Combining operational and environmental sustainability for an integrated flight scheduling and aircraft routing model of a full-cargo carrier

Master of Science Thesis

Ward Broeders

TU Delft

Combining operational and environmental sustainability for an integrated flight scheduling and aircraft routing model of a full-cargo carrier

Master of Science Thesis

by

Ward Broeders

to obtain the degree of Master of Science
at the Delft University of Technology,
to be defended publicly on January 20th, 2022.

Student number:	4454359	
Project duration:	February 2021 – January 2022	
Thesis committee:	Assoc. Prof. Dr. B.F. Lopes Dos Santos	TU Delft, Chair
	Dr. A. Bombelli	TU Delft, Supervisor
	Dr. J. Sun	TU Delft, Examiner

An electronic version of this thesis is available at http://repository.tudelft.nl/.

TU Delft

Acknowledgements

This thesis marks the final part of my master's degree in Air Transport & Operations at the faculty of Aerospace Engineering of the TU Delft. Over the past years, I have thoroughly enjoyed studying at the university and developing my engineering skills. For this final project, the goal was to develop a model that optimises a flight schedule, which is both financially feasible and aims to reduce the aircraft emissions. This thesis offered me a very interesting combination of optimisation problems, the air transport industry and sustainability. I hope that my thesis can contribute to future research in this field and help to more easily and efficiently increase the environmental sustainability of air cargo flight schedules.

Throughout the thesis I have had excellent supervision and support from Alessandro Bombelli. In our weekly meetings we always had interesting discussions where Alessandro gave great insights in ways to fix modelling issues I got stuck on. Unfortunately, most of these meetings had to be done remotely, but luckily in the summer months we could also speak in person. I believe that the support of Alessandro has greatly improved my enthusiasm on the topic and enhanced the quality of this thesis. Therefore, I sincerely want to thank you! Furthermore, I want to thank Paul Roling for joining all the milestone meetings and providing me feedback from an external perspective. Next, also special thanks to Bruno Santos and Junzi Sun for taking the time to join my thesis committee.

Finally, I would like to thank my family and friends for their support during the period of working on my thesis. First of all to my study friends Ide, Niels and Bas, with whom I often worked together during our master courses, internship and thesis. Also thanks to my roommates Didier and Stefan for their help reviewing my thesis. Thanks to my old roommates in Delft and my high school friends for the necessary distraction in between the thesis work. And finally many thanks to my girlfriend Jet for all her support.

Ward Broeders
Delft, January 2021

Contents

List of Figures vii

List of Tables ix

Nomenclature xi

Introduction xiii

I Scientific Paper 1

II Literature Study
previously graded under AE4020 29

1 Introduction 31

2 Research Framework 33
2.1 Problem Statement . . . 33
2.2 Research Questions . . . 33
2.3 Research Objective . . . 34

3 Airline Scheduling Problem 35
3.1 Network Setup . . . 35
3.1.1 Time-Space Network . . . 35
3.1.2 Connection-based Network . . . 37
3.1.3 String-based Network . . . 39
3.2 Objective and Constraints for Aircraft Routing Problems . . . 39
3.2.1 Objective Function . . . 40
3.2.2 Constraints . . . 40
3.3 Comparison of network types . . . 41

4 Air Cargo Modelling 43
4.1 Model Setup . . . 43
4.2 Objective and Constraints for Cargo Scheduling . . . 46
4.2.1 Objective Function . . . 46
4.2.2 Constraints . . . 47
4.3 Solution Methods . . . 48

5 Aircraft Emission Modelling 51
5.1 Aircraft Emissions . . . 51
5.2 Airport-based Emissions . . . 52
5.3 Flight Emissions . . . 53
5.4 Existing Emission Models . . . 54
5.5 Air Cargo Emissions . . . 54
5.6 Relevant Literature . . . 55

6 Conclusion 57

III Supporting work 59

A Input Data 61
A.1 Available airports and flight data . . . 61
A.2 Demand and request generation . . . 64

B Emission Model 67
B.1 Calculating aircraft mass and emissions 67
B.2 LTO emissions validation 68
B.3 Emission matrix. 69

C TSN Pre-processing 71

D Pareto front flight schedules 73
D.1 Network 1: EU 73
D.2 Network 2: EU-NA 74
D.3 Network 3: NA 75

Bibliography 77

List of Figures

3.1 Time line of a single airport in a time-space network [20] . . . 35
3.2 Time-space network with two airports and eight flights, with examples for node aggregation and island isolation [31] . . . 36
3.3 Connection-based network of 3 airports and 4 flights [48] . . . 38
3.4 (a) Arc aggregation of two pairs of two flights. (b) Node aggregation of arrival and departure node. (c) Preprocessing a larger network. All images taken from Clarke et al. [7] . . . 39

4.1 Spider graph for an express cargo airline with a single hub and eight base airports [33] . . . 44
4.2 A graphical representation of the cargo-flow time-space of Yan et al. [46], showing delivery, holding and demand arcs. . . . 45
4.3 Example of the time-space network with three cargo requests [9]. . . . 46
4.4 Graphical representation of the increase in complexity if multiple stops are allowed in a time-space network [46] . . . 48

5.1 Flight phases that are grouped into the LTO-cycle and cruise phase [15] . . . 52
5.2 Reference times and emissions for the standard ICAO LTO-cycle [23] . . . 52
5.3 Carbon efficiency per passenger as a function of flight distance [26] . . . 53

B.1 Iteration to find the aircraft take-off mass for two flights, using a B747-8F with a load factor of 1. The figures show how the aircraft mass changes during the flight. . . . 68
B.2 Analysis of OpenAP difference in LTO calculations. . . . 69

C.1 Example of a time-space network consisting of node set $\mathcal{N}$ and arc set $\mathcal{A}$ . . . 71
C.2 A time-space network pre-processed for example request r_1 . . . 72

List of Tables

3.1 Summary of used literature on FAP, ARP and Air Cargo Scheduling for three different network types . . . 42

4.1 Summary of the available literature on flight scheduling and cargo modelling, with the key characteristics of each model. . . . 50

5.1 Summary of the available literature on aircraft emission modelling for airport-base, flight-based and cargo airline emissions. . . . 56

A.1 Available airports, with the corresponding taxi times. . . . 61
A.2 Flight distances for Network 1: EU in kilometres. . . . 62
A.3 Flight times for Network 1: EU in hours. . . . 62
A.4 Flight distances for Network 2: EU-NA in kilometres. . . . 62
A.5 Flight time for Network 2: EU-NA in hours. . . . 63
A.6 Flight distances for Network 3: NA in kilometres. . . . 63
A.7 Flight times for Network 3: NA in hours. . . . 63
A.8 Weekly flight frequencies used as input for the demand generation for Network 3: NA. . . . 64
A.9 Demand matrix with the weekly cargo demand generated for Network 3: NA in tonnes. . . . 64
A.10 Request set generated using 100% of the cargo demand of Network 3: NA. . . . 65

B.1 Example of an emission matrix of a Boeing 747-8F for three origin-destination pairs. . . . 70

C.1 Amount of nodes before and after pre-processing for an instance with 3 aircraft and 25 cargo requests for all three networks. . . . 72
C.2 Amount of arcs before and after pre-processing for an instance with 3 aircraft and 25 cargo requests for all three networks. . . . 72

D.1 Flight schedule for aircraft 0 for Network 1: EU with 30 available cargo requests and $w_{CO_2} = 0$ . . 73
D.2 Flight schedule for aircraft 0 for Network 1: EU with 30 available cargo requests and $w_{CO_2} = 0.1$ 73
D.3 Request set available for Network 1: EU. . . . 74
D.4 Flight schedule for aircraft 0 for Network 2: EU-NA with 30 available cargo requests and $w_{CO_2} = 0$ 74
D.5 Flight schedule for aircraft 0 for Network 2: EU-NA with 30 available cargo requests and $w_{CO_2} = 0.55$ . . . 75
D.6 Request set available for Network 2: EU-NA. . . . 75
D.7 Flight schedule for aircraft 0 for Network 3: NA with 30 available cargo requests and $w_{CO_2} = 0$ . . 76
D.8 Flight schedule for aircraft 0 for Network 3: NA with 30 available cargo requests and $w_{CO_2} = 0.2$ 76
D.9 Request set available for Network 3: NA. . . . 76

Nomenclature

Abbreviations

ARP	Aircraft Rotation Problem
ASM	Available Seat Mile
BADA	Base of Aircraft Data
CN	Connection-based Network
CO	Carbon Monoxide
CO_2	Carbon Dioxide
CORSIA	Carbon Offsetting and Reduction Scheme for International Aviation
CSP	Crew Scheduling Problem
EEA	European Environment Agency
EMEP	European Monitoring and Evaluation Programme
ETS	Emission Trading System
EU-ETS	European Union Emissions Trading System
FAP	Fleet Assignment Problem
IATA	International Air Transport Association
ICAO	International Civil Aviation Organisation
LTO	Landing and Take-off
MILP	Mixed-Integer Linear Programming
MIP	Mixed-Integer Programming
NO_x	Nitrogen Oxides
OD	Origin-Destination
OpenAP	Open Aircraft Performance Model
SDP	Schedule Design Problem
SN	String-based Network
SO_x	Sulfur Oxides
TSN	Time-Space Network

Notation

$\%_{reduct}$	Amount of carbon emissions that has to be reduced in the revised schedule.	
Δt	Time step.	hr
$\mathscr{A}, \mathscr{A}^I$	Set of arcs and itinerary arcs.	
$\mathscr{A}^G, \mathscr{A}^F, \mathscr{A}^R, \mathscr{A}^N$	Set of ground arcs, flight arcs, request access/egress arcs and no-service arcs.	
C_F	Fixed aircraft operating cost parameter.	$€/hr$
C_{CO_2}	CO_2 cost parameter.	$€/tonne$
C_{fuel}	Fuel cost parameter.	$€/tonne$
$C_{handling}$	Handling cost parameter.	$€/tonne\ TOW$
Cap_1	Payload capacity at maximum fuel load.	kg
Cap_{avg}	Average aircraft capacity used to simulate demand.	kg
Cap^k_{ij}	Payload capacity of aircraft k on arc (i,j).	kg
Cap_{max}	Maximum payload capacity.	kg
E	Set of emission decision variables.	
E_0, E_{max}	CO_2 emissions at $LF = 0$ and LF_{max} for arc (i,j) with aircraft k.	$tonne$
e^k_{ij}	Continuous variable indicating emissions of aircraft k over arc (i,j).	$tonne$

$Emission_{max}$	Maximum emissions used for normalisation.	$tonne$
$freq_{ij}$	Weekly flight frequency on arc (i,j).	
$fuel_0, fuel_{max}$	Fuel weight at $LF = 0$ and LF_{max} for arc (i,j) with aircraft k.	kg
$i_k^+, i_k^-, i_r^+, i_r^-$	Origin and destination nodes of aircraft k and request r.	
k	Aircraft.	
$\mathcal{K}$	Set of aircraft	
l	Airport.	
$\mathcal{L}$	Set of airports.	
LF, LF_{max}	Load Factor and maximum allowed load factor for arc (i,j)	
LF_{avg}	Average load factor used to simulate demand.	
$m_{Fuel\ est}$	Estimation of the take-off fuel weight used for fuel flow calculations.	kg
m_{Fuel}	Calculated fuel weight after fuel flow calculations.	kg
$MTOW$	Maximum Take-Off Weight.	kg
$MZFW$	Maximum Zero Fuel Weight.	kg
$\mathcal{N}, \mathcal{N}^I, \mathcal{N}^R$	Set of nodes, itinerary nodes and request nodes.	
n_{discr}	Number of sections used in the discretisation of the CO_2 emissions.	
$n_{flights}$	Maximum number of flight arcs that a cargo request can be routed over.	
OEW	Operational Empty Weight.	kg
$OperCost_{ij}^k$	Operational costs for routing aircraft k over arc (i,j).	€
p_r	Cargo price.	€/kg
$Profit_{max}$	Maximum profit used for normalisation.	€
Q	Set of cargo routing decision variables.	
q_{ij}^r	Binary variable indicating if cargo request r is routed over arc (i,j).	
r	Cargo request.	
$\mathcal{R}$	Set of cargo requests.	
R_1	Maximum aircraft range at Cap_1.	km
$R_{max\ cap}$	Maximum aircraft range at Cap_{max}.	km
R_{max}	Maximum aircraft range.	km
$ratio_1, ratio_2, ratio_3$	Ratio used to distribute demand over primary, secondary and tertiary destinations.	
Rev_{ij}^r	Revenue earned from routing cargo request r over arc (i,j).	
S	Set of LF segment decision variables.	
s_r	Strategic weighting factor.	
s_{ij}^k	Binary variable indicating if the LF of aircraft k falls within this segment.	
$\mathcal{T}_{hor}$	Set of time steps.	
t_r^+, t_r^+	Release time and due time of cargo request r.	
$T_{block\ max}$	Maximum allowed block time during the simulation horizon.	hr
T_{hor}	Final time-stamp of the time horizon.	hr
t_{ij}	Flight time for flight arc $(i,j.)$	hr
t_{LTO}	LTO time.	hr
t_{TAT}	Turn-around time.	hr
TOW	Take-Off Weight.	kg
V_{cruise}	Cruise velocity	km/h
w_r	Weight of cargo request r.	kg
w_{CO_2}	Emission weighting factor.	
$w_{fuel,ij}^k$	Continuous variable to indicate fuel weight of aircraft k for arc (i,j).	kg
$w_{take-off,ij}^k$	Continuous variable to indicate take-off weight of aircraft k for arc (i,j).	kg
X	Set of aircraft routing decision variables.	
x_{ij}^k	Binary variable indicating if aircraft k is routed over arc (i,j).	
$\mathcal{Z}$	Objective function of total profits.	

Introduction

Just like other industries, the global aviation market is under large pressure to decrease the global emissions. In 2018, 2.4% of the world wide fossil fuel emissions are emitted by the aviation industry, with 8% of that related to dedicated freighter flights [19]. Even though the passenger airlines are heavily impacted by the COVID-19 pandemic, demand for full freighter airlines has only increased. Boeing expects that the world air cargo traffic keeps increasing with 4% per year [3], which will most likely result in even higher carbon emissions [22].

Some attempts are being made to reduce the carbon emissions of the aviation industry, like phasing out older aircraft, improving the carbon efficiency of new aircraft models and the introduction of alternative fuels. However, due to the long aircraft lifetime and relatively slow innovation in sustainable fuel techniques, improvements can only be expected on the long term. For reduction of aviation emissions on a shorter term, ICAO and the European Commission have instated two initiatives that directly target airlines financially to reduce and offset their carbon emissions: CORSIA and EU-ETS. These emission trading systems penalise airlines if they exceed the allowed amount of carbon emissions. This creates financial incentive for airlines to decrease the emissions of their operations. A way to do this on a shorter term than renewing the fleet is to alter the flight schedule. Therefore, this paper looks into what potential CO_2 reduction is available when also taking account carbon emissions in the flight scheduling process.

When designing a flight schedule, multiple sub-problems have to be solved. Based on the available demand between airports, trade-offs have to be made between the operational costs of certain flights and the revenue that can be earned from transporting cargo. In addition to this existing trade-off, which is already broadly described in literature, this paper also introduces the aircraft emissions into the decision making process. A novelty of the model developed in this paper is that the emissions are modelled with a direct dependency on the cargo that is on-board of the aircraft. Using this model, multiple experiments are performed to show the financial feasibility of schedule revisions that are necessary to decrease the emissions. However, this will come with a loss of profit, meaning that it is up to the airline to decide what schedule revisions are carried out. Overall, this paper is a starting point for further research in this field and can help airlines to identify profitable schedule revisions to reduce the carbon emissions of their operations.

This thesis report is organised as follows: In Part I, the scientific paper is presented. Part II contains the relevant Literature Study that supports the research. Finally, in Part III, the supporting work will go deeper into some parts of the model, input data and additional results.

I

Scientific Paper

Combining operational and environmental sustainability for an integrated flight scheduling and aircraft routing model of a full-cargo carrier

W.A. Broeders*

Delft University of Technology, Delft, The Netherlands

Abstract

This study proposes a flight scheduling model with an added aircraft emission model to solve the schedule design, aircraft routing and cargo routing problems for a full-cargo airline, where aircraft emissions are explicitly part of the decision-making process. Our model considers both operational sustainability (maximisation of profit) and environmental sustainability (minimisation of CO_2 emissions) in the objective function and can be used to identify trade-offs between the two potentially contrasting objectives. Aircraft emissions are modelled based on the aircraft type and load factor for each flight leg. Several experiments have been performed using 3 different sub-networks of a full-cargo airline as a reference, with instances of up to 8 airports, 3 aircraft and 25 cargo requests. The results show how different network characteristics and changes in cargo demand affect the profit decrease required to reduce emissions. On average, for a reduction of 25% of carbon emissions, profits in networks with short to medium-range flights decrease by roughly 14%. The expected loss of profit is larger and more inconsistent for networks that include long-range flights.

Keywords: Aircraft emission modelling, Flight scheduling, Aircraft routing, Cargo routing, Mixed-integer linear programming, Pareto front

1 Introduction

The global aviation industry is responsible for a large amount of carbon emissions. Although the contribution of 2.4% of the world wide fossil fuel emissions seems relatively small, this translates to 918 million metric tons of CO_2 [Graver et al., 2019]. Of all aviation emissions in 2018, 19% can be assigned to the transportation of cargo, with 8% related to dedicated freighter flights and the rest coming from freight transported in the cargo belly of passenger aircraft. The carbon emissions increased with 32% over the five year period prior to 2018 and forecasts by the International Civil Aviation Organisation ICAO show this growth continuing in the coming years [ICAO, 2016b]. This is further substantiated by the expectation that the world air cargo traffic keeps growing with 4% per year [Boeing, 2020]. Despite this, the International Air Transport Association IATA aims to achieve carbon neutral-growth from 2020 onward, with a following goal to halve the carbon emissions in 2050 of what they were in 2005 [IATA, 2020]. Initiatives like the introduction of alternative fuels and innovations in aircraft fuel efficiency are expected to only make a difference in the long term. Therefore, emission trading systems are set up on a European (EU-ETS) and global (CORSIA) scale to stabilise the increase of emissions until the previously mentioned improvements will become available. In these schemes, airlines have to offset the carbon emissions of their operations. This puts a direct financial pressure on the airlines themselves to reduce CO_2 emissions. With cargo-only carriers generally flying older aircraft, such airlines can choose to replace older aircraft types in their fleet by newer, more efficient models, which can significantly reduce the emissions of the airline [Brueckner and Abreu, 2017]. However, this is both very costly and takes multiple years to complete. A more short term option would be to increase the carbon efficiency of the operations by adapting the flight schedule and aircraft rotations.

The flight schedule design of a full freighter airline consists of multiple subproblems, where the best airport combinations, flight frequencies, aircraft types for each flight leg, and cargo transported on each flight leg are determined [Derigs et al., 2009]. In the aircraft routing part, a specific aircraft is assigned to each flight in the schedule to create flight rotations for all aircraft in the fleet. The goal is to maximise the available revenue of the network, while minimising the operating costs of aircraft rotations. Adding the reduction of aircraft emissions to the objectives of the flight planning can help to find schedule revisions that reduce the carbon emissions, while still being financially feasible. This paper proposes a mixed-integer linear programming (MILP) flight

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

scheduling model with an integrated aircraft emission model that can be used to find what schedule revisions are necessary to reduce the emissions of a dedicated cargo airline. This model aims to maximise the profit of the airline, which consists of revenue from transporting cargo and operating costs to fly the aircraft. An emission term is introduced in the objective function that has a negative impact on the profit. The emissions are mapped in an emission matrix that depends on each unique (flight leg, aircraft type, load factor) combination, hence influencing routing decisions both at aircraft and cargo level. A set of cargo requests is available for the airline to transport, which are generated based on Cargolux flight frequencies. A general overview of the different steps that lead to MILP optimisation model is shown in Figure 1.

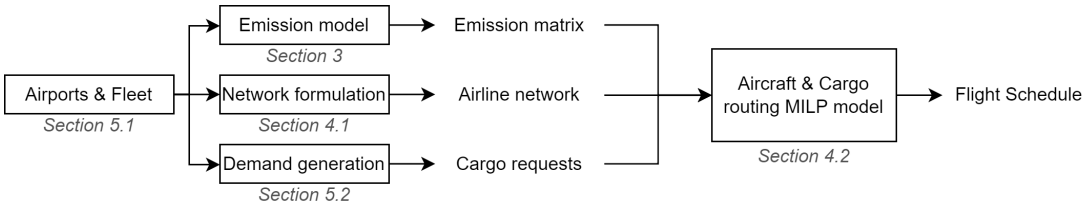

Figure 1: Flowchart showing the relation between the different segments of the flight schedule optimisation.

Multiple flight scheduling and aircraft routing models for full-cargo airlines are available in the academic literature. Also, a lot of research is available that calculates and analyses aircraft emissions around airports and for airline networks. However, to the best of our knowledge, no prior work exists that combines these aspects and, more specifically, computes aircraft emissions based on cargo routing decisions (via the load factor). Therefore, the first contribution of this work is the development of a model that simultaneously optimises (i) aircraft and cargo routing and (ii) the environmental impact of these rotations considering explicit dependency on cargo routing decisions. A further contribution is made by performing experiments with the proposed model, which are performed based on three sub-networks of a major full-cargo airline, where possible flight legs are determined from publicly available aviation data repositories. Due to the lack of available data, apart from global average load factors, fully synthetic cargo demand data is generated. The results presented in this paper confirm the financial feasibility of reducing emissions by flight schedule revisions, but also show that this comes with a loss of profit and it is up to the airline's decision-makers to decide where the sweet spot is.

This paper is structured as follows. Section 2 gives an overview of the related literature to this research. This is followed by the methodology of this study, which is split up in the description of the emission model in Section 3 and the flight scheduling model in Section 4. Section 5 describes the experimental setup and defines the four different experiments that are performed. The results of these experiments are discussed in Section 6. Section 7 presents the discussion, followed by the conclusion and recommendations in Section 8.

2 Literature Review

This section gives an overview of the relevant literature for this research. First, the different aspects of airline scheduling will be discussed, which is followed by a more detailed look into air cargo operations. After that, the available literature on aircraft emission models and research that combines flight scheduling for cargo airlines with aircraft emission modelling is analysed.

The airline planning problem can be split up into four sub-problems, namely the schedule design (SDP), fleet assignment (FAP), aircraft rotation (ARP) and crew scheduling problem (CSP). The first three of these problems are modelled in this study, as the CSP is considered less relevant in this stage of research. The SDP generates a timetable based on the available demand, which is further discussed in the air cargo modelling part. The model developed in this research solves the SDP simultaneously with the FAP and ARP due to the chosen network modelling approach. The FAP aims to maximise the profit of a given flight schedule by assigning aircraft types to each flight. This is followed by the ARP, which assigns specific aircraft to every flight, often also considering maintenance constraints. For the modelling of the FAP and ARP, three different methods are described in literature: Time-space networks, connection-based networks and string-based networks. A recent literature study that analyses the different parts of the airline planning problem is presented by [Zhou et al., 2020].

The time-space network is first proposed by [Hane et al., 1995] to solve the FAP. In such a network, a node represents an airport location at a certain time-stamp and is connected to other nodes by flight arcs, ground arcs and wrap-around arcs. [Sherali et al., 2006] uses this approach to solve the ARP, which is further extended to an aircraft maintenance routing problem by [Liang et al., 2011] and [Khaled et al., 2018] These studies all use three pre-processing steps that are described by [Hane et al., 1995] to reduce the number of arcs and nodes in the network and speed up the computational time.

