that combines airport selection, fleet routing and cargo routing, together with the use of a (random) mandatory flight list and timetable setting for all other optional flights.

This paper is structured as follows, in section 2 an overview of the academic literature concerning this field of research that was available at the time of writing is provided. Hereafter, in section 3, a detailed description of the research problem is given. In section 4 the methodology used to find a solution to the problem at hand is elaborated on. This is followed by section 5 which describes the instances that will be used to test the effectiveness of the model. In section 6 an in-depth analysis of all the obtained results is performed. Finally, in section 7 concluding remarks of the research and ideas for future research are discussed.

2 Literature Review

Designing a feasible and profitable flight schedule is a very complex task. Operational Research (OR) professionals have already been working on developing tools and methods to design optimal schedules since the 1950s. The design of the complete air freight transportation schedule can be divided into four major interdependent problems: *schedule planning*, *fleet assignment*, *rotation planning*, and *cargo routing*. Many researchers have developed models and methods to solve one or multiple of these sub-problems. This section will elaborate on the objectives and methods used in several of these studies. It should be noted that the planning horizon of this study is several weeks. This means that rotation planning is not yet important. Therefore, no attention is paid to this part of the design process.

2.1 Schedule Design

Most of the time, the schedule design is not started from scratch ([Lohatepanont and Barnhart, 2004], [Gopalakrishnan and Johnson, 2005], [Derigs et al., 2009]). Instead, an existing schedule is used as a basis, and changes are made, reflecting the difference in forecast demand. This is also known as *schedule development*. According to [Lohatepanont and Barnhart, 2004], there are a few reasons this is the industry's practice. First of all, it is operationally impractical and computationally tricky to build a schedule from scratch. Second, changing a network could require significant investments in infrastructure at certain airports. Finally, reliability and consistency are important to the customers of an airline. By using an existing schedule, the consistency of the network can more easily be retained. Furthermore, it makes the effort required for network planners to complete this step of the schedule design tractable.

[Derigs et al., 2009] referred to this approach as the *the pragmatic planning paradigm*. Their designed model uses a predefined master flight list that contains mandatory and optional flights. This model maximizes network-wide profit by simultaneously optimizing the selection of flight legs, fleet rotation, and cargo routing based on inputs such as the predefined master flight list and a forecast origin & destination-matrix or in short: O&D-matrix. This research was further extended by [Derigs and Friederichs, 2013]. Their optimization problem's objective was altered to a minimization problem, namely the minimization of the total operating cost. Moreover, the fleet assignment was added as a sub-problem to be solved by the system. The goal of their research was to give a proof of concept of the pragmatic planning approach. With this intention, they performed extensive tests on data that reflected the different types of cargo airlines. Finally, they were able to prove the computational tractability and effectiveness of their approach.

2.2 Fleet Assignment

Once the main schedule design is done, the fleet assignment problem can be considered. In literature, two classical methods are most frequently applied to model the flight network, namely the connection network (CN) and the time-space network (TSN). One of the advantages of the CN is that it can capture the flow of individual aircraft. This is for example useful if one is interested in ensuring equal wear and tear among aircraft in their fleet. Moreover, it can be used to solve the rotation planning problem. However, compared to the TSN the CN requires more nodes and arcs to model the network. Therefore, if the network size would increase, the CN would grow faster in size than the TSN would. As a consequence, finding a solution to an optimization problem using the CN would demand more time. As mentioned before, the aircraft rotation does not yet have to be considered. Therefore, it was decided to use the TSN for this research. For a more in-depth analysis of the differences between the CN and TSN, the reader is referred to [Zhou et al., 2020].

The TSN was first introduced by [Berge and Hopperstad, 1993]. This network is also known as the activity-on-edge network because the arcs represent actual movements in the network. In the past, this network representation has been used for various reasons. It has been used to model a network with varying departure times to determine the most profitable combination of departure times for each flight leg and the corresponding fleet assignment (see, e.g., [Rexing et al., 2000] or [Bélanger et al., 2006]). Moreover, the TSN has also been used to integrate the fleet assignment problem with timetable setting. In the research of [Yan et al., 2005] and [Tang et al., 2008] a case-study for a Taiwan airline was performed. In both pieces of research airport selection, fleet routing and timetable setting were combined. A family of heuristics was applied to solve the model. Especially the set-up of the TSN for these researches formed a source of inspiration for this research. This will become clear once the network representation of this study is explained (in section 4). Finally, it should be noted that although most of this literature is focused on passenger transportation, it is expected that these models can be easily modified to apply to cargo transportation networks as well.

2.3 Cargo Routing

Cargo routing is an essential part of the flight schedule design. [Derigs and Friederichs, 2013] explained this as follows: a cargo airline offers the conceptually simple service of offering insurance of timely delivery of goods for a certain price. In general, an airline forecasts how much demand it will have to transport within a certain period and how much revenue this will approximately generate. After that, it is up to the airline to transport this demand as cost-effectively through the network as possible, as this, in turn, maximizes the profit.

The cargo routing problem is often formulated as a multicommodity network flow problem (MNFP). In general, there are three methods to decompose the MNFP. Namely the node-arc formulation, the tree formulation and the path-flow formulation. [Jones et al., 1993] showed that when the MNFP was solved using the path formulation, substantially fewer master problem iterations were required than when the other two methods were used. Therefore, the use of the path-formulation will be further investigated. [Li et al., 2006] used the cargo routing problem to more accurately model the fleet assignment. To retain a tractable model, a two-step modelling approach was applied. First, all feasible paths that satisfied the aforementioned constraints were generated for all commodities. Thereafter, the MNFP was formulated with only the columns of these feasible paths. This two-step formulation is an intelligent approach to reduce the problem's size. Moreover, also [Derigs and Friederichs, 2013] emphasized the advantage of the path-flow model as it provides the opportunity to incorporate practical constraints on an itinerary's feasibility before the optimization model has to run.

2.4 Flight Schedule Recovery

It is impossible to completely avoid disruptions from happening. In the case of passenger transport, most disruptions occur on the supply side. This means that, for example, an aircraft has a mechanical failure or members of the crew are sick. In case of such disruptions, the schedule needs to be recovered as quickly as possible. Thereby, the focus is on minimizing the negative impact on operational costs and passengers. The loss of customers due to negative experiences with rerouting can have severe consequences for an airline. Therefore, it is important that the amount of change required to absorb the disruptions in the network is brought to a minimum. On the other hand, in the case of cargo transport, disruptions occur both on the supply and the demand side. The latter is a direct consequence of the unpredictability of demand for air cargo. To the best of our knowledge, only one study has focused its attention on how a cargo carrier can best recover from demand disruptions. In this study, [Delgado et al., 2020] introduced the Air Cargo Schedule Recovery Problem (ACSRP).

Their study aimed to develop a model that could redesign an operational flight schedule in reaction to certain demand disruptions. The disruptions they considered occurred maximal 3 days, or 72 hours, before the actual time of flight. The ACSRP aims to minimize the total operating cost while also considering a penalty for deviating from the original schedule. It was found that their model, when compared to a benchmark solution where cargo could only be re-routed, was able to achieve cost savings of about 10%. Although their research focuses on a different planning horizon and considers the cost of crew rescheduling, which does not have to be taken into account for this thesis, their model is a source of inspiration for the to-be-built model.

It can be seen that there already is a variety of research available. Several aspects of the aforementioned literature are combined for this research.

3 Problem Statement

3.1 Research Objective

The size of trade flows around the world is constantly prone to change. It is safe to say that accurately forecasting how much demand there will be for all the possible O&D pairs worldwide is impossible. Nonetheless, thanks to sophisticated forecasting models a reasonable estimate can be made a few weeks in advance. However, a flight schedule has to be made several months in advance. Therefore, when the schedule is made, many things can still happen that influence the forecast demand. Up and until now, to cope with changes in the forecast demand, the flight schedules of airlines have been adjusted manually by experienced planners. There are three main reasons why a model might be able to improve the quality of the solution that is implemented in the end. First, it is difficult to manually quantify the cost of several solutions. Therefore, it is even more difficult to determine which solution is the most cost-effective one. Second, the size of the air freight network is generally growing. Consequently, the solution space will grow even faster, making it inevitable that the optimal solution cannot be found manually. Third, some less obvious, but potentially cost-effective, solutions might exists which are overlooked by the planners.

Usually, an airline designs a flight schedule that can be repeated every week. This means that aircraft in general will start and end their rotation at the same location in the network. If the network has to be changed, this should be kept in mind as this will allow for a smooth transition between the existing schedule and the changed schedule. Therefore, the input for this research will be an existing flight schedule for one week that starts on Monday and ends on Sunday. The output schedule should comply with the start and end location of the aircraft in the existing schedule. Moreover, it is extremely difficult to accurately predict how much revenue will be generated by transporting a certain amount of demand. Therefore, this problem will be formulated as a cost minimization problem. The objective of the model is to determine a set of flights that can transport as much of the predicted demand as possible.

It should be noted that this research is conducted in cooperation with a large airline. As this airline will remain anonymous, we will from now on refer to this airline as “our airline”. The demand of our airline either needs to be delivered overnight or it can be delivered within several days. The first type of demand is also referred to as “express demand”. Our research is focused on this so-called express demand because it is especially challenging to design a profitable flight schedule for demand with tight delivery constraints.

3.2 Problem Setting

This section describes what the expected input and output of the model are. Furthermore, it will elaborate on the important assumptions and simplifications that were made to design the model. Finally, it will also highlight the operational constraints that are taken into account.

Input data

A summary of the main input data is provided below:

- The flight schedule as it is currently planned, from now on referred to as the base schedule. The user may decide how much of this schedule should remain the same.
- The potential O&D demand. The demand is divided into different requests. Each request represents an origin, a destination, a time from which it is available for pickup and a time before when it should be delivered and finally also the size of the request expressed in kgs. (The volume of a request is not taken into account for this research)
- The fleet and its respective characteristics such as the number of available aircraft per aircraft type, the respective maximum capacity, turnaround time, estimated usage cost, etc.
- Per aircraft type, a list of flight legs it is allowed to fly and the respective flight times.
- List of airports in the network and their respective specifics, such as the local timezone and opening and closing hours.
- Per O&D pair a list of maximum 20 shortest paths that can be used to transport the package from its origin to its destination. (This will be further explained in section 2.3)

Output data

A summary of the main output data is provided below:

- A flight schedule, showing the expected load factor and which requests are (partially) transported on each flight.
- An overview of how many requests are fully, partially or not at all transported and the respective combined weight.
- An overview of the aircraft operational details, such as the tot flight time, the total number of flight legs flown, total payload transported, estimated usage cost and fuel burn.

- An overview of certain Key Performance Indicators (KPIs) such as the total aircraft cost, fuel cost, total demand transported and the average cost per transported kg of demand.

Assumptions

There are several assumptions made for this research, which are expected to be in line with the industry practice. First of all, it is assumed that the same number of aircraft is available as are used in the base schedule. Furthermore, their starting and ending positions are assumed to be the same as provided in the current schedule. Generally speaking, an aircraft will start and end at the same airport, but exceptions may occur. Moreover, since rotation planning is not yet important, it is assumed that tail swapping is acceptable. This means that it only matters that an aircraft of a certain type starts and ends its rotation at a specific airport, but that it does not matter which aircraft of the available aircraft ends there. Furthermore, the cost components that are taken into account are the aircraft usage and the fuel burn cost. It is assumed that these costs give an acceptable indication of what the operational cost of the network will be. The aircraft usage cost is calculated by using the Aircraft Crew Maintenance and Insurance (ACMI) cost. It is assumed that these costs will cover the largest aircraft usage cost. The cost of keeping an aircraft on the ground was omitted since too little information was available about the cost of keeping an aircraft on the ground at different airports. The fuel costs can be decomposed into two main parts: the dry operating fuel, which is the fuel required to fly a certain leg with a certain aircraft, and the marginal extra fuel, which is proportional to the amount of payload brought on the aircraft. The fuel burn calculations which are made for each aircraft type will later on, in section 4.2, be explained in more detail. Moreover, it is assumed that unloading and loading of cargo can happen within the provided turnaround times. In addition, cargo is also allowed to be transloaded at any point in the network. This means that transloading is not restricted to the hubs of the network. Finally, it should be noted that each request may be transported over several paths in any way deemed best by the model.

Constraints

The model is subject to several constraints. These constraints ensure that the model returns a realistic solution that could be implemented. They can be divided into two main categories, aircraft related constraints and cargo related constraints. The first constraint is the conservation of aircraft. This constraint ensures that the number of aircraft arriving at an airport is equal to the number of aircraft departing that same airport. The second is the aircraft capacity constraint. This ensures that only less than (or an equal amount of) demand than the maximum available capacity is assigned to a flight. Thirdly is the cover constraint. As mentioned before certain flights have to be flown, this constraint ensures that only one aircraft is assigned to this particular flight. Fourth is the fleet size constraint. The number of aircraft of any type used in the solution may not exceed the number of available aircraft of that same type. Finally, there is the access and egress constraint. This constraint secures that the solution has the same number of aircraft starting and ending at different airports as was stated in the base schedule.

The constraints related to the cargo routing are described as follows: similar to aircraft, there is a conservation of cargo constraint. This ensures that all cargo is transported either via paths in the actual network or via a so-called no-service path. The specifics of the network layout will be explained in section 4.3. The second constraint is the cargo path flow demand constraint. This constraint ensures that the flow on all of the available paths for a request does not exceed the maximum demand of this request. The last cargo-related constrain is the feasible path constraint. This constraint ensures that a path can only be chosen by the model if all the flights in it are assigned to an aircraft.

4 Methodology

4.1 Demand Determination

One of the main inputs of the model is the predicted demand. It is out of the scope of this research to try to make a forecast of future demand. Especially due to the COVID-19 pandemic, predicting cargo demand has become even more difficult. Luckily, our airline was able to provide data on historic movements of cargo through their network several years ago. Furthermore, the airline also provided the base schedule for the respective period of data. This data was used to reconstruct an estimate of the O&D demand. As mentioned earlier, the focus of this research is on overnight delivery. To estimate when demand is generally available and when it is due, our airline provided a list of the airports in the network, which were either classified as “Primary” or a “Secondary” airports. For primary airports, the latest time of pickup is later in the evening when compared with the secondary airports. Similarly, the delivery time for primary airports is earlier in the morning than the delivery time for secondary airports. Although this is also a simplification of reality, it gives an appropriate distinction between the two types of airports. Secondary airports are generally the outer gateways of the network, which have less stringent timing constraints.

Given one week of cargo movements, the data was aggregated as follows: For each day of the week, it was determined how much cargo was transported between a certain O&D pair. Based on the classification of the origin and destination airport the *local* time when demand was available and before when it was to be delivered was determined. Hereafter, the local time was converted to *Coordinateed Universal Time*, or UTC. This was done such that the model used a normalized version of time for the delivery constraints. Similarly, the base schedule was converted to UTC as well. For each day of the week and for each O&D pair with a given demand for that day, a request was created. In Table 1 an example of three requests between Amsterdam and Madrid is shown. Please note, this table serves just as an example and does not represent actual data from our airline.

Table 1: Example of input demand between Amsterdam and Madrid. The column header "t av loc" and "t due loc" respectively mean the time when a request is available and when it is due. "av day" and "due day" are the day of the week the request becomes available and on which it has to be delivered. "t av UTC" and "'t due UTC" respectively mean the time when a request is available and when it is due defined in UTC. Note, this does not represent actual data from our airline.

req	**ORG**	**DST**	**t av loc**	**t due loc**	**av day**	**due day**	**t av UTC**	**t due UTC**	**weight [kg]**
1	AMS	MAD	2018-09-10 21:00	2018-09-11 7:00	Mon	Tue	2018-09-10 19:00	2018-09-11 5:00	1,000
2	AMS	MAD	2018-09-11 21:00	2018-09-12 7:00	Tue	Wed	2018-09-11 19:00	2018-09-12 5:00	1,500
3	AMS	MAD	2018-09-12 21:00	2018-09-13 7:00	Wed	Thu	2018-09-12 19:00	2018-09-13 5:00	1,250

4.2 Aircraft Performance Computations

As mentioned before, an important part of the costs that are taken into account are those associated with fuel. Fuel burn of an aircraft is dependent on the characteristics of the aircraft itself, but also on the operational conditions such as head or tailwind. Taking into account the influence of the latter is out of the scope of this research. However, a reasonable estimate of the expected fuel consumption can be made if several aircraft characteristics are known. Before these calculations are explained, it is useful to understand what a general payload-range diagram looks like. This is depicted in Figure 1. In case this diagram would have been drawn-up for a passenger or combi-flight, the "Mission Payload" would consist out of weight for passengers and weight for cargo.

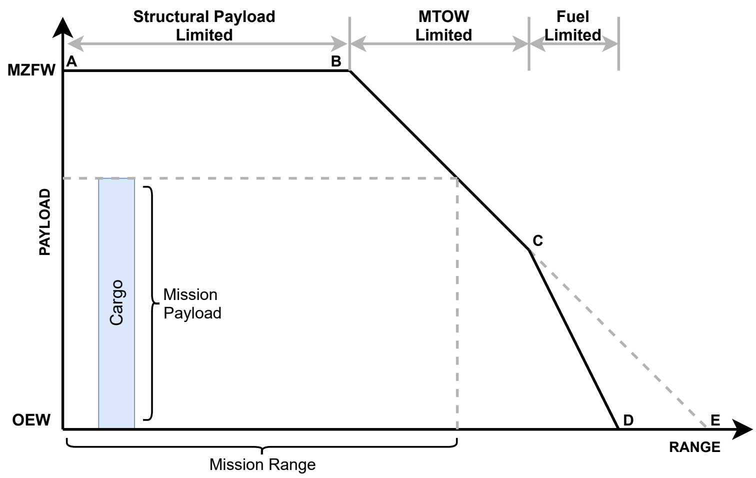

Figure 1: A schematic representation of a typical aircraft payload-range diagram. (Source: adapted from [Baxter et al., 2018].)

The x-axis of the figure represents the range, while the y-axis represents the payload. At the origin of the figure, one can see the OEW, which represents the operative empty weight (W_0) of the aircraft, which is the minimum take-off weight without any fuel. Point A represents the maximum zero fuel weight (MZFW) which is a theoretical point, in which the aircraft is loaded until maximum capacity, without any fuel. Going from point A to point B, one can add fuel without compensating for the amount of payload that can be brought. As at this point, the aircraft is limited by the maximum structural payload. At point B, the aircraft has reached the maximum take-off weight (W_{MTOW}). As can be seen in Equation 1 the MTOW consists of the maximum payload weight ($W_{p_{max}}$), the W_0 and the fuel required to fly the maximum range at maximum payload. This range is denoted by $R_{p_{max}}$, while the weight of the fuel is denoted by $W_{f_{MTOW}}$. When traversing from point B to point C, one has to sacrifice payload to bring enough fuel to achieve the extended range. Then, if one goes from point C to D, the range can be further enlarged. However, then also fuel is required to carry fuel, which

Table 2: Aircraft characteristics.

Aircraft Name	Abbr.	MTOW [kg]	OEW [kg]	MZFW [kg]	Max struc. PL [kg]	Max. R at max. PL [km]	TAT [hh:mm]
Boeing 737-400	73P	67,300	32,800	52,800	20,000	2,500	00:45
Boeing 757-200	75C	109,500	53,200	88,500	35,400	3,700	01:00
Airbus A300-600	ABY	170,500	83,000	130,000	47,000	4,800	01:15

explains the steeper slope of this part of the graph. For this research, this part of the graph is not taken into account. Instead, the line is treated as completely linear, which is represented by the segment between points C and E. Finally, the grey dotted lines provide an example. Given a certain amount of payload, the maximum mission range can be determined.

$$W_{MTOW} = W_{p_{max}} + W_0 + W_{f_{MTOW}} \tag{1}$$

The following computations have been inspired by [Wink, 2020]. However, other data on aircraft characteristics was available for this research. Therefore, the work has been adjusted to fit the needs of this research. The payload-range diagram, as shown above, can be derived from Breguet's range equation, which is shown in Equation 2. This equation represents the relation between the aircraft take-off weight and the range it will fly. It does so, by integrating the initial and final weight of the aircraft during the cruise phase.

$$R = \frac{VC_L}{C_D C_T} \ln\left(\frac{W_{st}}{W_{end}}\right) = \frac{VC_L}{C_D C_T} \ln\left(\frac{W_0 + W_p + W_f}{W_0 + W_p}\right) \tag{2}$$

In this equation R is the range, V is the cruise speed, C_L is the lift coefficient, C_D the drag coefficient, C_T the thrust coefficient, W_{st} and W_{end} are respectively the initial and final weight of the aircraft during cruise, W_p the weight of the payload and W_f the weight of the fuel. The values of the coefficients C_L, C_D and C_T are often not available to the general public. Moreover, they are also dependent on the flight operations, making it hard to estimate these values. However, by using point B of the payload-range diagram, the equation can be rewritten to variables of which information is available. As explained before, point B represent the W_{MTOW} and $R_{p_{max}}$. Rewriting the previous equation allows us to determine a constant C as shown in Equation 3.

$$C = \frac{VC_L}{C_D C_T} = \frac{R}{\ln\left(\frac{W_0+W_p+W_f}{W_0+W_p}\right)} = \frac{R_{p_{max}}}{\ln\left(\frac{W_{MTOW}}{W_0+W_p}\right)} \tag{3}$$

This constant is used for all the following fuel consumption calculations. There are three typical freighter jets used in this research. The data that is required for these calculations can be found in Table 2. From left to right the columns denote the following: aircraft name, abbreviation of the aircraft name, maximum take off weight, operating empty weight, maximum zero fuel weight, maximal structural payload, maximum range at maximum payload, turnaround time. By using Equation 1 and Equation 3, Equation 2 can be rewritten to Equation 4 as shown below.

$$R = \frac{VC_L}{C_D C_T} \ln\left(\frac{W_{MTOW}}{W_{MTOW} - W_{f_{MTOW}}}\right) = C \ln\left(\frac{W_{MTOW}}{W_{MTOW} - W_{f_{MTOW}}}\right) \tag{4}$$

The W_{MTOW} and $R_{p_{max}}$ are known, which means that this equation can be rewritten to the one below.

$$W_{f_{MTOW}} = W_{MTOW}\left(1 - \frac{1}{e^{\frac{R_{p_{max}}}{C}}}\right) = W_{MTOW}\left(1 - e^{-\frac{R_{p_{max}}}{C}}\right) \tag{5}$$

By filling in the values of the known variables, the weight required to fly $R_{p_{max}}$ can be calculated. Moreover, by assuming that the aircraft operates at W_{MTOW}, it can be seen that W_f is only dependent on the range. So, by using this information, the $W_{f_{MTOW}}$ can be calculated for a variety of ranges. As a result one can also calculate the theoretical W_p for these ranges by rewriting Equation 1 to Equation 6.

$$W_p = W_{MTOW} - W_0 - W_{f_{MTOW}} \tag{6}$$

If we now return to the payload-range diagram, it can be seen that in between points A and B no matter what range the aircraft is flying, the maximum structural payload cannot be exceeded. So if the range is smaller than $R_{p_{max}}$, this also means that less fuel is required to fly the aircraft. So, by rewriting Equation 2 and by using maximum payload for W_p, one can determine how much is required for the reduced range with Equation 7.

$$W_f = (W_0 + W_p)\left(e^{\frac{R}{C}} - 1\right) \tag{7}$$

Moreover, an aircraft does not always fly the maximum range, nor does it always fly with the maximum allowed payload. To determine how much fuel is required for a certain range given a certain payload, it is first determined what the required dry operating fuel is ($W_{f_{oew}}$). This represents the amount of fuel that is required to fly a certain range without payload. This is easily determined by assuming $W_p = 0$ in Equation 7. As can be seen, this value is again only dependent on range. The next step is to determine the marginal extra fuel required to bring any payload. This can be calculated by taking the derivative of Equation 7 with respect to W_p. This leads to Equation 8 which is shown below.

$$\frac{\partial W_f}{\partial W_p} = \frac{1}{\partial W_p}\left((W_0 + W_p)\left(e^{\frac{R}{C}} - 1\right)\right) = \left(e^{\frac{R}{C}} - 1\right) \tag{8}$$

It can be seen that the equation to calculate the dry operating fuel and the equation to calculate the marginal extra fuel are only dependent on the range. All of the aforementioned calculations were used to determine the payload range diagrams for the aircraft that are used in this research. Finally, it is also of great importance that an aircraft can divert from the original itinerary for whatever reason. Therefore, in addition to the initial fuel calculations, 5% extra reserve fuel is added. The resulting payload-range diagrams can be seen in Figure 2. The maximum payload capacity of aircraft k flying from airport l_i to airport l_j is denoted by $\kappa^k_{l_i l_j}$.

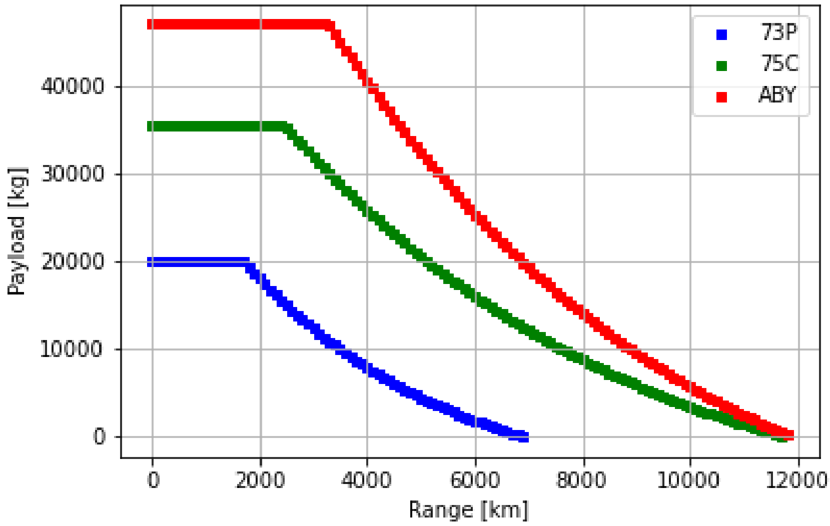

Figure 2: The payload range diagrams of the aircraft used in this research.

4.3 Network Representation

Before it is explained how the rest of the model is designed, it is important to understand the layout of the network. In case necessary, it will also be explained which sets are used in the model. The formal definition of all sets used in the model is provided in Table 5. A schematic representation of the network that is considered by the model is shown in Figure 3. This overview has a strong resemblance with the network that is presented in [Delgado et al., 2020], although there are some minor differences, such as the use of aircraft nodes. As can be seen in the legend, the network is composed of several types of nodes (N) and arcs (A). All nodes and arcs are used to create the TSN for the model, which is formulated as a directed graph $G = (N, A)$. The time horizon of the network is divided into equally sized periods of time $t \in T$. The set of airports in the network is represented by $l \in L$. Each airport for each point in time is represented by an activity node $i := (l, t)$. So an activity node represents both a place and a specific point in time. The airline operates a fleet of $k \in K$ aircraft. Each of the aircraft types in the fleet can serve a certain set of airports of $l_k \in L$. For each $k \in K$, there are two aircraft nodes: one for when the aircraft is available and one for when it is expected to be returned. These nodes are denoted by i^+_k and i^-_k respectively. The set of requests is denoted by $r \in R$. A request node consists of an origin node (i^+_r) and a destination node (i^-_r). The set of all nodes in the network is defined as $N := N^I \cup N^R \cup N^k$.

Five arc types connect the network. Flight arcs connect an airport node to another airport node later in time. To ensure that each aircraft can meet its next connection, the turnaround time of the aircraft is added to the flight time. This means that it depends on both the flight time and the turnaround time which activity nodes are connected with each other. Not all aircraft are allowed to fly to all airports in the network. Therefore, for each aircraft type, a subset of A^F is created that contains the set of legs (i, j) that each aircraft k can fly.

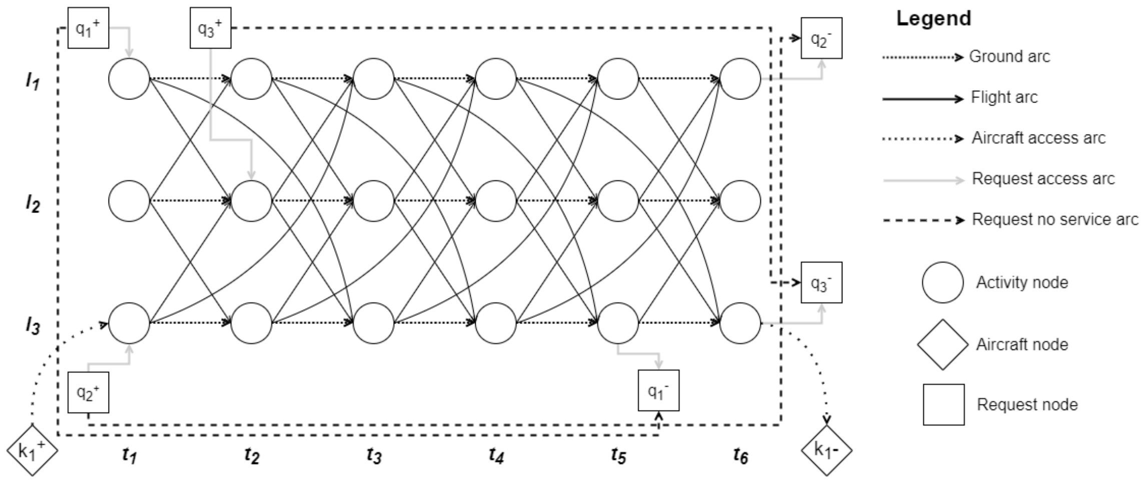

Figure 3: Schematic representation of the network that is considered by the model. (Source: adapted from [Delgado et al., 2020].)