The second approach is using a connection-based network, which was also first proposed to solve the FAP [Abara, 1989]. Instead of every airport only having one timeline with their available nodes, in a connection-based model every airport has a separate timeline for arrival and departure nodes. These two node types are connected using connection arcs. [Clarke et al., 1997] included maintenance and crew modelling in the FAP, with the addition of an extra pre-processing step to reduce the problem size. This research is extended to solve the ARP with maintenance constraints by [Haouari et al., 2013]. The advantage of the connection-based network lies in the ability to model the attractiveness of connections and transfers for passenger airlines. However, this is less suitable when cargo operations are modelled, as in this paper.

The final method is the string-based network, which is presented by [Pollack, 1974] to solve the ARP and later adjusted by [Barnhart et al., 1998] to both solve both the FAP and ARP. Flight strings are created, which describe a sequence of flights that can be flown sequentially by a single aircraft. [Liang and Chaovalitwongse, 2013] create their flight strings based on a time-space network, while [Barnhart et al., 1998] and [Froyland et al., 2014] use a connection-based network to find the feasible sequences of flights. The problem size rapidly increases due to the large amount of possible flight strings, which requires a more complex solving technique like column generation. The main benefit of this technique is the ability to solve larger networks or simulate longer term maintenance schedules.

The available research on flight planning for cargo airlines proposes integrated solutions for the SDP, FAP and ARP. While objective functions might slightly change according to the specific model, profit maximisation is the key driver. The model described by [Yan et al., 2006] selects the most profitable airports of the network and creates a flight planning, based on a time-space network. Demand arcs are added to show if certain cargo demand is served. To speed up the model, heuristics are used that vary the allowed number of intermediate stops in a single cargo routing. This model is extended to also optimise flight scheduling for freight airline alliances [Yan and Chen, 2008]. The string-based network is used by [Derigs et al., 2009], which create the flight strings based on a set of mandatory and optional flights. Maintenance modelling is added to this model in [Derigs and Friederichs, 2013]. [Delgado et al., 2020] investigates how cargo flight schedules can be redesigned as a consequence of demand disruptions (differences between expected and actual demand), in order to minimise the rescheduling costs. In the model, both aircraft and cargo requests are routed in the same time-space network. No-service arcs are added that directly connect the source and sink node of a cargo request, to allow the possibility not to transport some demand requests, resulting in a monetary penalty due to the missed delivery. Another approach to react to demand disruptions is given in [Delgado and Mora, 2021], where a combination of heuristics and column generation is used to deal with larger instances.

Several studies on aircraft emissions are available that can be categorised in research into (i) influence on local air quality and (ii) atmospheric climate effects. Aircraft pollution around airports is mainly caused by emissions during the different landing and take-off (LTO) phases. These are defined as all aircraft operations below 3000 ft and have different characteristics. During taxiing and idling the low thrust results in incomplete combustion, which leads to large amounts of CO and hydrocarbons. However, the high thrust during take-off and climb-out causes higher combustion temperatures, which result in a larger production of NO_x and SO_x [Pagoni and Psaraki, 2014]. The combustion of kerosene also leads to CO_2 emissions, with a relation of 3.149 kg CO_2 per kg kerosene [Carlier et al., 2006]. The LTO-emissions are generally calculated following the standard ICAO LTO-cycle, which provides a set of times and thrust setting per phase [ICAO, 2016a]. This cycle is also implemented in the LTO-emissions calculator from the European Environment Agency [EEA, 2019].

Most research into full flight emissions is limited to carbon emissions, as this is directly related to the amount of fuel that is burned. The effect of CO_2 on climate effects is also better understood than other emissions, such as nitrogen oxides and water vapour. The carbon emissions are dependent on multiple factors, such as the flight distance (discussed by [Jardine, 2005]), load factor and aircraft age [Brueckner and Abreu, 2017] and air traffic management influences [Miyoshi and Mason, 2009]. In addition, the type of network that is operated has an influence on the emissions of an airline. It was found in [Loo et al., 2014] that hub-and-spoke models decrease the CO_2 emissions per passenger km. However, they have a negative effect on the local air quality around the hub airports. Different emissions models are available, such as the Advanced Emissions Model[1] developed by Eurocontrol, which uses the Base of Aircraft Data (BADA) aircraft performance model. Publicly available models are Piano-X, the master flight emissions calculator of EMEP/EEA [EEA, 2019] and OpenAP, which is an aircraft performance model developed at the TU Delft [Sun et al., 2020] and is the model used in this paper.

Finally, three studies are available that combine aircraft emission modelling to air cargo modelling. The study of [Derigs and Illing, 2013] investigates how the network of a global airline changes due to the introduction of EU-ETS, showing a shift towards areas that do not fall under these European regulations. For the calculation of the emissions, a fixed emission rate per flight km is used, with added emissions if more payload is on board. [Chao, 2014] analyses how different freighter aircraft types are impacted by four ETS scenarios,

[1] https://www.eurocontrol.int/model/advancedemissionmodel, accessed on 19-04-2021

assuming a fixed fuel consumption throughout the whole flight. The third study does not consider the effects of emission trading systems, but proposes a multi-objective optimisation model to minimise the costs and carbon emissions of an air cargo alliance network [Yan et al., 2020]. Again fuel consumption is assumed to be constant during the cruise flight, with the LTO-emissions calculated following the ICAO LTO-cycle. The model is allowed to outsource operations to airlines in the alliance, which can results in lower emissions for the airline itself.

This study focuses on solving the flight scheduling, aircraft routing and cargo routing problems in an single model, with the aircraft emissions integrated in the objective function. For this, a time-space network is used, which also allows for the development of a flight schedule based on a given cargo demand. The emissions are calculated based on the amount of payload on board, which is used to analyse what schedule revisions are necessary to reduce the emissions.

3 Emission model

The emission model acts as a precursor to the flight scheduling model and generates an emission matrix that is used during the flight schedule design. The emission matrix holds emission data of all possible flights in the network for a range of load factors. During the flight schedule optimisation, emissions of a large number of flights with a variety of load factors have to be calculated. In order to speed up that process, the data-points of the emission matrix are used to approximate the emissions of each specific flight. This provides a good trade-off between having a continuous-like approach (more realistic, but computationally very difficult), and current approaches in the literature where emissions do not depend on the load factor at all. For the calculation of the emissions for the emission matrix, the emission model divides each flight into two sections: The combined taxi, landing and take-off phase (LTO-cycle) and the cruise or en-route phase. The CO_2 calculation for both phases is described in Section 3.1. The quality of the results is validated in Section 3.2 by comparing it to existing emission models.

3.1 Carbon Emissions Calculation

The CO_2 emissions are directly related to the combustion of kerosene, which means that the fuel flow can be used to find the amount of CO_2 emitted in the flight. The fuel used is multiplied to the emission coefficient (3.149 kg CO_2 per 1 kg kerosene [Carlier et al., 2006]). The emissions of the LTO phase and the cruise part of the flight are calculated separately and are described below.

3.1.1 LTO-emissions

The main reason that the emissions of the LTO phase are calculated separately is due to the standardised ICAO LTO-cycle [ICAO, 2016a]. This allows for a more reliable calculation of the emissions for this part of the flight. The cycle describes all operations that take place below an altitude of 3000 ft and gives standard times and thrust settings for each part of the LTO-cycle. Emission and fuel flow data of all commercial aircraft engines is available for the four thrust settings shown in Table 1 in the ICAO emission databank[2]. These values are independent of aircraft weight and load factor. The taxi times for each airport are retrieved from Eurocontrol[3]. If no data is available for the airport, the ICAO standard values of 7 and 19 minutes for taxi-in and taxi-out are used. The Open Aircraft Performance Model OpenAP is used for the calculation of the fuel flow and emissions for each of the LTO phases. This open-source performance model also uses values from the ICAO engine emission databank and allows for inputs of different thrust settings, altitudes and engine types [Sun, 2019]. The fuel usage and emissions of the LTO phases are found for each origin and destination airport combination, dependent on the local taxi times and altitude of the airport. The standard engine types are used for the two aircraft types in the fleet, which can be found in Table 2.

Table 1: Reference times and thrust settings for ICAO LTO-cycle [ICAO, 2016a]

Operating phase	Time [min]	Thrust setting [%]
Take-off	0.7	100
Climb	2.2	85
Approach & landing	4.0	30
Taxi & ground idle	Airport dependent 7.0 (in) & 19.0 (out)	7

[2]https://www.easa.europa.eu/domains/environment/icao-aircraft-engine-emissions-databank, accessed on 08-11-2021
[3]https://www.eurocontrol.int/publication/taxi-times-summer-2019, accessed on 12-04-2021

3.1.2 En-route emissions

The cruise part of the flight is less strictly defined than the LTO-cycle. First a flight trajectory is generated, where the flight length is assumed to be equal to the great-circle distance between the origin and destination. This is the shortest path between two locations on a sphere. A vertical trajectory is generated using the *TrajGen* function of OpenAP, which describes the altitude, airspeed and vertical velocity during the flight. The function creates the trajectories using a kinematic model which is based on statistical ADS-B data [Sun et al., 2020]. The trajectory is generated with a time step of 10 seconds. The default settings of the B747-400 are used, which represent the most occurring values of parameters such as calibrated airspeed during climb and descent, cruise Mach number and cruise altitude. The cruise range parameter is set equal to the great-circle distance of each airport pair, from which an average distance needed for climb and descent is subtracted. For flights shorter than 600 km, a lower cruise altitude of 20,000 ft is assumed. External influences such as wind, flight instructions from air traffic management and fuel saving procedures like a step climb during cruise are neglected. These factors are considered to be outside the scope of this study. Plots of the altitude, airspeed and vertical velocity for a flight from Hartsfield-Jackson Atlanta International airport (ATL) to Chicago O'Hare International airport (ORD) are shown in Figure 2. In order to prevent counting the LTO-emissions twice, the trajectory data-points below 3000 ft are not used in the emission calculation.

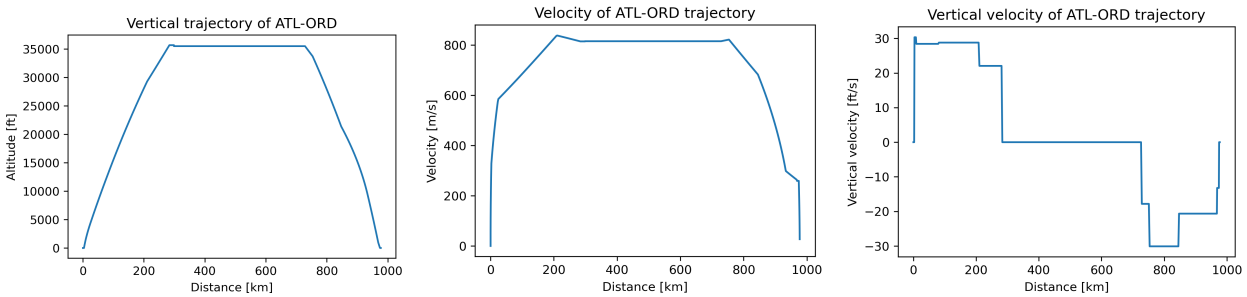

Figure 2: Output plots of altitude, velocity and vertical velocity for a generated trajectory between Atlanta (ATL) and Chicago 'O Hare airport (ORD)

The fuel use and emissions for each flight leg are calculated for 11 evenly spaced load factor values. The number of data-points can be varied depending on the desired accuracy, with a minimum of two load factors that are necessary for linear interpolation. The maximum load factor is based on the payload-range diagram of each aircraft. This diagram relates the payload capacity to the mission range, graphically shown in Figure 3. For a hypothetical range of 0 km, indicated in point A, the full payload capacity is available. This capacity is limited by the difference between the maximum zero fuel weight (MZFW) and the operational empty weight (OEW). For the ranges between the points A and B, fuel can be added without needing to reduce the payload weight. For longer ranges than $R_{max\ cap}$, a trade-off has to be made between payload capacity and fuel weight, to stay below the maximum take-off weight (MTOW). This decrease in payload capacity continues until point C, which represents the point where the maximum fuel limit is reached. From this point onward, no extra fuel can be added which decreases the payload capacity even quicker. Point D indicates the maximum mission range with no payload on board. The payload capacities and ranges indicated on the diagram are shown in Table 2 for the two most used aircraft by Cargolux.

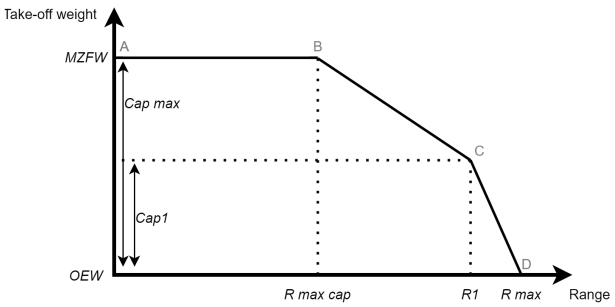

Figure 3: Payload-range diagram

The amount of fuel that is necessary for the flight can be determined for each load factor. The fuel use is among other dependent on the mass of the aircraft, and thus also on the amount of fuel that is loaded into the aircraft before the flight. Therefore, an estimation of the fuel needed for the flight is made ($m_{Fuel\ est}$), with an assumed

Table 2: Aircraft input values [Sun et al., 2020] [Boeing, 2012] [Boeing, 2002]

Aircraft type	Engine type	$MTOW$ [kg]	OEW [kg]	Cap_{max} [kg]	Cap_1 [kg]	$R_{max\ cap}$ [km]	R_1 [km]	R_{max} [km]
B747-8F	GEnx-2B67	447,700	197,000	134,000	70,000	7,778	13,890	16,112
B747-400F	RB211-524G	412,770	165,000	113,000	60,000	7,963	13,334	15,186

constant fuel flow of 3 kg/s. This gives the model an approximation of the take-off weight for this flight. For every time step, the fuel flow is calculated using the OpenAP *FuelFlow* function [Sun et al., 2020]. OpenAP uses a kinematic model to find the thrust needed to fly at the altitude, velocity and flight path angle provided by the generated trajectory. The aircraft mass for each time step is updated by subtracting the fuel used during that period. When the complete flight has been simulated, the sum of the fuel flow (m_{Fuel}) can be compared to the estimated fuel mass $m_{Fuel\ est}$. If the difference is larger than 2%, the fuel mass estimation is adjusted and the simulation is repeated. If the gap is smaller than 2%, this take-off mass and fuel flow are used in the en-route emissions calculation.

3.1.3 Full flight emissions

The integrated value of the en-route fuel flow is multiplied to the emission coefficient of CO_2 and added to the emissions of the LTO phase to find the total amount of CO_2 that is emitted on each flight. These values are stored in the emission matrix for all routes in the network. This is done for each available (aircraft type, load factor) combination. The matrix is constructed before the flight schedule optimisation to limit the computational time. An example of this output is plotted in Figure 4 for a flight between ATL and Luxembourg airport (LUX). The result is a slightly non-linear relation, with the emissions increasing faster for higher load factors. In Section 4.2 two methods are described that use these 11 known data-points to approximate the emissions for a specific load factor during the flight schedule optimisation.

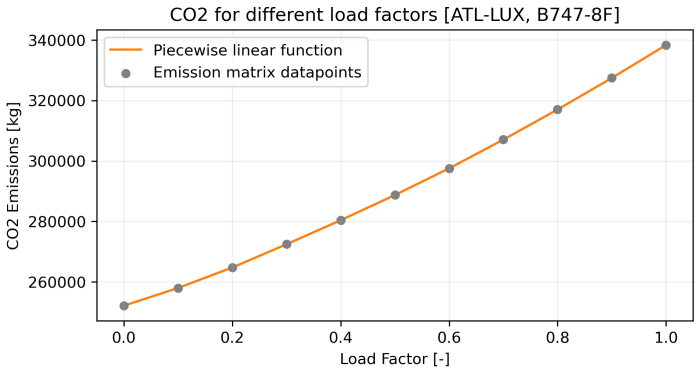

Figure 4: Full flight CO_2 emissions from the emission matrix for 11 load factor values, connected by a piecewise linear function.

3.2 Comparison with other emission models

Aircraft emissions for this study were computed using the OpenAP database, due to its flexibility and easiness of use and integration in the optimisation model. Because OpenAP is a recent addition to the existing literature on aircraft emission, the results obtained with OpenAP are validated using two other emission models. The EMEP/EEA model [EEA, 2019] was used to validate both LTO and full flight emissions, while Piano-X[4] was only used for the full flight emissions. The emissions are validated for an Airbus A380-800 equipped with GP7270 engines, as no Boeing 747 models are available in the Piano-X database. Although the A380 is not used in the remainder of this study, it displays much similarities to the both Boeing 747s, being a wide-body aircraft with four engines. Therefore, it gives a good indication on the accuracy of the results of the developed emission model.

[4]https://www.lissys.uk/PianoX.html, accessed on 08-11-2021

3.2.1 EMEP/EEA

This model is developed by the European Environment Agency and provides a separate LTO-emissions calculator[5] and a master emissions calculator[6] that can be used for emission calculations for a whole flight. EEA estimates the uncertainty of its model to be 5 to 10% for the LTO-emissions and 15-40% for the cruise phase[EEA, 2019]. The LTO-emissions are compared for a complete LTO-cycle at LUX. An average taxi-out time of 624 seconds and taxi-in time of 258 seconds are obtained from the Eurocontrol data. The comparison between the developed model and the EEA LTO-emissions calculator is shown in Table 3. The overall LTO emissions are 300 kg lower when using the OpenAP model (9,480 vs 9,772 kg, which corresponds to a relative difference of 3%). OpenAP finds lower values for the two taxi phases, while the emissions of the other three phases are slightly overestimated compared to the EEA calculator. A reason for these discrepancies might be that OpenAP approximates the relation between the thrust setting and fuel flow with a 3rd-degree polynomial [Sun, 2019]. This function is based on the values from the ICAO emissions databank[7], which only provides the fuel flow for the four thrust settings of the LTO-cycle. OpenAP uses this polynomial to also have the possibility to calculate the en-route fuel flows for different thrust settings. The EMEP/EEA approach uses the exact values from the ICAO emissions databank, which results in a difference between the two models. Furthermore, the altitude of the airport has been incorporated in the developed model, where the EEA model calculates all LTO emissions at sea level. The thinner air at higher altitude means that a larger fuel flow is necessary for a similar thrust force. Although LUX only lies at 350 meters altitude, assuming the airport at sea level would already result in a CO_2 reduction of around 250 kg.

Table 3: Validation of the LTO emissions of an A380-800 at Luxembourg Findel Airport (LUX). Comparison between the developed model, which uses OpenAP, and the LTO emissions calculator of EMEP/EEA

Operating phase	OpenAP [kg CO_2]	EMEP/EEA [kg CO_2]
Taxi-out	1,448	1,843
Take-off	1,440	1,396
Climb	3,697	3,607
Approach & landing	2,294	2,150
Taxi-in	599	775
LTO-cycle total	9,480	9,772

For the full flight emissions a short haul and a long haul flight are investigated, namely a flight from Glasgow Prestwick (PIK) to LUX (978 km) and a flight from ATL to LUX (7274 km). The EMEP/EEA Master emissions calculator does not allow for changes in the amount of payload, therefore only one value per flight is generated and shown in Table 4. The emissions for the PIK-LUX flight are just higher than the emissions for the flight with $LF = 1$ using OpenAP. However, the amount of CO_2 for the longer ATL-LUX flight that is found using the EMEP/EEA approach lies closer to the value calculated for $LF = 0.5$. The results are difficult to compare as no exact information is available on the take-off weight that is used for these flights. However, this gives a confirmation that the results are of a similar order of magnitude.

3.2.2 Piano-X

Second, the Piano-X aircraft emissions and performance model is considered. Only the full flight emissions are compared, as its definition of the LTO phases differs from to the ICAO LTO-cycle used in this study. The load factor is set at 0, 0.5 and 1 to show how the developed model relates to the results of Piano-X with different amounts of payload on board. Table 4 shows a comparison between the three emission models. The results for the short flight display an excellent agreement, with the largest error being just under 5% for the empty flight ($LF = 0$). The Piano-X model seems to be more sensitive to the changes in load factor, as the difference between flying with minimum and maximum payload being larger than for the OpenAP model. When considering the longer route, the difference increases as the OpenAP model expects larger emissions compared to Piano-X. A possible explanation is the assumption of a constant cruise altitude and speed that is used in the OpenAP model. The flight trajectory in Piano-X is optimised for the aircraft weight and also a step climb during cruise is implemented to save energy. Again, the impact of the load factor is slightly larger for the Piano-X model.

[5] https://www.eea.europa.eu/publications/emep-eea-guidebook-2019/part-b-sectoral-guidance-chapters/1-energy/1-a-combustion/1-a-3-a-aviation-1-annex5-LTO/view, accessed on 08-11-2021

[6] https://www.eea.europa.eu/publications/emep-eea-guidebook-2019/part-b-sectoral-guidance-chapters/1-energy/1-a-combustion/1-a-3-a-aviation-1/view, accessed on 08-11-2021

[7] https://www.easa.europa.eu/domains/environment/icao-aircraft-engine-emissions-databank, accessed on 08-11-2021

Table 4: Validation of the full flight emissions of an A380-800 for a short (PIK-LUX) and a long haul flight (ATL-LUX). Comparison between the OpenAP model, Piano-X and the EMEP/EEA Master emissions calculator.

Route	LF [-]	OpenAP [kg CO_2]	Piano-X [kg CO_2]	EMEP/EEA [kg CO_2]
PIK-LUX	0	48,155	45,998	
	0.5	51,707	52,547	
	1	58,124	59,884	60,468
ATL-LUX	0	281,980	249,585	
	0.5	322,325	298,878	
	1	377,931	357,329	339,803

4 Flight scheduling model

This section describes the flight scheduling model, which aims to find the most profitable flights and assigns an aircraft to each route. First, a time-space network is formulated in Section 4.1, which is used to route all aircraft and cargo requests between the different airports in the network. Section 4.2 describes the mathematical formulation of the MILP, together with two variants that are used to incorporate aircraft emissions (and hence, environmental sustainability) into the model.

4.1 TSN formulation

The routing of both the aircraft and cargo requests is carried out in a TSN that is inspired by [Delgado et al., 2020]. A TSN is a particular type of network that simultaneously maps events occurring in space and time. Each node in the network is uniquely defined by a tuple with a specific position (e.g. an airport) and a time-stamp. In a TSN, time is only defined as multiples of the chosen time step Δt. The overall planning horizon is defined between an initial time 0 and a final time T_{hor}, with a set of time-stamps $\mathcal{T}_{hor} = \{0, \Delta t, 2\Delta t, \cdots, T_{hor}\}$. The spatial component is represented by the set of airports $l \in \mathcal{L}$ that the aircraft can fly to. Hence, the aircraft network consists of a set of itinerary nodes ($\mathcal{N}^I$) for each unique combination of airport and time-stamp. These nodes are connected to each other by a set of ground arcs ($\mathcal{A}^G$) and flight arcs ($\mathcal{A}^F$). The ground arcs connect different nodes associated to the same airport, indicating that the aircraft is stationary on the ground. Flight arcs connect a node i to node j, where the airports associated to nodes i and j are different. The flight time is defined as $t_{ij} = distance_{ij}/V_{cruise} + t_{LTO} + t_{TAT}$. The great circle distance is divided by an assumed constant cruise speed of 900 km/h, with 30 minutes added for the LTO phase. A turnaround time (TAT) of 1 hour is already added to the flight time, which means that the aircraft can directly start the next flight using the same itinerary node. In the likely scenario that the flight time of a flight arc does not match a multiple of Δt, the arrival time is rounded up to the time-stamp of the next available node, hence overestimating the flight time. The combined set of ground arcs ($\mathcal{A}^G$) and flight arcs ($\mathcal{A}^F$) are referred to as itinerary arcs $\mathcal{A}^I = \mathcal{A}^G \cup \mathcal{A}^F$. Each aircraft k in the fleet $\mathcal{K}$ is assigned an origin and final airport, where it must be located at the start and end of the time horizon, respectively.

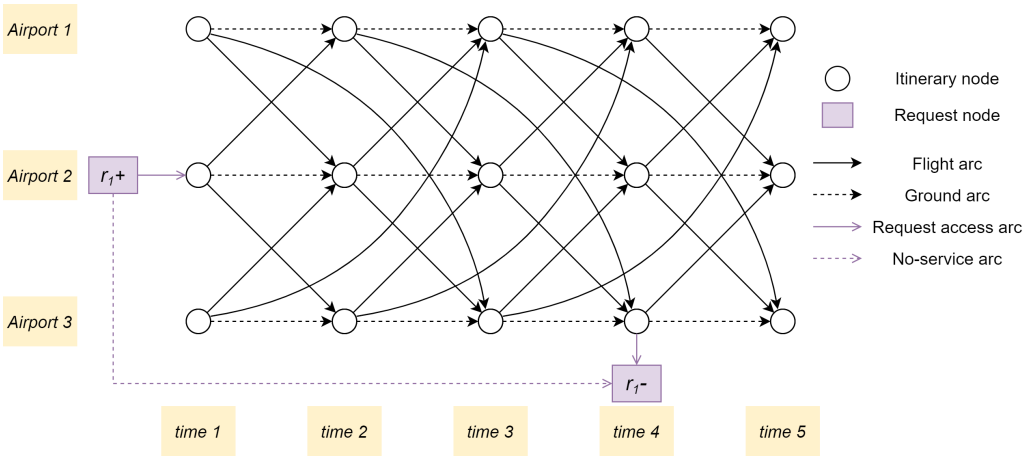

Figure 5: Example of a time-space network consisting of node set $\mathcal{N}$ and arc set $\mathcal{A}$

Each cargo request $r \in R$ is routed over the same TSN where the aircraft can move. Next to the itinerary nodes, two special request nodes are available for each cargo request: Origin node i_r^+ and destination node i_r^-. These request nodes represent, respectively, the (origin airport, release time) and (destination airport, due

time) tuples for request r. The request nodes are collected in the set $\mathcal{N}^R$ and are grouped with the itinerary nodes in the set $\mathcal{N} = \mathcal{N}^I \cup \mathcal{N}^R$. The request nodes are connected to the network of itinerary arcs and nodes by a set of request access/egress arcs: $\mathcal{A}^R = \mathcal{A}^{R+} \cup \mathcal{A}^{R-}$. Here $\mathcal{A}^{R+}$ define the access and $\mathcal{A}^{R-}$ define the egress arcs. Section 5.2 further elaborates on the generation of the cargo requests. Finally, each request node pair (i_r^+, i_r^-) is directly connected to each other via a no-service arc. The set of all no-service arcs is $\mathcal{A}^N$. A no-service arc is used to model the situation where a request r is not physically transported through the TSN, and thus represents a missed delivery, resulting in no revenue that is earned from this request. The whole set of arcs is $\mathcal{A} = \mathcal{A}^G \cup \mathcal{A}^F \cup \mathcal{A}^R \cup \mathcal{A}^N$. Cargo requests can use all arcs in $\mathcal{A}$ and nodes in $\mathcal{N}$, while aircraft can only use arcs $\mathcal{A}^I$ and nodes $\mathcal{N}^I$. An example of a TSN with a single cargo request r_1 is presented in Figure 5. The complete definitions of all sets used in the TSN are given in Table 5.

Table 5: Definitions of all sets used in the TSN.

Notation	Description	Definition
$\mathcal{K}$	Aircraft fleet	$\{k : k \in \mathcal{K}\}$
$\mathcal{R}$	Cargo requests	$\{r : r \in \mathcal{R}\}$
$\mathcal{L}$	Airports	$\{l : l \in \mathcal{L}\}$
$\mathcal{T}_{hor}$	Time steps	$\{t : t \in \mathcal{T}_{hor}\}$
$\mathcal{N}^I$	Itinerary nodes	$\{(l,t) : l \in \mathcal{L}, t \in \mathcal{T}_{hor}\}$
$\mathcal{N}^R$	Request nodes	$\{(l_r^+, t_r^+) : r \in \mathcal{R} \cup (l_r^-, t_r^-) : r \in \mathcal{R}\}$
$\mathcal{N}$	Nodes	$\{\mathcal{N}^I \cup \mathcal{N}^R\}$
$\mathcal{A}^F$	Flight arcs	$\{(i,j) : i,j \in \mathcal{N}^I \wedge l_i \neq l_j \wedge t_i < t_j\}$
$\mathcal{A}^G$	Ground arcs	$\{(i,j) : i,j \in \mathcal{N}^I \wedge l_i = l_j \wedge t_i + T = t_j\}$
$\mathcal{A}^R$	Request access/egress arcs	$\{(i_r^+, i) : r \in \mathcal{R} \wedge i \in \mathcal{N}^I \wedge l_i = l_r^+ \wedge t_r^+ \leq t_i \leq t_r^-\}$ $\cup \{(i, i_r^-) : r \in \mathcal{R} \wedge i \in \mathcal{N}^I \wedge l_i = l_r^- \wedge t_r^+ \leq t_i \leq t_r^-\}$
$\mathcal{A}^N$	No-service arcs	$\{(i_r^+, i_r^-) : r \in \mathcal{R}\}$
$\mathcal{A}^I$	Itinerary arcs	$\{\mathcal{A}^F \cup \mathcal{A}^G\}$
$\mathcal{A}$	Arcs	$\{\mathcal{A}^F \cup \mathcal{A}^G \cup \mathcal{A}^R \cup \mathcal{A}^N\}$

4.2 Mathematical formulation

The time-space networks that are discussed in Section 4.1 display all possible options that the model can choose from in the design of the flight schedule. In order to indicate which arcs are chosen for each aircraft and cargo request, two main decision variables are used: x_{ij}^k and q_{ij}^r. Both are binary decision variables, and indicate if an aircraft or cargo request is routed over arc (i,j), where i is the origin node and j is the destination node of the arc. A formal definition of the aircraft decision variable is X: $\{x_{ij}^k : k \in \mathcal{K}, (i,j) \in \mathcal{A}^I\}$, indicating that the aircraft can only be routed over itinerary arcs. Q gives the set of cargo routing decision variables: $\{q_{ij}^r : r \in R, (i,j) \in \mathcal{A}\}$.

Each aircraft k in the fleet $\mathcal{K}$ is characterised by a maximum payload capacity Cap_{ij}^k, which depends on the flight arc (i,j) and the aircraft type, due to payload-range restrictions. All aircraft need to comply with the maximum block time $T_{block\,max}$, with t_{ij} being the flight time between the airports associated with nodes i and j. Finally, $OperCost_{ij}^k$ characterises the operational costs for each aircraft. Note that $OperCost_{ij}^k$ is listed as a parameter, while it directly depends on decision variables X and Q and should formally be considered as such. More details are given in Section 4.2.1.

For each request $r \in \mathcal{R}$, a revenue Rev_{ij}^r is defined. Each request can be routed on a maximum number of flight legs equal to $n_{flights}$ to limit the amount of intermediate stops. Although this parameter is constant in this paper, it can be easily translated into a request-specific parameter to capture different cargo characteristics.

Next to the routing decision variables, the model uses another set of decision variables to map the CO_2 emissions produced by each aircraft $k \in \mathcal{K}$ when flying flight arc (i,j). This is given by the set E: $\{e_{ij}^k : k \in \mathcal{K}, (i,j) \in \mathcal{A}^F\}$. When assessing the environmental impact of the emissions, a fixed cost per unit weight of CO_2 is used, which is given by parameter C_{CO_2}.

In Table 6 and 7 the decision variables and parameters of the MILP are summarised.

Note that some of the described parameters are obtained as a function of other parameters that have not been presented yet for ease of notation. These will be thoroughly described in Section 4.2.1 and 4.2.2.

Table 6: Definitions of all decision variables used in the MILP.

Notation	Description
x_{ij}^k	Binary. Unitary if aircraft k traverses arc $(i,j) \in \mathcal{A}^I$
q_{ij}^r	Binary. Unitary if cargo request r traverses arc $(i,j) \in \mathcal{A}$
e_{ij}^k	Continuous. CO_2 emissions of aircraft k when traversing arc $(i,j) \in \mathcal{A}^I$

Table 7: Definitions of all parameters used in the MILP.

Notation	Description
$OperCost_{ij}^k$	Operational cost of aircraft k to fly arc (i,j)
Rev_{ij}^r	Revenue for request r for request access arc (i,j)
Cap_{ij}^k	Payload capacity of aircraft k on arc (i,j)
T_{block}	Maximum daily block time (i,j)
t_{ij}	Flight time for flight arc
n_{flight}	Maximum number of flight arcs that a cargo request can be routed over
C_{CO_2}	Cost per unit weight of CO_2 emissions

The flight schedule design and aircraft routing problem is formulated as a MILP model that aims to maximise the objective value, which is defined as follows:

$$\max \sum_{r\in\mathcal{R}} \sum_{(i,j)\in\mathcal{A}^{R+}} Rev_{ij}^r \cdot q_{ij}^r - \sum_{k\in\mathcal{K}} \sum_{(i,j)\in\mathcal{A}^F} OperCost_{ij}^k - \sum_{k\in\mathcal{K}} \sum_{(i,j)\in\mathcal{A}^F} e_{ij}^k \cdot C_{CO_2}, \tag{1}$$

subject to the following constraints:

$$\sum_{r\in\mathcal{R}} w_r q_{ij}^r \leq \sum_{k\in\mathcal{K}} Cap_{ij}^k x_{ij}^k, \qquad \forall (i,j) \in \mathcal{A}^F \tag{2}$$

$$\sum_{(i,j)\in\mathcal{A}^{I+}(i)} x_{ij}^k - \sum_{(j,i)\in\mathcal{A}^{I-}(i)} x_{ji}^k = \begin{cases} 1, & i = i_k^+ \\ -1, & i = i_k^- \\ 0, & otherwise \end{cases}, \qquad \forall i \in \mathcal{N}^I, k \in \mathcal{K} \tag{3}$$

$$\sum_{(i,j)\in\mathcal{A}^F} x_{ij}^k t_{ij} \leq T_{block\ max}, \qquad \forall k \in \mathcal{K} \tag{4}$$

$$\sum_{k\in\mathcal{K}} x_{ij}^k \leq 1, \qquad \forall (i,j) \in \mathcal{A}^F \tag{5}$$

$$\sum_{(i,j)\in\mathcal{A}^{+}(i)} q_{ij}^r - \sum_{(j,i)\in\mathcal{A}^{-}(i)} q_{ji}^r = \begin{cases} 1, & i = i_r^+ \\ -1, & i = i_r^- \\ 0, & otherwise \end{cases}, \qquad \forall i \in \mathcal{N}, r \in \mathcal{R} \tag{6}$$

$$\sum_{(i,j)\in\mathcal{A}^F} q_{ij}^r \leq n_{flights}, \qquad \forall r \in \mathcal{R} \tag{7}$$

$$q_{ij}^r \leq \sum_{k\in\mathcal{K}} x_{ij}^k, \qquad \forall (i,j) \in \mathcal{A}^F, \forall r \in \mathcal{R} \tag{8}$$

$$\sum_{(i,j)\in\mathcal{A}^I:\{t_i\leq t<t_j\}} x_{ij}^k = 1, \qquad \forall t \in \mathcal{T}_{hor}, k \in \mathcal{K} \tag{9}$$

$$\sum_{(i,j)\in\mathcal{A}:\{t_i\leq t<t_j\}} q_{ij}^r = 1, \qquad \forall t \in \{t_r^+ \leq t < t_r^-\}, r \in \mathcal{R} \tag{10}$$

$$x_{ij}^k \in \{0,1\}, \qquad \forall (i,j) \in \mathcal{A}^I, k \in \mathcal{K} \tag{11}$$

$$q_{ij}^r \in \{0,1\}, \qquad \forall (i,j) \in \mathcal{A}, r \in \mathcal{R} \tag{12}$$

$$e_{ij}^k \geq 0, \qquad \forall (i,j) \in \mathcal{A}^F, k \in \mathcal{K} \tag{13}$$

The objective function (Equation 1) consists of three terms that are used to determine the profit. The first term maps the revenue that is earned from transporting requests. The operational costs that are required are

mapped in the second term, and further explained in Section 4.2.1. The third term maps the costs due to CO_2 emissions, which is elaborated upon in Section 4.2.2. These terms are combined to find the most profitable flight schedule, while also taking account the corresponding aircraft emissions.

To find a feasible schedule, first a set of constraints is created which take care of the aircraft and cargo routing part. Constraints 2 assure that the payload per flight is smaller than or equal to the payload capacity of the aircraft operating that flight leg. Conservation of aircraft flow in the TSN is ensured by Constraints 3. Here $\mathcal{A}^{I+}(i)$ and $\mathcal{A}^{I-}(i)$ represent the set of all out- and ingoing arcs of node i. Subtracting all active incoming arcs from the outgoing arcs of a node should always result in zero, except for the starting and ending nodes of the aircraft (i_k^+ and i_k^-). Furthermore, Constraints 4 limit the operational time of each aircraft to the maximum allowed block time $T_{block\,max}$, which is set to 16 hours per day for this study. In addition, Constraints 5 enforce that at most one aircraft can operate each flight arc. Similarly to the aircraft flow conservation constraints, also cargo flow conservation is added by Constraints 6. This forces each cargo request r from its release source node i_r^+ to the sink node i_r^-. Constraints 7 limit the maximum number of flights of each cargo request to $n_{flights}$, which is set to 3 in this work.

On top of the routing and operational constraints mentioned above, an additional set of constraints is added to tighten the formulation. Cargo requests can only traverse flight arcs that are operated by an aircraft, which is checked using Constraints 8. Constraints 9 and 10 ensure that each aircraft and each cargo request is allocated to exactly one location at each time-stamp. For the aircraft this is done for the whole simulation horizon $\mathcal{T}_{hor}$, while the cargo only is checked for the period in which it is active, which is between t_r^+ and t_r^-. Finally, Constraints 11 and 12 define the binary nature of the x and q decision variables, while the emission variables e are continuous (Constraints 13).

4.2.1 Revenue and Operational cost

The first two terms of the objective function (Equation 1) are more elaborately described next.

$$Rev_{ij}^r = w_r p_r s_r \tag{14}$$

$$OperCost_{ij}^k = C_F t_{ij} x_{ij}^k + C_{fuel} w_{fuel,ij}^k + C_{handling} w_{take-off,ij}^k \tag{15}$$

$$w_{fuel,ij}^k \geq \left(fuel_0 + \frac{\sum_{r\in\mathcal{R}} w_r q_{ij}^r}{Cap_k} \cdot \frac{fuel_{max} - fuel_0}{LF_{max}} \right) - (1 - x_{ij}^k) \cdot M, \qquad \forall (i,j) \in \mathcal{A}^F, k \in \mathcal{K} \tag{16}$$

$$w_{take-off,ij}^k \geq \left(OEW + w_{fuel,ij}^k + \sum_{r\in\mathcal{R}} w_r q_{ij}^r \right) - (1 - x_{ij}^k) \cdot M, \qquad \forall (i,j) \in \mathcal{A}^F, k \in \mathcal{K} \tag{17}$$

$$w_{fuel,ij}^k \geq 0, \qquad \forall (i,j) \in \mathcal{A}^F, k \in \mathcal{K} \tag{18}$$

$$w_{take-off,ij}^k \geq 0, \qquad \forall (i,j) \in \mathcal{A}^F, k \in \mathcal{K} \tag{19}$$

The revenue for each request can be determined by multiplying the weight of the cargo request w_r by the freight price per kg p_r and a strategic weighting factor s_r (Equation 14). This strategic weighting factor can be used to vary the importance of specific requests. The revenue is calculated for each request access arc in the set $\mathcal{A}^{R+}$. If the request is not served, it is routed directly over a no-service arc and will therefore not yield any revenue. In the objective function, the revenues are multiplied to the corresponding cargo routing variable q and summed to find the total revenue.

The costs that are required to operate the aircraft are split up in three terms. First, the flight time is multiplied by a fixed cost term C_F, which includes maintenance cost, crew wages, asset depreciation costs, take-off and landing charges and other general operating costs. The hourly value for these fixed cost is retrieved from [van der Meulen et al., 2020] and stated with all other cost parameters in Table 8. Second, the fuel costs are considered. The fuel weight needed for each flight is determined using the emission model described before. A linear function between the fuel weight for the flight at zero and maximum payload ($fuel_0$ and $fuel_{max}$) is assumed. Constraints 16 show the linear interpolation between these two values to find the fuel weight for the load factor of the flight. This load factor is defined as the sum of all request weights on board of the flight divided by the payload capacity of aircraft k. A fuel price of 75 euro per barrel is used[8], which can be converted to 600 euro per tonne. $w_{fuel,ij}^k$ is implemented as an additional continuous decision variable (Equation 18), mapping the amount of fuel necessary as a function of decision variables x_{ij}^k and q_{ij}^k. Third, a term is added for the airport handling cost. This term is dependent on the take-off weight, which consists of the operational

[8] https://www.iata.org/en/publications/economics/fuel-monitor/, accessed on 03-11-2021

empty weight, the fuel weight and the weight of all cargo on board (Constraints 17). Similar to the fuel weight, $w^k_{take-off,ij}$ is an additional continuous decision variable, defined in Equation 19. Note that Constraints 16 and 17 are implemented using a big-M formulation that activates the constraint only if the associated x^k_{ij} is equal to 1.

Table 8: Used values for the different cost parameters used in the objective function.

Cost parameter		Value	
Cargo price	p_r	2	[€/kg]
Fixed operating cost	C_F	5,375	[€/hr]
Fuel cost	C_{fuel}	600	[€/tonne]
Handling cost	$C_{handling}$	10	[€/tonne TOW]
CO_2 cost	C_{CO_2}	50	[€/tonne]

4.2.2 Aircraft emissions

The final term of the objective function is the sum of all emission costs, which are subtracted from the profit term. These are calculated by multiplying the CO_2 emissions by the cost to emit one tonne of CO_2. This cost is based on the current prices for an allowance for one tonne in the EU Emission Trading System (EU-ETS)[9]. Although the pricing of these allowances is dependent on more factors, this fixed price is assumed for the CO_2 costs. The carbon emissions are calculated using the emission model, discussed in Section 3, as part of the pre-processing phase. Based on these known values that are stored in the emission matrix, an approximation of the carbon emissions of the flights that are considered in the optimisation is carried out. Two approximation methods are implemented: (i) Linear interpolation and (ii) discretisation of a piecewise linear function.

Linear interpolation is the most straightforward way to approximate the flight emissions. In Figure 6 the datapoints from the emission matrix of a ATL-LUX flight are plotted. Within this figure a linear function is shown, which connects the emissions of an instance with a load factor of 0 and 1. The emissions of both cases are referred to as E_0 and E_{max} respectively. The emissions are dependent on the load factor, which is calculated using cargo routing decision variable q. Therefore the linear interpolation is implemented as a constraint, shown by Constraints 20. The load factor is found by dividing the weight of all cargo requests on flight arc (i,j) by the cargo capacity of aircraft k. The values for emissions at a load factor of 0 and maximum capacity LF_{max} (E_0 and E_{max} respectively) are retrieved from the emission matrix. e^k_{ij} contains the carbon emissions of this flight and acts as a continuous decision variable that can directly be used in the objective function (Equation 1). The constraint is only activated if $x^k_{ij} = 1$, indicating that the flight arc is used in the flight schedule.

$$e^k_{ij} \geq \left(E_0 + \frac{\sum_{r \in \mathcal{R}} w_r q^r_{ij}}{Cap_k} \cdot \frac{E_{max} - E_0}{LF_{max}} \right) - (1 - x^k_{ij}) \cdot M, \qquad \forall (i,j) \in \mathcal{A}^F, k \in \mathcal{K} \tag{20}$$

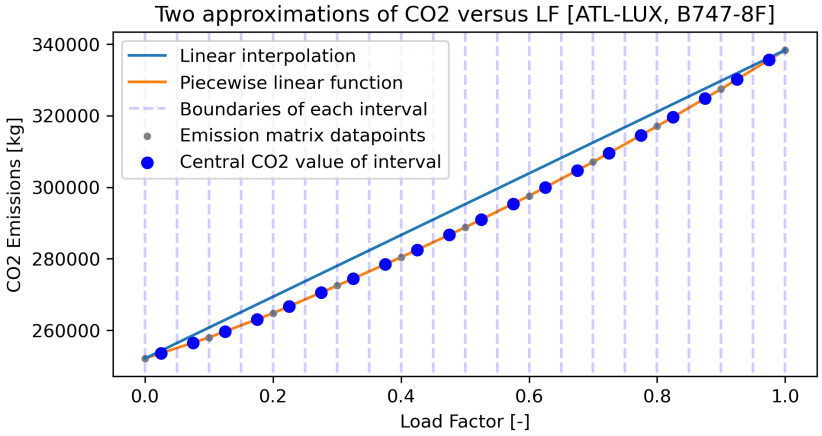

Figure 6: Graphical representation of the two methods used to approximate emissions based on load factor: Linear interpolation and discretisation based on a piecewise linear function.

[9]https://ember-climate.org/data/carbon-price-viewer/, accessed on 03-11-2021

A second method is proposed that can better capture the non-linear behaviours in the relationship between load factor and emissions. For this, the piecewise linear function is divided up into equally spaced segments, each with a lower and upper load factor bound. If the load factor of a certain flight falls between two of these bounds, the central value of that interval is chosen. An example with a discretisation of 20 sections ($n_{discr} = 20$) is shown in Figure 6. The load factor bounds are shown by the dotted vertical lines and the middle CO_2 values are depicted by the blue dots. An extra binary decision variable is created for each segment to indicate if it is used or not, to indicate if the (flight arc, load factor range) combination that the segment refers to is flown by aircraft k. The additional set of decision variables is defined by S: $\{s_{ij}^{k,n} : k \in \mathcal{K}, (i,j) \in \mathcal{A}^F, n \in n_{discr}\}$. The extra set of constraints that is necessary for the implementation of this second method is described below.

$$\sum_{n \in (0, n_{discr}-1)} s_{ij}^{k,n} = 1, \qquad \forall (i,j) \in \mathcal{A}^F, k \in \mathcal{K} \quad (21)$$

$$\sum_{n \in (0, n_{discr}-1)} \frac{n}{n_{discr}} s_{ij}^{k,n} \leq \frac{\sum_{r \in \mathcal{R}} w_r q_{ij}^r}{Cap_k} \leq \sum_{n \in (0, n_{discr}-1)} \frac{n+1}{n_{discr}} s_{ij}^{k,n}, \qquad \forall (i,j) \in \mathcal{A}^F, k \in \mathcal{K} \quad (22)$$

$$e_{ij}^k \geq \sum_{n \in (0, n_{discr}-1)} s_{ij}^{k,n} (co_2list_{ij}^{k,n}) - (1 - x_{ij}^k) \cdot M, \qquad \forall (i,j) \in \mathcal{A}^F, k \in \mathcal{K} \quad (23)$$

$$s_{ij}^{k,n} \in \{0,1\}, \qquad \forall n \in (0, n_{discr}-1), \quad (i,j) \in \mathcal{A}^F, k \in \mathcal{K} \quad (24)$$

Constraints 21 enforce that each unique (flight arc, aircraft) combination should be assigned a specific load factor segment. The load factor segment that is valid for this combination is computed using Constraints 22. The emissions associated to the middle-point load factor of the selected segment are computed with Constraints 23 which outputs the associated value of e_{ij}^k. Note that $co_2list_{ij}^{k,n}$ is a vector containing all CO_2 emission values of the middle-points of each segment of the discretisation. Again, only emissions of scheduled flights are added to the objective function by checking the associated aircraft routing variable x_{ij}^k using a big-M formulation. Finally, the binary nature of the decision variables is defined by Constraints 24.