Moreover, for each flight arc, it is determined which aircraft in the fleet is allowed to be assigned to it. This is saved in parameter K_{ij}. Ground arcs connect an airport node to the same airport but to the next point in time. Again, a ground arc is only accessible for an aircraft if it is allowed to fly there. As mentioned before, each aircraft in the considered fleet has a fixed starting and ending position. These positions are ensured by the fixed aircraft access arcs, denoted by set A^K. A distinction is made between access arcs that allow aircraft to start their rotation in the network (A^K_{org}) and access arcs that allow aircraft to finish their rotation in the network (A^K_{dst}). These arcs either emanate from an origin aircraft node (i^+_k) to an origin activity node i_{org_k} or return to an aircraft destination node (i^-_k) from a destination activity node i_{dst_k}. In case $(i, j) \in A^K_{org}$, the access arc is in the form of (i^+_k, i_{org_k}). In case $(i, j) \in A^K_{dst}$, the access arc is in the form of (i_{dst_k}, i^-_k). Furthermore, the parameters a_{org_k} and a_{dst_k} denote the number of aircraft that is required to be assigned to these aircraft access arcs. For example, if two aircraft of type k start their rotation from origin node *org*, then $a_{org_k} = 2$. Request access arcs allow a request to enter and exit the "physical" network. For each request two access arcs are drawn: one to enter the network and one to exit the network. These arcs are unique for each request. The request no-service arcs connect the origin node of a request with the destination node of that same request. Again, these arcs are unique for every request and are only accessible to that particular request. All ground and flight arcs are accessible to the requests. The complete set of arcs is defined as $A := A^F \cup A^G \cup A^K \cup A^N \cup A^R$.

The objective of the model is to determine a set of flights that *maximizes* the amount of demand transported by *minimizing* the total operational cost. The following cost components are taken into account: the ACMI cost associated with the flown flight legs, the cost for the fuel that is required to fly all flight legs empty, the cost for the marginal extra fuel that is required to transport payload, and last, the cost associated with not servicing demand. The parameter F^k_{ij} denotes the cost incurred when aircraft k would fly leg $(i, j) \in A^F$ *empty*. This parameter is composed out of two components, the ACMI cost to fly a leg and the fuel cost for flying that same leg empty. The ACMI component is calculated by multiplying the respective ACMI cost per hour for $k \in K$ (c_k) with the flight time of aircraft k for leg (i, j) (τ^k_{ij}). The fuel required for operating leg (i, j) empty ($W_{f_{OEW}}$) was estimated by using the calculations presented in section 4.2. The geodetic distance between any airport in the network is known, so by using Equation 7 the $W_{f_{OEW}}$ can be determined for each $(i, j) \in N^I$ and $k \in K$. It was assumed that 1 gallon of kerosene weight 3.45 *kg* and that 1 gallon of kerosene cost 1 *EUR*. It is worth noting that it is expected that this cost is lower than the actual market price. However, the authors of this research preferred to not overestimate the cost of fuel, as there are many more factors that influence the fuel consumption that are not taken into account. Moreover, if our airline intends to use the designed model, this value can easily be adjusted. So, the cost associated with flying empty are calculated as: $(W_{f_{OEW}}/3.45\ kg) * 1$. These costs are denoted by parameter F^k_{ij}.

The cost associated with transporting a kg of payload over arc $(i, j) \in A$ is represented by the parameter M_{ij}. Similarly to the cost for $W_{f_{OEW}}$, Equation 8 was used to determine the cost of the marginal extra fuel required to bring any kg of payload on leg $(i, j) \in A^F$. As can be seen, this parameter is not dependent on r. The costs associated with using the no-service arcs ($(i, j) \in A^N$) are set to a value that is at least as high as the cost to transport a request through the actual network. This is done to ensure that the model transports as much of the predicted demand as possible. This approach was chosen, such that the level of service provided to customers, which is a key driver for many airlines, is maximized. Finally, the cost associated with using ground arcs and aircraft and request access arcs are assumed to be zero, i.e. $F^k_{ij} = 0$ for $(i, j) \in A^G \cup A^K$ and $M_{ij} = 0$ for $(i, j) \in A^G \cup A^R$.

4.4 Cargo Routing

It has been proven that the path-flow model is a useful and efficient method of routing cargo through the network (see section 2.3). Therefore, it will be integrated into this model. In essence, the model works as follows: for every unique O&D pair it is determined what paths can be used to connect it via the network. A request is defined by a unique origin node and destination node. In general, a path is found if there is a set of arcs in the network that connect these nodes. A path-finding algorithm is used to find all available paths for each O&D pair. However, when the network grows in size, it becomes increasingly time-consuming to find all available paths. To reduce the effect of this problem, [Yan et al., 2005] implemented a meta-heuristic. This is a type of heuristic that can be used to guide the search process. In this case, the heuristic forced the path-finder to only search for specific paths. Four strategies were proposed that focused on the number of stops that were allowed for a path to be considered feasible. For a more detailed explanation of the strategies, the reader is referred to [Yan et al., 2005]. For this research, an adapted version of the proposed "mixed-stop heuristic" was implemented.

The mixed-stop-heuristic for this research is designed as follows: the number of stops that is allowed between a certain O&D pair is dependent on the *fastest total flight time* required to connect this O&D pair. The fastest total flight time is used because it is not assumed that all airports in the network are connected with direct flights. In case stop-overs are required to connect an O&D pair, then the total flight time that is taken into account is the sum of the flight time of the different legs. In Table 3 it can be reviewed how the different buckets of flight time are limited by the maximum allowed flight time of the total path and the maximum number of stops that are allowed in the path. It is important to understand that the goal of a meta-heuristic is not to exclude potentially realistic solutions but to exclude solutions that are infeasible or unrealistic.

Table 3: Definition of the buckets of the mixed-stop heuristic. "Flight time" refers tot the fastest connection provided by the network, "Max tot flight time" refers to the sum of the flight time of legs in the path.

Flight time	**Max tot flight time**	**Max nr of stops**
0-2 hrs	6 hrs	1 stop
2-4 hrs	10 hrs	2 stops
4-6 hrs	15 hrs	3 stops
>6 hrs	20 hrs	4 stops

The implementation of this heuristic showed improvements in terms of computing time required to find all potential paths for the different requests. However, the search for all these paths still required an undesirable amount of time (several hours). It was found that common search algorithms use brute force to find paths. What this means, is that all arcs are used to try to make a feasible path, even though it might be going in the wrong direction. To provide a tangible example: there are five airports in the network Amsterdam, Paris, Madrid, Hong Kong and New York. A request its O&D pair is Amsterdam-Madrid. A general pathfinder will consider *all* possible paths that start in Amsterdam and end in Madrid as feasible. Even a path going from Amsterdam to Hong Kong, back to Paris and finally to Madrid would be considered feasible. On top of this, in the TSN this path will be considered with all different arrival and departure options. Making numerous paths that only increase the size of the problem, without providing it valuable options. To overcome this problem the authors devised another meta-heuristic. For this heuristic, the *static network* of our airline was used. In this network, the airports of the network of our airline represent the nodes. Unlike with the TSN, all nodes in the static network are *not* associated with a point in time. The nodes are connected if a direct flight was flown in the past. This network was used to find the top 20 fastest (in terms of time) paths between all nodes (or airports) in the network. The time of one path was calculated by summing the flight time of all arcs (or flight legs) in the path. For each O&D pair, these paths, consisting out of a list of airports passed, with their respective total flight times were saved. Although most probably the path that is eventually assigned to a request is somewhere in the first few fastest path, the goal of the meta-heuristic is only to reduce the number of unrealistic paths. These saved paths were used to reduce the run-time required to find all the feasible paths from several hours to a few minutes. This was achieved as follows: each saved path consisted out of a list of airports in that path. These airports could then be used to create a subgraph of the whole TSN, that only contained those activity nodes that would be used to create a feasible path. Due to the mixed-stop-heuristic, the maximum number of airports in a path is limited to 6. One can imagine that finding all possible paths by brute force in a network of only 6 airports, is much faster than if the airport consists out of several tens of airports.

All the while taking into account the two aforementioned meta-heuristics, a cargo flow path p for request r is a sequence of connected activity nodes which connect the origin and destination node of that request. The set of all feasible flow paths is denoted by P. Furthermore, P_r is the set of paths available to request r. Finally, the cost associated with using each path can be determined by summing the marginal extra fuel cost M_{ij} for each leg (i, j) in path p. This leads to parameter V_p which represent the cost for each path $p \in P$. Moreover, set

P_{ij} is defined. This set denotes a set of paths p that use leg (i,j). Furthermore, for each path p it is determined how many flight arcs it contains. This is denoted by parameter n_p. Finally, for each path p a subset consisting of only flight arcs in that path is created. This is denoted by f_p. This subset was used to determine what the *minimum maximum* capacity of each of the flights in a path was. This forms an upper bound for how much demand can potentially be appointed to a path. If multiple aircraft can be assigned to one leg, the maximum of the two capacities is used. It is denoted by κ_p. In Table 4 an overview is provided of the parameters used in the model. Furthermore, Table 5 contains an overview of the sets used in the model and their mathematical formulation. And in Table 6 the decision variables of the model and their nature are presented.

The previously described path flow model is used to determine all potential paths each request can take. For some requests, there might be no feasible paths through the "physical" network within the given time window for delivery. The cause for this could for example be that the available aircraft type does not provide a direct connection between the O&D of the request and that the total flight time of the flights required to connect it, would exceed the delivery time window. Consequently, p_r would only contain the no-service path. Thus, it can be known before the optimization model is run, that these requests will not be serviced. Therefore, it was decided to remove these requests from the demand, before the problem was optimized. This directly reduced the number of decision variables in the problem. Additionally, the model only takes into account the no-service cost for those requests that could have been transported via the "physical" network.

Table 4: Parameters used in the model.

Parameter	**Definition**
$\kappa^k_{l_i l_j}$	Aircraft capacity of aircraft k on flight arc (i,j)
κ_p	Minimum maximum aircraft capacity offered on path p
τ^k_{ij}	Flight time required to fly leg (i,j) with aircraft k
F^k_{ij}	Cost of using aircraft k on leg (i,j)
M_{ij}	Marginal fuel cost on leg (i,j)
n_k	Number of aircraft of aircraft type k that are available
n_p	Number of flight arcs in path p
V_p	Cost to use flow path p
w_r	Weight of request r

4.5 Model Layout

It is essential to have a clear overview of what the complete model will look like in the end and how the sub-parts will interact with each other. Therefore, a schematic overview of the model its flow diagram is presented in Figure 4. As can be seen, the model is split up into two main parts. In one part of the model, a benchmark solution is determined, while in the other, the model determines the optimal solution. In the figure, the dark orange shaded blocks represent inputs, the yellow shaded blocks represent the models that will be used to determine the output blocks, which are shaded in green. Finally, the model's solution will be compared with the benchmark solution, which will produce the final output, which is shown in the blue shaded block.

Forecast demand and a base schedule should be provided to the model to determine the benchmark solution. This will then be used to model the cargo flow through the network. The flight schedule and resulting cargo flow are then used to estimate the operating cost and several other KPIs. Meanwhile, the same forecast demand and a mandatory flight list will be used as input for the other part of the model. The mandatory flight list is determined by randomly selecting a subset of flights of the base schedule. The cargo flow model and the schedule design model will then determine the optimal set of flights that should be added to this mandatory flight list. This newly designed schedule should serve as much of the predicted demand as possible. The results of the model solution can then be compared with that of the benchmark solution.

4.6 Mathematical formulation of the AFFSDP

The AFFSDP is formulated as a Mixed Integer Linear Program (MILP). Three decision variables (dv) are defined: $X := \{x^k_{ij} : \ \forall\ k \in K \wedge (i,j) \in A^I \cup A^K\}$, this is an integer variable that denotes whether aircraft k is assigned to arc (i,j). The upper bound of $x^k_{ij} = n_k$ where n_k denotes how many aircraft of type k are available. $Q := \{q_p : \ \forall\ p \in P\}$ which represents the amount of demand that is transported via path p. This dv is continuous. For every request there is a set of paths that only that request is allowed to be assigned to. Therefore, the natural upper bound of $q_p = w_r$ (w_r is the weight of request r). The last variable is

Table 5: Sets used in the model.

Sets	Definition	Formulation
A^F	Flight arcs	$\{(i,j) : i,j \in N^I \wedge l_i \neq l_j \wedge t_i < t_j\}$
A^F_k	Flight arcs aircraft k can use	$\{(i,j) \text{ s.t. } (i,j) \in \tau^k_{ij} \;\; \forall\, (i,j) \in A^F \wedge k \in K\}$
A^G	Ground arcs	$\{(i,j) : i,j \in N^I \wedge l_i = l_j \wedge t_i + 1 = t_j\}$
A^K	Aircraft access arcs	$\{(i^+_k, i) : k \in K \wedge i \in N^I \wedge l_i = l^+_k \wedge t^+_k \leq t_i \leq t^-_k\}$
		$\cup\{(i, i^-_k) : k \in K \wedge i \in N^I \wedge l_i = l^-_k \wedge t^+_k \leq t_i \leq t^-_k\}$
A^N	No-service arcs	$\{(i^+_r, i^-_r) : r \in R\}$
A^R	Request access arcs	$\{(i^+_r, i) : r \in R \wedge i \in N^I \wedge l_i = l^+_r \wedge t^+_r \leq t_i \leq t^-_r\}$
		$\cup\{(i, i^-_r) : r \in R \wedge i \in N^I \wedge l_i = l^-_r \wedge t^+_r \leq t_i \leq t^-_r\}$
A^I	“Physical” arcs	$A^F \cup A^G$
A	All arcs	$A^F \cup A^G \cup A^K \cup A^N \cup A^R$
f_p	Flight arcs in path p	$\{(i,j) \in p \text{ s.t. } (i,j) \in A^F \;\forall\, p \in P\}$
K	Aircraft types	-
K_{ij}	Aircraft types that can fly leg (i,j)	$\{k \in K \text{ s.t. } (i,j) \in A^F_k \;\; \forall\, (i,j) \in A^F\}$
L	Airports in the network	-
L_k	Airports serviced by aircraft type k	-
N^I	Itinerary nodes	$\{i : (l,t) \;\forall\, l \in L \wedge t \in T\}$
N^R	Request nodes	$\{i_r : (l^+, t^+) \;\forall\, r \in R\} \cup \{(l^-, t^-) \forall r \in R\}$
N^k	Aircraft nodes	$\{i_k : l^+_k \;\forall\, k \in K\} \cup \{l^-_k \forall k \in K\}$
N	Nodes	$N^I \cup N^R \cup N^k$
P	Flow paths	$\{p \in P_r \;\forall\, r \in R\}$
P_r	Flow paths available to request r	-
P_{ij}	Flow paths that use leg (i,j)	$\{p \in P \text{ s.t. } (i,j) \in p \;\forall\, (i,j) \in A^F\}$
T	Time periods	-

Table 6: Decision variables used in the model.

Decision Variable	Type	Definition
q_p	Continuous	Amount of flow assigned to path p
x^k_{ij}	Binary	Equals 1 if aircraft k is assigned to leg (i,j), 0 otherwise
z_p	Binary	Equals 1 if flow can be assigned to path p, 0 otherwise

$Z := \{z_p : \;\forall\, p \in P\}$. This dv is variable is binary and denotes whether path p can be used. It is equal to 1 if all flights in the path are assigned to an aircraft, otherwise it is 0. If $p = 0$, this means that q_p must equal zero as well.

In section 4.4 it was explained that the set of available flow paths is constructed by using two meta-heuristics. The model uses this set to find the optimal solution to the problem. The combined use of a meta-heuristic and an exact mathematical programming method is generally referred to as a matheuristic [Angelelli et al., 2020]. As will become clear in section 6, the matheuristic that was designed in this study showed that it can be used to solve large size problem instances within a reasonable computing time. The remainder of this section focuses on elaborating on the objective function and constraints of the model.

$$\text{Min} \sum_{k \in K} \sum_{(i,j) \in A^I} F^k_{ij} x^k_{ij} + \sum_{\text{p} \in P} V_p q_p \tag{9}$$

$$\text{subject to} \sum_{\text{p} \in P_{(i,j)}} q_p \leq \sum_{k \in K_{(i,j)}} \kappa^k_{l_i l_j} x^k_{ij} \quad \forall\, (i,j) \in A^I \tag{10}$$

$$\sum_{\text{k} \in K_{(i,j)}} x^k_{ij} \leq 1, \quad \forall (i,j) \;\in A^I \tag{11}$$

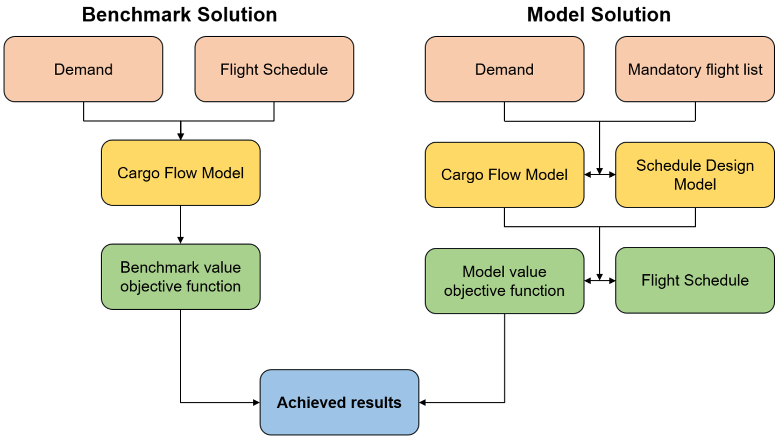

Figure 4: A schematic overview of how the model will be created and how the cost-effectiveness of it will be determined.

$$\sum_{(j,i)\in d_i^+} x_{ji}^k - \sum_{(i,j)\in d_i^-} x_{ij}^k = \begin{cases} -n_k & \text{if} \quad i = i_k^+ \\ n_k & \text{if} \quad i = i_k^-, \\ 0 & \text{if} \quad \text{i.o.c.} \end{cases} \qquad i \in N^I \wedge k \in K \tag{12}$$

$$\sum_{(k\in K_{ij})} x_{ij}^k = \begin{cases} a_{org_k} & \text{if} \quad (i,j) = A_{org}^K \\ a_{dst_k} & \text{if} \quad (i,j) = A_{dst}^K, \end{cases} \qquad (i,j) \in A^K \tag{13}$$

$$\sum_{\text{p}\in P^r} q_p = w_r \quad \forall\, r \in \text{R} \tag{14}$$

$$\sum_{(i,j)\in p} \sum_{k\in K_{ij}} x_{ij}^k \geq n_p z_p \;\forall\, p \in P \tag{15}$$

$$\sum_{(i,j)\in A^I : t_i \leq t \leq t_j} x_{ij}^k = n_k \quad \forall\, k \in K \wedge t \in T \tag{16}$$

$$x_{ij}^k \geq z_p \;\forall\, (i,j) \in f_p \wedge p \in P \tag{17}$$

$$q_p \leq min(\kappa_p,\; w_r) \;\forall\, p \in P_r \wedge r \in R \tag{18}$$

$$x_{ij}^k \quad \in \{0,1\} \;\forall\, (i,j) \in A, k \in K \tag{19}$$

$$q_p \quad \geq 0 \;\forall\, p \in P \tag{20}$$

$$z_p \quad \in \{0,1\} \;\forall\, p \in P \tag{21}$$

As can be seen in Equation 9, the objective function of this model is to minimize the total operating cost. There are two main components, the cost related to aircraft usage and the cost related to the transport of demand. The first constraint, shown in Equation 10, represents the aircraft capacity constraint. It ensures that for each leg the sum of flow on the paths (q_p) that use that leg is lower than or equal to the aircraft capacity of the aircraft that is assigned to each leg. The constraint shown in Equation 11 ensures that only one aircraft is assigned to each flight in the network. In Equation 12 the constraint that ensures the flow conservation of aircraft at each node in the network is shown. The aircraft nodes act as source and sink nodes, meaning that the number of aircraft leaving an aircraft origin node, should equal the total number of aircraft of that fleet type. Similarly, the number of aircraft entering an aircraft destination node should equal the total number of aircraft of that fleet type. For all other nodes in the network, the in-flow and out-flow of aircraft should be balanced and therefore must be 0. The constraint below (Equation 13) ensures that the number of aircraft of type k that start their rotation at a certain airport is equal to the number of aircraft that is supposed to end their rotation there (according to the base schedule). The next constraint, shown in Equation 14, forces the model to move all demand of each request from their origin node to their destination node, either via the “physical” network or via their no-service path. The constraint presented in Equation 15, implies the dependency of z_p on the flights in path p. It shows that z_p can only be 1 if the sum of aircraft assigned to flight legs in that path is equal to n_p. If this is not the case, the sum of x_{ij}^k will be lower than n_p directly forcing z_p to be equal to 0. The next two constraints are a special type of constraint named cuts. The difference between a constraint and a cut can be explained as follows; the model is properly defined with only the constraints shown in Equation 10 to Equation 15. Cuts are added to the model to tighten the bounds of the relaxed solution. This means that the model can find a solution more quickly when the cuts are added than when the cuts are omitted. In both cases, the model will converge to the same solution, but within a different time frame. The cut presented in Equation 16 takes advantage of the fact that at any given moment in time, the number of aircraft used by the model must be equal to the number of aircraft of that type in the fleet. The next cut, shown in Equation 17, exploits the fact that if even one of the flights in the set of flights of path p (f_p) is not flown, then z_p must equal 0 as well. A third cut was added, which is shown in Equation 18. This cut limits the amount of demand that can be assigned to q_p. This is either limited by the weight of the demand of request r or by the minimum-maximum capacity offered on the flight legs in the path (κ_p).

In addition to the cuts, there are two methods applied which allow the model to converge to an integer solution faster. It was found that decimal numbers are a common source for rounding errors in optimization problems [1]. The cost factor V_p consists of fractional numbers for all paths that are not no-service paths. To overcome this problem, all numbers in this set were multiplied by $10E+4$ and rounded to the nearest integer. It is important to use such a large number, as the difference in cost to transport one kg of demand for different paths was a very small value. Furthermore, for many requests, it happens that there are paths that are the same in terms of airports passed, but the moment in time at which these airports are passed differ. Therefore, the cost term V_p for these paths will be exactly the same. Symmetry in optimization models can be a cause of slow convergence because multiple solutions are equivalent. To overcome this problem, a slight distinction between each of the paths has to be made. This was done by implementing a path-preference system. For each path p available to request r it was determined how many periods t it arrived at the destination airport before the final delivery time of r. The difference was then subtracted from V_p. This leads to a preference for using paths that arrive earlier at the destination airports. V_p is in the order of hundreds or thousands. Therefore, the effect of subtracting a value from around 0 to 10 only makes a slight difference. Furthermore, in reality, it also makes sense to have this preference, as it would provide some contingency in case any flight is delayed.

To solve the MILP problem, the decision was made to make use of the IBM ILOG CPLEX 12.10 solver which was made available under an academic license. The LP-problem files were created using an HP notebook with 16GB of RAM. Solving the problem with the same notebook would take hours due to the high number of decision variables and constraints in the optimization problem. Therefore, a server that was made available by the TU Delft with a dual AMD EPYC 7551 server with 64 cores, 128 threads and 256 GB of RAM was used. This server was able to achieve acceptable results within only a few hours. Post-processing of the results could be done on the HP notebook as this did not require much time.

5 Description of the Test Instances

Several test instances have been used to analyze the effectiveness of the model. The aircraft types that form the basis of this analysis are the *ABY*, the *75C* and the *73P*. These aircraft types have been chosen for two main reasons. First of all, these aircraft are available in comparable numbers in the fleet of our airline. This allows for easier comparison between the different instances. The other reason is that there is a clear distinction in the maximum available payload capacity. As can be clearly seen in the payload range diagrams in Figure 2, the *73P* has the smallest capacity, thereafter follows the *75C* and the *ABY* is the largest of the three. This explains the choice of the aircraft types used for this research.

For each aircraft type, the following distinctions have been made. Either 50% of the flight arcs of the base schedule were fixed or 75% were fixed, or the actual schedule itself was used. Furthermore, for each of these instances a different demand scenario was determined, either *low*, *normal* or *high*. In the case of the low demand scenario, the O&D demand was determined, and the list of requests was created. Then, a random sample of 25% of the entries of this list was completely removed. In the case of the normal demand scenario, nothing was adjusted. In the case of the high demand scenario, a random sample of 25% of the entries of this list was determined and the demand for each of these entries was doubled. The use of these different demand scenarios is that it can be analyzed separately what the effect of fewer and more requests is, and the effect of having more or less demand in total but with the same amount of requests. The instances and their respective details are summarized in Table 7.

In general, a schedule repeats every week for one season long. Therefore, the time horizon of the base schedule of each of the aircraft types is one week as well. Furthermore, the time step t of the model was chosen to be 1 hour.

[1] https://www.ibm.com/docs/pl/icos/12.7.1.0?topic=problems-numeric-difficulties

Table 7: Definition of the instances that were used to test the model.

Instance	AC	Network description	Demand scenario
ABY-50-L	ABY	50% flight arcs fixed	Low
ABY-75-L	ABY	75% flight arcs fixed	Low
ABY-ACT-L	ABY	Base schedule	Low
ABY-50-N	ABY	50% flight arcs fixed	Normal
ABY-75-N	ABY	75% flight arcs fixed	Normal
ABY-ACT-N	ABY	Base schedule	Normal
ABY-50-H	ABY	50% flight arcs fixed	High
ABY-75-H	ABY	75% flight arcs fixed	High
ABY-ACT-H	ABY	Base schedule	High

Instance	AC	Network description	Demand scenario
75C-50-L	75C	50% flight arcs fixed	Low
75C-75-L	75C	75% flight arcs fixed	Low
75C-ACT-L	75C	Base schedule	Low
75C-50-N	75C	50% flight arcs fixed	Normal
75C-75-N	75C	75% flight arcs fixed	Normal
75C-ACT-N	75C	Base schedule	Normal
75C-50-H	75C	50% flight arcs fixed	High
75C-75-H	75C	75% flight arcs fixed	High
75C-ACT-H	75C	Base schedule	High

Instance	AC	Network description	Demand scenario
73P-50-L	73P	50% flight arcs fixed	Low
73P-75-L	73P	75% flight arcs fixed	Low
73P-ACT-L	73P	Base schedule	Low
73P-50-N	73P	50% flight arcs fixed	Normal
73P-75-N	73P	75% flight arcs fixed	Normal
73P-ACT-N	73P	Base schedule	Normal
73P-50-H	73P	50% flight arcs fixed	High
73P-75-H	73P	75% flight arcs fixed	High
73P-ACT-H	73P	Base schedule	High

6 Computational Results

The following section will elaborate on the results that were obtained by running the different instances described in the previous section. To ensure that that the data provided by our airline remains anonymous, all results have been normalized. For each part of the results that are discussed, it will be pointed out with respect to what value the data is normalized. For clarity, these cells are also highlighted in blue.

In Table 8 the details of the network of each of the instances are displayed. For each aircraft type, you can see the total number of airports, unique O&Ds it can fly to, flight legs in the network, fixed flight legs, paths that can be used to ship the request and at last the number of dvs. All data has been normalized w.r.t. instance 73P-ACT-L. Except for the number of dvs, as this does not reveal any sensitive data of our airline. The table provides useful insights into how the network of each of the different aircraft changes. For example, the number of airports in the network does not change much for the different aircraft types. But, the number of O&Ds flown does slightly increase with the increasing size of the aircraft. This also directly translates to the higher number of flight arcs in the network. Furthermore, it is interesting to note that the number of paths available to ship requests is the highest for the 73P. This might seem counter-intuitive, as the number of flight legs in the network is the lowest. But, the network serviced by the 73P is defined by more short-range flights than the networks of the ABY and the 75C. The average flight distance of the latter two is approximately 35% larger than the one of the 73P. Consequently, the O&Ds of requests serviced by the 73P are closer together, which generally means that the number of options to transport a request increases. In relation to this, it can be seen that the number of dvs for the 73P network is also the largest for each of the instances. As explained in section 4, there are three decision variables, two of which are related to the number of paths in the network and one which is related to the number of flight arcs. Therefore, the number of paths has a stronger impact on the size of the problem than the number of flight arcs does.