5 Experimental setup

After describing the emission model and flight scheduling model, this section will explain how the experiments are set up. First, Section 5.1 shows the three networks that are investigated in the experiments. Second, the simulation of the cargo demand and generation of cargo requests in these three networks is discussed in Section 5.2. Details on the different experiments that are performed are given in Section 5.3.

5.1 Description of the Networks

The ideal scenario would be to simulate the complete network of a full-cargo airline like Cargolux, however due to computational time limitations this is not possible. Therefore, a part of the Europe and North-American part of the Cargolux network is split up in three smaller networks that are analysed separately. First, a European network is created (EU), which consists of 7 airports and is characterised by relatively short flights. In the second network (EU-NA), three European airports are connected to five American airports. This allows for analysis of the combination of transatlantic flights and shorter continental flights. Finally, the airports on western side of North-America are joined in a third network (NA), which also contains the two largest airports in Mexico and Calgary in Canada. These are all medium range flights apart from the airport pair in Mexico, which is a shorter flight of just under 500 km. Figure 7 presents a graphical representation of the three networks.

Each network can be simulated with a fleet of 1, 2 and 3 aircraft, indicated by AC0, AC1 and AC2. The aircraft type of each aircraft is stated in Table 9, together with the airport locations and IATA codes of each network. In addition, the initial and final location of each aircraft is given. These airports act as the source and sink nodes in the time-space network as described in Section 4.1.

5.2 Demand and request generation

The next step is to define the demand between each airport in the networks. Very limited data on cargo demand is publicly available, therefore the demand was synthetically simulated to model the operations of a cargo airline. For each origin-destination pair, the flight frequency is found, based on publicly available data from flight tracking platforms. The demand between each city is based on the assumption that the cargo on a specific

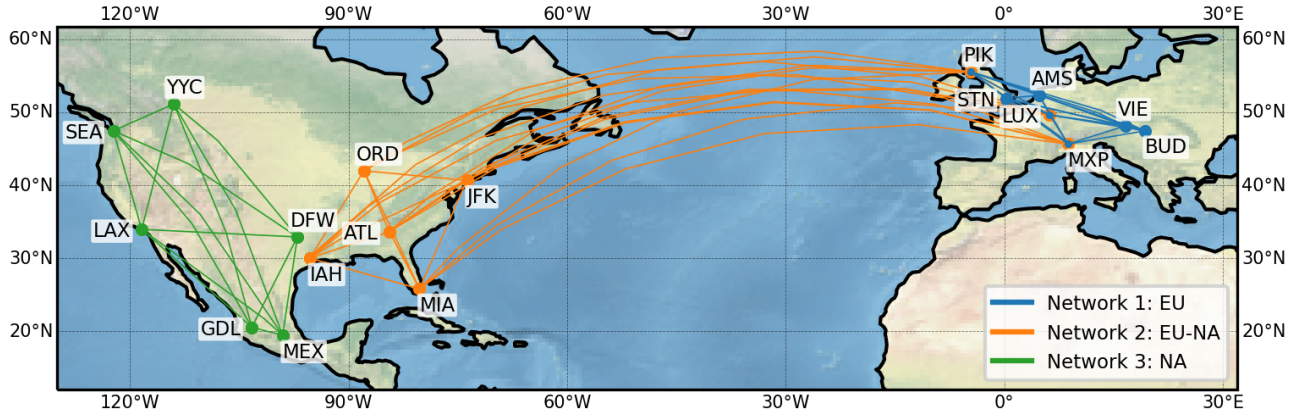

Figure 7: Graphical representation of the airports in the three networks that are investigated. The airports are connected by a line that shows the great-circle distance.

Table 9: Overview of the airports of the three networks that are investigated: EU, EU-NA and NA. The aircraft fleet is indicated with the initial and final airports per network.

	Airport	IATA-code		AC type	Initial	Final
EU	Luxembourg	LUX	AC0	B747-8F	LUX	LUX
	Amsterdam Schiphol	AMS	AC1	B747-8F	STN	PIK
	Glasgow Prestwick	PIK	AC2	B747-400F	MXP	LUX
	London Stansted	STN				
	Milan Malpensa	MXP				
	Vienna	VIE				
	Budapest Ferenc Liszt	BUD				
EU-NA	Luxembourg	LUX	AC0	B747-8F	LUX	ORD
	Glasgow Prestwick	PIK	AC1	B747-8F	ORD	LUX
	Milan Malpensa	MXP	AC2	B747-400F	MXP	ATL
	New York JFK	JFK				
	Miami	MIA				
	Chicago O'Hare	ORD				
	Hartsfield-Jackson Atlanta	ATL				
	George Bush Houston	IAH				
NA	Dallas Forth Worth	DFW	AC0	B747-8F	MEX	LAX
	Los Angeles	LAX	AC1	B747-8F	LAX	SEA
	Seattle-Tacoma	SEA	AC2	B747-400F	SEA	DFW
	Calgary	YYC				
	Guadalajara	GDL				
	Mexico-City	MEX				

flight is either aimed for its original destination (primary demand), one connection after the original destination (secondary demand) or two connections after the original destination (tertiary demand). The percentages for the primary, secondary and tertiary demand ($ratio_1, ratio_2, ratio_3$) can be manually adapted and sum up to 1.

The primary demand is determined for all airport pairs (i,j), using Equation 25, which is dependent on the number of flights between airports i and j ($freq_{ij}$) and on the ratio of the cargo that is assumed to be headed to airport j: $ratio_1$. These values are multiplied by the average load factor of Cargolux, which is reported to be 0.65 in 2019[10]. The average cargo capacity of an aircraft is based on a fleet of B747-400F and B747-8F aircraft. A similar procedure is followed to find the secondary demand, which is determined as the cargo that is first transported on flight (i,j) and then transferred to an available flight (j,g). This flight transports the cargo to its destination airport g. The cargo is equally divided over the different connections, based on the frequencies to each airport. Finally, the last part of the cargo from flight (i,j) is distributed in a similar manner over the available connections from airport g (Equation 27).

The departure times and dates are neglected in the determination of connections for the secondary and tertiary

[10]https://www.cargolux.com/media-room/media-releases/media-releases/Archives-2020/Cargolux-financial-results-for-2019, accessed on 09-08-2021

demand. It is assumed that the cargo can be transferred to any available connection from the airport that does not lead back to the origin airport of the cargo or an intermediate stop that the cargo has already visited. The demand ratios can be adjusted, given the characteristics of the considered airline. If most cargo is flown to its destination directly, $ratio_1$ can be increased, and if more connections are used in the network, $ratio_2$ and $ratio_3$ should be increased. A division of $\{ratio_1 = 50\%, ratio_2 = 30\%, ratio_3 = 20\%\}$ is used during this research.

$$\text{Primary demand } (ij) = freq_{ij} * ratio_1 * LF_{avg} * Cap_{avg} \qquad \forall i, j \in Airports \tag{25}$$

$$\text{Secondary demand } (ig) = freq_{ij} * ratio_2 * LF_{avg} * Cap_{avg} * \frac{freq_{jg}}{\sum conn_j}$$
$$\forall i, j, g \in Airports,\ g \neq i \tag{26}$$

$$\text{Tertiary demand } (ih) = freq_{ij} * ratio_3 * LF_{avg} * Cap_{avg} * \frac{freq_{jg}}{\sum conn_j} * \frac{freq_{gh}}{\sum conn_g}$$
$$\forall i, j, g, h \in Airports,\ g \neq i,\ h \neq i, j \tag{27}$$

The primary, secondary and tertiary demand between each airport couple is summed to create a demand matrix that is used in the generation of cargo requests. Not every experiment is simulated with the full demand of the Cargolux network to reduce the computational time, therefore the size of the demand matrix can be scaled down. This allows for creating more suitable request sets for each experiment, while still keeping the same demand distribution. The request generation creates new cargo requests until the demand between a certain airport pair is met. The weight of each request (w_r) is randomly generated between 15,000 and 30,000 kg. The due time t_r^- can vary randomly between the beginning of day 2 (after 24 hours) and the end of the time horizon. The release time t_r^+ should be at least 24 hours and at most 48 hours before the due time. The value of t_r^+ is varied randomly between the two bounds, as long as the value is not negative. This can be described by $t_r^+ : \{max\{0, t_r^- - 48\}, t_r^- - 24\}$. A strategic weight s_r is given to each request, which has a standard value of 1 and is set to 1.5 for transatlantic flights, to reward the model to fly these longer and more costly flights. Also, this represents the increased value of commodities being transported overseas.

5.3 Description of experiments

Using the three created airline networks and the generated demand, a set of experiments is performed. These experiments each aim to investigate different parts of the research objective. First, the two emission approximation methods are compared in Experiment 1. This is done to find which approach is most suitable for the implementation of the carbon emissions into the MILP in the remainder of the research. Second, an analysis on the relation between the profit and emissions term is performed by means of a Pareto front analysis. This helps to understand how the model behaves when the two terms are disconnected and weighted differently. A modification to the objective function that is necessary for this analysis is explained in Section 5.3.1. The third experiment aims to find what decrease in profit can be expected given a pre-determined reduction of carbon emissions. The used constraint for this experiment is briefly discussed in Section 5.3.2. The goal of this experiment is to find how financially feasible a fixed reduction of emissions would be for different networks and combinations of fleet size and cargo requests. Finally, a sensitivity analysis is performed in Experiment 4 to determine how the model reacts to changing input variables. The first of these sub-experiments intends to find the consistency of the results with different request sets available. This is done by feeding the same model ten different variations of a request set that is generated from the same demand matrix. Furthermore, the effects of a different aircraft type on the flight scheduling is analysed. A flight schedule is optimised for both the B747-8F and the B747-400F with the same initial conditions. An overview of all experiments is shown in Table 10, which also indicates the fleet size and number of cargo requests that are used in each instance.

Table 10: Summary of the four experiments that are performed.

	Subject	[Fleet size/Cargo requests]	Discussed in
E1	Comparison of emission approximation methods	[1/20]	Section 6.1
E2	Pareto front search	[1/30]	Section 6.2
E3	Constraining maximum allowed emissions	[2/20], [3/25]	Section 6.3
E4	Sensitivity analysis	[2/15], [1/25]	Section 6.4

The networks used for this study are set up with a time horizon T_{hor} of 72 hours and a time step Δt of 3 hours. This time step provides a good trade-off between scheduling accuracy and computational complexity. The MILP is written in Python 3.8 and solved using the Gurobi Optimiser[11]. The maximum run time is set to 2 hours for a single flight schedule optimisation.

[11] https://www.gurobi.com/, accessed on 18-11-2021

5.3.1 Objective function for Pareto front search

For experiment E2, a slight adjustment to the objective function is necessary. Because the objective function consists of two distinct parts, namely increasing profit and reducing emissions, the problem can be considered a multi-objective optimisation. In such a problem, a Pareto front search can be performed to analyse how the two objectives relate to each other. A point is considered Pareto optimal if one of the objectives cannot be improved, without the other objective resulting in a lower value. In order to find this front, a weighting factor and normalisation are applied. This emission weighting factor w_{CO_2} can range continuously from 0 to 1 and is multiplied to the emission objective, while the profit part is multiplied by the $(1 - w_{CO_2})$ term. Increasing w_{CO_2} will therefore move the objective more towards minimising the emissions at the cost of a decrease in profit. Furthermore, both terms of the objective function are normalised. First a run with $w_{CO_2} = 0$ is performed for each instance to find the maximum profit ($Profit_{max}$) that can be achieved. In this profit, the costs that the airline has to pay for the emitted CO_2 is neglected. This run also indicates the maximum emissions ($Emission_{max}$) in tonnes that can be expected for the instance. By dividing the first term by $Profit_{max}$ and the second by $Emission_{max}$, the terms now both vary between 0 and 1, giving them equal influence in the objective function. The adjusted objective function is shown in Equation 28.

$$\begin{aligned} \max \quad & \frac{(1 - w_{CO_2})}{Profit_{max}} \left(\sum_{r \in \mathcal{R}} \sum_{(i,j) \in \mathcal{A}^R} Rev_{ij}^r \cdot q_{ij}^r - \sum_{k \in \mathcal{K}} \sum_{(i,j) \in \mathcal{A}^F} OperCost_{ij}^k \right) \\ & - \frac{w_{CO_2}}{Emission_{max}} \left(\sum_{k \in \mathcal{K}} \sum_{(i,j) \in \mathcal{A}^F} e_{ij}^k \right), \qquad \text{with } 0 \leq w_{CO_2} \leq 1 \end{aligned} \tag{28}$$

5.3.2 Constraint to restrict allowed emissions

An extra constraint is added in E3 to investigate how the profit responds to a forced reduction of the carbon emissions. First, a model is optimised with the conventional objective function (Equation 1) to find the maximum amount of CO_2 that is produced in the optimal schedule. Input parameter $\%_{reduct}$ can be used to enter the percentage of carbon emissions that is wished to be reduced. Equation 29 is used to ensure that the total emissions are reduced by the desired percentage $\%_{reduct}$. In the experiment $\%_{reduct}$ is varied in the range [0,5,10,15,20,25].

$$\sum_{k \in \mathcal{K}} \sum_{(i,j) \in \mathcal{A}^F} e_{ij}^k \leq (100 - \%_{reduct}) \cdot Emission_{max} \tag{29}$$

6 Results

This section will present and discuss the results of the four experiments that are presented in Table 10. Sections 6.1 to 6.4 present the results of Experiments E1, E2, E3 and E4 respectively.

6.1 E1: Comparison of emission approximation methods

In Section 4.2.2, two methods are proposed to implement the CO_2 emissions into the MILP, without needing to calculate the emissions for each flight throughout the optimisation. Because both methods are an approximation of the actual emissions, this experiment aims to find the validity of this approximation and to determine which of the two methods is the most suitable for this research.

First, the average error between the approximated values and the actual CO_2 emissions is analysed. This is done for a long haul (ATL - LUX) and a short haul (PIK - LUX) flight, which are each simulated 1,000 times with random load factors. The actual emissions are calculated using the emission model and approximations are made using the linear and discretisation approach. The error is calculated as a relative percentile error: $(E_{approx} - E_{actual})/E_{actual} * 100\%$. The results for the two approaches are shown in a boxplot in Figure 8(a). The linearised approach overestimates the emissions both for short and long range flights, with a maximum error of 2.5%. This could be expected from Figure 6, as the linear function is constantly above the calculated data-points. The maximum error can be expected around a load factor of 50%. The errors for the discretisation approach have a mean of 0, with a tighter spread for the short range flight. This can be explained due to the larger impact that the load factor has on longer flights. The difference between flying the PIK-LUX flight empty and with maximum payload is around 8 tonnes of CO_2, while this value lies over 85 tonnes for the ATL-LUX flight. This results in a steeper curve and a higher expected error when dealing with longer flights. Figure 8(b)

shows the error distribution of both flights combined, which results in a shape resembling a normal distribution for the discretisation approach. This can be a consequence of the selection of the middle-point of the identified load factor segment the flight belongs to, as described in Section 4.2.2. The errors of the linear approach are more uniformly distributed, which can be related to the randomly generated load factors and the larger spread of the interquartile range in the boxplot.

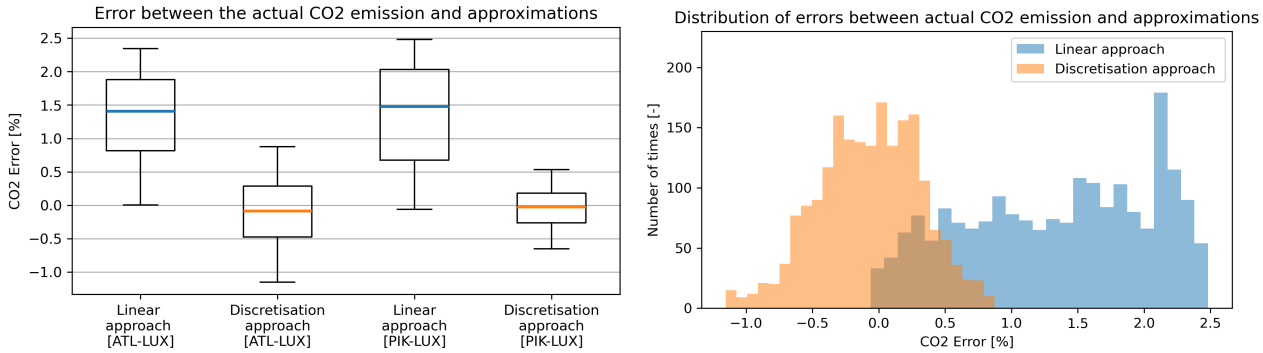

(a) Boxplot of the error distribution for the two approximation methods and two flights

(b) Combined error distribution of both approximation methods

Figure 8: Calculated errors between the actual and estimated emission values for 1000 runs with a random LF on a short range and long range flight.

Both methods were implemented in the MILP to analyse the behaviour of the approximation methods in the optimisation. The first observation was that the designed schedules for both methods are identical for all networks. This indicates that the approximation method has little to no influence on the aircraft and cargo routing when using a carbon price of €50 per tonne CO_2. Table 11 shows more detailed information on the calculated carbon emissions and the run times. The emissions for the linear approach are slightly higher than the discretisation method, which could be expected from the analysis of Figure 8. However, the differences seem negligible when compared to the increase in run time, which is almost an order of magnitude higher for the discretisation approach. This is mainly caused by the extra set of decision variables that is needed to identify the correct load factor segments (Section 4.2). If the computational time is not a limiting factor, the discretisation approach was proven to be slightly more accurate. However, for this research limiting the computational time as much as possible (provided a reasonable accuracy), was considered more important in order to analyse more instances in the performed experiments. Therefore, linear interpolation is used for the remainder of this study.

Table 11: Results of a run on each network, using a single aircraft and 20 available cargo requests.

	CO_2 [tonnes]		Run time [s]	
Network	Linear	Discretisation	Linear	Discretisation
Network 1: EU	363	360	43.7	480
Network 2: EU-NA	1,039	1,034	3.0	17.2
Network 3: NA	745	740	3.6	37.4

6.2 E2: Pareto front search

The second experiment investigates the relation between the profit and emissions terms, using the adjusted objective function (Equation 28). The model is run while varying the emission weighting factor w_{CO_2} between 0 and 1 in steps of 0.05. The combined objective value for this range of weighting factors is shown in Figure 9(a). Since the objective function is normalised, all curves start from 1 when w_{CO_2}=0. At this point, no weight is given to the emissions and the maximum available profit is found. With the weighting factor increasing, the objective value steadily decreases. If nothing would change to the flight schedule, this line would linearly decrease to a value of -1 at $w_{CO_2} = 1$. However, due to the two contrasting objectives, revisions to the flight schedules occur. These changes aim to decrease the CO_2 emissions without cutting too much of the profit.

When comparing the three different networks in Figure 9(a), the first difference that can be observed is the objective value at $w_{CO_2} = 1$. This value is negative for the EU-NA and NA networks, while being 0 for the EU network. This can be attributed to the initial and final location of aircraft AC0 being the same airport for

the EU network. Therefore no mandatory flights are operated. For the other two networks at least one flight is required to satisfy aircraft repositioning requirements, which results in CO_2 emissions and thus a negative objective value. In addition, no cargo requests are carried on these mandatory flights. This would only result in extra emissions due to the higher load factor, because the profit term is cancelled out by the $(1 - w_{CO_2})$ term. Another notable difference is the orange line of EU-NA that stays almost linear until $w_{CO_2} = 0.5$, where the objective value turns negative. This indicates that until that point no schedule revision is needed to improve the objective value. For the other two networks, revisions of the network are found for smaller values of w_{CO_2}, which can be seen from the lines slightly deviating upwards.

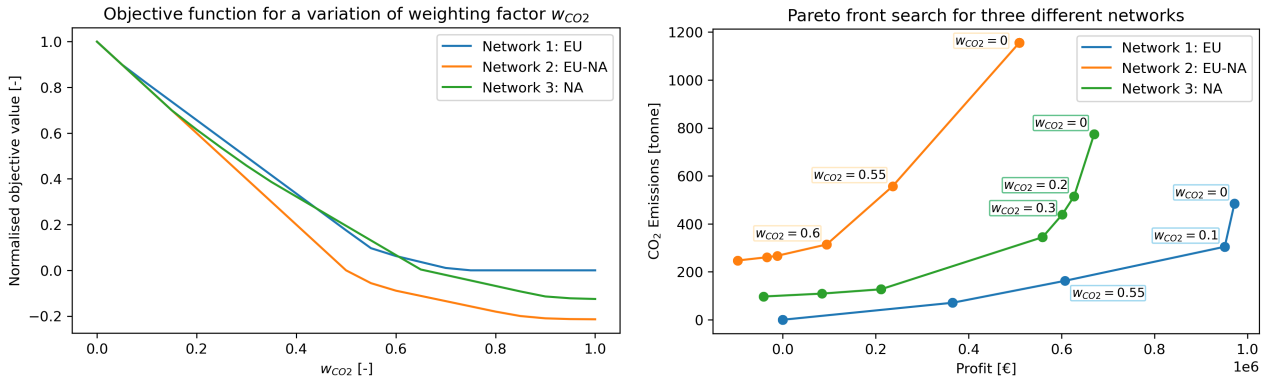

(a) Relation between w_{CO_2} and the normalised objective value.

(b) Pareto fronts showing the relation between the profit and emissions terms. The emission weighting factors of the three most profitable schedules are given.

Figure 9: Results of the model using the normalised objective function with the emission weighting factor w_{CO_2} for the three different networks. A single aircraft is used (AC0), with 30 cargo requests available.

The trade-off between the CO_2 emissions and the profit term is graphically shown by the Pareto fronts in Figure 9(b). Each point shows the amount of carbon emissions and the profit for an optimal schedule that was found during the variation of the emission weighting factor. The lines that connect the points represent the Pareto front, which indicates that the profit of a schedule can only be increased by also emitting more CO_2. Similarly, the emissions cannot be reduced, unless lower profits are accepted. Next to the three most profitable schedules of each network the w_{CO_2} value is shown. This is the smallest weighting factor that led to this revision of the schedule. These values correlate to the points in Figure 9(a) where a change in the slope can be observed. A first look at the three Pareto fronts in Figure 9(b) already highlights some of the characteristics of each network. The blue line of the EU network stays the most in the bottom-right part of the plot, indicating that large profits can be made with relatively little emissions. This is the result of the cargo prices not being dependent on the flight distance, which is beneficial for the short-haul flights within Europe. For EU-NA a large gap is identified between the two most profitable points on the Pareto front, which corresponds to the line being linear until $w_{CO_2} = 0.5$ in Figure 9(a). The NA network lies in between the other two, with a relatively steep slope for the first four schedule revisions. The small difference between the weighting factors of these schedules suggests that the most profitable schedule emits relatively much CO_2, and that there are many options in the network to reduce the emissions without losing a lot of profit. These observations are further discussed below using the corresponding flight schedules.

The two most profitable flight schedules for EU are shown in Figure 10. From the Pareto front it can be seen that the first revised European schedule ($w_{CO_2} = 0.1$) considerably reduces the carbon emissions, while almost maintaining the same profit as the original schedule. This is achieved by still flying the most profitable routes, while cancelling more polluting routes or routes with low load factors. A part of the schedule is unchanged, namely from the PIK-AMS flight 12 hours into the simulation horizon until the VIE-LUX flight that ends at $t = 30$. Most of these flights are performed with a relatively high load factor, which brings in a lot of revenue in an efficient way. In the revised schedule, the aircraft stays on the ground longer at LUX to wait for two new cargo requests that become available for the LUX-PIK flight at $t = 9$. In the original schedule these requests are transported using an extra round trip from LUX to PIK at hour 30, making those two flights redundant. Such revisions result in a profit decrease of just over 2%, while reducing CO_2 emissions by 180 tonnes. When analysing the Pareto front in Figure 9(b), the remainder of the blue line is not as steep as the line between 0 and 0.1. This means that a further reduction of the emission will require more severe schedule changes, with a larger impact on the profit of the airline.

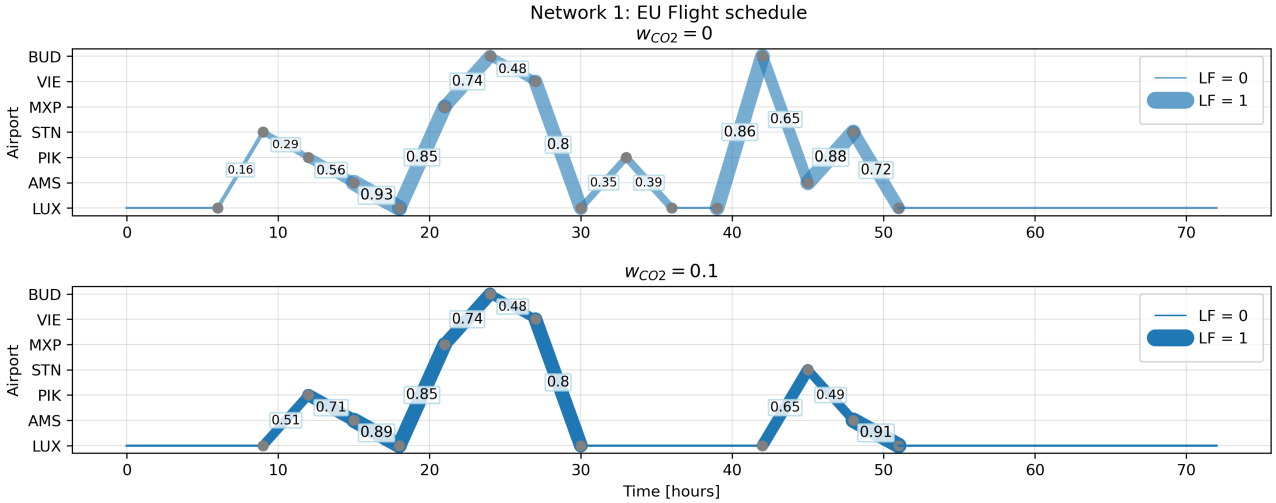

Figure 10: Flight schedules for the EU Network for aircraft AC0. Schedules are optimised for $w_{CO_2} = 0$ and 0.1. The load factor is given for each flight arc, which is also indicated by the thickness of the line.

As mentioned during the analysis of Figure 9, no revised schedule for the EU-NA network is found until $w_{CO_2} = 0.55$. A comparison between the initial and revised schedule is shown in Figure 11. A big difference between both schedules is the extra transatlantic round trip that is scheduled in the $w_{CO_2} = 0$ instance. In this work, the revenue per kg for transatlantic cargo demand is assumed to be 50% higher, which makes these trips very appealing. In order to capture both the demand from Europe to the US and back, three transatlantic flights must be performed to also reach the final aircraft location ORD. At the start of the initial schedule, the aircraft is scheduled to fly from PIK to ORD with a load factor of only 0.17, which results in a loss of around €15,000. After that, the aircraft picks up cargo in ATL, ORD and JFK to fly back to LUX at $t = 30$ at maximum payload capacity. The huge profit that is made on this flight is very hard to match by flying extra flights in Europe or in the US. This schedule only becomes non-optimal for an emission weighting factor of 0.55 and higher. Due to a transatlantic flight emitting around 250 tonnes of CO_2, compared to 50-100 tonnes for other flights in the network, these flights are not profitable anymore for such high values of w_{CO_2}. Another interesting observation is that both aircraft rotations end with an empty flight to ORD. This indicates that the mandatory final location of the aircraft is not ideal for this instance.

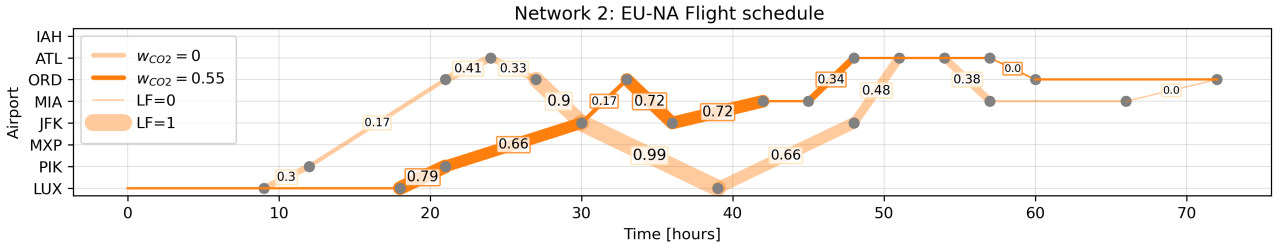

Figure 11: Flight schedules for the EU-NA Network for aircraft AC0. Schedules are optimised for $w_{CO_2} = 0$ and 0.55. The load factor is given for each flight arc, which is also indicated by the thickness of the line.

Finally, Figure 12 presents the two most profitable schedules for Network 3. The revised schedule reduces the amount of carbon emissions by more than 250 tonnes, while only making around 6% less profit. This is achieved by slightly increasing the average load factor of all flights and decreasing the average flight length. However, only picking the flights with a higher load factor does not directly end up in a better objective. This can be seen by comparing the first flights of both schedules. In the initial schedule the aircraft is first routed to LAX and GDL before heading back to MEX. Although the average load factor on these three flights is above 0.6, it is replaced by just flying back and forth between MEX to GDL, with an average load factor of only 0.34. In this case the revenue from the cargo requests in LAX does not weigh up to the extra emissions from the longer flight. Such trade-offs between the revenue from the cargo requests and the carbon emissions show how the model maximises the combined objective function. When looking back at the Pareto front in Figure 9(b) it can be seen that the first three revised schedules for the NA network all come up with a significant CO_2 reduction with limited impact to the profit. This network and the corresponding cargo requests allow for easier schedule revisions than the EU-NA network. The demand is more evenly spread out in Network 3 and the costly and long transatlantic flights of Network 2 reduce the amount of options to vary the schedules. This points out how

the relationship between the profit and the emissions term is influenced by the airports in the network and the available cargo demand.

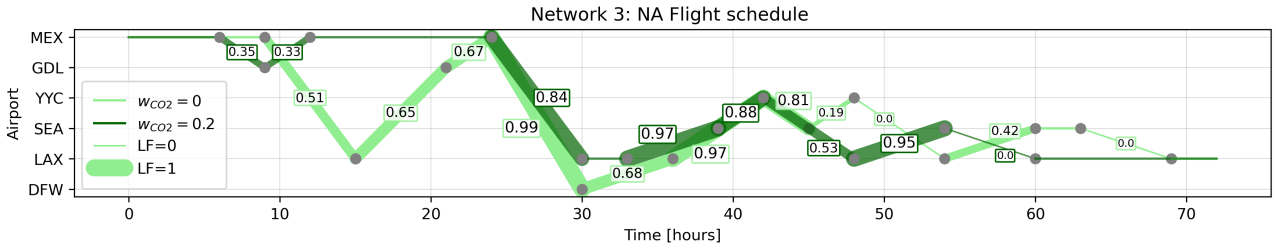

Figure 12: Flight schedules for the NA Network for aircraft AC0. Schedules are optimised for $w_{CO_2} = 0$ and 0.2. The load factor is given for each flight arc, which is also indicated by the thickness of the line.

6.3 E3: Constraining maximum allowed emissions

For the third experiment an extra constraint is added to the MILP formulation, as presented in Section 5.3.2. This constraint allows the model to find an optimal schedule which reduces the emissions by at least a predefined percentage $\%_{reduct}$, with respect to an initial schedule. The original objective function (Equation 1) is used, where the objective value is equal to the profit. This profit consists of the revenue, operational costs and the CO_2 costs, which is why in this section the amount of carbon emissions is discussed in terms of euros. The initial schedule is created by running the model with $\%_{reduct} = 0$. The maximum amount of CO_2 that is allowed is reduced in five steps of 5% to find how the revenue and cost parameters react.

The results of the EU network are presented in Table 12. The first three columns specify the input parameters of the instance, followed by the resulting revenue, cost and profit. As expected, aiming for lower carbon emissions decreases the profit of the network. Fewer cargo requests are served compared to the initial schedule ($\%_{reduct} = 0$), resulting in a reduction of the revenue. However, rerouting aircraft and cancelling flights also reduces the operational and CO_2 costs, which prevents the profit to decrease faster. The larger number of available cargo requests and the extra aircraft in the fleet result in higher profits for the network. In these cases, more demand is available, which makes it easier to find profitable flights. In both instance sets, a fairly large percentage of the profit is lost in order to comply with the reduction in carbon emissions, with mainly the larger values of $\%_{reduct}$ leading to bigger drops in the profit.

Table 12: Results for Network 1: EU.

Instance			Output			Objective value			
$\%_{reduct}$	Fleet	Requests	Revenue [€]	Operational cost [€]	CO_2 cost [€]	Profit [€]	Decrease [%]	Run time [s]	Gap [%]
0	2	20	773,968	192,672	15,249	566,047	0.0	225.9	0.0
5	2	20	739,846	175,532	14,147	550,167	2.8	134.3	0.0
10	2	20	693,984	155,521	12,712	525,750	7.1	374.0	0.0
15	2	20	693,984	155,521	12,712	525,750	7.1	144.3	0.0
20	2	20	600,592	141,854	12,083	446,654	21.1	872.6	0.0
25	2	20	588,592	142,027	10,862	435,703	23.0	877.1	0.0
0	3	25	1,090,308	257,737	20,459	812,111	0.0	154.6	0.0
5	3	25	1,023,996	232,393	18,434	773,169	4.8	700.9	0.0
10	3	25	1,023,996	233,016	18,237	772,744	4.8	1,820.2	0.0
15	3	25	966,962	212,530	17,063	737,369	9.2	2,249.4	0.0
20	3	25	922,810	207,888	16,114	698,808	14.0	1,688.7	0.0
25	3	25	872,096	188,789	14,735	668,572	17.7	1,421.0	0.0

The results of the second network have some other interesting aspects. Table 13 shows that for the smaller instance with 2 aircraft and 20 requests only one revised schedule is found. Even though the revenue of this new schedule has decreased by almost €300,000, only a 6.3% profit decrease is needed. Next to that, the amount of CO_2 is lowered by over 40%, which makes this schedule also optimal for the cases with higher values of $\%_{reduct}$. When focusing on the larger instance set, no optimal solution can be found for the CO_2 reductions of 10% and higher. The final column in the table shows the optimality gap when the branch-and-bound process was stopped due to run-time restrictions. The value indicates the gap between the best lower and upper bound solutions that the model has found. If both bounds are equal, a global optimum is found. If no optimal solution can be

found within the maximum run time of 2 hours, the gap and the best bound objective value are given. These results show the limitations of the model for larger instances. The complexity of the solution rapidly increases due to options for cargo to transfer between different aircraft in the fleet. The problem size is also increased by adding more airports and cargo requests, which need extra nodes and arcs in the TSN. Therefore, this model is not ideal for the simulation of larger networks or for optimising over longer time horizons.

Table 13: Results for Network 2: EU-NA.

Instance			Output			Objective value			
$\%_{reduct}$	Fleet	Requests	Revenue [€]	Operational cost [€]	CO_2 cost [€]	Profit [€]	Decrease [%]	Run time [s]	Gap [%]
0	2	20	1,071,448	649,995	75,777	345,676	0	19.0	0.0
5	2	20	717,301	352,760	40,798	323,743	6.3	126.9	0.0
10	2	20	717,301	352,760	40,798	323,743	6.3	140.7	0.0
15	2	20	717,301	352,760	40,798	323,743	6.3	120.1	0.0
20	2	20	717,301	352,760	40,798	323,743	6.3	60.9	0.0
25	2	20	717,301	352,760	40,798	323,743	6.3	90.7	0.0
0	3	25	1,506,257	892,026	106,567	507,664	0.0	2,762.7	0.0
5	3	25	1,405,257	809,794	98,630	496,834	2.1	6,654.9	0.0
10	3	25	1,345,711	781,947	95,818	467,946	7.8	7,200.1	11.7
15	3	25	1,133,918	645,783	72,827	415,308	18.2	7,200.1	21.9
20	3	25	1,133,918	645,783	72,827	415,308	18.2	7,200.0	19.7
25	3	25	1,133,918	645,783	72,827	415,308	18.2	7,200.1	15.8

Table 14 shows the results of the North-American network. For both instances the profit reduction stays limited to a maximum of 10% when reducing the emissions by 25%. This is a similar observation to the one that was made during the analysis of the Pareto front (Section 6.2). The emissions are reduced with relatively small adjustments to the schedule, without the need to cancel a lot of cargo requests. The demand is more evenly spread out over the airports compared to the EU network which is mostly focused on the Cargolux hub in LUX. The medium range flights also allow for more flexibility in the aircraft routing than the transatlantic flights in the EU-NA network.

Table 14: Results for Network 3: NA.

Instance			Output			Objective value			
$\%_{reduct}$	Fleet	Requests	Revenue [€]	Operational cost [€]	CO_2 cost [€]	Profit [€]	Decrease [%]	Run time [s]	Gap [%]
0	2	20	882,668	371,060	37,475	474,133	0.0	26.7	0.0
5	2	20	845,870	344,961	35,252	465,657	1.8	142.6	0.0
10	2	20	812,022	324,816	33,376	453,829	4.3	95.7	0.0
15	2	20	746,726	289,257	31,029	426,439	10.0	230.8	0.0
20	2	20	742,018	287,653	28,061	426,304	10.1	127.0	0.0
25	2	20	742,018	287,653	28,061	426,304	10.1	118.3	0.0
0	3	25	1,083,248	477,093	49,698	556,457	0.0	4,457.7	0.0
5	3	25	1,036,642	439,948	45,510	551,184	0.9	6,053.2	0.0
10	3	25	1,026,144	431,571	44,106	550,467	1.1	7,200.2	2.4
15	3	25	959,558	389,291	39,553	530,713	4.6	4,584.1	0.0
20	3	25	959,558	389,291	39,553	530,713	4.6	1,675.2	0.0
25	3	25	869,928	326,109	33,240	510,579	8.2	1,456.5	0.0

Finally, a more detailed breakdown of the profit of each network is shown in Figure 13. The results of the initial schedule and the 25% CO_2 reduction instance are plotted, with the yellow dotted line showing the profit. In the initial schedule, almost all cargo requests are served for each network, but the higher strategic factor for the transatlantic cargo requests results in a much higher revenue for the EU-NA network. However, this is cancelled out by much larger operational, fuel and CO_2 costs, resulting in the lowest profit of the three. The longer flight length clearly leads to higher costs for the airline, compared to the EU network, where only a small portion of that cost is needed to serve the same amount of cargo requests. A similar distribution can be seen for the instance where the carbon emissions have to be reduced by at least a quarter. Only the most profitable flights are chosen, which leads to both less revenue and less costs for the airline. Analysing more closely the portion of the costs that comes from the carbon emissions, it can be concluded that this is fairly small for all

instances. Increasing the penalty to emit CO_2, would create more financial incentive for airlines to change their flight schedules in order to reduce emissions.

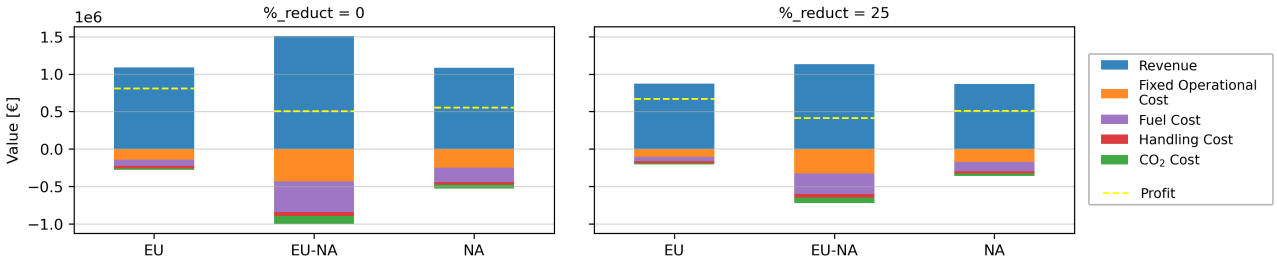

Figure 13: Distribution of the revenue and the different cost terms for the initial and a revised schedule with $\%_{reduct} = 25$. The schedules are found for all three networks, with a fleet of 3 aircraft and 25 cargo requests available.

6.4 E4: Sensitivity Analysis

The final experiment analyses the influence of changing the input parameters on the outcome of the model. First, the sensitivity of the models to changes in the set of cargo requests is investigated by running similar instances with newly generated request sets. Second, the influence of different aircraft types in the fleet is assessed.

6.4.1 Cargo Requests

In all previous experiments, only one set of cargo requests is generated and used to find the results. The demand distribution is based on the current flight network of Cargolux, which is kept constant. However, the request generation is also dependent on three parameters that can vary randomly between a given upper and lower bound, namely the weight, release time and due time of the request (Section 5.2). This means that even though the general division of the total cargo weight across the network is similar, changes can exist between different request sets. This section aims to find how the model behaves when receiving different request sets based on the same network and demand distribution. Each network is run ten times with a newly generated set of 15 requests and a schedule is optimised for a fleet of two aircraft (AC0 and AC1). The randomised request weight results in a difference in the maximum amount of revenue that is possible for the 15 requests. To create a good comparison between the different instances, a similar set-up as in the E3 has been used, where the emissions are constrained by a certain $\%_{reduct}$. The decrease in profit is shown for all three networks in Figure 14.

The thin lines in each indicate the results of one of the ten individual runs, which are combined to find the mean, shown by the thicker line. In addition, a 95% confidence interval around the mean is formed. A distinctive difference can be observed between Network 2 (EU-NA) and the two networks that operate only in Europe and North-America. Network 1 and 3 have similar characteristics, with a mean profit decrease of 14.1% and 12.9% respectively at $\%_{reduct} = 25$. When comparing this figure to the results of the European network in the previous section, it can be concluded that the profit decrease found in Table 12 lies on the higher end of the scale. The results from Network 3 in Table 14 are closer to the mean shown Figure 14. The different runs show that the results are influenced by changes in the request sets, but that this impact is much larger for the EU-NA network. This matches the analyses from the Pareto front, where it was found that in this network few options are available to revise the schedule. The larger uncertainty can also be attributed to the variation of the release and due times of the requests. For Network 2 this is an important factor in making the long transatlantic flights profitable. Because this experiment is performed using instances with only 15 cargo requests, it is a great advantage if multiple requests of a similar origin-destination pair are available at the same time period. When this is not the case, the aircraft are forced to fly from Europe to North-America or back with a very low load factor in order to reach the final scheduled destination. The low number of request also means that the available profit is lower in the first place, automatically resulting in larger relative differences.

Overall, the results of Network 1 and 3 are relatively consistent when designing the schedules for different request sets and give a reliable indication on the profit decrease that can be expected for a certain reduction of carbon emissions. For the longer range network (EU-NA) the results are more inconsistent. This shows that the potential to decrease emissions by changing the flight scheduling and aircraft routing is dependent on the characteristics of the network and the available cargo demand. The model proposed in this study can still help determining how to revise aircraft rotations to improve operations.

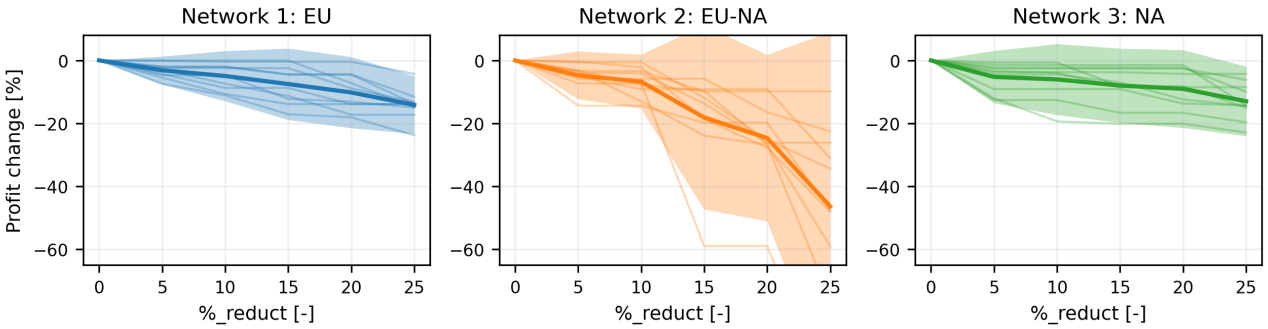

Figure 14: Progression of the objective function when constraining the carbon emissions. The thick line shows the mean of the ten runs. The larger contours gives a 95% confidence interval around the mean.

6.4.2 Aircraft Types

Next to the cargo inputs, also the available aircraft are varied in this study. The two most common aircraft types that Cargolux operates are used in this research, namely the Boeing 747-8F and the 747-400F. In E3, schedules for a non-homogeneous fleet with both aircraft types are optimised. However, using a mixed fleet it is hard to compare how the aircraft types relate to each other in terms of performance. In this section, a schedule is optimised for all three networks, with the similar set of 25 requests available for both aircraft types operated separately. The initial and final location of AC0 are used (Table 9), to keep all conditions equal. The aircraft parameters are given in Table 2, which shows the longer range and larger payload capacity of the newer B747-8F. The emission model is used to generate emission matrices for both aircraft.

Table 15 shows the output of the model, with the operational cost split out into the fixed operational cost, fuel cost and handling cost. For the European network, both aircraft are scheduled to serve all 25 cargo requests. However, from the difference in fixed operational costs it can be concluded that different schedules are produced. The fixed operating costs per hour (C_F) are assumed equal for both aircraft types, which means that the B747-400F needs more flights hours to transport the same amount of cargo. This is caused by the lower payload capacity compared to the B747-8F. Furthermore, the longer flights result in higher fuel costs and over 60 tonnes more CO_2. In return, the larger OEW of the B747-8F causes an increase in the average take-off weight, resulting in larger handling costs.

Larger profits can be made in all three networks when allocating the B747-8F, which makes this aircraft type more attractive for airlines. The aircraft is also more fuel efficient and emits less CO_2 per flight. However, the proposed schedule of the B747-400F for the EU-NA network ends up with lower costs for fuel and carbon emissions. Because the goal of the model is to maximise the profit, trade-offs between revenue and costs are made in the determination of an optimal schedule. The Boeing 747-8F still makes more profit, because it flies a slightly different route where it can make more revenue by transporting heavier cargo requests. If the price for carbon emissions or fuel would be increased, this results in a schedule that is less focused on the revenue.

Table 15: Results for both aircraft types on the three networks, with 25 requests available.