Table 8: Overview of the network details of each of the different test instances.

ABY	AP	OD	Legs	Fixed legs	Paths	DVs
50%, L	1.1	1.2	58.6	0.6	14.7	40,616
75%, L	1.1	1.2	58.6	0.9	14.7	40,616
act, L	1.1	1.2	1.2	1.2	1.1	1,886
50%, N	1.1	1.2	58.6	0.6	19.7	47,272
75%, N	1.1	1.2	58.6	0.9	19.7	47,272
act, N	1.1	1.2	1.2	1.2	1.5	2,382
50%, H	1.1	1.2	58.6	0.6	19.7	47,272
75%, H	1.1	1.2	58.6	0.9	19.7	47,272
act, H	1.1	1.2	1.2	1.2	1.5	2,382

75C	AP	OD	Legs	Fixed legs	Paths	DVs
50%, L	1.1	1.1	54.5	0.5	12.2	35,783
75%, L	1.1	1.1	54.5	0.7	12.2	35,783
act, L	1.1	1.1	1.0	1.0	0.9	1,551
50%, N	1.1	1.1	54.5	0.5	15.9	40,803
75%, N	1.1	1.1	54.5	0.7	15.9	40,803
act, N	1.1	1.1	1.0	1.0	1.2	1,935
50%, H	1.1	1.1	54.5	0.5	15.9	40,803
75%, H	1.1	1.1	54.5	0.7	15.9	40,803
act, H	1.1	1.1	1.0	1.0	1.2	1,935

73P	AP	OD	Legs	Fixed legs	Paths	DVs
50%, L	1.0	1.0	50.4	0.5	24.2	50,427
75%, L	1.0	1.0	50.4	0.7	24.2	50,427
act, L	1.0	1.0	1.0	1.0	1.0	1,697
50%, N	1.0	1.0	50.4	0.5	31.4	60,005
75%, N	1.0	1.0	50.4	0.7	31.4	60,005
act, N	1.0	1.0	1.0	1.0	1.4	2,171
50%, H	1.0	1.0	50.4	0.5	31.4	60,005
75%, H	1.0	1.0	50.4	0.7	31.4	60,005
act, H	1.0	1.0	1.0	1.0	1.4	2,171

The effect of an increasing number of dvs on the quality of the results from CPLEX can be seen in Table 9. In this table, one can see the objective value, the optimality gap and the solution time for each of the instances. It can be seen that in the case of the 73P, which had the most dvs, the optimal solution was never found within the given time limit of 3 hours. Furthermore, if for each aircraft type the results of the different demand inputs are compared, it can be seen that the optimality gap always increases from the actual instance to the one where only 50% of the network is fixed. This makes sense because the model is the most constrained in the case of the actual instances, and the least for the 50% instances. Finally, when one would look at the column of the objective value, it seems like that the actual model is performing the best. As will become clear in the rest of this section, this is not the case. Nonetheless, this seemingly illogical column can be explained as follows: the objective value is largely dominated by the no-service cost. As was explained in section 4.4, only those requests for which a path through the "physical" network exists, are considered by the model. Since the actual schedule provides fewer routing options (because it is completely fixed), this means that there are substantially fewer requests considered for the actual network. In turn, this means that the no-service cost will also be lower. Therefore, the objective value results in a lower value.

Table 9: Details of the CPLEX solutions.

ABY	**obj val** [€]	**opt gap** [%]	**sol t** [s]
50%, L	3.7E+08	2.78	10,814
75%, L	8E+08	1.00	584
act, L	8.1E+07	0.00	0
50%, N	6.8E+08	7.22	10,800
75%, N	1.2E+09	0.50	932
act, N	1.5E+08	0.00	0
50%, H	1.1E+09	5.57	10,813
75%, H	1.8E+09	1.02	10,805
act, H	6E+08	0.00	0

75C	**obj val** [€]	**opt gap** [%]	**sol t** [s]
50%, L	1.8E+08	1.00	8,791
75%, L	7.46E+08	0.92	92
act, L	1.74E+08	0.00	0
50%, N	4.29E+08	2.44	10,812
75%, N	1.12E+09	0.98	162
act, N	2.74E+08	0.00	0
50%, H	6.24E+08	3.58	10,813
75%, H	1.54E+09	0.97	147
act, H	6.32E+08	0.00	0

73P	**obj val** [€]	**opt gap** [%]	**sol t** [s]
50%, L	3.25E+08	13.61	10,813
75%, L	7.39E+08	2.95	10,815
act, L	5.37E+08	0.00	0
50%, N	7.28E+08	7.59	10,813
75%, N	1.21E+09	2.34	10,816
act, N	8.7E+08	0.00	0
50%, H	1.45E+09	4.84	10,820
75%, H	2.04E+09	1.46	10,821
act, H	1.52E+09	0.00	0

In Table 10, Table 11 and Table 12, a summary of the service analysis of each of the instances is displayed. In order, they display the instances with the low demand, normal demand and high demand. In order from left to right, the columns denote the following things, the demand, the average size of each request in the demand, the average weighted distance of each of the requests and the respective weighted detour factor. The average distance represents the average direct flight distance of each request. It is weighted in the sense, that for each request the direct flight distance was determined separately, and then the size of each request was used to determine the average weighted distance. Similarly, the weighted detour factor represents how many extra kilometres each request is flown in comparison to its direct flight distance, weighted w.r.t. the size of each request. From top to bottom the rows represent the following, all requests considered, all the requests that have been either fully serviced or partially serviced, and finally all the requests that have not been serviced at all. In Appendix A a more detailed version that also contains a distinction between fully serviced and partially serviced demand can be found. The key takeaways can be extracted from these summarized versions. Therefore, for conciseness here only these extracts are presented. Furthermore, it only makes sense to calculate the weighted detour factor, for those requests that were sent. Therefore the value in the rows of *all* and *no s* are denoted by *n.a.* (not applicable). Note, the data in Table 10, Table 11 and Table 12 was normalized w.r.t. instance 73P-ACT-L. The results of this instance are highlighted in blue.

First of all, in all of the instances, it can be seen that the weighted distance of the serviced demand is lower than that of all the demand considered. This makes sense as it is relatively more cost-effective for the model to transport more demand on multiple short routes than to transport less demand on a longer route because this results in lower no-service costs. In line with this observation, it can also be seen that the size of the average not-serviced request is much smaller than those of a serviced request. Additionally, also the average distance of the not-serviced requests is further than that of the serviced requests.

Furthermore, when one compares the average weighted distance of each of the instances, it becomes immediately clear that the average distance of serviced requests is the lowest for the 73P. This is in line with the expectation as this aircraft also flies the smallest average flight distance. However, as explained before, the closer requests are together, the more routing options there will be. This can be seen in the results of the service analysis of the instance where the model had either 75% or 50% of the flight network fixed for this aircraft, as for all of these instances the weighted detour factor of the 73P network is the highest. There was no freedom for the model to determine flights for the "actual" instances. Therefore, it is not fair to compare these results directly with that of the model. Furthermore, it can be seen that the detour factor of the instances of the 75C and ABY are usually quite close. This also makes sense as the network connected by these aircraft shows much resemblance in terms of average flight distance.

Table 10: Service analysis in case of low demand.

ABY, 50%, L

serv	dem	avg size	avg w dist	w det fac
all	2.279	1.494	1.239	n.a.
tot s	2.130	1.688	1.209	1.074
no s	0.147	0.545	1.670	n.a.

75C, 50%, L

serv	dem	avg size	avg w dist	w det fac
all	1.589	1.221	1.142	n.a.
tot s	1.522	1.312	1.136	1.082
no s	0.061	0.455	1.318	n.a.

73P, 50%, L

serv	dem	avg size	avg w dist	w det fac
all	1.731	1.221	0.978	n.a.
tot s	1.587	1.286	0.948	1.175
no s	0.105	0.571	1.341	n.a.

ABY, 75%, L

serv	dem	avg size	avg w dist	w det fac
all	2.279	1.494	1.239	n.a.
tot s	1.949	1.766	1.209	1.077
no s	0.315	0.740	1.406	n.a.

75C, 75%, L

serv	dem	avg size	avg w dist	w det fac
all	1.589	1.221	1.142	n.a.
tot s	1.278	1.351	1.115	1.069
no s	0.304	0.857	1.243	n.a.

73P, 75%, L

serv	dem	avg size	avg w dist	w det fac
all	1.731	1.221	0.978	n.a.
tot s	1.421	1.286	0.935	1.172
no s	0.265	0.844	1.184	n.a.

ABY, act, L

serv	dem	avg size	avg w dist	w det fac
all	1.572	1.468	1.067	n.a.
tot s	1.547	1.455	1.060	1.094
no s	0.010	0.740	1.800	n.a.

75C, act, L

serv	dem	avg size	avg w dist	w det fac
all	1.091	1.247	1.097	n.a.
tot s	1.024	1.221	1.069	1.099
no s	0.025	0.584	1.467	n.a.

73P, act, L

serv	dem	avg size	avg w dist	w det fac
all	1.000	1.000	1.000	n.a.
tot s	0.779	0.909	0.980	1.091
no s	0.092	0.662	1.283	n.a.

Table 11: Service analysis in case of normal demand.

ABY, 50%, N

serv	dem	avg size	avg w dist	w det fac
all	2.975	1.455	1.206	n.a.
tot s	2.680	1.675	1.155	1.085
no s	0.283	0.636	1.656	n.a.

75C, 50%, N

serv	dem	avg size	avg w dist	w det fac
all	2.128	1.234	1.124	n.a.
tot s	1.957	1.338	1.100	1.092
no s	0.161	0.597	1.415	n.a.

73P, 50%, N

serv	dem	avg size	avg w dist	w det fac
all	2.222	1.169	0.972	n.a.
tot s	1.907	1.260	0.941	1.170
no s	0.241	0.636	1.173	n.a.

ABY, 75%, N

serv	dem	avg size	avg w dist	w det fac
all	2.975	1.455	1.206	n.a.
tot s	2.486	1.740	1.140	1.078
no s	0.468	0.753	1.525	n.a.

75C, 75%, N

serv	dem	avg size	avg w dist	w det fac
all	2.128	1.234	1.124	n.a.
tot s	1.658	1.390	1.091	1.083
no s	0.454	0.844	1.233	n.a.

73P, 75%, N

serv	dem	avg size	avg w dist	w det fac
all	2.222	1.169	0.972	n.a.
tot s	1.714	1.208	0.930	1.170
no s	0.405	0.844	1.140	n.a.

ABY, act, N

serv	dem	avg size	avg w dist	w det fac
all	2.182	1.519	1.132	n.a.
tot s	2.133	1.519	1.120	1.099
no s	0.015	0.675	2.012	n.a.

75C, act, N

serv	dem	avg size	avg w dist	w det fac
all	1.415	1.208	1.128	n.a.
tot s	1.306	1.182	1.076	1.082
no s	0.047	0.714	1.777	n.a.

73P, act, N

serv	dem	avg size	avg w dist	w det fac
all	1.353	1.013	0.980	n.a.
tot s	0.992	0.896	0.953	1.073
no s	0.154	0.688	1.251	n.a.

Table 12: Service analysis in case of high demand.

ABY, 50%, H

serv	dem	avg size	avg w dist	w det fac
all	3.685	1.805	1.224	n.a.
tot s	3.213	2.091	1.169	1.082
no s	0.425	0.831	1.542	n.a.

75C, 50%, H

serv	dem	avg size	avg w dist	w det fac
all	2.600	1.506	1.111	n.a.
tot s	2.341	1.675	1.100	1.077
no s	0.239	0.727	1.229	n.a.

73P, 50%, H

serv	dem	avg size	avg w dist	w det fac
all	2.758	1.455	0.983	n.a.
tot s	2.133	1.545	0.918	1.133
no s	0.449	0.870	1.261	n.a.

ABY, 75%, H

serv	dem	avg size	avg w dist	w det fac
all	3.685	1.805	1.224	n.a.
tot s	2.922	2.156	1.138	1.077
no s	0.692	1.000	1.562	n.a.

75C, 75%, H

serv	dem	avg size	avg w dist	w det fac
all	2.600	1.506	1.111	n.a.
tot s	1.952	1.649	1.058	1.075
no s	0.613	1.117	1.271	n.a.

73P, 75%, H

serv	dem	avg size	avg w dist	w det fac
all	2.758	1.455	0.983	n.a.
tot s	1.897	1.442	0.908	1.133
no s	0.621	1.065	1.198	n.a.

ABY, act, H

serv	dem	avg size	avg w dist	w det fac
all	2.708	1.896	1.134	n.a.
tot s	2.468	1.805	1.101	1.095
no s	0.038	0.584	1.467	n.a.

75C, act, H

serv	dem	avg size	avg w dist	w det fac
all	1.778	1.519	1.127	n.a.
tot s	1.518	1.416	1.077	1.069
no s	0.075	0.740	1.738	n.a.

73P, act, H

serv	dem	avg size	avg w dist	w det fac
all	1.717	1.299	0.986	n.a.
tot s	1.081	1.013	0.949	1.063
no s	0.230	0.896	1.260	n.a.

Furthermore, it was also analyzed how the fuel consumption differed for each of the instances. The data for this analysis can be found in Table 13. The columns from left to right represent the total demand available, the total demand serviced and the amount of fuel used. For this analysis, the results have been normalized differently. The three tables on the left side show for each aircraft for a given demand scenario how the model behaves with an increased amount of freedom w.r.t. the actual scenario. The three tables on the right compare for each aircraft for each amount of freedom how the model behaves w.r.t. an increased amount of demand. There are two important things to focus on when reading these tables. First of all, for each of the aircraft types and each of the demand scenarios, it can be seen that when the freedom of the model increases, the amount of fuel required to service the whole network increases less than the amount of extra demand serviced. Furthermore, if one compares the 70% instances to the 50% instances w.r.t. how much additional demand is serviced and how much additional fuel is required, it can be seen that the demand serviced increases by 10-25%, while the fuel required either decreases or only increases by several percentages. The tables on the right highlight that in case the demand increases, the model is able to accommodate 20-50% more demand while the required amount of fuel only increases slightly.

Table 13: Comparison of fuel usage in the different test instances.

Instances with equal demand

ABY	Tot dem	Tot dem s	fuel used
50%, L	1.449	1.377	1.186
75%, L	1.449	1.259	1.213
act, L	1.000	1.000	1.000
50%, N	1.365	1.256	1.168
75%, N	1.365	1.165	1.197
act, N	1.000	1.000	1.000
50%, H	1.361	1.301	1.190
75%, H	1.361	1.183	1.221
act, H	1.000	1.000	1.000

73P	Tot dem	Tot dem s	fuel used
50%, L	1.732	2.035	1.642
75%, L	1.732	1.823	1.565
act, L	1.000	1.000	1.000
50%, N	1.645	1.919	1.658
75%, N	1.645	1.725	1.612
act, N	1.000	1.000	1.000
50%, H	1.606	1.971	1.701
75%, H	1.606	1.753	1.643
act, H	1.000	1.000	1.000

75C	Tot dem	Tot dem s	fuel used
50%, L	1.455	1.486	1.464
75%, L	1.455	1.247	1.361
act, L	1.000	1.000	1.000
50%, N	1.505	1.495	1.460
75%, N	1.505	1.268	1.381
act, N	1.000	1.000	1.000
50%, H	1.462	1.540	1.485
75%, H	1.462	1.285	1.357
act, H	1.000	1.000	1.000

Instances with equal freedom

ABY	Tot dem	Tot dem s	fuel used
50%, L	1.000	1.000	1.000
50%, N	1.306	1.258	0.985
50%, H	1.618	1.508	1.004
75%, L	1.000	1.000	1.000
75%, N	1.306	1.277	0.987
75%, H	1.618	1.500	1.007
act, L	1.000	1.000	1.000
act, N	1.386	1.379	1.000
act, H	1.722	1.596	1.000

73P	Tot dem	Tot dem s	fuel used
50%, L	1.000	1.000	1.000
50%, N	1.285	1.201	1.010
50%, H	1.595	1.343	1.036
75%, L	1.000	1.000	1.000
75%, N	1.285	1.205	1.030
75%, H	1.595	1.334	1.051
act, L	1.000	1.000	1.000
act, N	1.353	1.274	1.000
act, H	1.719	1.387	1.000

75C	Tot dem	Tot dem s	fuel used
50%, L	1.000	1.000	1.000
50%, N	1.340	1.284	0.998
50%, H	1.634	1.537	1.015
75%, L	1.000	1.000	1.000
75%, N	1.340	1.298	1.015
75%, H	1.634	1.528	0.997
act, L	1.000	1.000	1.000
act, N	1.296	1.276	1.000
act, H	1.626	1.483	1.000

It should be noted, however, that this can be explained by the results shown in Table 14. From left to right the columns denote the following: total demand considered, total demand serviced, number of legs in the final network, number of legs flown completely empty, the average leg load factor, the amount of fuel required, the fuel required for flying each leg in the network without payload, the fuel required to transport demand on all of the scheduled flights, the total available freight tonnes, the freight tonne kilometres, the ACMI cost per *kg* of transported demand, the ACMI plus fuel cost per *kg* of transported demand and finally the ratio of the amount of fuel required to transport that amount of demand. Note, that all results are again normalized w.r.t. 73P-ACT-L. Also, the average leg load factor is normalized, which explains why this value in some cases can be more than 1.00. In this table, you can see that for each aircraft type and their respective 50% and 75% instances, the average leg load factor increases. But if the respective number of flown legs in these instances is inspected, it can be seen that this number remains approximately the same. This means that the additional demand that is available when switching from the normal demand scenarios to the high scenarios is simply accommodated on flights that were also scheduled in the results of the 75% instances. Nonetheless, the difference between the results of the base schedule of each of the aircraft and the 75% instances, show a serious improvement.

Finally, it is also important to look at the column which shows the total cost per *kg* of transported demand. Here it can be seen that although the model in most cases performs better than the base schedule, it sometimes also performs slightly less. This might seem like an unfavourable result, but it is important to understand that the network of the model is able to transport much more demand. This has a huge influence on the potential revenue generated by the network. If all the extra demand can be accommodated for either less or just slightly more average cost per transported kg of demand, this is a desirable result. It should be noted, that in reality many cost components are now left out of the picture that would increase the expected operational cost (i.e. making network adjustments is frequently not free of charge). However, it is expected that these costs are of a different order than the potential extra profit that could be generated by transporting the additional demand.

Table 14: Overview of the general results and values of KPIs in the different test instances.

ABY	**dem**	**dem s**	**Nr legs**	**Nr empty legs**	**avg LF**	**fuel [kg]**	**dof [kg]**	**m fuel [kg]**	**ATK**	**FTK**	**€ACMI / kg dem**	**€(ACMI + fuel) / kg dem**	**Fuel kg / pl kg**
50%, L	2.281	2.732	1.468	1.848	0.813	2.693	2.694	1.932	4.272	3.290	1.263	1.013	0.986
75%, L	2.281	2.498	1.468	2.205	0.719	2.755	2.756	1.779	4.372	3.014	1.406	1.097	1.103
act, L	1.574	1.985	1.199	1.607	0.719	2.271	2.272	1.260	3.602	2.168	1.429	1.088	1.144
50%, N	2.979	3.438	1.462	1.616	1.031	2.653	2.653	2.358	4.206	3.986	1.000	0.841	0.772
75%, N	2.979	3.189	1.473	1.875	0.938	2.720	2.720	2.154	4.311	3.638	1.098	0.899	0.853
act, N	2.183	2.737	1.199	1.241	0.969	2.272	2.272	1.857	3.602	3.155	1.038	0.872	0.830
50%, H	3.689	4.121	1.468	1.554	1.188	2.703	2.703	2.867	4.284	4.817	0.842	0.758	0.656
75%, H	3.689	3.748	1.501	1.911	1.063	2.775	2.775	2.526	4.397	4.254	0.955	0.811	0.740
act, H	2.711	3.167	1.199	1.241	1.125	2.272	2.272	2.119	3.602	3.582	0.895	0.784	0.717

75C	**dem**	**dem s**	**Nr legs**	**Nr empty legs**	**avg LF**	**fuel [kg]**	**dof [kg]**	**m fuel [kg]**	**ATK**	**FTK**	**€ACMI / kg dem**	**€(ACMI + fuel) / kg dem**	**Fuel kg / pl kg**
50%, L	1.591	1.952	1.583	2.071	0.719	1.899	1.899	1.287	3.471	2.239	1.459	1.110	0.973
75%, L	1.591	1.638	1.431	2.045	0.625	1.765	1.766	1.038	3.229	1.817	1.602	1.189	1.078
act, L	1.094	1.313	1.025	1.304	0.750	1.297	1.298	0.816	2.374	1.430	1.451	1.106	0.988
50%, N	2.132	2.506	1.599	1.821	0.875	1.895	1.895	1.630	3.461	2.803	1.135	0.921	0.756
75%, N	2.132	2.126	1.499	1.875	0.781	1.791	1.792	1.340	3.273	2.326	1.263	0.991	0.843
act, N	1.417	1.676	1.025	0.893	0.906	1.298	1.298	1.037	2.374	1.812	1.143	0.921	0.774
50%, H	2.600	3.001	1.605	1.821	1.031	1.927	1.927	1.956	3.522	3.324	0.962	0.815	0.642
75%, H	2.600	2.504	1.479	1.848	0.938	1.761	1.761	1.543	3.216	2.654	1.060	0.859	0.703
act, H	1.779	1.948	1.025	0.893	1.031	1.298	1.298	1.210	2.374	2.084	0.977	0.824	0.666

73P	**dem**	**dem s**	**Nr legs**	**Nr empty legs**	**avg LF**	**fuel [kg]**	**dof [kg]**	**m fuel [kg]**	**ATK**	**FTK**	**€ACMI / kg dem**	**€(ACMI + fuel) / kg dem**	**Fuel kg / pl kg**
50%, L	1.732	2.035	1.675	1.661	1.344	1.642	1.642	2.080	1.640	2.065	0.820	0.872	0.807
75%, L	1.732	1.823	1.563	1.696	1.250	1.565	1.564	1.820	1.564	1.810	0.865	0.890	0.858
act, L	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000
50%, N	2.226	2.444	1.700	1.438	1.563	1.659	1.658	2.497	1.656	2.466	0.692	0.793	0.679
75%, N	2.226	2.197	1.625	1.589	1.469	1.612	1.611	2.205	1.610	2.182	0.737	0.819	0.734
act, N	1.353	1.274	1.000	0.705	1.250	1.000	1.000	1.223	1.000	1.219	0.789	0.855	0.785
50%, H	2.762	2.733	1.765	1.616	1.625	1.702	1.701	2.641	1.698	2.609	0.639	0.740	0.623
75%, H	2.762	2.432	1.683	1.679	1.500	1.644	1.643	2.321	1.641	2.294	0.684	0.767	0.676
act, H	1.719	1.387	1.000	0.705	1.313	1.000	1.000	1.327	1.000	1.314	0.722	0.815	0.721

The results previously discussed show that the model was able to determine a set of flights that could accommodate much more of the predicted demand than was possible using the base schedule. Moreover, in most of the cases, the model determined a combination of flights that could operate for less average cost than the base schedule would have done. Even though these results clearly show that there is potential for a tool that gives network planners suggestions on how to adjust the network according to the re-predicted demand, there are also some important remarks that have to be kept in mind when interpreting these results. First of all, the delivery constraints of the requests are simplified in several ways. For all airports in the network different pick-up and delivery times are considered. This may depend on many different aspects, such as the delivery day of the week, night curfews, or certain agreements with airports. These detailed characteristics of the network were not available. Therefore, it was assumed that all demand was available from a certain time onward and should be delivered before a certain time. In reality, not all demand will be available at the beginning of the evening. For example, in the base schedule, two flights to one destination may be planned. The first to accommodate the material that is available at the start of the evening, the other to transport the demand that only came in after the first flight had already departed. In the model, all this material could potentially be transported by one flight at the beginning of the evening making the second flight redundant.

Furthermore, as was also explained in section 3.2, it was assumed that unloading and loading of the aircraft can always be done within the given turnaround times. Additionally, transloading of requests is also assumed to be possible within these turnaround times. In reality, the TATs of aircraft and the transloading time of requests is highly dependent on the facility where an aircraft arrives or departs. Moreover, it is also dependent on the time of arrival or departure. For example, during peak hours it might take longer to sort all material. Experienced network planners could take into account such special cases by applying some contingency measures in the form of more ground time between certain connections. Again, the model has a benefit compared to the base schedule as it only uses the minimum TAT.

Finally, it is important to consider what input demand is used. It was explained in section 3.1 that this research focuses on the express business of our airline. Therefore, the total express demand was used as an input. In reality, this demand can be divided in any way over multiple aircraft types. In our model, however, only a subset of the fleet was available to transport this demand. This means that relatively more express demand could have been made available to this subfleet than in reality would have been the case. This is important to take into account when comparing the figures for total demand serviced by the actual network and those of the model with more freedom.

7 Conclusion

The air freight market is expected to continue to grow in the future. This growth comes together with certain challenges such as coping with demand changes. Currently, experienced network planners have to decide how the network will be adjusted based on this re-predicted demand. However, due to the increasing size of not only the trade flow but also the size of the network itself, it becomes increasingly difficult to determine the most cost-effective way to do this. Therefore, the goal of this research was to design a model that could help network planners with deciding on how to adjust or develop their network. It should use an existing schedule as an input, together with a predicted demand, and based on this, it should determine a set of flights that could accommodate as much of the demand as possible.

The result of this research is a model that combines airport selection, fleet routing and cargo routing, together with the use of a (random) mandatory flight list and timetable setting for all other optional flights. To the best of our knowledge, such a model has not yet been developed. The problem consists of several sub-problems that have been studied in the past for a variety of reasons. Therefore, an extensive literature study has been performed to acquire an understanding of what models are beneficial in what kind of case studies. This knowledge was then used to decide how the to be designed model should be structured. Additionally, different methods were studied that could potentially be used to decrease the total run time of the model.

In the final design of the model two meta-heuristics have been used. These heuristics reduced the pre-processing time of the model from several hours to several tens of seconds. In addition, the computational performance of the model was improved by reducing the number of fractional numbers in the lp-problem. Such numbers are a known source for rounding errors and instability of optimization models. It was also found that symmetry could be a cause of slow convergence. Symmetry exists if there are multiple solutions with the same value. Initially, the costs for paths that passed the same set of airports were equal, which meant that there were many similar solutions the model could find. Therefore, a path-preference model was designed that greatly reduced the amount of symmetry.

The results of the research showed that the model was able to uplift up to 50% more demand than the base schedule could. Moreover, the estimated average cost to transport one kg of demand in most cases reduced as well. Transporting demand is what directly generates revenue for an airline. Therefore, these results clearly show that there is potential for using a model like this. However, as also mentioned in section 6, certain assumptions give the model more opportunities than there in reality are. For example, not all demand will be available at the beginning of the evening. Or transloading of requests could take much longer than the TAT of an aircraft.

However, these findings also naturally lead to interesting future research directions that could improve the designed model. Currently, a random set of flights from the base schedule is fixed. This random set could be replaced by a list of mandatory flights that the network planners could create. With the current setup, flights may be fixed which are known to have zero demand given the re-predicted demand. The model is created such that it allows for easy implementation of this predetermined list.

Additionally, the model could potentially show further improvements if the time step was reduced from 1 hour to for example only 30 minutes. Currently, if the expected flight time plus turnaround time of a leg is 2 hours and 5 minutes, the model will round this up to 3 hours to connect the associated activity nodes. This is done to ensure that the conservation of the flow of aircraft is always adhered to. It should be noted, however, that reducing the step size is expected to have a huge effect on the problem size. As the path-flow model has many more opportunities to route the requests through the network. Therefore, a trade-off would have to be made between computational performance and the desired timetable preciseness.

Moreover, an important aspect of this research is that for each instance only one aircraft type was used. In reality, many more different aircraft are used together to accommodate a certain demand. Therefore, it could be very interesting to analyze what results come out of the model when different aircraft types are mixed. The effect of increasing the size of the problem on computation time could then be further investigated. There are several possibilities to attenuate the speed at which the model grows. For example, the path-flow model currently considers the top 20 fastest paths between any O&D pair, but this could be reduced. Or the minimum amount of demand for an O&D pair to be considered as a request could be increased, such that the number of requests does not grow too rapidly.

References

[Amaruchkul et al., 2007] Amaruchkul, K., Cooper, W. L., and Gupta, D. (2007). Single-leg air-cargo revenue management. *Transportation Science*, 41(4):457–469.

[Angelelli et al., 2020] Angelelli, E., Archetti, C., and Peirano, L. (2020). A matheuristic for the air transportation freight forwarder service problem. *Computers Operations Research*, 123:105002.

[Baxter et al., 2018] Baxter, G., Srisaeng, P., and Wild, G. (2018). The air cargo carrying potential of the airbus a350-900xwb and boeing 787-9 aircraft on their ultra-long-haul flights: A case study for flights from san francisco to singapore. *Transport and Telecommunication Journal*, 19:301–314.

[Berge and Hopperstad, 1993] Berge, M. E. and Hopperstad, C. A. (1993). Demand driven dispatch: A method for dynamic aircraft capacity assignment, models and algorithms. *Operations Research*, 41(1):153–168.

[Boeing, 2018] Boeing (2018). World air cargo forecast 2018 - 2037. Available at "https://file.veryzhun.com/buckets/carnoc/keys/3fa55da709101d0d937e78732a88cd9d.pdf".

[Bélanger et al., 2006] Bélanger, N., Desaulniers, G., Soumis, F., and Desrosiers, J. (2006). Periodic airline fleet assignment with time windows, spacing constraints, and time dependent revenues. *European Journal of Operational Research*, 175(3):1754–1766.

[Delgado et al., 2020] Delgado, F., Sirhan, C., Katscher, M., and Larrain, H. (2020). Recovering from demand disruptions on an air cargo network. *Journal of Air Transport Management*, 85:101799.

[Derigs and Friederichs, 2013] Derigs, U. and Friederichs, S. (2013). Air cargo scheduling: integrated models and solution procedures. *OR Spectrum*, 35(2):325 – 362.

[Derigs et al., 2009] Derigs, U., Friederichs, S., and Schäfer, S. (2009). A new approach for air cargo network planning. *Transportation Science*, 43:370–380.

[Gopalakrishnan and Johnson, 2005] Gopalakrishnan, B. and Johnson, E. (2005). Airline crew scheduling: State-of-the-art. *Annals of Operations Research*, 140:305–.

[Huang and Lu, 2015] Huang, K. and Lu, H. (2015). A linear programming-based method for the network revenue management problem of air cargo. *Transportation Research Procedia*, 7:459 – 473. 21st International Symposium on Transportation and Traffic Theory Kobe, Japan, 5-7 August, 2015.

[Jones et al., 1993] Jones, K. L., Lustig, I. J., Farvolden, J. M., and Powell, W. B. (1993). Multicommodity network flows: The impact of formulation on decomposition. *Mathematical Programming*, 62(1-3):95117.

[Li et al., 2006] Li, D., Huang, H.-C., Morton, A. D., and Chew, E.-P. (2006). Simultaneous fleet assignment and cargo routing using benders decomposition. *OR Spectrum*, 28(3):319335.

[Lohatepanont and Barnhart, 2004] Lohatepanont, M. and Barnhart, C. (2004). Airline schedule planning: Integrated models and algorithms for schedule design and fleet assignment. *Transportation Science*, 38:19–32.

[ReportLinker, 2021] ReportLinker (2021). Air cargo market forecast to 2028 - covid-19 impact and global analysis by type, service, and end user. Available at "https://www.reportlinker.com/p06130831/Air-Cargo-Market-Forecast-to-COVID-19-Impact-and-Global-Analysis-By-Type-Service-and-End-User.html".

[Rexing et al., 2000] Rexing, B., Barnhart, C., Kniker, T., Jarrah, A., and Krishnamurthy, N. (2000). Airline fleet assignment with time windows. *Transportation Science*, 34(1):120.

[Sandhu and Klabjan, 2006] Sandhu, R. and Klabjan, D. (2006). Fleeting with passenger and cargo origin-destination booking control. *Transportation Science*, 40:517 – 528.

[Tang et al., 2008] Tang, C.-H., Yan, S., and Chen, Y.-H. (2008). An integrated model and solution algorithms for passenger, cargo, and combi flight scheduling. *Transportation Research Part E: Logistics and Transportation Review*, 44(6):1004 – 1024.

[Wink, 2020] Wink, R. (2020). *A Novel Aircraft Market Assessment Model under Flexible Passenger Demand*. PhD thesis, Delft University of Technology, Delft.

[Yan et al., 2005] Yan, S., Chen, S.-C., and Chen, C.-H. (2005). Air cargo fleet routing and timetable setting with multiple on-time demands. *Transportation Research Part E: Logistics and Transportation Review*, 42(5):409 – 430.

[Zhou et al., 2020] Zhou, L., Liang, Z., Chou, C.-A., and Chaovalitwongse, W. (2020). Airline planning and scheduling: Models and solution methodologies. *Frontiers of Engineering Management*, 7:1–26.

Appendices

A Full Service Analsyis

In this section, a more detailed overview of the service analysis is provided. Compared to the one shown in the paper, this contains two additional rows that are denoted by "fully s" and "part s". The first refers to requests that were completely serviced, while the second refers to requests that are only partially serviced. In the case of the latter, this means that the no-service arc has been used to transport some of the demand. Furthermore, there is also one additional column, namely the "nr req". It represents the total number of requests that fall under the umbrella of a certain row. This distinction was made to be able to achieve a better understanding of what demand was or was not transported. Moreover, it could be used to validate whether the model is performing the way it was intended to behave. For example, it was analyzed how the results of the model varied if the cost for no-service were unequal. One of the experiments that was performed was tailoring the no-service cost to the maximum possible cost for shipping each request through the network. This was done by analyzing the potential cost that would be inquired if only that request would be shipped by a certain path. Also, the cost for all individual flights in this path was taken into account. The reason for this is that the model should always choose to ship a request via the network even if it results in unfavorable load factor performance. Naturally, this meant that the no-service cost for requests that required more individual flights and that were further apart were the highest. In these tables, it was visible that the model preferred to service requests of which the average weighted distance was further apart than was the case when the no-service costs were equal. Furthermore, it was found that the number of requests that were only partially serviced was extremely low, which is an interesting result. After further investigating this finding, it was determined that the relatively tight delivery constraint caused the model to choose a combination of flights that could be performed within this time window which in turn minimized the possible no-service cost. For future research, it could be interesting to analyze how the model would behave with more long-distance flights and a less dense network. The data in Table 15, Table 16 and Table 17 have been normalized w.r.t. instance 73P-ACT-L, which is highlighted in blue.

Table 15: Detailed service analysis for all instances with low demand.

ABY, 50%, L

serv	nr req	dem	avg size	avg w dist	w det fac
all	1.540	2.279	1.494	1.239	n.a.
tot s	1.268	2.130	1.688	1.209	1.074
fully s	1.258	2.077	1.662	1.186	1.074
part s	0.010	0.053	5.364	2.111	1.082
no s	0.272	0.147	0.545	1.670	n.a.

75C, 50%, L

serv	nr req	dem	avg size	avg w dist	w det fac
all	1.305	1.589	1.221	1.142	n.a.
tot s	1.169	1.522	1.312	1.136	1.082
fully s	1.159	1.490	1.299	1.142	1.077
part s	0.010	0.032	3.247	0.860	1.321
no s	0.136	0.061	0.455	1.318	n.a.

73P, 50%, L

serv	nr req	dem	avg size	avg w dist	w det fac
all	1.430	1.731	1.221	0.978	n.a.
tot s	1.245	1.587	1.286	0.948	1.175
fully s	1.123	1.386	1.247	0.936	1.154
part s	0.123	0.200	1.649	1.023	1.321
no s	0.185	0.105	0.571	1.341	n.a.

ABY, 75%, L

serv	nr req	dem	avg size	avg w dist	w det fac
all	1.540	2.279	1.494	1.239	n.a.
tot s	1.109	1.949	1.766	1.209	1.077
fully s	1.086	1.848	1.714	1.165	1.078
part s	0.023	0.099	4.312	2.020	1.062
no s	0.430	0.315	0.740	1.406	n.a.

75C, 75%, L

serv	nr req	dem	avg size	avg w dist	w det fac
all	1.305	1.589	1.221	1.142	n.a.
tot s	0.950	1.278	1.351	1.115	1.069
fully s	0.934	1.236	1.338	1.114	1.071
part s	0.017	0.041	2.494	1.179	1.005
no s	0.354	0.304	0.857	1.243	n.a.

73P, 75%, L

serv	nr req	dem	avg size	avg w dist	w det fac
all	1.430	1.731	1.221	0.978	n.a.
tot s	1.116	1.421	1.286	0.935	1.172
fully s	1.000	1.231	1.234	0.931	1.156
part s	0.116	0.189	1.649	0.960	1.275
no s	0.315	0.265	0.844	1.184	n.a.

ABY, act, L

serv	nr req	dem	avg size	avg w dist	w det fac
all	1.079	1.572	1.468	1.067	n.a.
tot s	1.066	1.547	1.455	1.060	1.094
fully s	1.033	1.467	1.429	1.036	1.086
part s	0.033	0.079	2.416	1.517	1.242
no s	0.013	0.010	0.740	1.800	n.a.

75C, act, L

serv	nr req	dem	avg size	avg w dist	w det fac
all	0.884	1.091	1.247	1.097	n.a.
tot s	0.841	1.024	1.221	1.069	1.099
fully s	0.798	0.887	1.117	1.073	1.077
part s	0.043	0.137	3.208	1.041	1.239
no s	0.043	0.025	0.584	1.467	n.a.

73P, act, L

serv	nr req	dem	avg size	avg w dist	w det fac
all	1.000	1.000	1.000	1.000	n.a.
tot s	0.861	0.779	0.909	0.980	1.091
fully s	0.732	0.507	0.701	0.997	1.053
part s	0.129	0.272	2.117	0.945	1.163
no s	0.139	0.092	0.662	1.283	n.a.

Table 16: Detailed service analysis for all instances with normal demand.

ABY, 50%, N

serv	nr req	dem	avg size	avg w dist	w det fac
all	2.056	2.975	1.455	1.206	n.a.
tot s	1.613	2.680	1.675	1.155	1.085
fully s	1.570	2.498	1.597	1.090	1.086
part s	0.043	0.181	4.247	2.059	1.065
no s	0.444	0.283	0.636	1.656	n.a.

75C, 50%, N

serv	nr req	dem	avg size	avg w dist	w det fac
all	1.742	2.128	1.234	1.124	n.a.
tot s	1.470	1.957	1.338	1.100	1.092
fully s	1.430	1.877	1.325	1.100	1.077
part s	0.040	0.076	1.922	1.095	1.465
no s	0.272	0.161	0.597	1.415	n.a.

73P, 50%, N

serv	nr req	dem	avg size	avg w dist	w det fac
all	1.907	2.222	1.169	0.972	n.a.
tot s	1.530	1.907	1.260	0.941	1.170
fully s	1.344	1.610	1.208	0.935	1.136
part s	0.185	0.294	1.597	0.972	1.356
no s	0.377	0.241	0.636	1.173	n.a.

ABY, 75%, N

serv	nr req	dem	avg size	avg w dist	w det fac
all	2.056	2.975	1.455	1.206	n.a.
tot s	1.437	2.486	1.740	1.140	1.078
fully s	1.397	2.267	1.636	1.052	1.080
part s	0.040	0.219	5.558	2.044	1.056
no s	0.619	0.468	0.753	1.525	n.a.

75C, 75%, N

serv	nr req	dem	avg size	avg w dist	w det fac
all	1.742	2.128	1.234	1.124	n.a.
tot s	1.199	1.658	1.390	1.091	1.083
fully s	1.156	1.567	1.364	1.090	1.072
part s	0.043	0.090	2.104	1.101	1.280
no s	0.543	0.454	0.844	1.233	n.a.

73P, 75%, N

serv	nr req	dem	avg size	avg w dist	w det fac
all	1.907	2.222	1.169	0.972	n.a.
tot s	1.424	1.714	1.208	0.930	1.170
fully s	1.268	1.435	1.143	0.932	1.134
part s	0.156	0.277	1.792	0.915	1.361
no s	0.483	0.405	0.844	1.140	n.a.

ABY, act, N

serv	nr req	dem	avg size	avg w dist	w det fac
all	1.440	2.182	1.519	1.132	n.a.
tot s	1.417	2.133	1.519	1.120	1.099
fully s	1.364	1.924	1.416	1.044	1.092
part s	0.053	0.209	3.974	1.819	1.164
no s	0.023	0.015	0.675	2.012	n.a.

75C, act, N

serv	nr req	dem	avg size	avg w dist	w det fac
all	1.179	1.415	1.208	1.128	n.a.
tot s	1.113	1.306	1.182	1.076	1.082
fully s	1.060	1.159	1.104	1.081	1.064
part s	0.053	0.147	2.805	1.032	1.223
no s	0.066	0.047	0.714	1.777	n.a.

73P, act, N

serv	nr req	dem	avg size	avg w dist	w det fac
all	1.334	1.353	1.013	0.980	n.a.
tot s	1.109	0.992	0.896	0.953	1.073
fully s	0.940	0.656	0.701	0.956	1.045
part s	0.169	0.336	2.013	0.945	1.128
no s	0.225	0.154	0.688	1.251	n.a.

Table 17: Detailed service analysis for all instances with high demand.

ABY, 50%, H

serv	nr req	dem	avg size	avg w dist	w det fac
all	2.056	3.685	1.805	1.224	n.a.
tot s	1.546	3.213	2.091	1.169	1.082
fully s	1.493	2.960	2.000	1.110	1.085
part s	0.053	0.252	4.792	1.868	1.042
no s	0.510	0.425	0.831	1.542	n.a.

75C, 50%, H

serv	nr req	dem	avg size	avg w dist	w det fac
all	1.742	2.600	1.506	1.111	n.a.
tot s	1.411	2.341	1.675	1.100	1.077
fully s	1.361	2.193	1.623	1.110	1.068
part s	0.050	0.146	2.974	0.953	1.220
no s	0.331	0.239	0.727	1.229	n.a.

73P, 50%, H

serv	nr req	dem	avg size	avg w dist	w det fac
all	1.907	2.758	1.455	0.983	n.a.
tot s	1.387	2.133	1.545	0.918	1.133
fully s	1.189	1.649	1.390	0.895	1.116
part s	0.199	0.482	2.442	0.999	1.192
no s	0.520	0.449	0.870	1.261	n.a.

ABY, 75%, H

serv	nr req	dem	avg size	avg w dist	w det fac
all	2.056	3.685	1.805	1.224	n.a.
tot s	1.364	2.922	2.156	1.138	1.077
fully s	1.281	2.590	2.039	1.074	1.070
part s	0.083	0.331	4.039	1.640	1.138
no s	0.692	0.692	1.000	1.562	0

75C, 75%, H

serv	nr req	dem	avg size	avg w dist	w det fac
all	1.742	2.600	1.506	1.111	n.a.
tot s	1.189	1.952	1.649	1.058	1.075
fully s	1.132	1.776	1.584	1.067	1.066
part s	0.056	0.175	3.143	0.965	1.165
no s	0.553	0.613	1.117	1.271	n.a.

73P, 75%, H

serv	nr req	dem	avg size	avg w dist	w det fac
all	1.907	2.758	1.455	0.983	n.a.
tot s	1.325	1.897	1.442	0.908	1.133
fully s	1.129	1.442	1.286	0.859	1.096
part s	0.195	0.453	2.338	1.059	1.251
no s	0.583	0.621	1.065	1.198	n.a.

ABY, act, H

serv	nr req	dem	avg size	avg w dist	w det fac
all	1.440	2.708	1.896	1.134	n.a.
tot s	1.374	2.468	1.805	1.101	1.095
fully s	1.288	2.063	1.610	1.064	1.091
part s	0.086	0.405	4.740	1.286	1.115
no s	0.066	0.038	0.584	1.467	n.a.

75C, act, H

serv	nr req	dem	avg size	avg w dist	w det fac
all	1.179	1.778	1.519	1.127	n.a.
tot s	1.076	1.518	1.416	1.077	1.069
fully s	0.983	1.215	1.247	1.082	1.043
part s	0.093	0.304	3.312	1.054	1.169
no s	0.103	0.075	0.740	1.738	n.a.

73P, act, H

serv	nr req	dem	avg size	avg w dist	w det fac
all	1.334	1.717	1.299	0.986	n.a.
tot s	1.076	1.081	1.013	0.949	1.063
fully s	0.884	0.674	0.766	0.942	1.039
part s	0.192	0.407	2.130	0.960	1.104
no s	0.258	0.230	0.896	1.260	n.a.

II

Literature Study

A decision-making tool to improve an existing flight schedule of a full cargo airline

Literature Study

Anne Bart Beijneveld

TU Delft

A decision-making tool to improve an existing flight schedule of a full cargo airline

Author: A.B.M. Beijneveld

Supervisor: Dr. A. Bombelli

This literature study is submitted in fulfillment for the course AE4020

March 2021

Signed:
Date: Tuesday 23rd March, 2021

Contents

List of Figures

List of Tables

Acronyms

ACSRP Air Cargo Schedule Recovery Problem.

ACTN Air cargo transportation network.

AFT Available freight tonne.

AFTK Available freight tonne kilometer.

ALLFV Average Leg Load Factor (*volume*).

ALLFW Average Leg Load Factor (*weight*).

B&B Branch and bound.

B&P Branch and price.

B&P& Branch and price and cut.

CN Connection network.

DST Destination.

DV Decision variables.

ESSND Express Shipment Service Network Design.

Fl Nr Flight number.

H&S Hub-and-spoke.

KPI key performance indicators.

MILP Mixed-integer linear programming.

MNFP Multicommodity network flow problem.

MP Master problem.

O&D Origin and destination.

OR Operational research.

ORG Origin.

PP Point-to-point.

RMP Restricted master problem.

TOC Total operating cost.

TSN Time-space network.

Acronyms

1 Introduction

In the past few decades, technological developments in the air transportation industry have made air transport cheaper, faster, and safer. Thereby, it has become indispensable for the global economy. Boeing has predicted an average growth rate of 4.2 percent for the air freight industry due to the world trade growth (Boeing (2018)). Tree types of carriers are responsible for arranging transport of the total demand, namely the combination carriers, the cargo-only carriers, and the integrators. The combination carriers carry both passengers and cargo, while the cargo-only carriers and the integrators carry cargo only. Integrators are a particular form of cargo-only carriers; where cargo-only carriers offer airport-to-airport services, the integrators offer door-to-door services. The cargo-only carriers' network is explicitly designed for the air freight demand, while the network of the combination carriers is designed for passenger demand only. Thereby, the cargo-only carriers have the advantage of being able to offer a more flexible network. For a more detailed analysis of the different business models, the reader is referred to Popescu et al. (2011).

Designing a feasible and profitable flight schedule is a very complex task, especially for cargo-only carriers that have to cope with trade imbalances. These imbalances make it difficult to design the network such that aircraft do not fly empty legs. On top of this, accurately predicting air cargo demand is complex. This unpredictability is a consequence of several things. First of all, capacity is two-dimensional, meaning that both the weight and the volume of goods influence how much can be transported in total. Furthermore, an estimation is made of the weight and volume of the goods to be shipped when the booking is made, but this often differs from the actual size on the day of transport (Huang & Lu (2015)). Second, cargo bookings are often placed relatively close to the day of the actual flight, making it challenging to ensure that enough capacity is available on each route (Sandhu & Klabjan (2006)). Finally, it also frequently occurs that bookings do not show up at all (Amaruchkul et al. (2007)).

As a consequence of this unreliability, carriers are frequently forced to reschedule their flights. In many airlines, this process is performed manually by experienced planners. For some smaller networks, it might be possible to find an optimal solution manually. However, for larger networks, the solution space may become rather large. In this case, the chances of determining the most cost-effective solution manually are minimal. Therefore, the objective of the research is defined as follows:

> *To develop a tool that can assist a full-cargo airline in deciding how to adjust their flight schedule such that it can transport the re-predicted demand most cost-effectively.*

The model shall be formulated as a mixed-integer linear programming (MILP) model. Several solution techniques will be compared to determine which one is most effective in solving the model. This research's first challenge is to determine how different models can be integrated to model the network. The second challenge is to formulate a method that enables the model to find an optimized aircraft routing that allows the network to accommodate all of the predicted demand. Furthermore, it is well-known that the problem sizes can become rather large. Therefore, the final challenge is determining the most suitable solution technique that can handle such large problem instances most effectively. Several techniques might have to be combined to accomplish this goal.

In the project plan, a general overview of the available literature was provided. To guide the search for this literature, several research questions had been formulated. The purpose of this report is to provide a more in-depth analysis of this literature such that the research questions can be answered. Furthermore, it will provide guidelines to develop the final model. Moreover, it will provide the reader with a clear overview of the state-of-the-art literature. The remainder of the report is structured as follows. In chapter 2 the research outline is explained. After that, in chapter 3 an overview of the air cargo transport supply chain is provided. Then, in chapter 4 an elaborate analysis of how flight schedules are designed is given. Hereafter, it is discussed in chapter 5 what different solution techniques have been applied to solve the flight schedule design problems. Finally, in chapter 6, it will be explained how all the previous chapters' information will be combined into one model.

2 Research Outline

This chapter discusses the outline of the research. First, in section 2.1 the research problem is defined. An explanation of the research objective follows this in section 2.2. After that, in section 2.3 the research questions that were formulated are presented. Finally, the project planning is briefly elaborated upon in section 2.4.

2.1 Research Problem

For every airline, flight schedule design is the central element of the planning process. Several problems have to be solved either simultaneously or sequentially: determine which markets should be connected by direct or indirect flights, determine how frequently these markets should be connected, determine how much capacity should be offered on these routes, and determine how the available aircraft are routed (Derigs et al. (2009)). These planning decisions have to be made several months in advance based on the then-predicted demand. However, as mentioned earlier, it is hard to predict cargo demand accurately. Nonetheless, thanks to sophisticated forecasting models, a reasonable estimate can be made several months in advance. Moreover, the closer the day of operations becomes, the more accurate the forecast usually gets. It is constantly verified whether the current schedule would be able to accommodate all of this expected demand. If this is not the case, then the flight schedule has to be adjusted accordingly.

In practice, reevaluating a schedule and adjusting it according to specific requirements is called *schedule development*. This has been studied both for passengers and cargo by Lohatepanont & Barnhart (2004) and Derigs et al. (2009) respectively. Both of the models that were developed for these researches require a list of mandatory and optional flights as input. The network planners must prepare this flight list. As a consequence, the input of the planners influences the optimality of the final solution. Moreover, the objective of the models is to determine the most profitable combination of flights. This research aims to minimize the total operating cost required to transport all of the forecast demand. So, this research does not consider revenue generated by delivering cargo. Furthermore, the goal is to build a model that is independent of the level of expertise of planners. Therefore, no input from planners will be required.

A model that required similar inputs as are expected to be used for this research is the one developed by Delgado et al. (2020). This model determined how a flight schedule could best be adjusted in response to last-minute demand changes while also taking into account crew rescheduling cost. The authors defined this as the *air cargo schedule recovery problem*. Although specific short-term recovery issues, such as no-shows and crew rescheduling, do not apply to this research (due to the planning horizon of several weeks), this research can serve as a source of inspiration for the to-be-built model. Furthermore, their model incorporates penalties for deviating from the original schedule. This is logical for their model as it is developed to cope with last-minute changes. Due to the planning horizon of several weeks, there is more freedom to reschedule aircraft in this research. To the best of our knowledge, a tool that can assist a full-cargo airline in deciding how to adjust their flight schedule such that it can transport the re-predicted demand most cost-effectively has not yet been developed. It is expected that such a tool will be valuable for airlines.

2.2 Research Objective

The previous section explained what this research is expected to contribute to the body of knowledge. Formally, the objective of the research is defined as:

> *To develop a tool that can assist a full-cargo airline in deciding on how to adapt their flight schedule such that it can transport the re-predicted demand most cost-effectively.*

Up and until now, to cope with changes in the forecast demand, the flight schedules of airlines have been adjusted manually by experienced planners. However, for certain networks, the solution space is so ample that it is challenging to find an optimal solution manually. The goal is to develop a tool that can determine rescheduling options that reduce airlines' operational costs more than was achieved with manual solutions. The model will be tested with data that a large cargo-only carrier will provide. The data will consist of historic flight schedules, demand forecast, and cargo movements through the network. The historical data can serve as a benchmark solution. Then, it can be determined what the quality of the solution from the model is by comparing it with the benchmark.

2.3 Research Questions

It was determined that a decision-making tool that can optimize an existing flight schedule based on a particular demand has not yet been developed. The aim of the authors is not only to develop such a tool but also to determine what the potential cost benefits for an airline would be if such a tool would be used. Therefore, the main research question of this thesis has been defined as follows:

What cost improvements could potentially be achieved by using a decision-making tool that can determine how flight rotations in an existing flight schedule of a full-cargo airline should be adjusted based on the expected demand?

The main research question is split into several sub-questions in order to answer it. The answers to these sub-questions will, together, provide the answer to the main question posed above. Below each sub-question, a short explanation is provided of why it is a relevant question for the main research.

1. *How should the uncertainties in the forecast demand be taken into account when redesigning an existing schedule?*
   Demand for air freight is highly unpredictable. Only moments before take-off it is known how much cargo will be loaded onto the aircraft. Therefore, it will be important to learn how planners incorporate these uncertainties when designing the base schedule. Understanding how this is done will greatly improve the feasibility of the proposed solution.

2. *What are the most important constraints that should be taken into account to be able to find a feasible fleet rotation?*
   Designing a schedule that is both feasible and profitable is extremely difficult. An in-depth analysis should be conducted such that it can be determined what the common pitfalls and challenges are to design such a schedule.

3. *What are the cost components that influence the choice of a path between a certain O-D pair the most?*
   It is expected that a thorough understanding of the cost build-up is important to be able to determine which rotations can best be added or deleted from the system. The reason for this is that independent of which network representation is chosen, the edge weights have to be determined. One or several of these cost components could potentially be used for this. The choice of edge weight is expected to have a large influence on the decision-making tool's outcome.

4. *What solution techniques are suitable to solve the flight schedule optimization problem?*
   The goal is to develop a tool that can assist network planners with designing a network that can transport the complete demand as cost-effectively as possible. The solution space of this problem can become rather large. Therefore, it will be essential to determine what kind of methods can solve such a problem in a timely manner. However, since the run time for this research is not too stringent, an acceptable run time is less important than finding an optimal solution. Therefore, the focus of determining the most appropriate solution techniques should be on analyzing their potential to find (near) optimal solutions.

2.4 Research Planning

The complete thesis can be divided into four phases which are separated by certain meetings. The first phase consists of performing a literature study and writing a project plan. Once these two reports are handed-in, a Kick-off meeting will be held. Hereafter, the second phase of the thesis starts, which is also known as the initial phase. In this phase, the first part of the research will be initiated, meaning that data will be gathered, and the first steps of developing a model are taken. After about three months, a Mid-term meeting will be planned. During this meeting, the progress and project management of the researcher will be evaluated by the supervisors. The supervisors' feedback can then be used to start on the final phase, which will also take about three months. In this phase, validation and verification of the study will be done. Once the researcher has almost finished these tasks, a Green Light Meeting will be planned. During this meeting, the final results of the study will be presented. Once the study achieves the green light, the last points of improvement can be incorporated in the report, and the researcher will defend her thesis about two weeks later. An overview of the entire time horizon is shown in Figure 2.1.

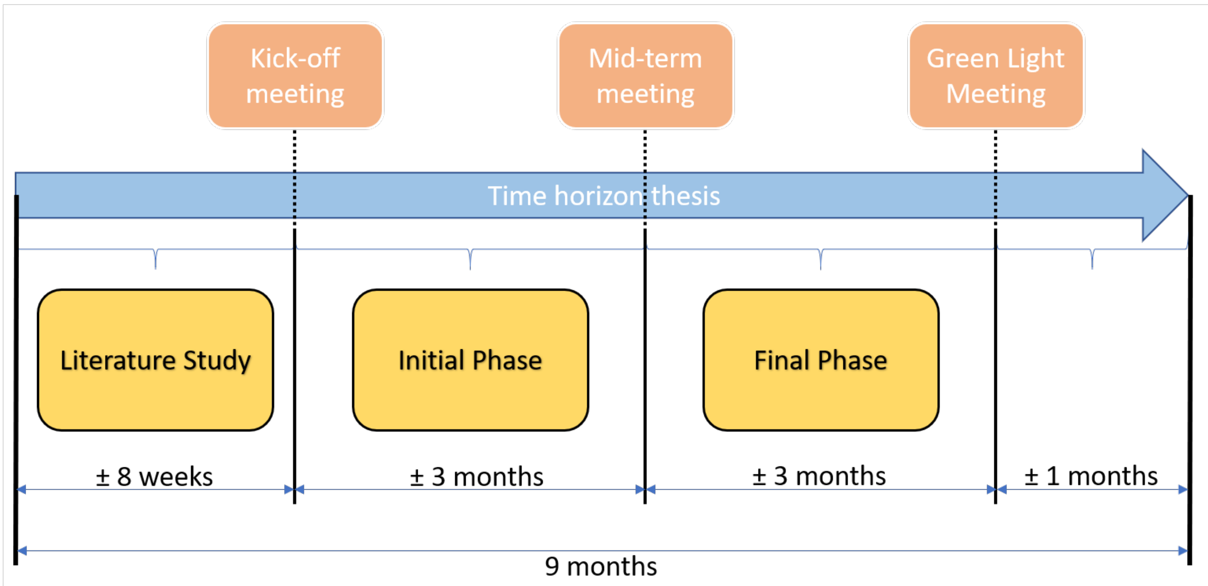

Figure 2.1: Overview of the time horizon for the thesis procedure for the department of Air Transport Operations of the Aerospace Engineering faculty of the TU Delft.