Instance		Output						Objective
Network	Aircraft type	Requests served	Revenue [€]	Fixed oper cost [€]	Fuel cost [€]	Handling cost [€]	CO_2 cost [€]	Profit [€]
EU	B747-8F	25	1,090,308	157,500	88,437	35,593	22,574	786,202
	B747-400F	25	1,090,308	163,974	100,602	32,598	25,722	767,408
EU-NA	B747-8F	14	904,317	248,694	234,686	30,296	61,623	329,016
	B747-400F	14	822,019	224,282	215,591	23,937	56,605	301,601
NA	B747-8F	20	905,166	200,833	150,843	33,341	39,330	480,817
	B747-400F	19	874,424	200,833	155,578	29,202	40,724	448,084

Finally, Figure 15 gives a visual representation of the two flight schedules that are created for Network 3. The schedules for both aircraft are identical for the first part. The two values on top of each flight arc indicate the load factor of the different aircraft. From this, the effect of a larger payload capacity can be clearly seen, as the load factor of the B747-8F is always lower. A problem arises when the load factor of the B747-8F is higher

than 0.84, which is equal to the maximum payload capacity of the Boeing 747-400F. This occurs at $t = 45$, for the flight from LAX to YYC. The B747-400F is now forced to fly a slightly different route, with the result that one request less can be served. This is creates the largest difference in profit between the two aircraft.

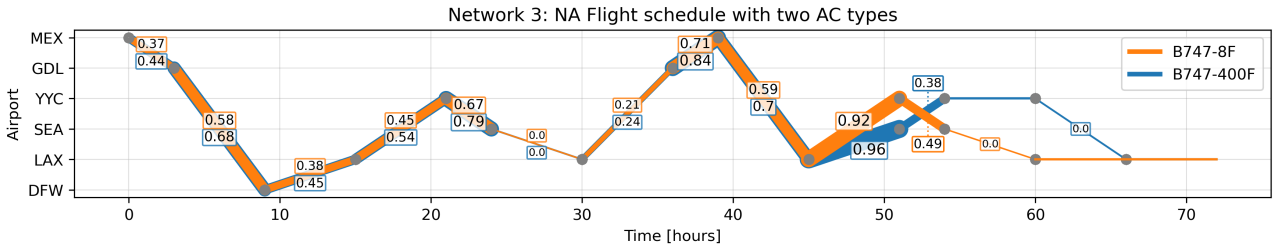

Figure 15: Flight schedule of the NA Network with 25 cargo requests, optimised for a B747-400F and B747-8F aircraft. The load factor is given for both aircraft on each flight arc, which is also indicated by the thickness of the line.

The results of this analysis show that aircraft type that is used in the fleet directly influences the flight schedule and thus also on the corresponding profit and emissions. This is already visible in the schedule design for two very similar aircraft types that both belong to the Boeing 747 family. From the results presented in this section, the expected conclusion can be drawn that the newer B747-8F is the more efficient aircraft of the two, and allows for higher profits and fewer CO_2 emissions.

7 Discussion

This section briefly discusses the limitations of the developed model and the effects of certain assumptions and simplifications on the results.

Computational and modelling limitations As highlighted in Section 6.3, the model is not always able to find an optimal result within the set maximum run time of 2 hours. Increasing the number of airports in the network, the fleet size and/or the amount of time steps in the simulation horizon has a clear effect on the solution complexity and therefore also the computational time. This creates challenges when optimising operations for the whole fleet and network of a full-cargo carrier which, in the case of Cargolux, corresponds to roughly 30 aircraft and 100 airports. Even though the problem size is reduced in a pre-processing step that eliminates unnecessary decision variables, no optimal flight schedules can be found for combinations with more airports, aircraft and cargo requests than a certain threshold. The same is valid if a smaller simulation time step is required, although for cargo operations generally a larger time step is allowed, compared to passenger operations. Nevertheless, the chosen time step of 3 hours already results in unnecessary long waiting times for short-haul flights that take far shorter than 3 hours. In addition, flights that are just longer than 3 or 6 hours are put in a disadvantage to flights that 'cost' one time step less to perform in the model, but much less in reality.

Concerning the two approximation methods used to implement carbon emissions into the objective function, the simpler linear approximation approach was chosen for the specific instances of this paper due to its computational efficiency. In case of severely non-linear behaviours, the piecewise linear approximation should be more appropriate. Nevertheless, further modelling improvements are possible for this approach. In the current framework, CO_2 emissions referring to the middle-point of the selected load factor segment are selected. This is an approximation within an already approximated function. An alternative would be to implement a linear combination of the CO_2 emissions of the left and right bounds of the selected segment, given the computed load factor, to determine a more precise value for emissions.

Synthetic data Second, the experiments performed in this study are all based on synthetic data, which is inspired by actual flight schedules and reported average load factors due to the lack of available real data. Therefore, the quantitative results of this study are not directly representative of the real operations of the airline that is considered in this paper. However, the results do give an impression of the capabilities of the proposed flight scheduling model. The model can also be easily adjusted to run experiments with real life input values and demand data. Furthermore, the decrease of the emissions in E3 (Section 6.3) is determined relative to a schedule that has already been optimised using the developed model, which includes an emissions term. It is possible that a smaller profit decrease is required if the schedules are revised based on an initial schedule that has been designed without considering the aircraft emissions.

Fleet initial and final location Next to the demand, also the fleet inputs are varied in the case studies. All aircraft are given fixed initial and final locations to keep the initial conditions of different runs similar. From the flight schedules that were generated, it could be seen that these fixed positions sometimes led to flights that are flown empty or with a very low load factor in order to comply with the repositioning requirements. Optimising the initial and final aircraft position was not investigated in this study, just as maintenance constraints that might justify such repositioning requirements. However, such features can easily be implemented in the model.

Emissions calculation Finally, the emission model uses simplifications in the calculation of the carbon emissions for each flight. The vertical flight path is not optimised for the aircraft weight during the flight and it is assumed that the aircraft can follow the great-circle arc to reach its destination. If it is desired to model the emissions with greater accuracy, such assumptions should be reconsidered. Furthermore, only a single aircraft type is used for the validation of the emission model. The results of this validation were found to be accurate enough for this study, however more extensive validation is necessary to improve the knowledge of how the developed model relates to existing emission models.

8 Conclusions and Recommendations

This paper proposes an integrated flight scheduling, aircraft routing and cargo routing model with an aircraft emission model to produce a more sustainable flight planning for a full-cargo airline, where trade-offs are made between operational (profit maximisation) and environmental (CO_2 emissions minimisation) sustainability. The emissions are introduced into the objective function with a direct dependency on the cargo that is on-board the aircraft, using the current EU-ETS price of €50 per tonne of CO_2. The aim of this study was to find how flight schedules can be revised in order to decrease the carbon emissions of a dedicated cargo airline, while monitoring the effect on the profit.

The emission model that is used as input to the optimisation model is developed using an open source aircraft performance model, and is validated by comparing the results to two existing emission models. The emissions are implemented into the MILP with a direct relation to the amount of cargo scheduled for each flight leg, using two different approaches. The first conclusion is that for this study a linear function that maps the emissions as a function of the load factor was the most suitable method due to computational performances. The other method, which uses discretisation of a piecewise linear function, does lead to slightly more accurate results. However, this increase in accuracy does not outweigh the much larger computational time. In case computational time is not a limiting factor, or if the simulated aircraft types and/or configurations lead to a more non-linear relationship between load factor and emissions, this second approach that was provided is more suitable.

The performance of the integrated model was assessed using three air cargo networks, varying between 6 and 8 airports, created to simulate part of the operations in Europe and North-America of Cargolux, a major European full-cargo airline. A Pareto front search was carried out to determine the interdependence between operational and environmental sustainability. From this it was concluded that the relationship between the profit and emissions term is heavily influenced by the characteristics of the network. The long transatlantic flights give less opportunity to implement schedule revisions than networks with shorter flight distances and more evenly spread out demand. Increasing the load factor can be used to increase the carbon efficiency. Also, reducing the number of long-haul flights from the schedule is effective in reducing the CO_2 emissions.

Another set of experiments was run with the goal of modifying an initial schedule by reducing emissions by a percentile range from 5 to 25%. For the largest instances of 3 aircraft and 25 cargo requests, this led to a profit decrease of 17.7% for a reduction of 25% of the CO_2 for the European network. The transatlantic Europe-North America network led to a 18.2% reduction of the profits for the same carbon reduction. However, the model could not find a global optimum within the 2 hour run time, showing the computational limitations for larger networks. The North-American network performed best and only lost 8.2% of the profit for the same 25% CO_2 reduction.

The consistency and reliability of this experiment was tested by running a set of ten similar instances with newly generated demand requests based on the same demand distribution. This showed consistent results for the EU and NA networks. However, the EU-NA network was more sensitive to request changes, leading to less reliable results. Lastly, it was found that the allocation of the B747-8F leads to higher profits than the B747-400F, due to its larger payload capacity and higher fuel efficiency. From this it is concluded that replacing older aircraft in the fleet can result in higher profits and thus more opportunity to decrease aircraft emissions.

A major takeaway of this research, especially in an era where all airlines and transport providers are advertising (environmental) sustainability and promising carbon reduction or neutrality, is that implementing these promises always comes with a cost. While we can only base our conclusions on the specific instances we tested,

it is evident how reducing CO_2 emissions has an effect (decrease) on profit. For our test cases, the two objectives proved to be contrasting and hence trade-offs must be made. At the same time, we developed a framework where these trade-offs can be easily mapped and assessed, hence providing a useful decision-making tool for airlines. In particular, one of our experiments specifically maps the loss of profit given a pre-determined percentile reduction in carbon emissions, to mimic current airline environmental sustainability goals. While having some modelling limitations, airlines could use this approach to assess if the carbon reduction they want to comply with is economically sustainable or if that goal needs to be revised in order to contain profit losses.

Finally, while this research provides a first solid step towards a better understanding of the interdependencies between operational and environmental sustainability for a full-cargo airline, it can be improved in many ways. First, the computational time should be reduced in order to better simulate larger networks, shorter time steps, and/or longer time horizons. While our results are still valid, they refer to a subset of a full-cargo airline both in terms of fleet size and airport network. Full freighter rotations are generally geographically specific and periodic (and hence different sub-networks operate almost independently). Nevertheless, modelling the whole set of operations is more accurate and might lead to similar or relatively different results. Improving the algorithmic performance of the model can be achieved, for example, by implementing a string-based network for the cargo routing, that is solved using meta-heuristics or column generation. Second, the emission model can be improved by generating a more realistic trajectory that includes air traffic management constraints and a function to optimise the airspeed and altitude during cruise. Furthermore, other emissions such as nitrogen oxides, water vapour, and CO should be added to the emission model, since they contribute to atmospheric pollution. This can help to better understand the atmospheric and local air quality effects of the newly designed schedules. Third, provided data availability, future research should be based on real demand and cost data, to better support conclusions on the feasibility and potential of carbon emissions reductions. Next to a more realistic financial analysis, operational features could be added, like maintenance modelling, crew scheduling, and optimisation of the fleet size and re-positioning.

References

[Abara, 1989] Abara, J. (1989). Applying Integer Linear Programming to the Fleet Assignment Problem. *Interfaces*, 19(4):20–28.

[Barnhart et al., 1998] Barnhart, C., Boland, N. L., Clarke, L. W., Johnson, E. L., Nemhauser, G. L., and Shenoi, R. G. (1998). Flight string models for aircraft fleeting and routing. *Transportation Science*, 32(3):208–220.

[Boeing, 2002] Boeing (2002). 747-400 airplane characteristics for airport planning. Technical report.

[Boeing, 2012] Boeing (2012). 747-8 airplane characteristics for airport planning. Technical report.

[Boeing, 2020] Boeing (2020). Boeing World Air Cargo Forecast (WACF). Technical report.

[Brueckner and Abreu, 2017] Brueckner, J. K. and Abreu, C. (2017). Airline fuel usage and carbon emissions: Determining factors. *Journal of Air Transport Management*, 62:10–17.

[Carlier et al., 2006] Carlier, S., Smith, J., and Jelinek, F. (2006). GAES-Future Engine Technology Environmental Impact. Technical report, EUROCONTROL, Bretigny sur Orge.

[Chao, 2014] Chao, C. C. (2014). Assessment of carbon emission costs for air cargo transportation. *Transportation Research Part D: Transport and Environment*, 33:186–195.

[Clarke et al., 1997] Clarke, L., Johnson, E., Nemhauser, G., and Zhu, Z. (1997). The aircraft rotation problem. *Annals of Operations Research*, 69:33–46.

[Delgado and Mora, 2021] Delgado, F. and Mora, J. (2021). A matheuristic approach to the air-cargo recovery problem under demand disruption. *Journal of Air Transport Management*, 90:Article 101939.

[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:Article 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. *Source: Transportation Science*, 43(3):370–380.

[Derigs and Illing, 2013] Derigs, U. and Illing, S. (2013). Does EU ETS instigate Air Cargo network reconfiguration? A model-based analysis. *European Journal of Operational Research*, 225(3):518–527.

[EEA, 2019] EEA (2019). EMEP/EEA Air Pollutant Emission Inventory Guidebook, Part B: Aviation. Technical report, European Environment Agency, Copenhagen.

[Froyland et al., 2014] Froyland, G., Maher, S. J., and Wu, C. L. (2014). The recoverable robust tail assignment problem. *Transportation Science*, 48(3):351–372.

[Graver et al., 2019] Graver, B., Zhang, K., and Rutherford, D. (2019). CO2 emissions from commercial aviation, 2018. Technical report, International Council on Clean Transportation.

[Hane et al., 1995] Hane, C. A., Barnhart, C., Johnson, E. L., Marsten, R. E., Nemhauser, G. L., and Sigismondi, G. (1995). The fleet assignment problem: Solving a large-scale integer program. *Mathematical Programming*, 70(1-3):211–232.

[Haouari et al., 2013] Haouari, M., Shao, S., and Sherali, H. D. (2013). A lifted compact formulation for the daily aircraft maintenance routing problem. *Transportation Science*, 47(4):508–525.

[IATA, 2020] IATA (2020). Aviation & Climate Change. Technical report.

[ICAO, 2016a] ICAO (2016a). Airport Air Quality Manual. Technical report.

[ICAO, 2016b] ICAO (2016b). Environmental Trends in Aviation to 2050. Technical report.

[Jardine, 2005] Jardine, C. (2005). Calculating the Environmental Impact of Aviation Emissions. Technical report, Oxford University Centre for the Environment, Oxford.

[Khaled et al., 2018] Khaled, O., Minoux, M., Mousseau, V., Michel, S., and Ceugniet, X. (2018). A compact optimization model for the tail assignment problem. *European Journal of Operational Research*, 264(2):548–557.

[Liang and Chaovalitwongse, 2013] Liang, Z. and Chaovalitwongse, W. A. (2013). A network-based model for the integrated weekly aircraft maintenance routing and fleet assignment problem. *Transportation Science*, 47(4):493–507.

[Liang et al., 2011] Liang, Z., Chaovalitwongse, W. A., Huang, H. C., and Johnson, E. L. (2011). On a new rotation tour network model for aircraft maintenance routing problem. *Transportation Science*, 45(1):109–120.

[Loo et al., 2014] Loo, B. P., Li, L., Psaraki, V., and Pagoni, I. (2014). CO2 emissions associated with hubbing activities in air transport: An international comparison. *Journal of Transport Geography*, 34:185–193.

[Miyoshi and Mason, 2009] Miyoshi, C. and Mason, K. J. (2009). The carbon emissions of selected airlines and aircraft types in three geographic markets. *Journal of Air Transport Management*, 15(3):138–147.

[Pagoni and Psaraki, 2014] Pagoni, I. and Psaraki, V. (2014). A tool for calculating aircraft emissions and its application to Greek airspace. *Transportation Planning and Technology*, 37(2):138–153.

[Pollack, 1974] Pollack, M. (1974). Some aspects of the aircraft scheduling problem. *Transportation Research*, 8(3):233–243.

[Sherali et al., 2006] Sherali, H. D., Bish, E. K., and Zhu, X. (2006). Airline fleet assignment concepts, models, and algorithms. *European Journal of Operational Research*, 172(1):1–30.

[Sun, 2019] Sun, J. (2019). *Open Aircraft Performance Modeling Based on an Analysis of Aircraft Surveillance Data*. PhD thesis, TU Delft.

[Sun et al., 2020] Sun, J., Hoekstra, J. M., and Ellerbroek, J. (2020). OpenAP: An Open-Source Aircraft Performance Model for Air Transportation Studies and Simulations. *Aerospace*, 7(104).

[van der Meulen et al., 2020] van der Meulen, S., Grijspaardt, T., Mars, W., van der Geest, W., Roest-Crollius, A., and Kiel, J. (2020). Cost Figures for Freight Transport – final report. Technical report, Panteia.

[Yan and Chen, 2008] Yan, S. and Chen, C. H. (2008). Optimal flight scheduling models for cargo airlines under alliances. *Journal of Scheduling*, 11(3):175–186.

[Yan et al., 2006] Yan, S., Chen, S. C., and Chen, C. H. (2006). Air cargo fleet routing and timetable setting with multiple on-time demands. *Transportation Research Part E: Logistics and Transportation Review*, 42(5):409–430.

[Yan et al., 2020] Yan, Y., Zhang, J., and Tang, Q. (2020). Aircraft fleet route optimization based on cost and low carbon emission in aviation line alliance network. *Journal of Intelligent and Fuzzy Systems*, 39(1):1163–1182.

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

II

Literature Study
previously graded under AE4020

1

Introduction

Just like all other industries and transport sectors, the airline industry is under large national and international pressure to decrease the emissions of greenhouse gasses and other hazardous emissions that influence the air quality around airports. Graver et al. [19] determined that the global CO_2 contribution of the aviation industry in 2018 was 2.4% of the world wide fossil fuel emissions, which translates to 918 million metric tons of CO_2. The CO_2 emissions increased with 32% over the five year period prior to 2018 [19] and forecasts by the International Civil Aviation Organisation ICAO [22] show this growth continuing. Innovations in aircraft efficiency and the use of alternative fuels are expected to slow down the increase in emissions compared to the increase in aviation demand, but international goals on sustainability will still be hard to meet. The International Air Transport Association IATA aims to achieve carbon neutral-growth from 2020 onward and to reduce the carbon emissions in 2050 to 50% of what they were in 2005 [25]. Emission trading systems like CORSIA[1] and EU-ETS[2] are created to makes this carbon-neutral growth achievable, by offsetting the additional emissions of airlines compared to 2019. Such goals and systems make the need for airlines to reduce their emissions more urgent.

When looking more specifically to full cargo airlines, it is estimated that they contributed to 8% of all aviation CO_2 emissions in 2018 [19]. The world cargo forecast from Boeing shows an average yearly growth in cargo demand of 4.3% from the years 2009 to 2019 and an average growth of 4% per year is expected for the coming 20 years [3]. Just over 50% of all global cargo is carried on average by full freighters, with the other cargo carried in the belly of passenger aircraft. This represents the importance of the full freighter airlines in the air cargo industry. The routing of cargo aircraft is mostly focused on maximising profits or minimising the cost to operate the network, with no real attention being payed on how much different flight types add to emissions of the network. The research on the sustainability of air cargo networks is very limited, which means that there is a lot of room for improvement in this field. With most cargo airlines generally operating older, less efficient aircraft[3], making the routing of aircraft more efficient could be a relatively easy and low-cost way to reduce the aircraft emissions. This research aims to combine an emission model with an air cargo routing model to investigate the possibility of reducing the airlines emissions due to changes in the network.

This report is a Literature Review of the existing research on air cargo operations, aircraft scheduling models and aircraft emissions. The aim is to gain this knowledge for the master thesis, to get an overview of the research in these fields that has been performed and to identify a research gap in the literature. First, the research framework of the thesis is discussed in Chapter 2. This is followed by Chapter 3, which analyses the three modelling setups that are used to simulate aircraft scheduling problems, which is then applied to cargo airlines and cargo routing in Chapter 4. Chapter 5 reviews the literature on aircraft emission modelling and the different emission models that are available.

[1] https://www.icao.int/environmental-protection/CORSIA/Pages/default.aspx, accessed on 23-03-2021
[2] https://ec.europa.eu/clima/policies/ets_en, accessed on 23-03-2021
[3] https://simpleflying.com/cargo-operators-older-planes/, accessed on 23-03-2021

2

Research Framework

2.1. Problem Statement

The problem that is considered for this project is the routing of cargo and fleet for a full freighter cargo airline in a more sustainable way. This is a difficult combination, as reducing the emissions of a network will most likely go against the objective of the cargo airline which is to maximise the profit. In this problem it is assumed that nothing can be changed to the aircraft fleet and other ground operations to decrease emissions, and the solution should solely be focused on revisions of the network and flight schedule.

The fleet routing, flight scheduling and cargo routing problems are mostly focused at maximising the profit of the operations of the airline, by serving cargo requests between origin and destination airports. These requests do not have to be served by a single flight, but can also be transferred between aircraft to reach the destination in a sequence of one or more flights. By doing this, cargo demand on a less busy OD-pair can also be served without the need for a direct connection between these two airports. An existing network of the scheduled cargo carrier Cargolux will be taken as a reference, which restricts the fleet and airports that are served. The existing network can impose extra constraints like flying to some airports with a minimum frequency or aircraft beginning and ending their rotation at a hub.

By also introducing the aircraft emissions in the objective function, the schedule can be changed to minimise the overall objective. If certain types of flights contribute a lot to the networks emission, these might be cancelled and replaced by more sustainable routes. The cargo and fuel loads are also incorporated in the emission model to investigate their influence on the amount of greenhouse gasses emitted. A choice should be made how these emissions are integrated in the objective function and at which scale they are considered. A revised schedule could reduce the overall flight emissions, but increase the emissions around a specific airport due to a large number of flights. This trade-off should be considered to find how the sustainability of the network can be improved the most. Finally, the revisions to the schedule should not have a too large impact on the costs and profits of the network, to make sure it is financially feasible.

2.2. Research Questions

The main research question that is aimed to be answered is stated below:

> "How can the aircraft emissions of a scheduled full cargo airline be reduced by introducing an emission model to a MILP air cargo scheduling model, while also minimising the cost for the airline?"

The main research question can be be further detailed into lower level questions, which are stated below:

1. What should an air cargo scheduling model look like in order to minimise both network operating costs and the airline emissions?

    (a) What method should be used to model the air cargo scheduling model?

   (b) What input data should be used to realistically model the network of a scheduled full cargo airline like Cargolux?
   (c) How should an emission model be altered to implement it into a MILP cargo scheduling model?
   (d) How can the cost of aircraft emissions be compared to the operational costs and profits of the airline?
   (e) How can local airport-based emissions be compared to the global effects of greenhouse gasses?

2. How should the network and schedule of a cargo airline be changed to minimise the airline emissions?
   (a) How do factors like load factor, taxi time and flight distance influence fuel usage and aircraft emissions?
   (b) What types of flights in the flight network of a scheduled cargo airline like Cargolux contribute the most to the emissions of the airline?
   (c) What is the difference in emissions of a revised schedule that minimises emissions compared to an original schedule?
   (d) What are the extra costs that would arise from a more sustainable aircraft routing schedule?

2.3. Research Objective

The main research objective of this thesis is defined below:

> To investigate the effect of aircraft routing on the aircraft emissions of a scheduled full cargo airline like Cargolux, by implementing an aircraft emission model that is dependent on the route, aircraft type and load factor, into a MILP flight scheduling model which generates aircraft rotations to serve cargo requests, while both maximising the total operating profit and minimising the flight emissions of the airline.