3 The Air Cargo Transport Supply Chain

Much of the research conducted in the area of air transportation networks has been focused on passenger transportation. There are many similarities between the transportation of passengers and freight. However, some important discrepancies also exist. First, these differences are highlighted in section 3.1. Thereafter, the main business models of full cargo airlines and integrators are discussed in section 3.2. Finally, in section 3.3, the air freight network's key performance indicators are presented.

3.1 Difference between passenger and freight transport

Over the past decades, a wide variety of topics within the air transportation industry have been studied, most often focusing on the transportation of passengers. However, in recent years it has become increasingly popular to investigate the counterpart of passenger transport, namely cargo transport. First of all, the definition of the two options should be clear. Airfreight or cargo is usually defined as anything other than passengers and luggage belonging to passengers. Where passengers in general travel only with commercial passenger aircraft, freight can be transported in several ways. Either it is transported in the belly of a passenger aircraft or in dedicated freighters, which is a type of aircraft that is solely used to transport freight. Freighter aircraft fly both scheduled and unscheduled. It is essential to understand the differences between passenger and freight transport. Then, it can be determined how the models developed for passenger transport could be adjusted to be applicable for cargo transport.

The first notable difference between passenger and cargo transport is the difference in control of the flow through the operator's network. Airlines that transport passengers have no control over the final itinerary a passenger decides to follow. On the other hand, airlines that transport freight have the complete freedom to choose any itinerary for packages as long as they are delivered to the final destination on time. Such a network, where the set of locations and routing options are centrally planned, is what O'Kelly (1998) described as "delivery systems". The opposite of this system, which is typically seen for passenger airlines, is what O'Kelly called a "user attracting" system. Decision-makers benefit from the centrally planned system, as it gives them, for example, the freedom to decide where the larger sorting centers are placed. These sorting centers are essential to process the consolidated flows as efficiently as possible. On the other hand, decision-makers have to make a more educated guess of how customers will use the facilities for the user-attracting system. The delivery system and the user attracted system lie at the poles of a continuum, with most airlines operating a combination of the two.

A second significant difference is that the flow of passengers is more or less symmetric and that of freight asymmetric. This is because passengers usually make a round trip, while freight usually has a one-way trip from the origin to its final destination. This would not be a problem if trade flows around the world would have been of equal size. However, the trade flows are greatly imbalanced. For example, in 2019, 2.7 million tonnes of freight were transported from East Asia to North America, while only 1.7 million tonnes were transported from North America to East Asia (Boeing (2020)). To cope with these imbalances, full-cargo airlines have introduced triangular and circular routes in their network (Gardiner et al. (2005)). This means that if a flight route from A to B exist, that one may not assume that the return flight from B to A exists as well. In the case of passenger transport, this usually is safe to assume. Therefore, the most realistic way to model an air cargo transportation network (ACTN) is by use of a directed graph. Another reason to introduce circular routes into the network is for maintenance reasons, as maintenance of full freighters is usually performed at fixed intervals at the hub of an airline (Bombelli et al. (2020)).

Another interesting point is that when designing an ACTN, one does not have to consider passenger preferences. For example, Paleari et al. (2010) compared different networks to determine which provided the best service to passengers. To do so, they determined certain so-called 'key performance indicators (KPIs) such as the average travel time between any OD pair in 1, 2, 3, or more than 3 steps, the corresponding average waiting times, and the routing factor (the ratio between in-flight distance and potential direct flight distance). KPIs like these do not directly matter for cargo transport. On the contrary, it is not relevant as long as the cargo is delivered at the final destination before the latest allowable time of arrival (Sandhu & Klabjan (2006); Amaruchkul et al. (2007)). The KPIs that can be used to judge the performance of an ACTN will be discussed in section 3.3. As another example, passengers prefer to travel during the day, while cargo usually has peak operations during the night (O'Kelly (2014)).

3.2 Overview of the different air freight carriers

The air freight market can be roughly divided into two types of carriers: integrators and non-integrators. Integrators organize everything from the start of the journey at the shipper to the end of the journey at the consignee. On the other hand, non-integrators provide airport to airport delivery of freight. Both types of carriers will be more thoroughly analyzed in this section. An schematic representation of possible air cargo delivery options is shown in Figure 3.1.

Integrators are entirely integrated across all transport modes. Currently, only three companies - FedEx, UPS, and DHL - are almost completely responsible for serving this part of the market. A market that a small group of sellers dominates is called an oligopoly. The main challenge for the analysis of an oligopoly is the interdependence of the competitors, which means that the assumptions made about the most likely reactions or actions of competitors influence the optimal strategic behavior of each of the other competitors (Lipczynski et al. (2017)). Furthermore, an integrator is generally the owner of all assets that are involved in transporting freight from the shipper to the consignee. This includes physical assets such as aircraft, trucks, labor assets, and information assets (Forster & Regan (2001)).

Non-integrators can be divided into combination carriers and all-cargo carriers. Combination carriers are carriers that combine the transport of passengers with that of cargo. Some of these carriers limit this part of their business to the transport of express packages and mail. However, there are also examples of combination carriers that have both passenger aircraft and dedicated freight aircraft. Examples of combination carriers are AirFrance KLM, Lufthansa, and Emirates. All-cargo airlines only operate dedicated freight aircraft. The two types of services these airlines provide are scheduled and unscheduled services, whereas the latter is often referred to as charter services. These carriers mainly operate in markets where the competition from integrators and combination carriers is limited, such as Africa, Latin America, and the Middle East (Group (2009)). Examples of all-cargo carriers are Cargolux, Atlas Air, and Nippon Cargo Airlines. An essential difference between all-cargo carriers and combination carriers is that all-cargo carriers' network is explicitly designed for the air freight demand. In contrast, the network of the combination carriers is designed for passenger demand only. Thereby, the all-cargo carriers have the advantage of being able to offer a more flexible network. For a more detailed analysis of the different business models, the reader is referred to Popescu et al. (2011).

Finally, it should be noted that integrators can also be seen as all-cargo carriers. Therefore there is also an arrow drawn from shipper to all-cargo carriers and from there to the consignee. This research aims to develop a tool that is useful for every type of all-cargo carrier.

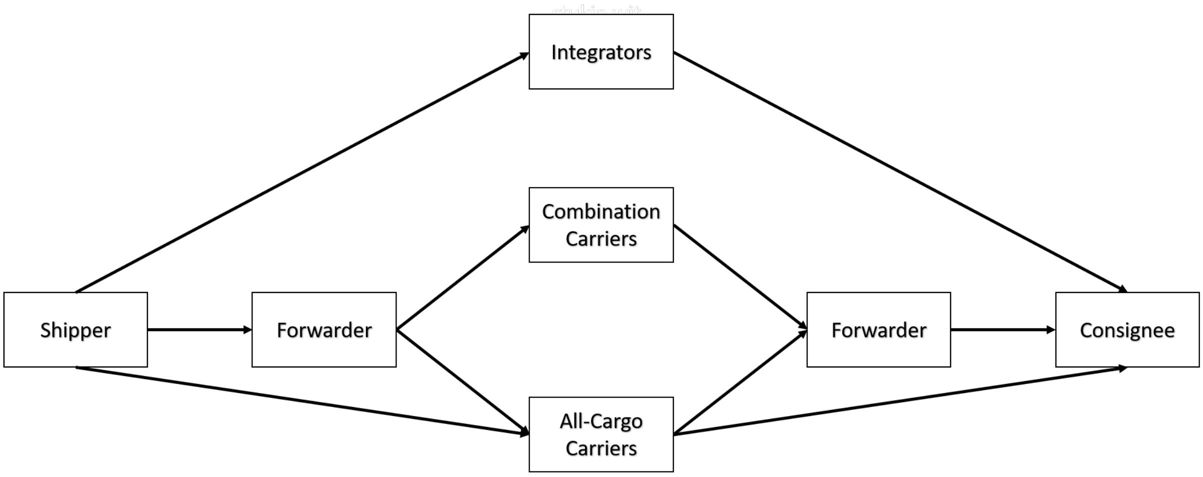

Figure 3.1: A schematic representation of possible air cargo delivery options. Source: Adapted from Forster & Regan (2001).

3.3 Key Performance Indicators

A wide variety of key performance indicators exist in the air transportation industry. They can be used to compare the performance of different companies with each other but also to evaluate the performance of a network. For this study, it will be beneficial to use the KPIs to compare the tool's quality of different solutions to each other. Moreover, it should be noted that the KPIs for passenger transport differ from those for cargo transport, as was already mentioned in the previous chapter. This research is focused on determining an improved flight schedule. Therefore, the discussion of KPIs in this chapter is limited to those that are especially useful for evaluating an airline's network performance. In Table 3.1 an overview is provided of KPIs that can be used to evaluate the performance of an ACTN. A short explanation of each of these KPIs is provided below the table. All of the KPIs were extracted from a list of KPIs used in a study that compared 47 air cargo carriers and clustered them based on management strategies (Dewulf et al. (2014)).

Total Operating Cost (TOC): The total operating cost of an ACTN is composed of a large number of different factors, such as fixed aircraft cost, fuel cost, labor cost, ground-handling cost, etc. Retrieving accurate data about all these factors is highly elusive. However, certain cost components can be measured to a certain extent, such as fixed aircraft cost, the block hour cost of aircraft, and fuel cost. It is assumed that these cost components can provide a good indication of the quality of a solution in terms of cost.

Total Hours Flown: This KPI measures the total flight hours all aircraft made within a certain period. In air travel, the total number of hours flown is a much more relevant indicator than the

Key Performance Indicator	Unit
Total Operating Cost (TOC)	USD
Total Hours Flown	Hrs
Total Tonnes Carried	Tonnes
Total Fuel Consumption	Gallons
Available Freight Tonne (AFT)	Tonnes
Available Freight Tonne Kilometer (AFTK)	Tonnes · km
Average Leg Load Factor *(weight)* (ALLFW)	[-]
Average Leg Load Factor *(volume)* (ALLFV)	[-]
On-time performance	%
Operating Cost per transported tonne	USD / tonnes
Hub performance	Tonnes / Hub
Average length of transported tonne	Km
Average transportation time of transported tonne	Hrs

Table 3.1: Key performance indicators that can be used to evaluate the performance of an ACTN.

total distance flown. This is because flying in East or West direction can have a significant influence on the total flight time. For example, flying from New York to Amsterdam will take about 7 hours, while the return trip will take about 8 hours[1]. Therefore, the total number of flight hours is a better indicator of how much is flown than the number of kilometers.

Total Tonnes Carried: This KPI represents the total amount of tonnes the network should, in theory, be able to transport. This could differ per solution based on how the model is set-up.

Total Fuel Consumption: This KPI represents the total amount of fuel expected to be burned if the network is operated as modeled by the tool.

Available Freight Tonne (AFT): The AFT represents the total capacity that is offered on the network in terms of weight.

Available Freight Tonne Kilometer (AFTK): The AFTK can be calculated by multiplying an aircraft's available capacity on a particular flight leg by the flight distance of that same flight leg. The total AFTK of the complete network can be calculated by summing this value for each of the flown flight legs within a certain period.

Average Leg Load Factor *(weight)* (ALLFW): This KPI can be calculated by determining each of the flight legs' expected load factors in the network. This can be done by dividing the load expected to be transported over the available capacity in terms of weight.

[1]Estimated flight times were retrieved from https://www.flightconnections.com on 02/01/2021

Average Leg Load Factor *(volume)* (ALLFV): The ALLFV is calculated in the same way as ALLFW, but then the expected volume to be transported should be divided over the available capacity in terms of volume.

On-time performance: This KPI represents the number of tonnes expected to be delivered before the latest allowable time of arrival of the total number of tonnes expected to be carried. Depending on if on-time performance is modeled as a hard constraint or as a soft constraint, this KPI can differ per solution.

Hub performance: Another interesting thing to compare between different solutions is how well the hubs perform. In theory, most of the transport flows should be modeled through the hubs, as this would save both time and cost due to efficient sorting machines. Therefore, it could be interesting to compare how much certain hubs are used in terms of arriving and departing freight tonnes between different solutions.

Average length of transported tonne: This KPI represents the average length a freight tonne travels through the network. This can indicate how efficiently freight is transported through the network. The fewer kilometers freight has to be transported through the network, the less fuel will be required.

Average transportation time of transported tonne: This KPI represents the average time it takes to send freight from its origin to its final destination. This can indicate how efficiently freight is transported through the network. The less time freight spends in the network, the better the assurance of timely delivery is.

4 Flight Schedule Design

Designing a feasible and profitable flight schedule is a very complex task. Operational Research (OR) professionals have already been working on developing tools and methods to design optimal schedules since the 1950s. The design of the complete air freight transportation schedule can be divided into four major interdependent problems: *schedule planning*, *fleet assignment*, *rotation planning*, and *cargo routing*. This chapter is organized as follows, section 4.1 explains how a flight schedule is created in general. After that, an overview of the available research for solving the sub-problems of the flight schedule design problem is provided in section 4.2. This is followed by a discussion in section 4.3. Finally, in section 4.4 an analysis of how an air cargo flight schedule can be recovered is provided.

4.1 Flight Schedule Design in General

As mentioned before, four main problems have to be solved to design a flight schedule for air freight transportation. When a passenger transportation network is considered, the crew scheduling problem should also be considered. This is, however, less difficult for freight transport, as this only requires two pilots and no cabin crew. Therefore, this problem is not taken into account for this study.

The first step of *schedule planning* is usually started about 12 months in advance. This can be divided into three smaller steps. First, it should be determined which city pairs or markets the airline will serve. Thereafter, it has to be determined how frequently each of the markets is served. Finally, the timetable itself should be created. To complete these steps successfully, it is crucial to understand how demand and supply interact with each other. Lohatepanont & Barnhart (2004) explained this interaction as follows:

A market is defined by one origin and one destination precisely. This means that New York - Amsterdam is one market, and Amsterdam New York is another. The airline has to determine what the maximum demand is that it can capture for each distinct market. This is also known as the unconstrained market demand. Hereafter, the airline has to determine the unconstrained itinerary demand. This means that each of the flight legs is not yet capacity constrained. Then, it should be analyzed how the airline could supply this demand. For this purpose, an appropriate network structure should be thought of. Generally, this network combines the well-known point-to-point (PP) and hub-and-spoke (H&S) network structures. For a more in-depth analysis of these two networks, the reader is referred to Cook & Goodwin (2008). Interestingly, the problem can also be solved in reverse: first, decide on the network structure and thereafter determine how much demand can be captured. Finally, it should be noted that additional demand can be stimulated by increasing the number of flights and vice-versa. This clearly shows the strong interaction between demand and supply.

Although it is essential to understand this interaction, it should be noted that most of the time, schedule planning does not start from scratch (Lohatepanont & Barnhart (2004), Gopalakrishnan & Johnson (2005), Derigs et al. (2009)). Instead, an existing schedule is used as a basis, and changes are made, reflecting the difference in forecast demand. This is also known as *schedule development*. According to Lohatepanont & Barnhart (2004), there are a few reasons this is the industry's practice. First of all, it is operationally impractical and computationally tricky to build a schedule from scratch. Second, changing a network could require significant investments in infrastructure at certain airports. Finally, reliability and consistency are important to the customers of an airline. By using an existing schedule, consistency of the network can more easily be retained. Furthermore, it makes the effort required for network planners to complete this step of the schedule design tractably.

Once the *schedule planning* or *schedule development* is done, the fleet assignment problem may be considered. One fleet consists of one or several aircraft of one specific model or type of aircraft. Most airlines operate several fleets. Therefore, the fleet assignment problem is solved to match the demand for a certain flight leg as close as possible with the aircraft's capacity. If the aircraft is too small, demand is spilled. On the other hand, if the aircraft is too big, the aircraft flies partially empty, meaning it is less cost-effectively transporting demand.

The purpose of the next step, *rotation planning*, is to assign each specific aircraft in a fleet to a sequence of flight legs, such that a feasible rotation for each aircraft is found. In this part of the planning also maintenance constraints can be included.

As a final step, *cargo routing* should be considered. The difficulty here lies in the fact that cargo does not have a preference for a certain itinerary. Cargo can take any route through the network, as long as it is delivered to the final destination on time. Therefore, it should be carefully analyzed whether all expected OD demand can be transported through the final network. If this is not the case, the network needs to be redesigned. This can be seen as a two-phase planning process of schedule construction and schedule evaluation. Systems that constructed the schedule in this way emerged in the 1960s (Brough (1966), Tobin & Butfield (1970), Loughran (1972)). A schematic overview of the steps described above and how the concept mentioned above is applied to it is shown in Figure 4.1.

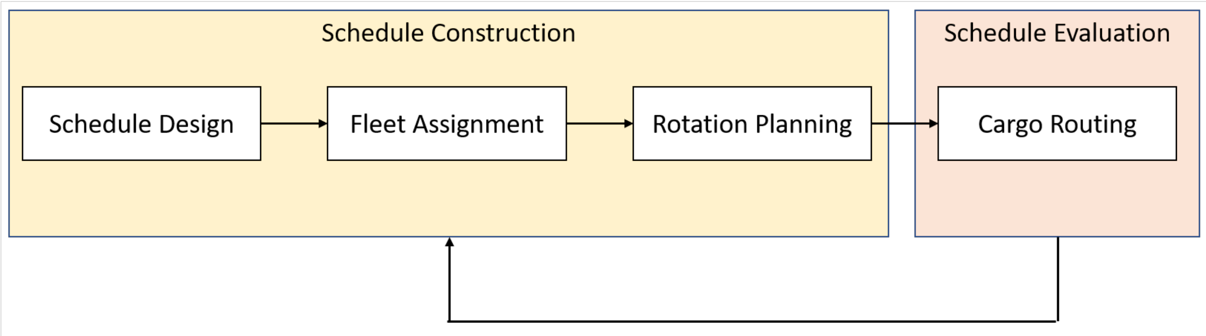

Figure 4.1: A schematic overview of how the concept of schedule construction and schedule evaluation may be applied to the steps that have to be undertaken to design a flight schedule. Source: Adapted from Derigs & Friederichs (2013)

4.2 Applied methods to design a flight schedule for a full-cargo airline

In the previous section, the design process of a flight schedule as a whole was discussed. However, as mentioned earlier, this process is often divided into smaller sub-problems. Many researchers have developed models and methods to solve one or multiple of these sub-problems. This section will elaborate on the objectives and methods used in several of these studies.

4.2.1 Schedule Design

The profitability of an airline is critically influenced by the flights it offers. Therefore, Lohatepanont & Barnhart (2004) focused their attention on the sub-problems of schedule design and fleet assignment. For this purpose, they have developed a model that simultaneously optimizes the selection of flight legs to be flown and the assignment of aircraft types to it. They do not design the network from scratch but use a master flight list as input. This list consists of mandatory and optional flights. Another important input for their model is the average unconstrained itinerary demands. As output, the model then provides a list of recommended flights to include in the schedule and the associated fleet assignment that combined yield an optimized schedule in terms of cost. Although this study focused on designing a network for passenger transport, this approach was in later research also adopted for the design of an ACTN. For example, Derigs et al. (2009) used this approach as well. They refer to this approach as *the pragmatic planning paradigm*. It is pragmatic because it leaves the generation of the predefined master flight list in the hands of experienced planners and only supports their decision-making process by a model-based optimization system. Their designed model maximizes network-wide profit by simultaneously optimizing the selection of flight legs, fleet rotation, and cargo routing based on inputs such as the predefined master flight list and a forecast O&D-matrix. This research was further extended by Derigs & Friederichs (2013). Their optimization problem's objective was altered to a minimization problem, namely the minimization of the total operating cost. Moreover, the fleet assignment, which will be explained in the next section, was added as a sub-problem to be solved by the system. The goal of their research was to give a proof of concept of the pragmatic planning approach. With this intention, they performed extensive tests on data that reflected the different types of cargo airlines. Finally, they were able to prove the computational tractability and effectiveness of their approach.

Although the pragmatic planning approach's effectiveness has been proven, studies that developed methods to design the network from scratch can also be found. For example, Yan et al. (2005) developed a model that combined airport selection, fleet routing, and timetable setting. The objective of their model was to maximize the operating profit given a particular forecast demand. The airline they used for reference only operated one type of aircraft. Therefore their model was limited to one fleet type. This aside, they did compare interesting modeling heuristics. The developed heuristics consider the maximum number of stops cargo is allowed to make before it arrives at its final destination. In the heuristics they developed cargo was allowed either to make no stops at all, just one stop, or unlimited stops. Finally, they also developed a mixed-stop heuristic where the maximum number of stops was determined based on the OD distance. This meant that for short-haul O&D demand, only direct flights were allowed, for middle-haul one-stop flights and long-haul unlimited stops. Especially this latter heuristic seems very interesting for the modeling of the ACTN as it reduces the problem scale, but it does not compromise the realisticness of the network.

4.2.2 Fleet assignment

As explained in section 4.1, once the airport selection and flight route selection have been performed, the fleet assignment problem can be considered. In literature, two classical methods are most frequently applied to model the flight network, namely the connection network (CN) and the time-space network (TSN). These networks can be used to determine the feasible paths the available aircraft can follow through the network. Therefore, these networks are frequently used to create the final flight schedule. Both of these networks will be shortly discussed in the following subsections. It should be noted that although most literature is focused on passenger transportation, it is expected that these models can be easily modified to apply to cargo transportation networks as well.

Connection Network

The connection network resembles the network that is also frequently seen in vehicle routing problems. Abara (1989) was the first to introduce it into OR for the aviation industry. In this case, a node represented either an arriving or a departing flight, which is why it is also known as an activity-on-node network. An arc could represent three things: a leg arc, a connection arc, or an origination/termination arc. A leg arc was used to represent a possible connection between two flights. A connection arc existed between an arrival and departure node if the time between those nodes exceeded the minimum turnaround time. An origination/termination arc represented the origination/termination of a flight sequence at the start/end of the day. In the original model, the balance of incoming and outgoing aircraft was ensured by balancing the number of originating and terminating sequences. Instead, source and sink nodes could also have been used to balance the network. Moreover, two timelines were created to represent each airport's departure and arrival timeline in the network. The fleet assignment model can then be used to create a sub-network for each fleet. When the sub-networks are combined, cover constraints are used to ensure that each flight is assigned to precisely one fleet type. A representation of this network is shown in Figure 4.2.

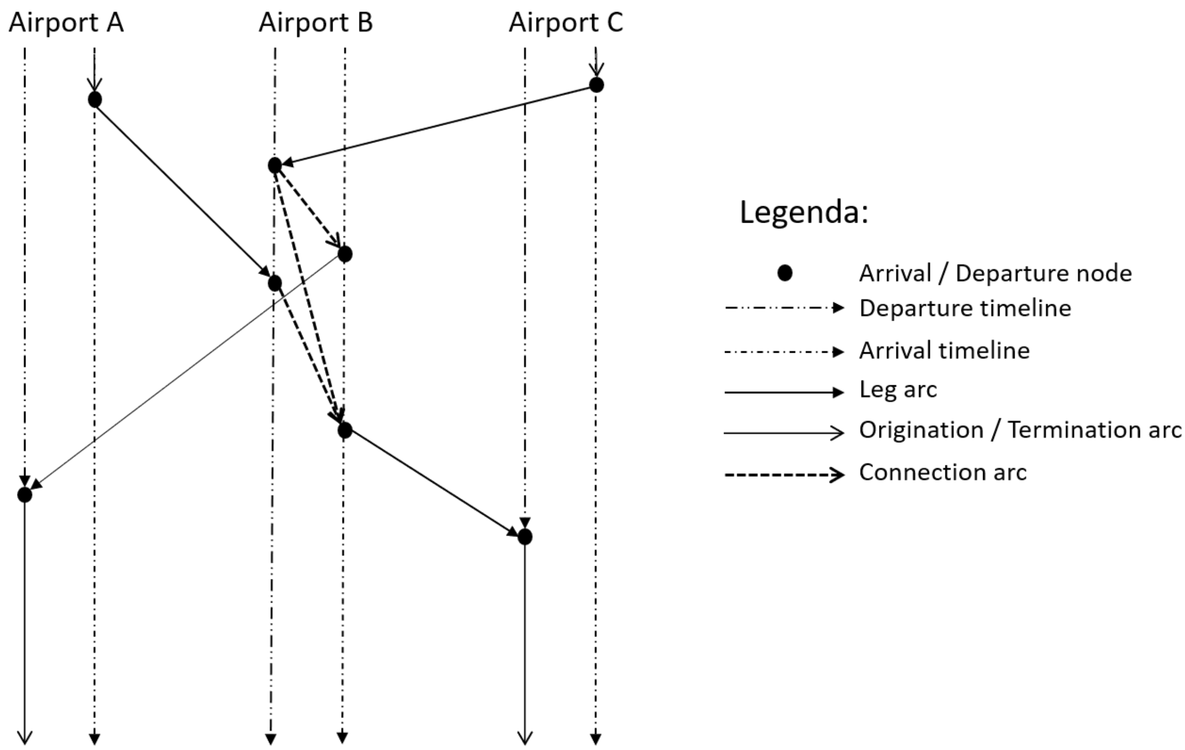

Figure 4.2: The connection network that was proposed by Abara (1989). Source: Adapted from Zhou et al. (2020)

Zhou et al. (2020) provides the following basic mathematical formulation for the connection-based fleet assignment model:

Sets

K	set of fleet, indexed by k
L	set of legs, indexed by l, i or j
L^+	$= L \cup \{0\}$. The index $i = 0$ denotes the original arc, and index $j = 0$ denotes the terminal arc. Given a leg connection $i \to j$, $i, j \in L^+$, if $i = 0$ then j is the first leg of a daily aircraft route; if $j = 0$, then i is the last leg of a daily aircraft route.
S	set of stations, indexed by s
L_s^A	set of legs arriving at station s
L_s^D	set of legs departing from station s

Constants

M_k	number of available aircraft of fleet k

Parameters

c_k	cost of each aircraft in fleet k
p_{jk}	benefit of operating leg j by fleet k

Variables

x_{ijk}	$\in \{0, 1\}$. $x_{ijk} = 1$, if fleet k covers the connection $i \to j$, $i, j \in L^+$ and $x_{ijk} = 0$, otherwise.

Basic Fleet Assignment Model based on the connection network

$$\text{Max} \sum_{i \in L^+} \sum_{j \in L} \sum_{k \in K} p_{jk} x_{ijk} - \sum_{j \in L} \sum_{k \in K} c_k x_{0jk} \tag{4.1}$$

$$\text{subject to} \sum_{i \in L^+} \sum_{k \in K} x_{ijk} = 1 \qquad \forall j \in L \tag{4.2}$$

$$\sum_{i \in L^+} x_{ilk} - \sum_{j \in L^+} x_{ljk} = 0 \qquad \forall l \in L, \forall k \in K \tag{4.3}$$

$$\sum_{l \in L_s^D} x_{0lk} - \sum_{l \in L_s^A} x_{l0k} = 0 \qquad \forall s \in S, \forall k \in K \tag{4.4}$$

$$\sum_{l \in L} x_{0lk} \leq M_k \qquad \forall k \in K \tag{4.5}$$

$$x_{ijk} \in \{0, 1\} \qquad \forall i, j \in L^+, \forall k \in K \tag{4.6}$$

As can be seen in Equation 4.1 this is a maximization problem. In the first part, the profit gained from flying all legs and connections is summed, and in the second part, the cost incurred by using aircraft is calculated. The cover constraint that ensures each flight is assigned to one fleet type is represented by Equation 4.2. The constraint shown in Equation 4.3

ensures that the network flow balance is maintained. Equation 4.4 is the schedule balance constraint that ensures that an equal amount of aircraft arrive and depart from each airport. As a result, the schedule can be repeated every day. The last constraint (Equation 4.6) is the aircraft count constrained, which ensures that the number of aircraft used in the schedule does not exceed the number of available aircraft for each of the fleet types.

One of the advantages of this network representation is that maintenance restrictions can be easily included in the model via feasible maintenance paths (see, e.g., Abara (1989), or Barnhart et al. (1998)). Moreover, costs can be allocated to the different arcs in the network. This provides richer modeling possibilities for this network when compared to the TSN (Barnhart et al. (1998)). In line with this, Sherali et al. (2013) exclaimed that this network provides easy accommodation of integrated operational considerations. Furthermore, this network can capture the flow of individual aircraft. This is a useful aspect for airlines that would like to ensure equal wear and tear among the aircraft in their fleet. Friederichs (2010) developed a model that incorporated equal aircraft usage constraints. At last, this network is used to both analyze how the risk of disruptions can be reduced and determine how disruptions can best be remedied (see, e.g., Rosenberger et al. (2004), Hu et al. (2017)).

Time-Space Network

The time-space network was first introduced by Berge & Hopperstad (1993). This network is also known as the activity-on-edge network because the arcs represent actual movements in the network. The nodes in the network represent the arrival or departure times of flights. Furthermore, in this network, an arc could represent either a leg arc or a ground arc. A leg arc represented a flight, and a ground arc an aircraft remaining on the ground. Moreover, for each airport, one timeline was created that represented the planning horizon. A representation of this network is shown in Figure 4.3.

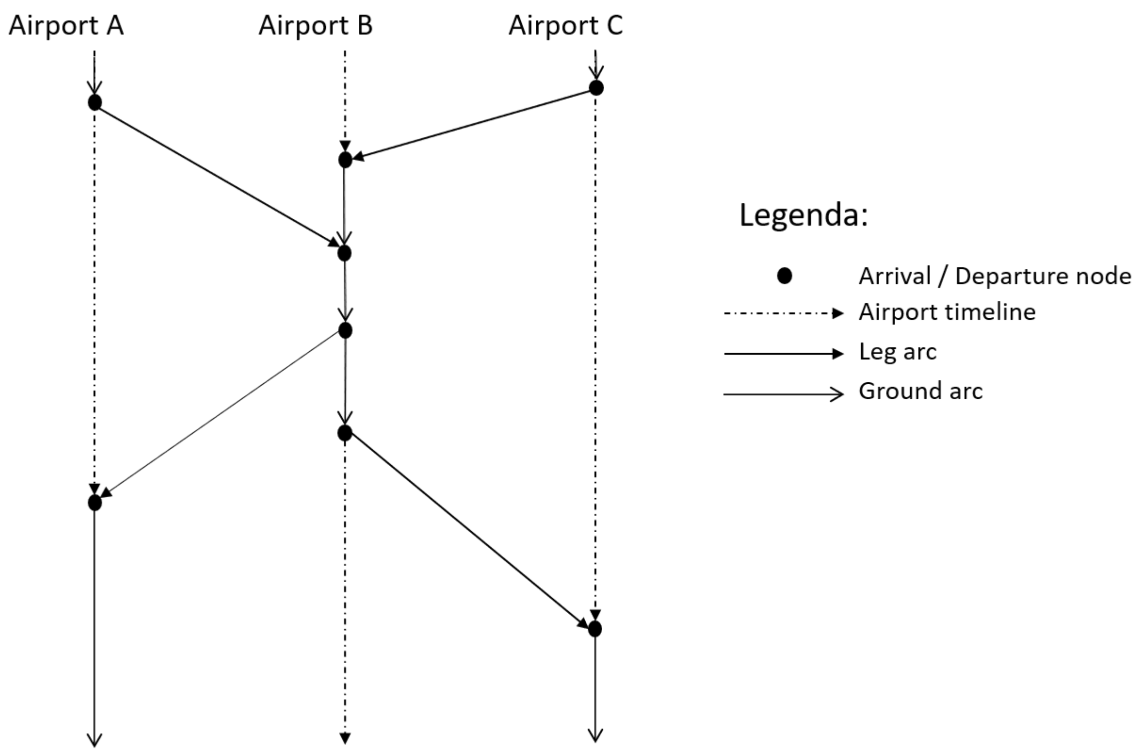

Figure 4.3: The time-space network that was proposed by Berge & Hopperstad (1993). Source: Adapted from Zhou et al. (2020)

The model formulation of the CN is extended by Zhou et al. (2020) in the following way to formulate the TSN.

Sets (cont.)

N_k	set of nodes in the network for fleet *k*, indexed by *n*
L_{n^+}	set of legs inbound to node *n*
L_{n^-}	set of legs outbound from node *n*
L^P	set of legs crossing the aircraft count time. The count time can be any time point of a day.
N_k^P	set of nodes in the network for fleet *k*, and the ground arcs into these nodes cross the aircraft count time

Variables (cont.)

x_{ijk}	$\in \{0, 1\}$. x_{ijk} = 1, if fleet *k* covers the leg *l*; x_{ijk} = 0, otherwise.
y_{n^+}	number of aircraft on the ground arc into node *n*, where n $\in N_k$, k $\in$ K
y_{n^-}	number of aircraft on the ground arc out node *n*, where n $\in N_k$, k $\in$ K

Basic Fleet Assignment Model based on the time-space network

$$
\begin{aligned}
\text{Max} \sum_{l \in L} \sum_{k \in K} p_{lk} x_{lk} & & & (4.7)\\
\text{subject to} \sum_{k \in K} x_{lk} = 1 & & \forall l \in L & (4.8)\\
\sum_{l \in L_{n^+}} x_{lk} + y_{n^+} - \sum_{l \in L_{n^-}} x_{lk} - y_n = 0 & & \forall n \in N_k, \forall k \in K & (4.9)\\
\sum_{l \in L^P} x_{lk} + \sum_{n \in N_k^P} y_{n^+} \leq M_k & & \forall k \in K & (4.10)\\
x_{lk} \in \{0, 1\} & & \forall l \in L, \forall k \in K & (4.11)\\
y_{n^+}, y_{n^-} \geq 0 & & \forall n \in N_k, \forall k \in K & (4.12)
\end{aligned}
$$

The objective function shown in Equation 4.7 maximizes the total profit. Equation 4.8 represents the cover constraint for the TSN. The flow balance constraint is shown in Equation 4.9. The aircraft counting constraints is represented by Equation 4.10.

Furthermore, when Figure 4.2 is compared to Figure 4.3 it can already be seen that one of the advantages of the TSN is that it requires fewer arcs to be built. Therefore, when the network under consideration is relatively large, it will cost less time to model than the CN for the same network. By using the TSN, it will not be possible to track individual aircraft. However, for certain applications, it is not required that individual aircraft can be tracked. For instance, the TSN was used to model a network with varying departure times to determine the most profitable combination of departure times for each flight leg and the corresponding fleet assignment (see, e.g., Rexing et al. (2000) or Bélanger et al. (2006)). The main differences between the two networks are presented in Table 4.1.

Table 4.1: A comparison between the connection and time-space network. (Source: Zhou et al. (2020))

	Connection Network	Time-Space Network
Cost	Assignment cost + connection cost	assignment cost
# Nodes	$O(\|L\|)$	$O(\|L\|)$
# Arcs	$O(\|L\|^2)$	$O(\|L\|)$
# Variables	$O(\|L\|^2\|K\|)$	$O(\|L\|\|K\|)$
# Constraints	$O(\|L\|\|K\|)$	$O(\|L\|\|K\|)$
Connection Information	Known	Unknown

Moreover, the TSN has also been used to integrate the fleet assignment problem with timetable setting (see e.g. Yan et al. (2005), Tang et al. (2008) or Sherali et al. (2013)). The research of Tang et al. (2008) is used as example. For their modeling purpose, the TSN had a slightly different representation, as can be seen in Figure 4.4. Now, a node for every point in time at every airport is created. Three arcs are used to connect nodes, flight leg arcs, ground arcs, and cycle arcs. The flight leg arc represents a connection between two airports. In theory, every possible flight should be represented in the network; however, only a few connections are drawn to prevent cluttering of the figure. Each arc contains the departure and arrival time, the departure and arrival airport, and the operating cost. The arc flow's upper bound is one, which implies only one aircraft can be assigned to that particular flight. The ground arc represents an aircraft remaining on the ground. The cycle arc ensures continuity between two consecutive planning periods. Furthermore, this research integrated passenger, cargo, and combi flight scheduling into one model. Therefore, for each of the different aircraft, a different sub-network was created.

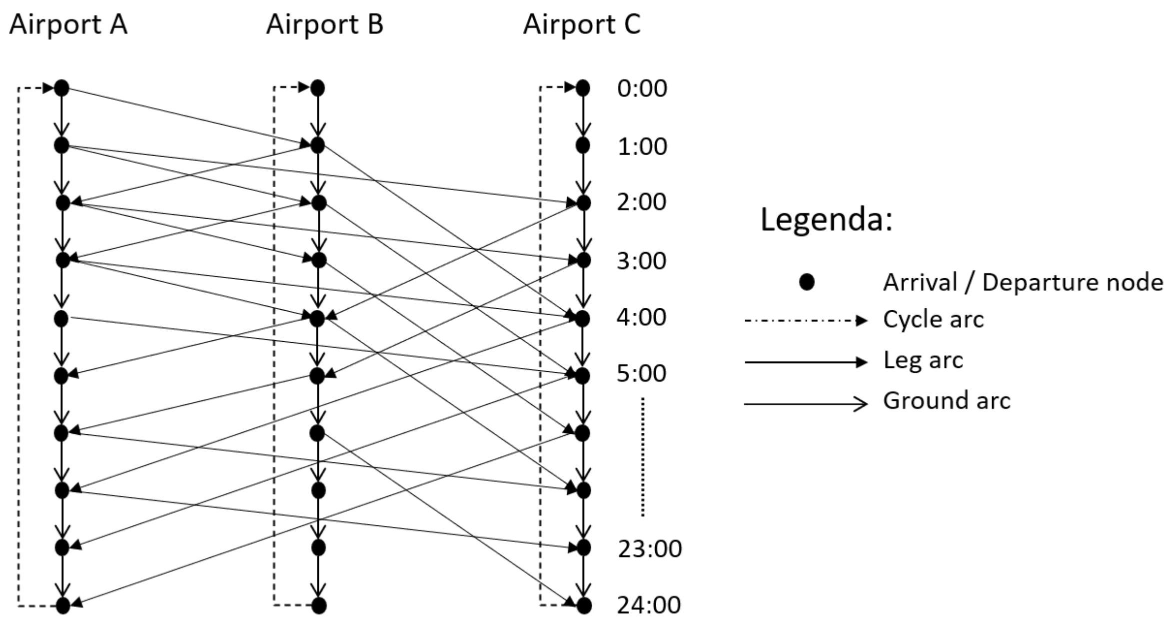

Figure 4.4: The time-space network used to integrate the fleet assignment problem with timetable setting. Source: Adapted from Tang et al. (2008)

The model formulation is very much similar to the ones mentioned before. The objective is to minimize the system cost. Furthermore, the model is subject to constraints such as the flow conservation constraints, the available aircraft constraint, the flight leg service constraint, and the aircraft capacity constraints. Several others are important for their research but are not relevant to this research, so they are not discussed.

To conclude, both the CN and TSN have their advantages and disadvantages, which were summarized in Table 4.1. The CN grows faster in size with an increasing network size but does capture the flow of individual aircraft. The TSN cannot provide information on the latter, but the problem size grows much slower. Since this research's scheduling time horizon is several weeks, it is not required yet to determine the flow of individual aircraft. Instead, at this point in the schedule development, it is interesting to analyze how the aircraft routing and timetable setting can be further optimized. Therefore, it was decided that a TSN will be used to model the network for this research.

4.2.3 Aircraft Rotation

Once the schedule design and the fleet assignment are determined, the aircraft rotation problem may be considered. The goal is to find a feasible sequence of flights for each aircraft such that all flight legs in the network are flown by one aircraft, while the number of available aircraft in each fleet is not exceeded (Clarke et al. (1996)). In this part of the design problem, it is often ensured that each aircraft follows a route that also satisfies that aircraft its maintenance requirements. Two typical methods can be distinguished in literature. Either maintenance checks are performed based on a fixed amount of passed calendar days, or maintenance checks are performed after a certain number of flight hours have been flown (Zhou et al. (2020)). However, this study will not go into this level of detail of the flight schedule design. The goal of this study is to determine a feasible flight schedule until the fleet assignment is determined. For an in-depth analysis of the different available aircraft routing models, the reader is referred to Zhou et al. (2020).

4.2.4 Cargo Routing

One of the first studies to consider air cargo routing was conducted by Antes et al. (1998). They developed a model that evaluated an airline's flight schedule with respect to cost, revenue, and contribution to profit. Thereafter, air cargo routing was also used to determine the maximal contribution to profit for each flight leg for revenue management purposes by Bartodziej & Derigs (2004). This was done by taking into account yield values for a set of possible O&D pairs. Furthermore, in Derigs & Friederichs (2013) it was emphasized that the cargo routing problem is an essential part of the flight schedule design since the choice of routing directly influences the profitability of the schedule. The reason for this is simple: a cargo airline offers the conceptually simple service of offering insurance of timely delivery of goods for a certain price. In general, an airline forecasts how much demand it will have to transport within a certain period and how much revenue this will approximately generate. After that, it is up to the airline to transport this demand as cost-effectively through the network as possible, as this, in turn, maximizes the profit. In this research, the cargo routing problem is formulated as a path-flow model. A path represents a sequence of flight legs that connects an O&D pair. For all these O&D itineraries, decision variables are introduced. This approach has the advantage that practical constraints on an itinerary's feasibility can be incorporated before the optimization model has to run. The set of all cargo flow paths is represented by *P* and the set of cargo flow paths that use leg *l* is represented by P_l. This is mathematically formulated as shown in Equation 4.13.

$$P_l :\; = \{p \in P \| \exists k \in \{1, \ldots, n_p\} : l_k^p = l\} \tag{4.13}$$

Furthermore, the amount of cargo that is transported over each path $p \in P$ is measured by x_p^{flow}. The rest of the path-flow based cargo routing model (CRP-P) is formulated as follows:

The CRP-P Model

$$\text{Max} \sum_{p \in P} m_p x_p^{flow} \tag{4.14}$$

$$\text{subject to} \sum_{p \in P_{od}} x_p^{flow} \leq d_{od} \qquad \forall od \in OD \tag{4.15}$$

$$\sum_{p \in P_l} x_p^{flow} \leq w_l^{max} \qquad \forall l \in L \tag{4.16}$$

$$\sum_{p \in P_l} vol_{od_p} x_p^{flow} \leq v_l^{max} \qquad \forall l \in L \tag{4.17}$$

$$x_p^{flow} \geq 0 \qquad \forall p \in P \tag{4.18}$$

The objective, shown in Equation 4.14, of this model is to maximize the contribution to profit. This profit is calculated by multiplying the total flow on a path by a certain freight rate. In Equation 4.15 the constraint that restricts the amount of cargo that can be transported on each O&D pair in P_{od} is shown. The next two constraint, represented by Equation 4.16 and Equation 4.17, ensure that the leg capacity of each flight leg in terms of weight and volume are respected.

Furthermore, the cargo routing problem studied by Li et al. (2006) was also formulated as a path-flow model. For their research, several side constraints had to be taken into account, such as a minimum transfer time for cargo. Also, cargo could only be transferred to another aircraft at hubs and at no other stations. Moreover, delivery time windows had to be taken into account. To retain a tractable model, a two-step modeling approach was applied. First, all feasible paths that satisfied the aforementioned constraints were generated for all commodities.

Thereafter, the multicommodity network flow problem (MNFP) was formulated with only the columns of these feasible paths. The rows of the problem are created by capacity and demand constraints. This two-step formulation seems especially useful for this research as the network can become rather large. Therefore, it is an intelligent approach to reduce the problem's size by selecting only feasible paths instead of considering all possible routes.

Finally, it should be noted that there are two other standard methods to decompose the MNFP, namely the node-arc and tree formulation. However, Jones et al. (1993) showed that when the MNFP was solved using the path formulation, substantially fewer master problem iterations were required than when the other two methods were used. Therefore, the path-formulation will also be used for this research.

4.3 Applied methods to design a flight schedule for an integrator

The previous section elaborated on what different methods could be used to design a profitable flight schedule for a full-cargo airline. The integrators can be considered a special form of full-cargo airlines, as explained in section 3.2. Therefore, this section will provide an overview of the available literature that focused on designing models to design (parts of) a flight schedule for an integrator.

Kuby & Gray (1993) introduced a network planning problem called the hub network design problem with stopovers and feeders. In this study, they compared the effectiveness of a H&S network with stopovers and feeders with that of a H&S network with only direct flights into a hub. Based on the network of FedEx in the western United States, they determined that the first configuration allows for a more efficient network design.

Barnhart & Schneur (1996) introduced the Express Shipment Service Network Design (ESSND) problem and present a column generation approach for its solution. Solving this problem aims to design a network of flights that enables overnight transportation of packages at minimum cost. They were able to determine the service design, fleet size, and fleet composition simultaneously. However, in their model, they assumed several operational restrictions that simplified the problem. Examples are, only a single hub is allowed, and each gateway could only be served by one aircraft type. A few years later, Kim et al. (1999) were able to solve the ESSND problem with flexible hub assignment. To achieve this, they first had to add valid inequalities, which strengthened their linear programming relaxation. Thereafter, they applied a series of innovative problem reduction methods, which reduced their problem size tremendously. They formulated the ESSND as follows:

ESSND Model Formulation

$$\min \sum_{f \in F} \sum_{r \in R^f} d_r^f y_r^f \tag{4.19}$$

$$\text{subject to} \sum_{k \in K} x_{ij}^k \leq \sum_{f \in F} \sum_{r \in R^f} \delta_{ij}^{fr} u_r^f y_r^f \qquad (i,j) \in A \tag{4.20}$$

$$\sum_{j:(i,j) \in A} x_{ij}^k - \sum_{j:(j,i) \in A} x_{ji}^k = \begin{cases} b^k & \text{if } i = O(k) \\ -b^k & \text{if } i = D(k), \\ 0 & \text{otherwise} \end{cases} \qquad i \in N, k \in K \tag{4.21}$$

$$\sum_{r \in R^f} \beta_i^r y_r^f = 0 \qquad i \in N, f \in F \tag{4.22}$$

$$\sum_{r \in R^f} y_r^f \leq n_f \qquad f \in F \tag{4.23}$$

$$\sum_{f \in F} \sum_{r \in R^f} \delta_h^r y_r^f \leq a_h \qquad h \in H \tag{4.24}$$

$$x_{ij}^k \geq 0 \qquad (i,j) \in A, k \in K \tag{4.25}$$

$$y_r^f \in \mathbb{Z}_+, \qquad r \in R^f, f \in F \tag{4.26}$$

The constraint shown in Equation 4.20 restricts the amount of flow on any arc in the network to the maximum capacity that is assigned to it. In Equation 4.21 the flow balance constraint is depicted that ensures conservation of flow for all O&D demand. Then, Equation 4.22 represents the aircraft balance constraint that forces an equal number of aircraft of each fleet type to arrive and depart from each airport. The constraint shown in Equation 4.23 ensures that not more than the number of available aircraft in each fleet type is used. In Equation 4.24 the constraint for landing capacity for aircraft at hubs is shown. Finally, they assumed that the sorting capacities of the hubs were not exceeded. By applying their model to different data sets of an express operator, they determined that annual cost savings of tens of millions of dollars could be achieved.

Armacost et al. (2002) had a different approach to resolve the problem of poor lower bounds on the optimal integer solution of the ESSND problem. They introduced composite variables that represent a combination of aircraft routes that ensure a feasible flow for all packages between any origin and destination. The resulting formulation for the network design problem had strong LP relaxations. The composite variable formulation for Next-Day-Air Network Design was formulated as follows:

Composite Variable Formulation for Next-Day-Air Network Design

---

Sets

- A set of all arcs in time-space network
- C set of all composites
- C_P set of composites constructed from pickup routes
- C_D set of composites constructed from delivery routes
- G set of gateways (airports)
- H set of hubs (airports)
- F set of fleet types
- R^f set of routes flown by fleet type, $f \in F$

Indicators

- δ_c^{gh} = 1 if composite, c covers the demand between gateway g and hub h, 0 otherwise.
- δ_{ij}^r = 1 if route r covers includes arc (i,j), 0 otherwise.
- δ_h^p = 1 if path p passes through hub h, 0 otherwise.

Data

- γ_c^f = number of aircraft of type f included in composite c.
- $\gamma_c^f(\bar{g})$ = number of aircraft of type f included in composite c originating at airport g.
- $\gamma_c^f(\underline{g})$ = number of aircraft of type f included in composite c terminating at airport g.
- d_c = cost of all aircraft routes in composite c.
- b_P^{gh} = pickup demand volume from gateway g to hub h.
- b_D^{gh} = delivery demand volume from hub h to gateway g.

Decision Variables

- v_c = 1 if composite c is selected, 0 otherwise.

$$\min \sum_{c \in C} d_c v_c \tag{4.27}$$

$$\text{subject to} \sum_{c \in C_P} \delta_c^{gh} v_c \geq 1 \qquad \text{for all } (\mathrm{g}, h) : b_P^{gh} > 0 \tag{4.28}$$

$$\sum_{c \in C_D} \delta_c^{gh} v_c \geq 1 \qquad \text{for all } (\mathrm{g}, h) : b_D^{gh} > 0 \tag{4.29}$$

$$\sum_{c \in C_P} \gamma_c^f(\bar{g}) v_c - \sum_{c \in C_D} \gamma_c^f(\underline{g}) v_c = 0 \qquad \text{for all } g \in G, f \in F \tag{4.30}$$

$$\sum_{c \in C_p} \gamma_c^f(\underline{h}) v_c - \sum_{c \in C_D} \gamma_c^f(\bar{h}) v_c = 0 \qquad \text{for all } h \in H, f \in F \tag{4.31}$$

$$\sum_{c \in C_p} \gamma_c^f v_c \leq n_f \qquad \text{for all } f \in F \tag{4.32}$$

$$\sum_{f \in F} \sum_{c \in C_P} \gamma_c^f(\underline{h}) v_c \leq a_h \qquad \text{for all } h \in H \tag{4.33}$$

$$v_c \in \{0, 1\} \qquad \text{for all } c \in C \tag{4.34}$$

As can be seen in Equation 4.27 the objective is to minimize the cost associated with the routes in the composites. The constraints shown in Equation 4.28 and Equation 4.29 are cover constraints that ensure that all pickup and delivery demand is covered. Equation 4.30 and Equation 4.31 represent the constraints that force a balance of arriving and departing aircraft of all types at all airports in the network. The next constraint, shown in Equation 4.32, limits the number of aircraft that can be used to the maximum available for each fleet type. Finally, the constraint shown in Equation 4.33 ensures that the number of arriving and departing aircraft at each airport in the network does not exceed the allowed maximum. The deployment of this approach to the UPS Next Day Air delivery network led to hundreds of millions of dollars in cost savings (Armacost et al. (2004)). Furthermore, Fleuren (2013) showed that TNT could have saved 132 million dollars if this approach would be implemented with only a few minor changes that capture the TNT Express operation's specifics.

The aforementioned examples clearly show the potential value of this approach. Therefore, it will be analyzed how this model can be adjusted such that it applies to the problem at hand as well. One of the differences that should be more closely studied, for example, is that they differentiate between pickup and delivery routes. In this research, this differentiation will not be made. Luckily, a few other researchers explored the use of composite variables and the ESSND problem, as will become clear when the rest of this section is read.

Similar to the ESSND problem is the movement scheduling problem, which was solved by Louwerse et al. (2014). They split the problem up into three parts. First, they determined which direct routes were required (either due to time constraints or due to enough demand). Thereafter, the routes from depots to hubs and vice versa were determined (DH/HD-problem). They modeled the DH/HD-problem as a set partitioning problem and solved it using a column generation algorithm. Finally, the scheduling of routes between hubs (HH-problem) was determined. This problem was formulated as a networking loading problem and solved with a local search algorithm. Meuffels (2015) used a similar approach to tackle congestion of a single hub in a service network. For this purpose, the author developed a model that could design a multi-hub express network. The resulting network consisted of scheduled movements, vehicles required for each movement, and the routing of flows through the network. The solution approach was tested on three data instances of an express service provider. The results demonstrated that a cost reduction of 18.6% of the total transportation cost could be realized.

The second model that focused on the ESSND problem with flexible hub assignment was developed by Shen (2004). To solve the problem, she first implemented a disaggregated information-enhanced column generation approach that reduced the number of variables from hundreds of thousands to only thousands. Thereafter, the author introduced a new model, referred to as the gateway cover and flow formulation. By applying this solution approach to the network design problem of a large express service provider, it was shown that tens of millions of dollars could annually be saved if the model were implemented. This model was used as a basis by Quesada Pérez et al. (2018). They strengthened the model by adding three families of valid inequalities. Thereafter, they introduced route covers to reduce the number of variables and constraints further. This new model was named the Route and Hub model with Cuts and Covers. FedEx Express Europe built test scenarios to test the model. It was found that the model was able to achieve an average improvement of 20% in terms of cost. This model was extended by Quesada Pérez et al. (2020) such that it could also design complex routes. For this research, the added value of five different complex routes was analyzed. Based on test data provided by FedEx Europe, it was found that if all route types were included in the model, cost savings of almost 5% could be realized. The reader is referred to Quesada Pérez et al. (2020) for a more detailed explanation of the complex routes that were considered.

4.4 Flight Schedule Recovery

The previous sections explained how the complex task of designing a flight schedule often is decomposed into smaller sub-problems and how these sub-problems can be solved. If no disruptions occur during the day of operations, this schedule is executed as planned. Unfortunately, it is impossible to completely avoid disruptions from happening. In the case of passenger transport, most disruptions occur on the supply side. This means that, for example, an aircraft has a mechanical failure or members of the crew are sick. On the other hand, in the case of cargo transport, disruptions occur both on the supply and the demand side. The latter is a direct consequence of the unpredictability of demand for air cargo, which has been discussed in section 3.1.

The focus of most research for flight schedule recovery for passenger transport is therefore focused on determining the best way to return to the original schedule as quickly as possible. However, for this research, the goal is not to return to the original schedule as fast as possible. Instead, the goal is to optimize the flight schedule given a change in forecast demand. This can be compared with a disruption in the demand. To the best of our knowledge, only one study has focused its attention on how a cargo carrier can best recover from demand disruptions. In this study, Delgado et al. (2020) introduced the Air Cargo Schedule Recovery Problem (ACSRP). Although there are some discrepancies between the ACSRP and the problem studied in this thesis, there are also some interesting similarities. Therefore, the goal of their research and their model formulation are discussed in more detail below.

Their study aimed to develop a model that could redesign an operational flight schedule in reaction to certain demand disruptions. The disruptions they considered occurred maximal 3 days, or 72 hours, before the actual time of flight. The ACSRP aims to minimize the total operating cost while also considering a penalty for deviating from the original schedule. Furthermore, they assume an airline operates a fleet $k \in K$. Each fleet has a specific capacity κ_k and serves a set of airport $l \in L$. Each aircraft is assigned a location at the beginning and the end of the recovery horizon, l_k^+ and $l_k^- \in L$ respectively. Moreover, there is a set of requests $r \in R$ that resembles the demand. Each request has a specific weight w_r, a strategic weight factor s_r, an O&D airport (l_r^+ and $l_r^- \in L$) and a release and due time t_r^+ and $t_r^- \in T$.

A directed time-space network, $G = (N, A)$, was created to formulate the ACSRP. This network has two types of nodes, namely itinerary nodes and request nodes. Each itinerary node represents a combination of a point in time and an airport, while each request node contains an origin and destination node $i_r^+ := (l_r^+, t_r^+)$ and $i_r^- := (l_r^-, t_r^-)$. Furthermore, four different arcs are used to connect nodes in the network. Flight arcs connect one airport to another one in the future. Only the arcs that are feasible in terms of travel time and aircraft permits are drawn. Both the aircraft and requests can be transported over these arcs. The second type of arcs is the ground arcs. These arcs connect one node in time with the next on the same timeline of an airport. These arcs can be used for transporting aircraft and requests too. The third type of arc is the no-service arc. This arc represents a direct connection between an O&D pair of a request. This arc does not represent an actual connection. The fourth type of arc is the request access (and egress) arc. This arc represents the connection between a request node and an itinerary node. These are only drawn after and before the release and due time respectively have passed. The no-service and request access arcs can only be used to transport requests.

Furthermore, they tested three different crew management policies. These policies affect how the cost of deviating from the original schedule is accounted for. For each of the policies, they proposed a certain penalty metric. They modeled the ACSRP as a MILP. The objective function and constraints of the model are presented below.