The research objective can be split up in more detailed and tangible sub-goals. The first part of the project mainly is about rewriting an existing air cargo scheduling model from literature that can be used for this research. When this model can be used to realistically simulate the aircraft and cargo routing of a full freighter cargo airline like Cargolux, the goal is to add the emissions part. Existing models and theories will be used for the emissions modelling, which will be extensively discussed in this review. With the complete model working, experiments with different weights for the emission costs in the objective function are performed. The goal is to find which types of flights in the schedule are responsible for the most emission and to determine if and how the network can be adapted in order to reduce emissions. The operational costs of the revised network should be compared to an original network, to investigate the financial feasibility of this schedule redesign. This combination of objectives has not been researched yet and would create an useful addition to the existing literature in the air cargo scheduling field.

3

Airline Scheduling Problem

The airline planning problem consists of four sub-phases, which each deal with a specific part of the airline schedule [48]. The schedule design problem (SDP) considers which airports will be served and can determine a preliminary timetable. The SDP in this project is mostly focused at cargo scheduling and will be further discussed in Chapter 4. This chapter mainly focuses on the flight planning parts, consisting of the fleet assignment problem (FAP) and aircraft rotation problem (ARP). The overall goal of these problems is to assign aircraft to the available flights and create aircraft routings for each aircraft in the fleet. The fourth and final subproblem is the crew scheduling problem (CSP), which does not have priority in this project. A general overview of the objective function and constraints of the flight schedule planning is given in section 3.2, but first the different network designs are presented in section 3.1. Finally, section 3.3 compares the advantages and disadvantages of the different network types.

3.1. Network Setup

The three methods that are used the most to model an aircraft routing network are described in this section: First time-space networks are discussed (subsection 3.1.1), then connection-based networks (subsection 3.1.2) and finally subsection 3.1.3 discusses string-based networks [48]. An overview of the used literature for the different network types and the scheduling applications can be found in Table 3.1.

3.1.1. Time-Space Network

The time-space network was first used to model the fleet assignment problem by Hane et al. in 1995 [20], where an airport is represented by a single time line. A node on this time line is created for each arrival or departure and a flight is defined as a connection between a departure node and an arrival node on two different airports. A decision variable is used to designate that a fleet is assigned to a specific flight. The end of the time line is connected to the beginning to create a cycle, to for example model a daily or weekly schedule. An example for a time line of a single airport is shown in Figure 3.1 [20], where every flight to and from this airport is assigned to a specific fleet.

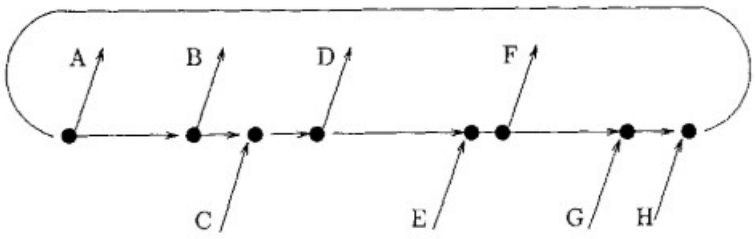

Figure 3.1: Time line of a single airport in a time-space network [20]

This method can also be used to model an ARP. Such a model often uses three types of arcs to connect the nodes: Flight arcs, ground arcs and wrap-around or overnight arcs [28, 31]. The flight arcs serve as a direct flight between two airports for a single aircraft. In this flight arc, the time that is needed for the turnaround at

the airport is added to the flight time, to ensure the connection to the next flight of the aircraft is feasible. A ground arc connects two time steps at the same airport and thus means that the aircraft remains unused on the airport for this time period. Multiple aircraft can stay on a ground arc at the same time. The overnight or wrap-around arc connects the last node of the planning period to the first time stamp of the new period. An example is shown in Figure 3.2a from [31], where two airports are connected by eight flight arcs and ground arcs connect the event nodes at the airport. The time increases to the right and the overnight arcs are used to create a daily schedule.

For problems with more flights, the time-space network can get large very fast. In order to reduce the number of nodes and ground arcs and thus to reduce the size of the problem, Hane et al. [20] introduced three pre-processing steps, namely node aggregation, island isolation and the elimination of missed connections. The technique of node aggregation is to combine nodes if the connections between two flights are still valid when the exact time of the event node changes. Two or more flights that arrive consecutively at an airport can share their arrival node, as these flights can still be connected to any flight departing from this node. The same can be done for consecutive departing flights and flights that subsequently arrive and depart from an airport, as described for an ARP by Sherali, Bish and Zhu [40]. Liang et al [31] shows a schematic representation of the node aggregation in Figure 3.2b, where the number of nodes is reduced from 16 to 7 nodes. For example, the first two departing flights from airport A can be merged into a single node, because no flights arrive or depart between those two flights and the conservation of aircraft is still valid. Also at station B the arrival node at 18:30 and the departure node at 19:00 are merged into an aggregated node, because it does not matter when for the model when exactly the arrival and departure take place, as long as the connection is valid.

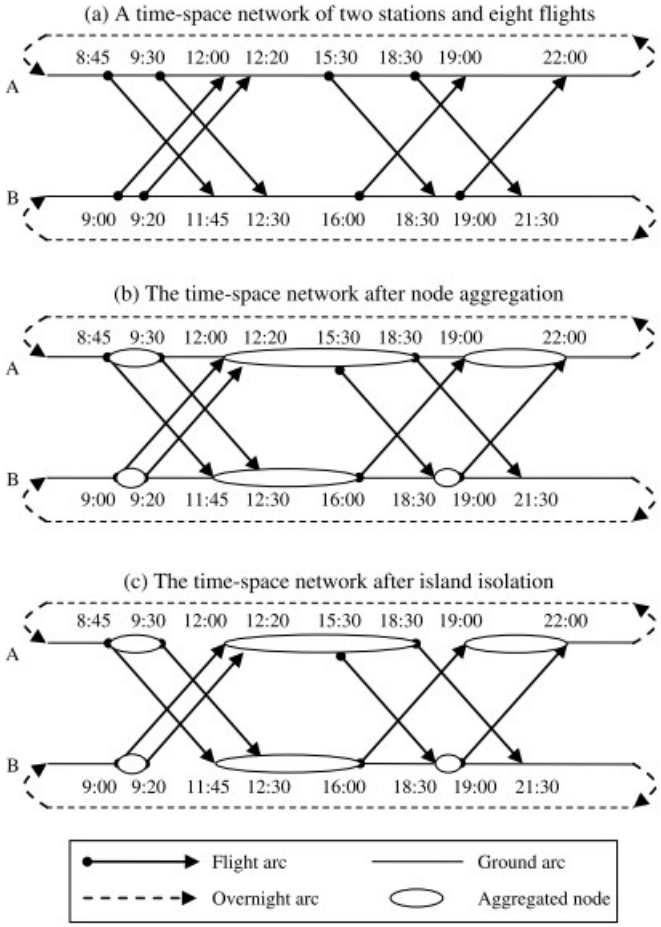

Figure 3.2: Time-space network with two airports and eight flights, with examples for node aggregation and island isolation [31]

An extra step of preprocessing that Hane et al. [20] propose is island isolation, which means to remove unused ground arcs and thus creating periods of time, 'islands', where one or more aircraft are on the ground

at that airport. This is mostly true for smaller airports with less traffic, where this might lead to a large reduction of arcs, for example in spoke airports of a hub-and-spoke network [40]. This is vizualised by Liang et al. [31] in Figure 3.2c, where the ground arcs are removed for the periods where no aircraft are present at the airports. This reduces the number of ground arcs from 9 to 6 for this specific case. For islands which consist of one arrival and one departure, the departure flight always has to be flown by the aircraft that also served the arriving flight, as no other aircraft are on the ground at that airport. Therefore, the decision on the aircraft that flies the outbound flight is already made and the decision variable now defines a flight path of one or more consecutive flights that are connected [20]. When considering different aircraft types with different turnaround times, it can occur that the time between the two connecting flights is too short to be served by a certain aircraft type with a long turnaround time, for example a wide-body aircraft. Hane et al. [20] and Sherali et al [40] describe that as a third preprocessing step these missed connections should be eliminated. The aircraft types that miss the connection cannot serve these flights and thus this set of flights does not have to bee considered in network of this specific aircraft.

Implementation of these preprocessing steps can aid in strongly reducing the problem size, as proven by Hane et al. [20]. For this study the first two preprocessing methods could be of use to decrease the problem size and computational time. During operations of a scheduled cargo carrier like Cargolux, a lot of time no aircraft are on the ground at smaller airports that are not a hub, which results in many empty ground arcs which can be removed by island isolation. Also combining consequent arrival and departure nodes at an airport into a single aggregated node is possible to further reduce the problem size. Removing missed connections will probably less effective due to the fairly homogeneous fleet of a carrier like Cargolux, which only uses Boeing 747s[1]. Although there are small differences in the versions of the aircraft, this most likely will not affect the turnaround time, making the removal of missed connections less likely.

Extra arc types can be added to deal with the uncertainty and reliability of aircraft, for example the reserve arcs as proposed by Burke et al. [5]. If an aircraft is assigned to such an arc, this means that it can be used as a back-up aircraft at this airport. The probability that an aircraft can depart on time is modelled using stochastic distributions and the spare aircraft that is scheduled on the reserve arc can be used to replace any delayed or cancelled flights. Adding aircraft reliability into the aircraft routing modelling lies outside the scope of this research and nominal operations of the flights will be assumed. Yan and Young [45] proposed arcs for an intermediate stop (one-stop arcs) as a replacement for non-stop flights and rental arcs to increase the capacity of a fleet. This is also deemed outside the scope of this research.

The combination of aircraft routing and air cargo scheduling is often performed using time-space networks and results in the implementation of extra arc types, like the no-service arcs and request arcs proposed by Delgado et al. [9] and delivery, holding and demand arcs in the study of Yan et al. [46]. This will further be discussed in Chapter 4.

3.1.2. Connection-based Network

Just like the time-space network, the connection-based network first was used for a fleet scheduling problem. In 1989, Abara [1] proposed this network type for the modelling of the American Airlines fleet schedule. The main difference to the time-space network is the fact that each airport now has a separate timeline for arrivals and departures. This is schematically shown in Figure 3.3 [48], showing two timelines per airport, which increase downwards.

The flight arcs or leg arcs still represent a flight between two airports, connecting a departure and arrival node. Connection arcs are used to connect an arrival event on the arrival timeline to the next departure event on the departure timeline. The wrap-around arcs connect the last event at an airport with the first event of the next day, meaning that aircraft which are allocated to such an arc stay the night at this airport [48]. The overnight can also be modelled using originating and terminating arcs, which represent the starting and ending situations at the airports [29]. Unlike for the time-space network, for the connection network a preliminary flight schedule should be available before assigning the flights and connections to specific aircraft. This can be a time-space network with some of the flight options filtered out, to give the model room to shift flights and create feasible connections. An advantage of the connection network over the time-space network

[1] https://www.airfleets.net/flottecie/Cargolux.htm, accessed on 29-03-2021

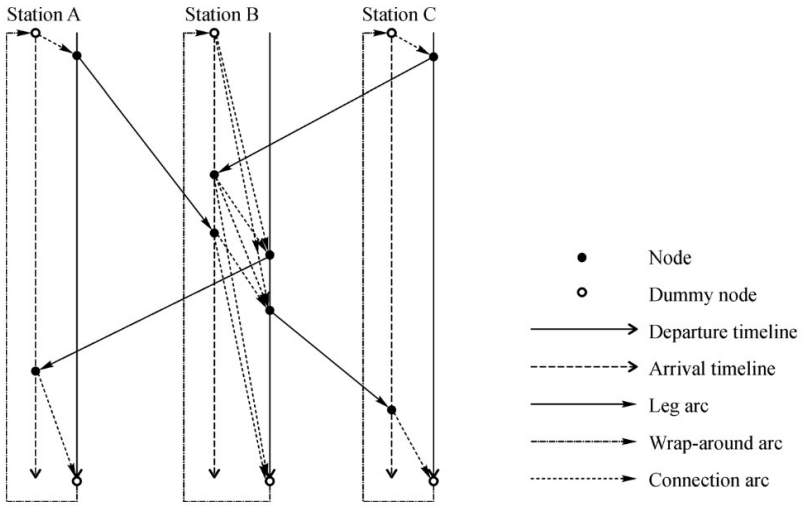

Figure 3.3: Connection-based network of 3 airports and 4 flights [48]

is the possibility to easily track the flights and flight path of specific aircraft, with arcs representing connections between two flights [38].

The connection network is often modelled for passenger airlines, where through-values are used to note if a flight with an intermediate stop appeals to passengers who otherwise would pick a direct flight. This means that the stop cannot add too much flight time and thus has to lie more or less on the route. Clarke et al. [7] propose that passengers rather stay in the aircraft on an intermediate stop than that they have to switch to another aircraft. Connections that have a ground time longer than 1.5 hours are given a penalty, as these flights are not attractive to passengers. Haouari, Shao and Sherali [21] give the through values a negative cost as extra revenue can be generated by these connections. For an air cargo network these connection times or transfers between aircraft are less important, as for cargo it does not matter what route it takes or how long a certain connection is. The only thing that is important is that the cargo is at its final destination before the due time. This makes penalties for long ground times or negative costs for good connections obsolete.

Due to the extra connection arcs that have to be modelled, the problem size increases faster than a time-space network as the number of legs increase [48]. To cope with this increased problem size, again preprocessing steps can be taken. Next to node aggregation, also arc aggregation is proposed by Clarke et al. [7]. A sequence of flights that is flown by the same aircraft can be aggregated into a super-arc, reducing the number of arcs and possibilities to make connections. This is shown in Figure 3.4a, with a network of 4 flights connecting a hub to a spoke airport. Flight 1 is an outgoing flight, which is connected to flight 2 to fly the aircraft back to the hub airport. By aggregating flight arcs 1 and 2, these two flights have to be served by the same aircraft and no extra attention has to be paid to the connection between these two flights at the spoke airport. The same has been done with flights 3 and 4. Figure 3.4b shows the node aggregation of the arrival and departure node of the hub airport, creating a singe super-node for all operations at the airport. This removes the ground arc that connected the two nodes. In this network, the only choices that have to be made is which sequence of flights each aircraft has to serve. For example after completion of flight 2, the aircraft can then be connected to either sequence 3-4 or again to sequence 1-2. The aircraft that performed sequence 3-4 in the first rotation will then be allocated to the other sequence. The goal of this preprocessing is to create a single super-node for each airport with available possibilities to connections. At a larger airport with more flights, this does require some extra steps. In Figure 3.4c, a schematic representation is showed, before the aggregation into a super-node. If ground arc y would be removed, this would give the possibility to connect the arriving flights 7 and 8 to the departure of flight 1. However, flight 1 has already departed once flights 7 and 8 arrive, meaning this is an impossible connection. Therefore, the through values are given a penalty of negative infinity to restrict this connection to be made [7].

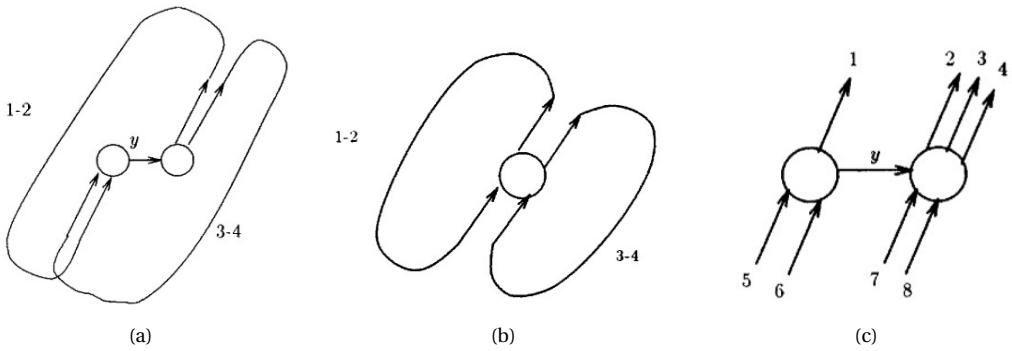

Figure 3.4: (a) Arc aggregation of two pairs of two flights. (b) Node aggregation of arrival and departure node. (c) Preprocessing a larger network. All images taken from Clarke et al. [7]

3.1.3. String-based Network

The third and final network type that is used in aircraft routing is the string-based network. A string is determined in 1974 by Pollack [37] as a sequence of flights served sequentially by a specific aircraft. Pollack proposes a method to generate strings from a flight schedule and determine the minimum aircraft fleet size to serve the schedule. The model assigns one flight per string, which is followed by the model trying to connect two strings together. This is only possible if the final arrival airport of the first string matches the first departure airport of the second string and if the departure time of the second string is later than the arrival time of the first. The objective of this model is to minimise the number of strings necessary to cover all the flights. The problem is solved using a heuristic that shifts departure times between the earliest and latest allowed departure times to decrease the number of strings necessary.

Barnhart et al. [2] are the first to implement the string network into an integrated aircraft fleeting and routing problem. They determine an augmented string as a sequence of linked flights, with a maintenance slot added at the end of the string. The objective is to allocate exactly one aircraft from the fleet to each flight segment, while minimising the cost of the sum of all strings. In order to model the string-based network, the base of either a time-space network or a connection-based network is used. Liang et al. [30] modify a time-space network to create flight strings, while Barnhart et al. [2] and Froyland, Maher and Wu [18] connect their flights using a connection-based network. Like described in subsection 3.1.1 and subsection 3.1.2, the advantage of a time-space network is that it contains less arcs and thus limits the problem size. However this does come at the cost of less extensive modelling possibilities and information about the connections in the created schedule. In order to also have these possibilities in a time-space network, Liang et al. propose to add penalty, zero cost and through arcs to have extra connection and maintenance modelling options [30]. To deal with the larger size of the string network problems, often column generation is used [2, 18]. In this technique the linear programming problem is divided into a master problem and multiple subproblems which are then solved iteratively. The integer programming problem is then solved using a branch-and-price algorithm [48].

The main use of the string-based model is to add the maintenance planning to the aircraft rotation schedule. The schedule is often designed for a longer time horizon, like a weekly maintenance schedule discussed by Liang et al. [30], and for requirements like equal fleet utilization over a longer period of time [2]. The only available use of the string network in cargo based operations is from Derigs and Friederichs [10]. The time-space network is extended to model the maintenance constraints by creating augmented flight strings as proposed by Barnhart et al. [2]. Adding maintenance constraints into this research will not have priority and will be first considered beyond the scope of this research. With maintenance constraints not essential for the project, the disadvantages of a string network, like a larger problem size and the use of column generation, quickly grow larger than the advantages of creating a more extensive solution.

3.2. Objective and Constraints for Aircraft Routing Problems

To get a better oversight of the way that aircraft routing and scheduling models are set up, this section discusses the general structure of such models.

3.2.1. Objective Function

In modelling of aircraft planning problems, a decision variable is used in the objective function to determine if a certain leg, connection or string is flown by a specific aircraft. For flight leg l in a set of all flight legs L and aircraft k in a set of K aircraft in the fleet, x_{lk} would be the decision variable that notes if flight l is operated by aircraft k in the time-space network. When modelling with connections, the decision variable indicates the connection between flight i and j, which are part of set L; x_{ijk}. The same reasoning is used with string-based networks to form x_{rk} for string $r \in R$ [48]. Equation 3.1 shows that the binary decision variable can only vary between 0 and 1, depending on the choice of flights for each aircraft.

$$x_{lk},\ x_{ijk} \text{ or } x_{rk} \quad \begin{cases} 1 & \text{if leg } l\text{, connection } ij \text{ or string } r \text{ is operated by aircraft } k \\ 0 & \text{else} \end{cases} \tag{3.1}$$

When purely looking at the aircraft rotation planning, often the objective of the model is to minimise the total schedule cost. In such cases already a preliminary flight schedule has been made, which consists of the flight legs that have to be flown and basic flight times. The main goal is to choose which leg, connection or string is flown by which aircraft in order to minimise the cost of the schedule, while every flight in the schedule is assigned. From Khaled et al. [28], an example of objective functions for a time-space network is shown in Equation 3.2. The cost c_{lk} can very per flight leg l and aircraft k, which gives the model the opportunity to determine the lowest schedule cost for the given flight schedule.

$$\text{Minimise} \qquad \sum_{l \in L} \sum_{k \in K} c_{lk} x_{lk} \tag{3.2}$$

If the fleet assignment problem has also been done, each flight leg has been allocated a specific aircraft type. This means that the cost for each flight leg is now c_l, with the aircraft type already fixed. Zhou et al. [48] give an example for a string based network where the schedule cost is minimised (Equation 3.3). The model has to choose between all strings, which are sequences of flight legs operated by the same aircraft type. The cost of each string r is now determined as c_r, which leaves the objective function with the simplified decision variable x_r.

$$\text{Minimise} \qquad \sum_{r \in R} c_r x_r \tag{3.3}$$

An objective proposed by Clarke et al. for a connection-based network is to maximise the sum of all through values [7]. As earlier described in subsection 3.1.2, the through value of a connection between two flight is dependent on multiple factors including connection time and how the flight with intermediate stop relates to the direct flight time. When the main interest of the model lies in creating a maintenance routing planning, the model can also be setup as a feasibility problem. This is done by simply setting the objective function to zero [48]. minimising the costs of the aircraft maintenance rotation is possible by adding these costs to the objective function. In air cargo scheduling, the objective can also include parts like no-service costs for not serving a certain request [9] or maximising the total revenue of the cargo routing part [10]. This is further discussed in Chapter 4.

3.2.2. Constraints

In order to come up with a feasible solution, constraints are used during the optimisation. The constraints require that variables meet certain conditions, which are needed to create a feasible and usable schedule. Some constraints that are used in a typical aircraft routing problem are explained below.

The cover constraint can be found in nearly every aircraft scheduling problem. This constraint makes sure that each leg, connection or string is covered exactly once. This prevents that multiple aircraft are assigned to the same part of the schedule. An example for such a constraint for a time-space network is shown in Equation 3.4 [48]. For each flight leg l, the sum of the decision variables x_{lk} of each aircraft k should be exactly 1, meaning that the leg is assigned to a specific aircraft (ARP) or aircraft type (FAP). Zhou et al. also show cover constraints for different flight networks [48]. Note that for this constraint, already a flight schedule should be available where it is required that all legs are flown. This is almost always the case for a scheduled passenger network, however this is not always valid for cargo networks, which are more dependent on individual requests.

$$\sum_{k \in K} x_{lk} = 1 \qquad \forall l \in L \tag{3.4}$$

A schedule also needs to guarantee that the flow of aircraft in and out of airport is in equilibrium. This is done using flow balance constraints or equipment continuity constraints, which make sure that the aircraft that are scheduled to depart from a certain airport also are located at that airport at the planned departure time [28]. For a time-space network this constraint can be checked by comparing the sum of aircraft entering and exiting a specific node, as shown in Equation 3.5 [48]. In this constraint L_{n+} and L_{n-} are the sets of legs respectively entering and exiting node n. To count the number of aircraft on the ground that enter and exit this node, y_{n+} and y_{n-} are used. The constraint now forces an equal aircraft flow balance for each node. Similar constraints can be created for connection- and string-based networks [48].

$$\sum_{l \in L_{n+}} x_{lk} + y_{n+} - \sum_{l \in L_{n-}} x_{lk} - y_{n-} = 0 \qquad \forall n \in N_k, \forall k \in K \tag{3.5}$$

The model cannot assign more aircraft in the schedule than the number of available aircraft in the fleet M_k. Therefore an aircraft count is implemented that assures that the sum of aircraft in the air and on the ground at a certain count time P is smaller or equal than the number of aircraft available. Zhou et al. [48] give examples for the three different model networks. A constraint for a time-space network of a fleet assignment problem is shown in Equation 3.6. For a connection based network no specific count time is used, but this is checked at the first leg of the daily route of all aircraft, as shown in Equation 3.7. The sum of all original flight legs should be smaller or equal than the number of available aircraft. Finally, Equation 3.8 shows an example for a constraint for an ARP of a string network. Because the string network often runs over a multiple day time period, variable γ_r is introduced that counts the number of times that string r crosses the count time P. Aircraft that are on the ground at the count are summed like in the time-space network.

$$\sum_{l \in L^p} x_{lk} + \sum_{n \in N_k^P} y_{n+} \leq M_k \qquad \forall k \in K \tag{3.6}$$

$$\sum_{l \in L} x_{0l} \leq M \tag{3.7}$$

$$\sum_{r \in R} \gamma_r x_r + \sum_{g \in G^p} y_g \leq M \tag{3.8}$$

Extra constraints can be implemented, like initial condition constraints that require an aircraft to depart its first flight from a specific airport, for example the hub of the airline [28]. Similar constraints can be added for the final flight of a day or of a time horizon. Airports can also introduce minimum or maximum frequency constraints, which require the schedule to serve a specific airport a number of times during the modelling horizon [10]. For non-homogeneous fleets constraints might be used to deal with the compatibility of some aircraft types to serve certain flight legs, for example due to the maximum range or runway length at an airport [10]. Turnaround time or turn time constraints can also be used to create feasible connections between two flights. If the time between the arrival and departure of two connecting flights is too short, there is no time for things like boarding and refuelling, making the connection infeasible [28]. In some models this turnaround time is already added flight time, which eliminates the need for an extra constraint. Furthermore, maintenance constraints can be introduced to force a maintenance check after a certain amount of flight hours, flight cycles or a specific elapsed time since the latest maintenance check. The aircraft can then be constrained to undergo maintenance at a certain airport, where extra constraints can be added for maximum maintenance capacity at this airport [48]. Adding extra constraints will make the solution more realistic, however this does increase the complexity of a problem. The objective for this research lies more on the combination of the air cargo scheduling model with an emission model, therefore first a working scheduling model will be implemented and merged with the emissions model before the possible addition of extra constraints. Constraints more specific for air cargo modelling will be further discussed in Chapter 4.