The Air Cargo Schedule Recovery Problem Formulation

Sets

A^F	Set of fight arcs
A^G	Set of ground arcs
A^N	Set of no-service arcs
A^R	Set of request access arcs
A	Set of arcs
D	Set of days in the recovery horizon
$\delta^+(i), \delta^-(i)$	Set of arcs that emanate from (or end at) node *i*
$\delta^+_\cdot(i), \delta^-_\cdot(i)$	Set of arcs in set $A^\cdot$ that emanate from (or end at) node *i*
K	Set of aircraft
L	Set of airports
N	Set of nodes
N^I	Set of itinerary nodes
N^R	Set of request nodes
R	Set of requests
T	Set of periods of time in the recovery horizon

Data

C_F^k	Cost of operating aircraft *k* per unit of time.
C_V	Cost due to additional fuel consumption per unit of time and weight.
F_{ij}^k	Cost of aircraft *k* using arc (i,j).
κ_k	Capacity of aircraft *k*.
p_r	Fare charged by the airline per unit of weight for request *r*.
t_r^+, t_r^-	Release (or due) time for request *r*.
t_{ij}	Flight time (in number of periods) between nodes *i* and *j*.
t_{lm}	Flight time (in hours) between airports *l* and *m*.
τ_k	The minimum turn around time of aircraft $k \in K$.
τ_r	The minimum transfer time of request $r \in R$.
V_{ij}^r	Cost of request *r* using arc (i,j).

Indicators

α_{ij}^k	= 1 if aircraft $k \in K$ can fly from node $i \in N^I$ to node $j \in N^I$, 0 otherwise.
γ_{il}	= 1 if node $i \in N$ corresponds to airport $l \in L$, 0 otherwise.
z_{lm}^k	The number an additional flight between *l* and *m* performed by *k*.

Decision Variables

x_{ij}^k	= 1 if aircraft $k \in K$ traverses a given arc $(i,j) \in A^I$ as part of its schedule, 0 otherwise.
q_{ij}^r	= 1 if a request $r \in R$ traverses a given arc $(i,j) \in A^I$, 0 otherwise.

$$\min \sum_{k \in K} \sum_{(i,j) \in A^I} F_{ij}^k x_{ij}^k + \sum_{r \in R} \sum_{(i,j) \in A} V_{ij}^r q_{ij}^r + P(X,Q) \tag{4.35}$$

$$\text{subject to} \sum_{r \in R} w_r q_{ij}^r \leq \sum_{k \in K} \kappa_k x_{ij}^k \qquad (i,j) \in A^F \tag{4.36}$$

$$\tag{4.37}$$

$$\sum_{(i,j)\in\delta_i^+(i)} x_{ij}^k - \sum_{(j,i)\in\delta_I^-(i)} x_{ji}^k = \begin{cases} 1, & i = i_k^+ \\ -1, & i = i_k^-, \\ 0, & \text{i.o.c.} \end{cases} \qquad i \in N^I, k \in K \tag{4.38}$$

$$\sum_{(i,j)\in\delta^+(i)} q_{ij}^r - \sum_{(j,i)\in\delta_I^-(i)} q_{ji}^r = \begin{cases} 1, & i = i_r^+ \\ -1, & i = i_r^-, \\ 0, & \text{i.o.c.} \end{cases} \qquad i \in N, r \in R \tag{4.39}$$

$$\sum_{(j,i)\in\delta_F^-(i)} x_{ji}^k \leq x_{(l_i,t')(l_i,t'+1)}^k, \quad i \in N^I, k \in K, t' \in \{t_i, \ldots, \min\{t_i + \tau_k, |T|\} - 1\} \tag{4.40}$$

$$\sum_{(j,i)\in\delta_F^-(i)} q_{ji}^r \leq q_{(l_i,t')(l_i,'+1)}^r, \quad i \in N^I, r \in R, t' \in \{t_i, \ldots, \min\{t_i + \tau_k, |T|\} - 1\} \tag{4.41}$$

$$x_{ij}^k \in \{0,1\}, \qquad (i,j) \in A^I, k \in K \tag{4.42}$$

$$q_{ij}^r \in \{0,1\}, \qquad (i,j) \in A, r \in R \tag{4.43}$$

$$\sum_{(i,j)\in A^I: t_i \leq t < t_j} x_{ij}^k = 1, \qquad k \in K, t \in T \tag{4.44}$$

$$\sum_{(i,j)\in A: t_i \leq t < t_j} q_{ij}^r = 1, \qquad r \in R, t \in t' \in \left\{t_r^+, \ldots, t_r^- - 1\right\} \tag{4.45}$$

$$q_{ij}^r \leq \sum_{k\in K} x_{ij}^k, \qquad (i,j) \in A^F, r \in R \tag{4.46}$$

As mentioned before, the objective of this model is to minimize the total operating cost. The costs are divided up into three components, as can be seen in Equation 4.35. The components are the cost incurred by operating aircraft, the cost incurred by transporting a request over a certain arc, and lastly, the penalty cost associated with deviating from the original schedule. The first constraint, shown in Equation 4.36, forces requests to be moved over arcs with enough capacity. The second two constraints, presented in Equation 4.38 and Equation 4.39, are flow conservation constraints that ensure a balance of requests and aircraft in the nodes in the network. Thereafter, in Equation 4.40 and Equation 4.41, the constraints that impose that each aircraft and request spends the minimum turnaround or transfer time at each airport it arrives at. By imposing the constraints in Equation 4.42 and Equation 4.43 it is ensured that x_{ij}^k and q_{ij}^r can only take on binary values. The last three equations represent so-called cuts. These cuts are used to tighten the relaxation of the model. What this means is later explained in section 5.4. For now, it is important to understand that the first cut, shown in Equation 4.44, takes advantage of the fact that at any given moment in time, every aircraft can only be in one place in the network. The second cut, presented in Equation 4.45, does the same thing, but then for request in between their release and due times. The last cut, shown in Equation 4.46, specifies that a request can only traverse a certain flight arc if an aircraft is assigned to it.

The above-described model was tested by using 24 original schedules with up to 30 orders and 8 airports. Moreover, for each of the test instances, they considered 4 different disruption scenarios. It was found that their model, when compared to a benchmark solution where cargo could only be re-routed, was able to achieve cost savings of about 10%. Although their research focuses on a different planning horizon and considers the cost of crew rescheduling, which does not have to be taken into account for this thesis, their model is a source of inspiration for the to-be-built model.

5 Solution Techniques

In this chapter, several solution techniques that can be used to solve a MILP are discussed. A MILP is a problem in which several variables are constrained to be an integer, while others can be non-integers. Furthermore, the objective function is linear and is also subject to linear constraints. In the past few years, many different methods have been studied. However, this chapter will focus on the most frequently used methods, namely column generation, branch & bound, branch & price, and branch & price & cut. The main theory of each technique is first shortly explained. Thereafter, examples of studies that applied the considered technique and their results are discussed. Once the main techniques have been discussed, it will be explained which solver will be used to solve the MILP in section 5.5.

5.1 Column Generation

Column generation is a technique that can be used to solve MILP problems. The formulation was first introduced by Ford & Fulkerson (1958). The main advantage of this method is that not all possible solutions need to be considered. The so-called Master Problem (MP) is the problem that contains all the decision variables (DVs) of the problem. However, sometimes not all DVs are known, or it is known that many of the DVs do not have to be considered because they will have a value of 0. The column generation technique uses this in the following way, first an initial (non-optimal) solution is formulated. Then, it is analyzed whether adding a DV to the initial solution improves it. If this is the case, a column is generated for this DV. In this way, only the DVs that positively influence the optimal solution have to be considered to solve the MILP. This subset of DVs is referred to as the Restricted Master Problem (RMP). Since the RMP is much smaller than the MP, the computational time required to find an optimal solution is significantly reduced.

The problem studied in Bartodziej & Derigs (2004) is when to decide to accept a booking request given that it is known how much demand is expected and that there is a certain fixed set of accepted requests. The model used is a path flow formulation of a special multi-commodity flow problem and it is solved by applying column generation. The test data consisted of three real-world planning problems from a cargo airline. The test data consisted of three real-world planning problems from a cargo airline, ranging from 10 to 79 airports, 624 to 1592 legs, and 1338 to 3459 O&D pairs. Each of the instances is used to compare four different solution approaches, referred to as S1, S2, S3, and S4. In solution approaches S1 and S2, the master problem has to be solved after each iteration. The other two solution approaches, S3 and S4, aimed at reducing the number of dual variables that have to be considered. Consequently, these algorithms had the advantage of lower run times but also had the disadvantage of achieving lower values for the optimal solutions. The run time of S1 and S2 ranged from 150 min to 300 min, while S3 and S4 ranged from 20 to 150 min. The more significant number of LPs that S1 and S2 have to solve compared to S3 and S4 can explain this. Furthermore, they determined the maximum value the profit could reach by using a different approach and set this as a reference value. The contribution to this reference value by S1 and S2 for the larger instances was approximately 95%, while that of S3 and S4 ranged from 90% to 94%, which is a significant difference. The authors concluded that the running time for solving the LPs becomes unacceptably high if the size of the instances further increases. This would be a problem for time-critical environments. In that case, additional means, such as limiting the number of allowable legs per path, have to be applied to reduce the size of the master problem.

The column generation technique was also used by Derigs et al. (2009) to determine the best combination of a list of mandatory and optional flights such that the network-wide profit was maximized. They combined column generation with the shortest path algorithm that solved the sub-problems. A data generator was developed that designed five realistic problem instances that varied with respect to the number of prescheduled and optional flights and the number of O&Ds considered. For better comparison, each instance was only changed with respect to one parameter. The number of prescheduled flights was either 130 or 200, the number of optional flights switched between 30, 60, and 100, and the number of O&Ds was either 1,500 or 3,000. Moreover, the number of mandatory flights remained 40 for all of the instances, and the number of available aircraft was always 18. The instances were used to compare two models, referred to as *INT-A* and *INT-B*. The models are mostly identical, except that model *INT-B* calculates the fixed aircraft cost differently, due to which it has fewer constraints to consider than model *INT-A*. It was concluded that an increase in the number of optional flights resulted in a strong growth of the number of columns generated and the time required to find a solution. To solve the smallest instance with 30 optional flights, model *INT-A* had to consider 30,300 columns, which took about 1.1 min, while model *INT-B* had to consider 61,600 columns, which took 1.2 minutes. To solve the same instance, only with 100 optional flights, model *INT-A* had to consider 110,800 columns, which took about 20 min, while model *INT-B* had to consider 1,062,000 columns, which took 240 minutes. Instance IV and V were comparable to instance I, here only the number of prescheduled flights and the number of O&Ds were increased. For both models, the MIP-size and running times were comparable to those of instance I. These parameters have little impact on the results because prescheduled flights do not have to be considered for the rotation planning. Furthermore, column generation was used to determine feasible O&Ds, which is why an increase in the number of O&Ds also had little impact on the outcome. Both models' main limitation is that fleet assignment is not considered as they assumed only one type of aircraft was available.

5.2 Branch and Bound

Branching is a frequently used method to solve mathematical optimization problems. In general, the goal for solving such problems is to find a combination of certain variables that either maximize or minimize a certain objective function. A rooted tree representing the complete subsets of possible solutions can be created by branching on each of the different variable options. However, for larger problems creating and analyzing all of the different options becomes a substantial computational burden. To overcome this hurdle, the branch-and-bound (B&B) technique may be applied. Instead of applying brute force and creating all possible branches, the branch is checked against the optimal solution's estimated lower and upper bound. A branch is added to the tree if and only if it produces a solution in between these bounds. This process is repeated until there are no branches left to consider.

The problem considered in Rexing et al. (2000) is a combination of fleet assignment and scheduling their departures. Two algorithmic approaches were developed to solve the model. One that performed well in terms of speed and simplicity, which is referred to as the direct solution technique (DST). The other one, known as the iterative solution technique (IST), minimized memory usage. Three data sets from a major U.S. airline were used to test both of the algorithms their performance. The number of flights ranged from 1621 to 2037, while the number of fleets considered ranged from 7 to 11. For each of these sets, they varied the size of the time windows and arc copy interval to compare the effect on the LP matrix size and the run time. A time-window defines how much time a departure may be shifted in either direction of its original scheduled departure time. The arc copy interval defines at what interval departure times are allowed to occur within a time window. The problem sizes were measured in non-zero elements. For the IST model, the problem size ranged from 39,600 to 62,700, while that of the DST model ranged from 38,800 to 320,600 (it should be noted that for the latter problem instance, no solution could be achieved due to insufficient memory). Although the problem sizes considered differed significantly, the DST model solved almost all problem instances the fastest. The most significant instance that the DST model solved was 258,000 and required 400 min. The largest instance solved by the IST model was 62,700 and required 250 min. The only instance where the IST model outperformed the DST model was the instance in which a narrow time window with a narrow arc copy interval was considered. In this case, the DST was slowed down by many unnecessary flight copies, while the IST only had to consider copies for a few selected flights.

It should be noted that the B&B technique heavily relies on a proper estimation of the lower and upper bound of the optimal solution. In general, poorly defined bounds lead to an inefficient and larger search space for the optimal solution. This effect is clearly shown by Armacost et al. (2002), who designed the composite variable formulation to solve the ESSND. They explained that the conventional formulations gave poor bounds on the optimal integer solution for this specific network design problem for two reasons. First, due to certain constraints, fractional solutions for the aircraft decision variables are induced. This is because the model rather uses fractional planes than incurring costs for unused capacity in aircraft. The second reason is that aircraft balancing constraints amplify the previously mentioned reason. These constraints cause otherwise isolated aircraft to be connected to the rest of the network. To overcome this, composite variables were introduced, as was also earlier explained in section 4.3. A composite variable represents a combination of aircraft routes that ensure a feasible flow for all packages between any origin and destination. In a relatively small example, they compare the solutions produced by three different model formulations. The first has the least tight bounds, the third the strongest. The instance considered consisted of one hub and five gateways, three fleet types, and timing restrictions for pickup and delivery. The first model required 781 nodes in the B&B tree to come to a solution that was 63% off from the optimal LP relaxed objective value. The second model required 111 nodes and was only 19% off. The third model only required one node and returned the optimal solution. However, it should be noted that, in general, it is not guaranteed that the third model immediately returns the integer solution at the root node. Nonetheless, this example clearly shows that the use of stronger bounds allows for faster generation of very good approximations of the optimal solutions via B&B.

5.3 Branch and Price

Branch-and-price (B&P) is a method to solve (M)ILP with many variables. It is a hybrid method of the aforementioned solution techniques, namely B&B and column generation. A general explanation of the technique is given in Barnhart et al. (1996). The authors describe that, just like with column generation, this technique omits columns from the LP relaxation because many variables will have a value of zero in the optimal solution. Then, it is checked whether the LP solution is optimal by solving the so-called pricing problem. This problem is solved to find columns that can be added to the basis such that the optimal solution previously found is improved. If these columns exist, they are added to the RMP. Solving the RMP and solving the pricing problem is repeated until no more columns are generated by solving the pricing problem.

The B&P technique has been implemented by Barnhart et al. (1998). They studied the combined fleeting and routing problem for aircraft. Their model creates feasible sequences of flights that respect maintenance requirements called strings. They test their B&P algorithm on data provided by a long-haul airline with a planning horizon of a week. This schedule contained 1124 flights, visiting 40 cities, with 9 fleet types, containing in total 89 aircraft. It required about 4.5 hours to find the relaxed LP solution with over 88,000 generated columns. Furthermore, they evaluated the effect on solution time if a certain tolerance value was taken into account. This meant that a column was only generated at a node if that node's solution value exceeded the root node by the set tolerance value. Tolerance values between 0.25% and 1% were considered. If the tolerance value was set at 0.25%, the LP lower bound was 1.5% off from the optimal integer solution found and the run time was 5 hours and 39 minutes.

Decreasing the tolerance value to 1% led to a little reduced run time, namely 5 hours and 27 minutes. The authors do not mention the gap between the optimal solution and the lower bound.

The fleeting and routing problem has also been studied by Haouari et al. (2011). Although, this research focused specifically on solving the problem for TunisAir. Provided with a flight schedule, the goal was to determine a minimum cost route assignment that satisfied maintenance constrained. The test data consisted of six instances, in which the number of rotations ranged from 192 to 507, the number of flight legs from 426 to 1050, the number of aircraft from 29 to 34, and the number of fleet types from 6 to 8. For each of the instances, they analyzed how much their B&P algorithm improved the solution produced by TunisAir. Moreover, they analyzed run time for each of the instances and the number of nodes and columns generated. On average, the solution was improved by 1.2%, with a maximum reaching 2.2%. Furthermore, the total run time ranged from several seconds to a little more than an hour. The number of nodes created ranged from 15 to 27, while the numbers of columns generated ranged from 9,100 to 77,400. It was observed that finding the LP relaxed solution required most of the total required run time.

5.4 Branch and Price and Cut

Another method to find an integer solution for a relaxed (MI)LP is to use the cutting-plane method. The method works as follows, first, it is tested whether the optimum solution is an integer solution. If this is not the case, then a linear inequality can be added to the (MI)LP that separates the optimum from the feasible set of solutions. This linear inequality is what is referred to as a *cut*. As a result, the previous optimum is no longer feasible. This process is repeated until an integer solution to the (MI)LP is found. This method can be added to the B&P technique to further strengthen the relaxation (Desrosiers & Lübbecke (2011)).

The planning goal in Derigs & Friederichs (2013) is to develop a model that can identify the best possible combination of mandatory and optional flights while also taking into account the available capacity of external flights cargo handling cost, and constraints and maintenance constraints. A set of complex MIPs that represented the different sub-problems was developed and consequently integrated. A branch-and-price-and-cut algorithm was developed to solve the model. To test the model, they have created three realistic problem instances representing different planning scenarios for four different types of freight airlines. Moreover, for each of the instances, five different data sets were generated. So, in total 60 instances were tested. The smallest test instance consisted of approximately 70 mandatory flights, 30 optional flights 800 O&Ds, and on average 250 legs. On the other hand, the largest instance consisted of approximately 360 mandatory flights, 720 optional flights, 6,000 O&Ds en on average 2,500 legs. Their algorithm was able to produce high-quality integer results for all instances within a reasonable time. The solution gap ranged from 0 to a maximum of 0.66%. Moreover, the run time ranged from only several seconds to a maximum of 4 hours. The authors attributed these excellent results to the use of implied bound cuts. These cuts were introduced because they were better at accounting for fixed flight costs than the general cuts.

5.5 Available solver software

There is a wide variety of both commercial and non-commercial solvers available. Each of these solvers can apply different types of algorithms. In Table 5.1 an overview of the available commercial MILP solvers is given. To determine which solver is most suitable for this study, the previously discussed solution techniques were reviewed. The goal was to determine which studies had similar objectives as this study and which study considered similar problem size instances. Moreover, it was analyzed which solvers were used in these studies.

The first study that stood out, was the one conducted by Armacost et al. (2002). As previously explained, the authors solved the ESSND problem by using a composite variable formulation. Although the problem instances considered in this study were relatively small, this is still an interesting method, as in later studies it was also successfully applied to larger instances (see, e.g., Armacost et al. (2004) or Fleuren (2013)). To find their results, they used a B&B technique combined with the CPLEX solver.

The second study that was particularly interesting is the one by Derigs & Friederichs (2013). From a predetermined set of flights, they found the optimal combination of flights by integrating the fleet assignment and cargo routing problems. In this study, both these problems also have to be solved. Moreover, the larger instances that were considered are comparable with the problem instances that will be considered in this study. They applied the branch & price & cut technique in combination with CPLEX to find the optimal solutions.

It immediately becomes clear that both studies have used CPLEX to solve their optimization problems. This directly confirms that this solver will be able to apply both the techniques that are thought to be the most interesting for this study. Furthermore, since this study's planning horizon is about 6 - 8 weeks, the importance of a short run time is not the most important constraint. Therefore, the run time of several minutes to several hours of the previously discussed studies seems acceptable. Moreover, because the authors have experience with CPLEX, this makes it the convincing winner for the choice of MILP solver. To further support this choice, it was also found that CPLEX performs better than, for example GuRoBi in case of high dimensional problems (Anand et al. (2017)). Since the final goal of this thesis is to analyze and optimize the global network of a large cargo-only carrier, it is expected that the data used will be highly dimensional.

Table 5.1: MILP optimization software packages. (Source: adapted form Kumar & Mageshvaran (2020)

Software name	**Founders**	**Algorithms utilized**	**Parameters included**	**Features**	**Specifications**	**Interfaces, modelling languages**
CPLEX Mittelmann (2014) (IBM ILOG CPLEX Optimization Studio)	Bixby the founder of CPLEX, retained and provided by IBM	Branch and cut algorithm and Dynamic search algorithm	Mipemphasis meta parameter	Capable of calculating multiple optimal solutions and the solutions have stored in a solution pool	Version: 12.8.0, Website: http:// www.ibm. com/analytics/ cplex-optimizer	License: proprietary. C, C++, Java, .NET, MATLAB, Python, Microsoft Excel
GUROBI	Zonghao Gu, Edward Rothberg, and Robert Bixby	Include cutting planes algorithm, heuristics and search techniques	MIP-Focus meta parameter	New MILP solver that is designed with modern multicore processing technology to obtain an optimal solution	Version: 3.0, Website: www.gurobi.com. License: proprietary	Object-oriented interfaces for C++, Java, .NE T, and Python
LINDO Nash (1991) (Linear, Interactive, and Discrete Optimizer)	LINDO SYSTEMS INC	It offers different forms of cutting planes algorithms and different node selection rules	LINDO also comprises a mipemphasismet a parameter that has used for adjusting algorithm parameters	Significantly Faster on Large Quadratic Models. Improved Handling of Models with Discontinuou s Functions	Version:10.0, Website: www.lindo.com	License: proprietary. C, Visual Basic, MATLAB, Microsoft Excel
MOSEK	Mosek ApS, a Danish company	Branch and bound, branch and cut, and state-of-theart interiorpoint optimizer algorithm	Parameters include optimizer choice for solving linear problems, turning pre-solve parameter value and feasibility of tolerances value	MOSEK interior-point optimizer can reliably detect a primal and dual infeasible status of solutions	Version:9 beta, Website: www.mosek.com	License: proprietary. C, C++, Java and Python. Mosek is accessible for use by customers through a GAMS interface on the NEOS Server

6 Use Case

This chapter will discuss how the knowledge acquired by carrying out this literature study will be combined. As mentioned in chapter 2 the objective of this research is to develop a tool that can assist a full-cargo airline in deciding on how to adapt their flight schedule such that it can transport the re-predicted demand most cost-effectively. First, the general layout of the model is presented in section 6.1. Thereafter, the inputs of the model will be discussed in more detail in section 6.2. This will be followed by an explanation of which models are expected to be used in section 6.3. Finally, the assumptions that will be made to design and formulate the model are elaborated upon in section 6.4. It should be noted that the test data for this thesis is provided by a large cargo-only carrier that will remain anonymous. For ease of writing, the airline will be referred to as our airline for the remainder of this chapter.

6.1 Model Layout

It is essential to have a clear overview of what the complete model will look like in the end and how the sub-parts will interact with each other. Therefore, a schematic overview of the model its flow diagram is presented in Figure 6.1. As can be seen, the model is split up into two main parts. In one part of the model, a benchmark solution is determined, while in the other, the model determines the optimal solution. In the figure, the dark orange shaded blocks represent inputs, the yellow shaded blocks represent the models that will be used to determine the output blocks, which are shaded in green. Finally, the model's solution will be compared with the benchmark solution, which will produce the final output, which is shown in the blue shaded block.

Forecast demand and a flight schedule should be provided to the model to determine the benchmark solution. This will then be used to model the cargo flow through the network. The flight schedule and resulting cargo flow are then used to estimate the operating cost. Meanwhile, the same forecast demand will be used as an input for the other part of the model. In this part, a completely new flight schedule is designed. This will be accomplished by integrating a schedule design model with a cargo flow model. The output of these two integrated models will be a new flight schedule and an estimate for the expected operating cost. Thereafter, it can be determined whether the model has indeed improved the benchmark solution or if the model was unsuccessful in improving it. Furthermore, it should also be examined what the influence of time discretization on different time intervals is. On one side, the problem size can become intractable if the time step is too small. On the other side, a too-large step would make it inaccurate and less realistic. It has been decided that an initial time interval of one hour will be used. It will be analyzed what the effect of decreasing the time interval on the performance of the model has in terms of run time and optimality of the solution.

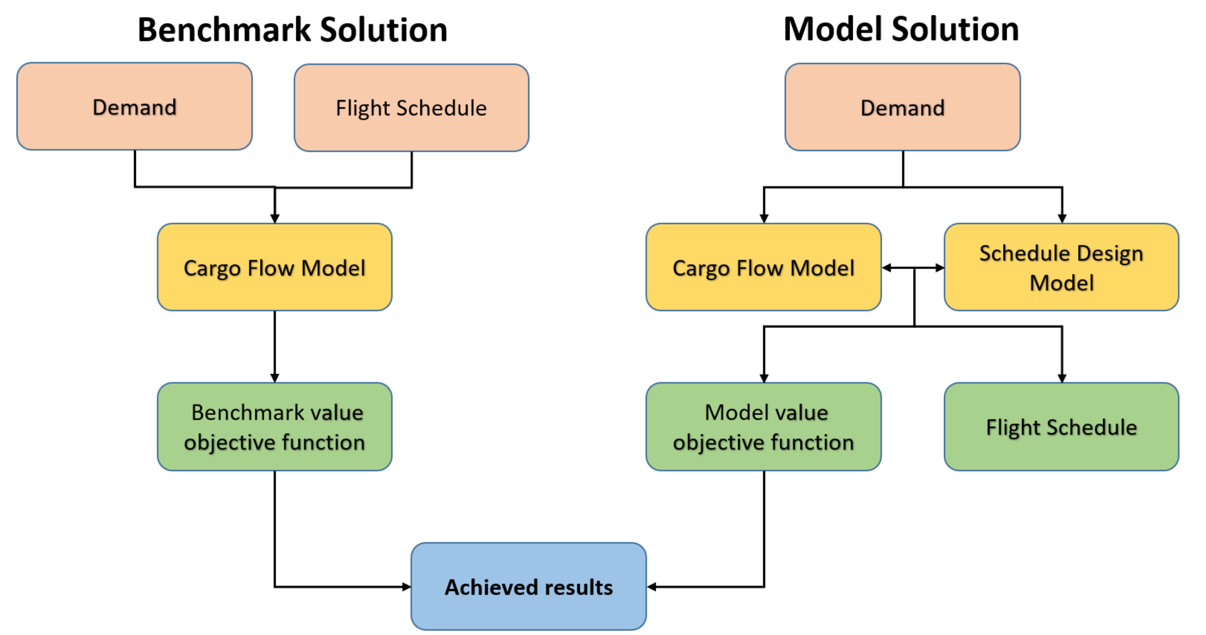

Figure 6.1: A schematic overview of how the model will be created and how the cost-effectiveness of it will be determined.

6.2 Model Inputs

The inputs shown in Figure 6.1 are the inputs that will change when different scenarios are tested. Apart from these inputs, several other inputs are static in the sense that they do not change (much) for different scenarios. Historic schedules will be used as test instances, as these will give a clear indication of what has happened without the model and what could have happened if the model would have been used. Each of the inputs that the model will use will be shortly discussed now.

Flight Schedule As mentioned before, to determine the benchmark solution for a scenario, the planned flight schedule has to be fed to the model. This flight schedule will be presented as is shown in Table 6.1. The headers have the following meaning. Rot is the rotation of which the flight leg is part. Fl nr means flight number, ORG means origin, and DST means destination. Furthermore, Dep and Arr stand for departure and arrival, respectively. Blk hrs represent the expected block time of the flight. This represents the total time between when the aircraft engines turn on at the origin location and when the engines turn off when the final destination has been reached. This means it includes not only flight time but also expected taxi time. The last column represents the aircraft fleet type that is expected to fly that particular flight.

Table 6.1: Example of a general flight schedule.

Rot	Fl Nr	ORG	DST	Dep Date	Dep Time		Arr Date	Arr Time		Blk hrs	A/C type
					UTC	Local		UTC	Local		
1	1234	AAA	BBB	dd/mm/yyyy	hh:mm	hh:mm	dd/mm/yyyy	hh:mm	hh:mm	hh:mm	B747

Demand The forecast demand plays a crucial role in this thesis as it serves both parts of the model as an input. The historical data that will be used for this analysis will look something like the table that is shown in Table 6.2. As can be seen, for each O&D pair, a particular demand will be given, which will be expressed in tonnes. Furthermore, a release and due date of the demand will be provided as well.

Table 6.2: Example of what the demand for a certain O&D pair will look like.

ORG	DST	Demand		Release Date	Due Date
		[kg]	[m3]		
AAA	BBB	*kg*	*vol*	dd/mm/yyyy	dd/mm/yyyy

Airports The model will also be provided with a list of all airports in the network of our airline. For each airport, it shall be checked if the model needs to take into account certain opening hours or curfews. However, at this point, these are still unknown. It is expected that the list of airports will look something like the table shown in Table 6.3. At this stage, it is still imperative to keep the model as flexible as possible. Although certain curfews or opening hours might not be known yet, they might influence the determined solution's feasibility for the flight schedule. Therefore, it should remain possible to incorporate these at a later stage.

Table 6.3: Example of what the airport list for the complete network would look like.

Airport	Long	Lat	Throughpout	Curfews	Opening hours	Transfer time
AAA	*deg*	*deg*	dd/mm/yyyy	hh:mm - hh:mm	hh:mm - hh:mm	hh:mm

Fleet Our airline operates several fleet types in its network. Moreover, certain aircraft are only allowed to fly within certain regions. This should be taken into account in the model. Furthermore, for each fleet type, the following things will be determined: a standard turnaround time, a payload range diagram, and an estimate of the average block-hour cost. The expected flight range of an aircraft strongly affects how much payload it can carry. Therefore, an attempt will be made to create a general payload range diagram for each fleet type based on historical data. Furthermore, the average block-hour cost represents costs such as leasing cost, insurance, cost of permits, cost of crew, etc. It is assumed that calculating the cost based o the block-hour cost of the used aircraft is a good indication for the total operating costs of the network.

Network As a final input for the model, a connection matrix is built. This matrix contains all the expected block-hour connection times for all of the possible O&D pairs in the network. For each of the fleet types, a different connection matrix will be built. This matrix can then also directly constraint unfeasible itineraries for certain aircraft.

6.3 Models used

As was shown in Figure 6.1 two main models will be used, namely the schedule design model and the cargo flow model. When the benchmark solution is determined the cargo flow model runs on its own. It is expected that the objective function of the cargo flow model will be to minimize cost. For example, a penalty could be incurred for not moving a certain part of the demand. Another option would be to determine the added marginal costs of bringing extra cargo on a flight. In order to do this, it could be derived from aircraft characteristics what the effect of the added weight on the fuel consumption is.

Similarly, when the cargo flow model and the schedule design model are integrated to determine the model solution, the objective will be to minimize the overall cost. The operating cost of aircraft can be calculated by summing the cost of all arcs used and their respective use cost for each specific fleet type. Both the models will be constraint by the following constraints:

- Flow balance of aircraft at every node in the network is ensured
- Flow balance of demand at every node in the network is ensured
- The number of aircraft used in the schedule does not exceed the available number of aircraft
- Demand can only be assigned to arcs that has an aircraft assigned to it
- Not more demand than the available aircraft capacity can traverse an arc
- Demand is only available after its release date (and time) and should be delivered before its due date (and time)
- Only feasible flow paths are considered for the cargo flow model
- The sorting capacity of each airport in the system cannot be exceeded
- After landing, each aircraft should remain grounded for at least the minimum TAT
- After landing, any part of the demand can only be transferred to flights that depart after the minimum transfer time has passed
- The maximum payload – range of each aircraft cannot be exceeded
- Each aircraft can only be at one place in the network at any point in time
- Any part of the demand can only be at one place in the network at any point in time

6.4 Assumptions

Finally, it is important to be clear which assumptions are made to develop the model. It is not expected that these assumptions strongly influence the model's outcomes, which is important because the model has to determine a realistic and feasible flight schedule as a solution.

1. The forecast demand is static. This means that once the forecast is fed into the model, it is assumed to remain the same.
2. Next day delivery applies to all of the demand.
3. Demand for each O&D pair can be divided in any way over multiple itineraries.
4. Demand can be transferred at any airport in the network.
5. Arrival and departure times of aircraft are deterministic.
6. Unloading and loading of an aircraft can always be done within the specified turnaround times of that aircraft.
7. For each airport, a standard transfer time is accounted for when transferring cargo irrespective of it being a stop-over or a final destination.
8. The maximum allowed arrival or departure rate at any airport is not taken into account.
9. Calculating the total cost of operating all scheduled aircraft by using their average block hour cost is a good estimate for the total operating cost of the network.
10. Maintenance constraints do not have to be considered at this point in the planning horizon.

A Gantt Chart

Thesis - Anne Bart Beijneve...	start	end
Literature Study and Project Plann...	**01/12/20**	**26/02/21**
Research relevent literature	01/12	12/02
Write Research Methodologies paper	21/12	20/01
Write Literature Study	04/01	12/02
Review RM paper	25/01	29/01
Review LS	22/02	26/02
Hand-in LS	26/02	26/02
Initial Phase	**01/03/21**	**28/05/21**
Analyze the static historical perform...	01/03	12/03
Analyze historical flow	15/03	19/03
Create network	22/03	09/04
Model constraints	12/04	23/04
Model the flow through the network	26/04	07/05
Reporting of results	10/05	21/05
Prepare Midterm Meeting	24/05	27/05
Midterm Meeting	28/05	28/05
Final Phase	**31/05/21**	**26/07/21**
Compare different solution policies	31/05	18/06
Perform sensitivity analysis	21/06	02/07
Reporting of results	05/07	16/07
Prepare Green Light Meeting	19/07	23/07
Green Light Meeting	26/07	26/07
Last Mile	**27/07/21**	**17/08/21**
Incorporate comments Green Light ...	27/07	02/08
Finalize thesis and paper	03/08	09/08
Prepare Thesis defence	10/08	16/08
Thesis defence	17/08	17/08

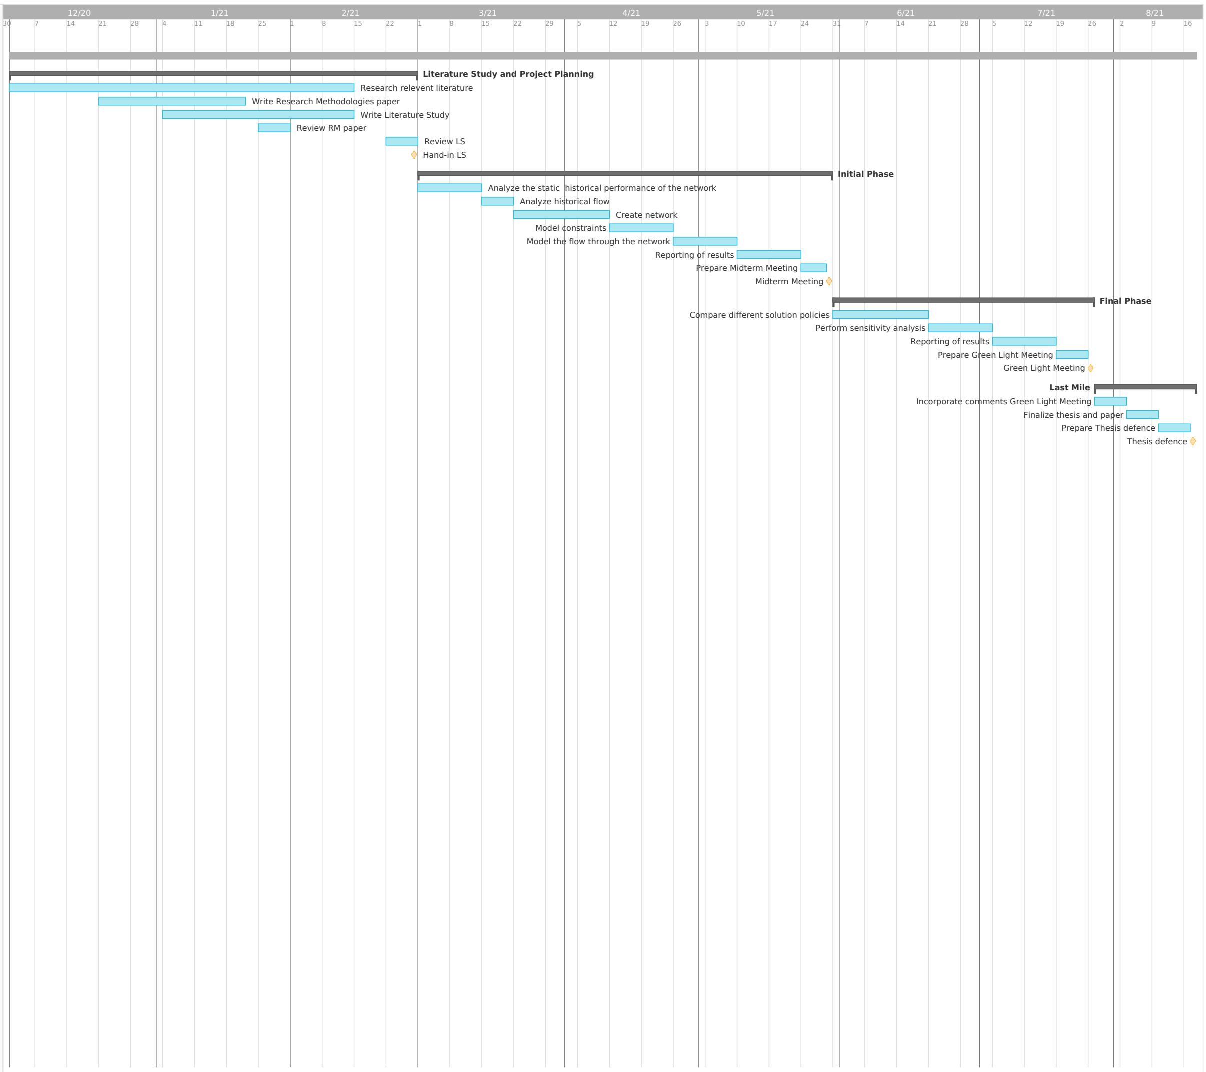

References

Abara, J. (1989), ‘APPLYING INTEGER LINEAR-PROGRAMMING TO THE FLEET ASSIGNMENT PROBLEM’, *INTERFACES* **19**(4), 20–28.

Amaruchkul, K., Cooper, W. L. & Gupta, D. (2007), ‘Single-leg air-cargo revenue management’, *Transportation Science* **41**(4), 457–469.

Anand, R., Aggarwal, D. & Kumar, V. (2017), ‘A comparative analysis of optimization solvers’, *Journal of Statistics and Management Systems* **20**(4), 623–635.

Antes, J., Campen, L., Derigs, U., Titze, C. & Wolle, G.-D. (1998), ‘Synopse: a model-based decision support system for the evaluation of flight schedules for cargo airlines’, *Decision Support Systems* **22**(4), 307 – 323.
URL: *http://www.sciencedirect.com/science/article/pii/S016792369800027X*

Armacost, A., Barnhart, C. & Ware, K. (2002), ‘Composite variable formulations for express shipment service network design’, *Transportation Science* **36**, 1–20.

Armacost, A., Barnhart, C., Ware, K. & Wilson, A. (2004), ‘Ups optimizes its air network’, *Interfaces* **34**, 15–25.

Barnhart, C., Boland, N., Clarke, L., Johnson, E., Nemhauser, G. & Shenoi, R. (1998), ‘Flight string models for aircraft fleeting and routing’, *Transportation Science* **32**, 208–220.

Barnhart, C., Johnson, E., Nemhauser, G., Savelsbergh, M. & Vance, P. (1996), ‘Branch-and-price: Column generation for solving huge integer programs’, *Operations Research* **46**.

Barnhart, C. & Schneur, R. R. (1996), ‘Air network design for express shipment service’, *Operations Research* **44**(6), 852 – 863.

Bartodziej, P. & Derigs, U. (2004), *On an Experimental Algorithm for Revenue Management for Cargo Airlines*, Lecture Notes in Computer Science, pp. 57 – 71.

Berge, M. E. & Hopperstad, C. A. (1993), ‘Demand driven dispatch: A method for dynamic aircraft capacity assignment, models and algorithms’, *Operations Research* **41**(1), 153–168.

Boeing (2018), ‘World air cargo forecast 2018 - 2037’, How it was published. Available at ”https://file.veryzhun.com/buckets/carnoc/keys/3fa55da709101d0d937e78732a88cd9d.pdf”.

Boeing (2020), ‘World air cargo forecast 2020 - 2039’. Available at ”http://www.boeing.com/commercial/market/cargo-forecast”.

Bombelli, A., Santos, B. & Tavasszy, L. (2020), 'Analysis of the air cargo transport network using a complex network theory perspective', *Transportation Research Part E: Logistics and Transportation Review* **138**.
URL: *http://www.sciencedirect.com/science/article/pii/S1366554520306104*

Brough, B.M., ed. (1966), *The Integration of Management Information System Planning - The BEA-IMPACT-Concept*, Vol. 6, AGIFORS.

Bélanger, N., Desaulniers, G., Soumis, F. & Desrosiers, J. (2006), 'Periodic airline fleet assignment with time windows, spacing constraints, and time dependent revenues', *European Journal of Operational Research* **175**(3), 1754–1766.
URL: *https://www.sciencedirect.com/science/article/pii/S0377221705004741*

Clarke, L., Hane, C., Johnson, E. & Nemhauser, G. (1996), 'Maintenance and crew considerations in fleet assignment', *Transportation Science* **30**, 249–260.

Cook, G. & Goodwin, J. (2008), 'Airline networks: A comparison of hub-and-spoke and point-to-point systems', *The Journal of Aviation/Aerospace Education & Research* **17**, 51–60.

Delgado, F., Sirhan, C., Katscher, M. & Larrain, H. (2020), 'Recovering from demand disruptions on an air cargo network', *Journal of Air Transport Management* **85**, 101799.
URL: *http://www.sciencedirect.com/science/article/pii/S0969699719302479*

Derigs, U. & Friederichs, S. (2013), 'Air cargo scheduling: integrated models and solution procedures', *OR Spectrum* **35**(2), 325 – 362.

Derigs, U., Friederichs, S. & Schäfer, S. (2009), 'A new approach for air cargo network planning', *Transportation Science* **43**, 370–380.

Desrosiers, J. & Lübbecke, M. (2011), *Branch-Price-and-Cut Algorithms*, John Wiley & Sons, Inc., pp. 1 – 18.

Dewulf, W., Meersman, H. & Voorde, E. (2014), 'From carpet sellers to cargo stars: Analyzing strategies of air cargo carriers', *Journal of Air Transport Studies* **5**, 96–119.

Fleuren, H. (2013), 'Supply chain wide optimization at tnt express', *Interfaces* .

Ford, L. R. & Fulkerson, D. R. (1958), 'A suggested computation for maximal multi-commodity network flows', *Management Science* **5**(1), 97 – 101.

Forster, P. W. & Regan, A. C. (2001), 'Electronic integration in the air cargo industry: An information processing model of on-time performance', *Transportation Journal* **40**(4), 46–61.

Friederichs, S. (2010), Air Cargo Schedule Planning, PhD thesis, University of Cologne, Cologne.

Gardiner, J., Ison, S. & Humphreys, I. (2005), 'Factors influencing cargo airlines' choice of airport: An international survey', *Journal of Air Transport Management* **11**(6), 393–399.

Gopalakrishnan, B. & Johnson, E. (2005), 'Airline crew scheduling: State-of-the-art', *Annals of Operations Research* **140**, 305–.

Group, W. B. (2009), *Air Freight: A Market Study with Implications for Landlocked Countries*, The World Bank, Washington, DC.

Haouari, M., Sherali, H. D., Mansour, F. Z. & Aissaoui, N. (2011), 'Exact approaches for integrated aircraft fleeting and routing at tunisair', *Computational Optimization and Applications* **49**(2), 213–239.

Hu, Y., Liao, H., Zhang, S. & Song, Y. (2017), 'Multiple objective solution approaches for aircraft rerouting under the disruption of multi-aircraft', *Expert Systems with Applications* **83**, 283–299.
URL: *https://www.sciencedirect.com/science/article/pii/S095741741730249X*

Huang, K. & Lu, H. (2015), 'A linear programming-based method for the network revenue management problem of air cargo', *Transportation Research Procedia* **7**, 459 – 473. 21st International Symposium on Transportation and Traffic Theory Kobe, Japan, 5-7 August, 2015.

Jones, K. L., Lustig, I. J., Farvolden, J. M. & Powell, W. B. (1993), 'Multicommodity network flows: The impact of formulation on decomposition', *Mathematical Programming* **62**(1-3), 95–117.

Kim, D., Barnhart, C., Ware, K. A. & Reinhardt, G. (1999), 'Multimodal express package delivery: A service network design application', *Transp. Sci.* **33**, 391–407.

Kuby, M. J. & Gray, R. G. (1993), 'The hub network design problem with stopovers and feeders: The case of federal express', *Transportation Research Part A: Policy and Practice* **27**(1), 1 – 12.
URL: *http://www.sciencedirect.com/science/article/pii/096585649390012A*

Kumar, P. H. & Mageshvaran, D. R. (2020), 'Methods and solvers used for solving mixed integer linear programming and mixed nonlinear programming problems: A review', *INTERNATIONAL JOURNAL OF SCIENTIFIC & TECHNOLOGY RESEARCH* **9**.
URL: *http://www.sciencedirect.com/science/article/pii/S0969699719302479*

Li, D., Huang, H.-C., Morton, A. D. & Chew, E.-P. (2006), 'Simultaneous fleet assignment and cargo routing using benders decomposition', *OR Spectrum* **28**(3), 319–335.
URL: *https://dx.doi.org/10.1007/s00291-006-0041-8*

Lipczynski, J., Wilson, J. O. S. & Goddard, J. A. (2017), *Industrial Organization: Competition, Strategy, Policy*, Pearson, Harlow, England.

Lohatepanont, M. & Barnhart, C. (2004), 'Airline schedule planning: Integrated models and algorithms for schedule design and fleet assignment', *Transportation Science* **38**, 19–32.

Loughran, B.P.., ed. (1972), *An Airline Schedule Construction Model*, Vol. 12, AGIFORS.

Louwerse, I., Mijnarends, J., Meuffels, I., Huisman, D. & Fleuren, H. (2014), Scheduling movements in the network of an express service provider, PhD thesis, Tilburg University.

Meuffels, W. (2015), The design of road and air networks for express service providers, PhD thesis, Tilburg University.

Mittelmann, H. (2014), 'Recent benchmarks of optimization software', *22nd European Conference on Operational Research* **4**, 1–314.

Nash, J. C. (1991), 'Optimizing add-ins: The educated guess. what's best!', *PC Magazine* **10**, 130–132.

O'Kelly, M. E. (1998), 'A geographer's analysis of hub-and-spoke networks', *Journal of Transport Geography* **6**(3), 171 – 186.
URL: *http://www.sciencedirect.com/science/article/pii/S0966692398000106*

O'Kelly, M. E. (2014), 'Air freight hubs in the FedEx system: Analysis of fuel use', *Journal of Air Transport Management* **36**, 1–12.
URL: *http://www.sciencedirect.com/science/article/pii/S0969699713001464*

Paleari, S., Redondi, R. & Malighetti, P. (2010), 'A comparative study of airport connectivity in china, europe and us: Which network provides the best service to passengers?', *Transportation Research Part E: Logistics and Transportation Review* **46**(2), 198 – 210. Selected Papers from the 12th ATRS Conference.
URL: *http://www.sciencedirect.com/science/article/pii/S1366554509001094*

Popescu, A., Keskinocak, P. & al Mutawaly, I. (2011), The air cargo industry, *in* L. Hoel, G. Giuliano & M. E. Meyer, eds, 'Intermodal Transportation: Moving Freight in a Global Economy', Eno Foundation for Transportation, Washington, DC, USA, chapter 7, pp. 209–237.

Quesada Pérez, J. M., Lange, J.-C. & Tancrez, J.-S. (2018), 'A multi-hub express shipment service network design model with flexible hub assignment', *Transportation Research Part E: Logistics and Transportation Review* **120**, 116–131.
URL: *https://www.sciencedirect.com/science/article/pii/S136655451830557X*

Quesada Pérez, J. M., Tancrez, J.-S. & Lange, J.-C. (2020), 'Express shipment service network design with complex routes', *Computers & Operations Research* **114**, 104810.
URL: *https://www.sciencedirect.com/science/article/pii/S0305054819302527*

Rexing, B., Barnhart, C., Kniker, T., Jarrah, A. & Krishnamurthy, N. (2000), 'Airline fleet assignment with time windows', *Transportation Science* **34**(1), 1–20.

Rosenberger, J., Johnson, E. & Nemhauser, G. (2004), 'A robust fleet-assignment model with hub isolation and short cycles', *Transportation Science* **38**, 357–368.

Sandhu, R. & Klabjan, D. (2006), 'Fleeting with passenger and cargo origin-destination booking control', *Transportation Science* **40**, 517 – 528.

Shen, S. (2004), Logistics service network design : models, algorithms, and applications, PhD thesis, Massachusetts Institute of Technology.

Sherali, H., Bae, K.-H. & Haouari, M. (2013), 'An integrated approach for airline flight selection and timing, fleet assignment, and aircraft routing', *Transportation Science* **47**, 455–476.

Tang, C.-H., Yan, S. & Chen, Y.-H. (2008), 'An integrated model and solution algorithms for passenger, cargo, and combi flight scheduling', *Transportation Research Part E: Logistics and Transportation Review* **44**(6), 1004 – 1024.

Tobin, N. & Butfield, T., eds (1970), *OR Models and their Implementation in the Airline's Planning Cycle*, Vol. 10, AGIFORS.

Yan, S., Chen, S.-C. & Chen, C.-H. (2005), 'Air cargo fleet routing and timetable setting with multiple on-time demands', *Transportation Research Part E: Logistics and Transportation Review* **42**(5), 409 – 430.

Zhou, L., Liang, Z., Chou, C.-A. & Chaovalitwongse, W. (2020), 'Airline planning and scheduling: Models and solution methodologies', *Frontiers of Engineering Management* **7**, 1–26.

III

Supporting work

A

Verification & Validation

A very important aspect of every scientific research is the verification and validation of the work that has been done. Verification is usually done mostly during the research. It is used to verify the correct behavior of the model. This is for example very useful for identifying potential bugs in your code. On the other hand, validation is used to validate if the model does what it was intended to do. To perform the process of verification and validation, two main questions can be posted:

- **Verification:** Are we designing the model in the right way?
- **Validation:** Are we designing the right model?

Verification

It was known beforehand that the final network would become rather large. This made it even more important to ensure the model behaved correctly while changes can still easily be made. Therefore, before the authors worked with a full-scale network, extensive tests were performed with smaller fabricated networks that consisted of several airports, a few periods and 1 or 2 aircraft. It should be noted that at this point it was not yet decided to use a path flow model. So cargo could flow in any feasible way through the network. Moreover, the maximum capacity of the aircraft was not yet dependent on range. So in this first stage of the design process, it was analyzed whether the constraints that should ensure the conservation of flow for both aircraft and demand were performing as intended. This was done by creating a dummy request list and analyzing the outcome of the model. Once this was confirmed, the constraint for maximum aircraft capacity was added. The correct behavior of this constraint was verified by increasing the demand of the requests to more than the aircraft could carry. The results showed that the aircraft indeed carried only its maximum allowed payload.

Once it was decided to design the path-flow model, it was carefully verified whether indeed the shortest paths were found and if the model only used airports that could be part of the itinerary for a certain O&D pair. An additional challenge for the design of the path-flow model was that it should exclude recursive paths from the path options. A recursive path is a path that would, for example, connect the O&D pair AMS-MAD by flying AMS-PAR-AMS-MAD. The difficulty lay within the activity node formulation. A node consisted of both a place and a time. To the model, each node was unique. Therefore, starting at AMS at time period 0 and going back to AMS at i.e. time period 6, was at first not recognized by the model as going back to an airport that had previously been passed. Furthermore, also ground arcs are part of any potentially feasible path. This meant that a path could not simply be eliminated if the airport was already in any of the previous activity nodes. This problem was resolved as follows: per potential path, an airport list was created. The airport of each activity node in the path was added to this list. Then, it was analyzed per airport if it occurred in the list only consecutively or if there was any part in the list where there was at least one other airport mentioned in between two places where the airport under investigation was mentioned. In the case of the first, the path was added to the list of path options, otherwise, it was omitted. It was known that the number of paths could increase quickly. Therefore it was even more important to ensure that the path-flow model would only consider those paths that are actually feasible and logical. With another dummy demand, it was analyzed if the paths chosen by the model were in line with the intended constraints of the model.

Table A.1: Example test demand

req	(ORG, t_av)	(DST, t_due)	DEM
1	(AMS,0)	(MAD,3)	10
2	(MAD,3)	(PAR,5)	10
3	(PAR,5)	(BER,7)	10
4	(BER,7)	(AMS,9)	10
5	(MIL,0)	(BRU,3)	15
6	(BRU,3)	(CGN,5)	15
7	(CGN,5)	(ZRH,7)	15
8	(ZRH,7)	(MIL,9)	15
9	(LON,0)	(OSL,3)	20
10	(OSL,3)	(CPH,5)	20
11	(CPH,5)	(DUB,7)	20
12	(DUB,7)	(LON,9)	20

Table A.2: Example rotations

Rotation 1:	
ORG	**DST**
(AMS,0)	(MAD,3)
(MAD,3)	(PAR,5)
(PAR,5)	(BER,7)
(BER,7)	(AMS,9)

Rotation 2:	
ORG	**DST**
(MIL,0)	(BRU,3)
(BRU,3)	(CGN,5)
(CGN,5)	(ZRH,7)
(ZRH,7)	(MIL,9)

Rotation 3:	
ORG	**DST**
(LON,0)	(OSL,3)
(OSL,3)	(CPH,5)
(CPH,5)	(DUB,7)
(DUB,7)	(LON,9)

Finally, there were three important things to analyze concerning routes that were chosen for aircraft. First of all, the routes should be chosen such that they transported most of the demand. Second, the routes of the model should be partially fixed in compliance with the actual schedule. Third, the starting and ending locations of the airports should also be fixed. A test demand was created that should force the model to choose a certain rotation to maximize the amount of demand that is transported. This was done by determining a request list such that a certain rotation was chosen that would yield the highest amount of demand transported. An example of this test demand can be seen in Table A.1. For clarity, the expected choice of rotations can be found in Table A.2. It is expected that in the case of the shown request list, rotation 3 is chosen. In case the request list was altered, such that the demand for the first 4 requests was triple its previous size, this should lead to the model determining the first rotation as the optimal one. This was indeed the case. Hereafter, it was verified whether if certain legs were fixed, that these legs were indeed included in the model results. This as well was the case. The final step of verification was easily performed by similarly fixing the aircraft access routes as was done for the previous step.

Validation

As mentioned before, the goal of validating your model, is to determine whether you have designed the right model. So, does the model effectuate the original purpose? As this process requires investigation of the functioning of the model as a whole, this is something that is usually done once the model is finished. For this research the objective was to maximize the amount of demand that could be transported by a certain amount of available aircraft. The model was set up as a cost minimization model with two main components: aircraft cost and transportation cost of demand, which also included the no-service cost. The latter also included the no-service cost. The model should reflect the goal of the objective. Therefore, it was analyzed whether the two main components were correctly taken into account. The model is forced to move demand as the cost of not moving it are extremely high. Though, it should be noted that careful attention was paid to not make this cost unnecessarily high. A too large value would nullify the rest of the cost. Moreover, it would also negatively affect the optimization process as a too large value could cause rounding errors. Such errors can be the source of loss of precision in the computation of a feasible solution [1]. A part of the validation process involved analyzing whether the cost were indeed high enough such that any kg of payload would be shipped through the network in case there was capacity on any aircraft available. Furthermore, it was explained that maximization of the level of service provided to all customers is a key driver for the design of the network

[1] https://orinanobworld.blogspot.com/2011/07/perils-of-big-m.html

of our airline. For this research it was assumed that all customers are treated equally and have equally sized request. Therefore, the cost for not servicing demand is equal for all demand. Several test were performed with varying no-service cost across the demand. As expected, the model would serve more requests that had a higher no-service cost, even though this led to a lower over all amount of demand transported. It was found that an equal no-service cost led to the best overall results.

Furthermore, it was essential to evaluate the aircraft cost. Did the model indeed choose the set of flights such that it not only transported the most demand, but also for as little cost as possible? This was validated by modifying the cost for different flight legs. To explain how this was done, the same example as previously shown is taken. There is a certain amount of demand, shown in Table A.3. Note, the demand is now equal. In Table A.4 the cost related to flying the available flight legs in the network are shown. All rotations would lead to the same amount of demand transported. Therefore, the model can only minimize the cost by choosing the flights that minimize the aircraft usage cost. The results showed that the model choose Rotation 1, which confirmed that the model was performing the way it was intended to.

Table A.3: Example demand for validating the correct behavior of the model

req	**(ORG, t_av)**	**(DST, t_due)**	**DEM**
1	(AMS,0)	(MAD,3)	20
2	(MAD,3)	(PAR,5)	20
3	(PAR,5)	(BER,7)	20
4	(BER,7)	(AMS,9)	20
5	(MIL,0)	(BRU,3)	20
6	(BRU,3)	(CGN,5)	20
7	(CGN,5)	(ZRH,7)	20
8	(ZRH,7)	(MIL,9)	20
9	(LON,0)	(OSL,3)	20
10	(OSL,3)	(CPH,5)	20
11	(CPH,5)	(DUB,7)	20
12	(DUB,7)	(LON,9)	20

Table A.4: Example cost for a different flight legs in a dummy network.

Flight leg	**cost**
AMS-MAD	100
MAD-PAR	100
PAR-BER	100
BER-AMS	100
MIL-BRU	200
BRU-CGN	200
CGN-ZRH	200
ZRH-MIL	200
LON-OSL	300
OSL-CPH	300
CPH-DUB	300
DUB-LON	300