3.3. Comparison of network types

In this chapter the three main network setups for the modelling of aircraft routing problems are discussed to gain enough knowledge to decide which network would be suitable for this research. The used research is

summarized in Table 3.1, with distinctions for the three network types, namely time-space, connection-based and string-based networks, and whether the network is used to model a fleet scheduling, aircraft routing or air cargo scheduling problem. Some papers treat multiple applications or network types.

Table 3.1: Summary of used literature on FAP, ARP and Air Cargo Scheduling for three different network types

Time-Space Network		Fleet Scheduling	Aircraft Routing	Air Cargo Scheduling
Hane et al. (1995)	[20]	✓		
Yan and Young (1996)	[45]	✓		
Yan et al. (2006)	[46]		✓	✓
Burke et al. (2010)	[5]		✓	
Liang et al. (2011)	[31]		✓	
Derigs and Friederichs (2013)	[10]	✓	✓	✓
Khaled et al. (2018)	[28]		✓	
Delgado et al. (2020)	[9]		✓	✓
Zhou et al. (2020)	[48]	✓	✓	
Connection-Based Network				
Abara (1989)	[1]	✓		
Clarke et al. (1997)	[7]		✓	
Sherali et al. (2006)	[40]	✓		
Haouari et al. (2013)	[21]		✓	
Safaei and Jardine (2018)	[38]		✓	
Zhou et al. (2020)	[48]	✓	✓	
String-Based Network				
Pollack (1974)	[37]	✓		
Barnhart et al. (1998)	[2]	✓	✓	
Derigs and Friederichs & Schäfer (2009)	[12]		✓	✓
Derigs and Friederichs (2013)	[10]	✓	✓	✓
Liang et al. (2013)	[30]	✓	✓	
Froyland, Maher & Whu (2014)	[18]		✓	
Zhou et al. (2020)	[48]	✓	✓	

The basics of the model objective functions and constraints are quite similar when comparing the three network types. Constraints for the flight coverage, flow balance and the number of aircraft can be found in each network type. Most models aim to minimise the schedule cost, with exceptions like maintenance feasibility objectives for some string-based models and through value optimisation for some connection networks. The difference between the three network types is the way that the schedules are created. In a typical time-space network, event nodes are created for each timestep and connected by either flight arcs, ground arcs or overnight arcs. This creates a fairly simple network setup, which is convenient to model a complete flight schedule from scratch. A connection network needs more preprocessing and a preliminary set of flight times which can then be connected using connection arcs. The advantage of this method is that individual aircraft can be more easily followed throughout the flight schedule, which is not possible in the time-space network. A major disadvantage is the problem size of a connection network, which grows quickly when more flight legs or aircraft are available. This is also a problem for the string-based network, with the number of possible strings growing exponentially if new flight legs or aircraft are added to the problem. Therefore, string-based models always have to be solved using column generation. The main use and research performed using string-based networks is the modelling of aircraft maintenance routing, which is not essential for this research. Looking more closely at the research performed on connection-based networks, it can also be concluded that almost no papers exist that combine this network type with air cargo scheduling (Table 3.1). This is due to the fact that the path that cargo takes is less important than for passengers, as the main focus does not lie on short connection times but on the delivery of cargo before the due time. More research is available on the combination of time-space networks and routing for cargo airlines, as can be seen in Table 3.1. Therefore, this network type will be further discussed in combination with cargo routing and modelling in the next chapter.

4

Air Cargo Modelling

With the theoretical basis of the aircraft routing problem and parts of the fleet scheduling problem covered, this chapter focuses on the application of these techniques on air cargo modelling. These applications cover a combination of scheduling flights, routing aircraft and routing cargo. This research only focuses on full cargo operations, which means no bellyhold cargo in passenger aircraft is considered. Feng et al. [17] review the available literature of all air cargo operations, covering the perspective of airlines, airports and the freight forwarders. This can be used to get a wider overview of the whole air cargo industry. The number of researchers that really combine the cargo routing to the airline scheduling problems is fairly limited and can be found in Table 4.1. This chapter will discuss the different problems that were solved in these papers, first by reviewing the different model setups (section 4.1), followed by the objective function and most important constraints of the models in section 4.2.

4.1. Model Setup

In Table 4.1 it is clear that the most used network type in cargo modelling is the time-space network (TSN), which was earlier discussed in subsection 3.1.1. However, the first research that was published on air cargo scheduling by Marsten and Muller [33] uses a network type that is designed more like a connection network. They design a hub-and-spoke flight schedule for an overnight express cargo airline. For this purpose, all aircraft fly in the evening from their base airport to the hub, where cargo from all airports is sorted and put back onto the outbound flights to their base airports. By combining multiple stops on the inbound and outbound flights towards the hub, not every airport needs to be served by a single aircraft, but a spider graph network can be created. An example with a single hub and eight base airports is shown in Figure 4.1, where can be seen that multiple airports are connected by a single multi-stop flight to and from the hub. The model created by Marsten and Muller maximises the profit by combining the origin-destination (OD) pairs and assigning the most suitable aircraft type towards each flight. The model is extended for multiple hub operations, where paths are limited to at most two hubs per overnight schedule. Also a daytime system has been added, where the aircraft fly towards the hub in the morning and return in the evening before the overnight operations begin. This model is designed specifically for express delivery operations with one or two cycles per day and cargo that needs to be delivered the next day. This is different to the scheduled carrier airline network this research is most interested in, where operations are more continuous throughout the day, less focused on big hubs and with a more relaxed delivery time.

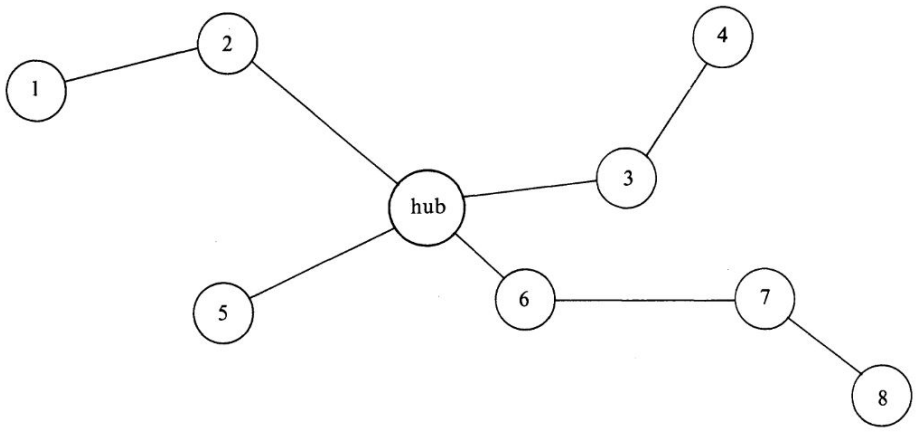

Figure 4.1: Spider graph for an express cargo airline with a single hub and eight base airports [33]

Yan et al. [46] present a method to model aircraft fleet routing and the determination of a timetable for a full freighter airline, based on a classic time-space network. Their model combines a fleet-flow and a cargo-flow time-space network to model both the fleeting problem and the cargo routing problem. The aircraft scheduling is done using the same three arc types as discussed in subsection 3.1.1, namely flight leg arcs, ground arcs and cycle arcs. For the cargo scheduling three new arc types are proposed: Delivery arcs, holding arcs and demand arcs. The delivery arcs represent cargo being allocated to a flight from one airport to another, similar to flight leg arcs. Costs for this arc type are determined by the handling cost, which is relative to the cargo weight and flight. The holding arc is similar to a ground arc, indicating that the cargo stays on the ground at this airport during a certain time. Holding costs can be induced for storing the cargo at an intermediate airport. These costs are mostly disregarded if the cargo is still on the origin airport or already at the destination airport. Finally, the demand arcs represent a negative cost if the cargo demand is served. The arc connects the destination to the origin if the model chooses to service this specific demand. The objective of the model is to minimise the network cost, which means that not all cargo demand needs to be transported if this does not bring any profit.

Figure 4.2 shows the three different arcs of the cargo routing model, with a timestep of 4 hours between each node. The cargo-flow network consists of multiple layers, each for a specific OD-time pair from the OD-table [46]. Yan and Chen [44] further elaborated on this paper by coordinating operations for freight airline alliances, which can further reduce operating cost and lead to higher profits. In this research only a single airline will be considered, however if the results are promising, it could be further examined how a cargo airline alliance could contribute to reducing emissions.

Two combinations of a string-based air cargo network are presented by Derigs, Friederichs and Schäfer [12] and Derigs and Friederichs [10]. Their first paper from 2009 creates a string network based on a connection network. In this model it is possible that a planner selects a set of optional flights, which will then be formed into feasible rotations. The aircraft rotation problem is then solved to connect sequences of flights together that are flown by the same aircraft. For this only the connections are of importance, not the exact departure or arrival times as described in the previous chapter. The model aims to minimise the number of aircraft needed to serve the selected flights. The cargo is then routed over the flights by creating cargo itineraries or strings which connect an OD pair. For larger problems not all OD-itineraries can be created and column generation has to be applied. The goal of the cargo routing is to maximise the profit of transporting the demand. A second option is also possible where the whole flight selection, aircraft rotation planning and cargo routing are done in a integrated model. The same steps as described before are taken, with column generation used to decrease the problem size and complexity of the problem [12].

Note that for this problem, a predefined set of mandatory and optional flights should already be available, which are than linked to create flight sequences. Furthermore, also bellyhold capacity of passenger aircraft is available in the model, which goes beyond the scope of this research.

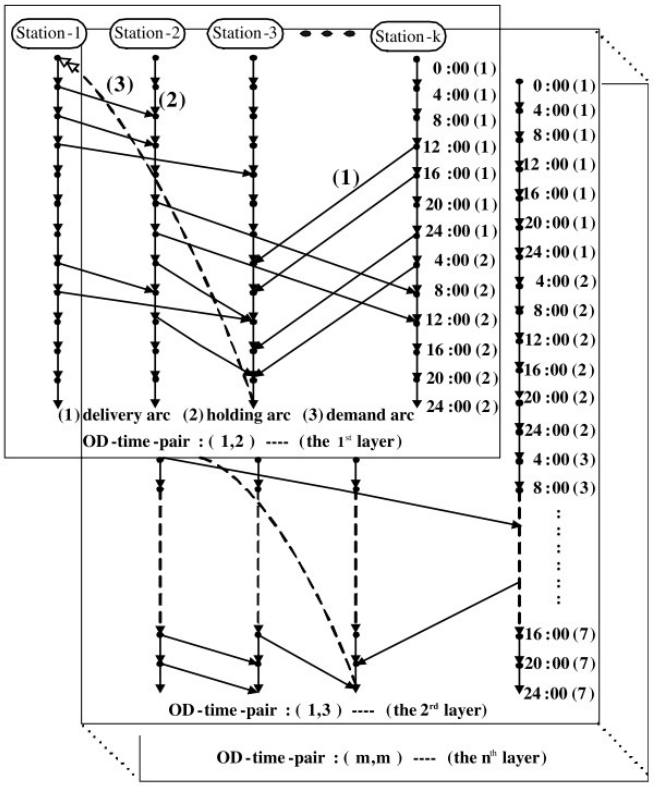

Figure 4.2: A graphical representation of the cargo-flow time-space of Yan et al. [46], showing delivery, holding and demand arcs.

In a paper from 2013, Derigs and Friederichs [10] present an extended version of the model described above. However, this model is based on a time-space or timeline network instead of modelling the connections between flights. A weekly schedule is created with timesteps of only 1 minute, creating a large number of possible routings. The nodes are connected with flight, ground and wrap-around arcs like in a typical time-space network. A string network is then created to include maintenance opportunities in the aircraft rotations [10]. The cargo routing problem is solved using the same OD-itineraries as described above [12]. Delgado and Mora [8] also solve the problem using a mix of a time-space and string network, in order to find a replacement schedule for demand disruptions in a short computational time. The flights are routed on flight and ground arcs, while cargo pickup and delivery nodes at an airport are grouped for a specific aircraft.

In 2020 Delgado et al. [9] implemented a pure time-space network to find an optimal schedule after demand disruptions. This network combines the flight and cargo schedules by implementing two types of nodes, namely itinerary nodes which represent all possible airport-time pairings, and request nodes which hold the origin and destination nodes of the cargo requests that are made. The itinerary nodes are connected by flight and ground arcs like in a normal time-space network, with request access and no-service arcs added for the cargo routing. A graphical representation is shown in Figure 4.3. The request access arcs connect the request node to the itinerary node, from where the cargo request can be routed on the same flight and ground arcs as the aircraft. The no-service arc directly connects the origin and destination request node and is activated if this cargo request is not served in the model, adding costs to the objective function for not receiving the fare of the customer.

The OD-pair for the request is randomly generated from a set of airports and like the weight, release and due time of the cargo. These values can vary between a selected range of cargo sizes and times in the time horizon. This information is stored in the request nodes [9]. The model uses a timestep of 20 minutes, to quickly find a replacement schedule for the 3 day time horizon. This is a larger timestep than the model of Derigs and Friederichs [10], who use a time discretisation of 1 minute, which gives a very detailed schedule of all flights. However, this does increase the problem size, as for the determination of a weekly schedule the time horizon consists of 10,080 time steps that have to be considered. Larger timesteps in the range of multiple hours can

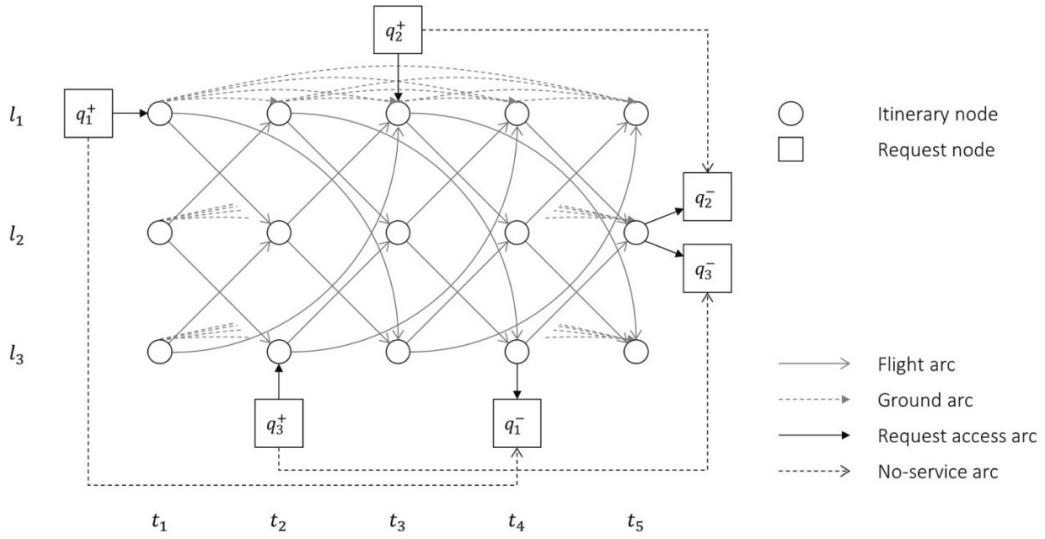

Figure 4.3: Example of the time-space network with three cargo requests [9].

also be used like Yan et al. [46], if a less detailed schedule is required.

4.2. Objective and Constraints for Cargo Scheduling

Just like for the models for aircraft routing discussed in the last chapter, this section will discuss the objective functions and constraints of the air cargo scheduling models.

4.2.1. Objective Function

Where the objective function of aircraft routing problems discussed in section 3.2 mostly consist just of a decision variable and a cost function for the flight, the objective functions for cargo scheduling problems often are more elaborate. Next to the flight costs, also costs and profits for carrying cargo are important. An example of such a function is given in Equation 4.1 [11, 12], with the first term representing the cargo profit which should be maximised. The second and third term should be minimised as they represent the total costs of the selected flights in the schedule and fixed aircraft costs which depend on the fleet size. Here the decision variable for cargo flow f_p over path p is multiplied by the profit margin m_p of the specific path that is flown. The costs c_f for all flights that are selected is determined by multiplying it with decision variable y_f. The number of aircraft needed to fly the schedule is equal to the number of rotations x_r, which is multiplied to the fixed costs per aircraft c^A. Minimising the aircraft fleet would be useful in case of aircraft leasing, however with a fixed fleet this does add value. The three different sub-problems can be solved in separate steps or in an integrated model, where all factors are combined in a single objective function. Derichs, Friederichs and Schäfer [12] and Derichs and Friederichs [10] present ways to take both approaches.

$$\text{Maximise} \qquad \sum_{p \in P} m_p f_p - \sum_{f \in F} c_f y_f - \sum_{r \in R} c^A x_r \tag{4.1}$$

The objective function can further be expanded, like the addition of a third term that subtracts the fixed costs of each aircraft in order to minimise the fleet size that is needed [12]. Yan et al. [46] also add costs for airports that are chosen in the schedule to compensate for the cost of ground infrastructure and personnel. Delgado et al. [9] introduce penalty cost for crew reassignment due to schedule disruptions. In other research crew scheduling or crew costs are mostly neglected. Changes to the flight schedule, like flight cancellations due to demand disruptions, are penalized by Delgado and Mora [8] in order to keep the original schedule intact the most as possible. They also propose a policy to split a cargo request and transport them on two different flights or itineraries, that could potentially increase the profits of the airline. However this is not preferred by the customer and leads to a more complex model [8].

The addition of an extra term or penalty function in the objective function could also be applied to implement

the emissions factor of each flight. A similar method is proposed by Derigs and Illing [11], who add a variable for extra EU ETS allowances that need to be bought. EU ETS is a European emission trading system that forces airlines to buy allowances if they emit more than the allowed amount of CO_2, this will further be discussed in Chapter 5[1]. In the model presented by Derigs and Illing only these extra costs for the whole airline are taken into account. However, there is no research in the emissions of specific flight legs, aircraft with different load factors or different parts of the flight, which is a gap that this research will focus on.

4.2.2. Constraints

For the aircraft scheduling part of the models mostly the same constraints are used, like the aircraft flow and balance constraints and the aircraft count constraints (section 3.2). The cover constraint is used less in the models using time-space networks, because often no preliminary flight schedule is available for the air cargo modelling, in contrast to the modelling of passenger flight schedules. The string-based networks do use a cover constraint to make sure every string is covered by an aircraft [10, 12].

The cargo routing problem is mostly constrained by the available capacity of each leg or aircraft. In some research this is only done by checking the cargo weight [8, 9, 33, 46], while others also include the cargo volume [10, 12]. Also the demand on a certain OD-pair should always be larger or equal than the actual cargo flow that is being transported over them. Equation 4.2, Equation 4.3 and Equation 4.4 show how Derigs and Friederichs set these constraints [10]. x_p^{flow} is the decision variable that states the weight of the cargo that is carried on cargo flow path p. P_l contains all cargo flow paths that are routed over flight leg l. The sum of the weight of the cargo routed over these paths should be smaller than the maximum weight capacity on leg l, given by $w_l^{\max}$. The same is done for the sum of volumes of the cargo path that should be smaller than the volume capacity $v_l^{\max}$. The set P_{od} holds all cargo flow paths that run between a certain OD-pair. The sum of the weight over all these paths should be smaller than the demand for this OD-pair.

$$\sum_{p \in P_l} x_p^{\text{flow}} \leq w_l^{\max} \qquad \forall l \in L \tag{4.2}$$

$$\sum_{p \in P_l} \text{vol}_p x_p^{\text{flow}} \leq v_l^{\max} \qquad \forall l \in L \tag{4.3}$$

$$\sum_{p \in P_{\text{od}}} x_p^{\text{flow}} \leq d_{\text{od}} \qquad \forall \text{od} \in OD \tag{4.4}$$

The constraints for aircraft flow balance can also be implemented for cargo flow, like shown by Delgado et al. [9] in Equation 4.5. In this case $\delta^+(i)$ and $\delta^-(i)$ are the sets of arcs that start and end at node i respectively and i_r^+ and i_r^- are the initial and final nodes of request r. If node i is the starting node of request r this should result in a positive flow of 1. Similarly the ending node acts as a sink resulting in a flow of -1. If the node is neither the start nor the end node of the request, the balance should be 0. This is because the cargo both enters and exits this node as an intermediate stop, or because the node is not used by the cargo request at all.

$$\sum_{(i,j) \in \delta^+(i)} q_{ij}^r - \sum_{(j,i) \in \delta^-(i)} q_{ji}^r = \begin{cases} 1, & i = i_r^+ \\ -1, & i = i_r^- \\ 0, & \text{else} \end{cases} \qquad i \in N, r \in R \tag{4.5}$$

If the model needs to be constrained more for the research purposes, more constraints can be added, like picking up the cargo only after it is available and delivering it before its due time [10]. Furthermore, only cargo transfers can take place at an airport with the right facilities [10] or a fixed ground time can be made mandatory for cargo to allow for cargo transfers [9]. A constraint can be made to fix the position of aircraft at the start and end of the time horizon [8]. Marsten and Muller [33] propose a fuel constraint that limits the total fuel usage of the whole fleet for each overnight operation. Yan et al. [46] use a constraint to determine if an airport is used and thus extra airport costs should be added. Finally, constraints can be used to make sure that cargo and aircraft cannot be assigned to multiple places at once [9], that capacity of airports or flight routes is not exceeded [46] and that cargo can only be routed over a flight arc if an aircraft is assigned to that arc. During the model specification it should be determined which constraints will be used and which will be relaxed to fit the schedule requirements.

[1] https://ec.europa.eu/clima/policies/ets_en, accessed on 23-03-2021

4.3. Solution Methods

The different models that are reviewed use different techniques to come to a solution, mainly due to the characteristics of the model, as also described in Table 4.1. The optimisation problem described in this chapter and Chapter 3 are called Mixed-Integer Programming models (MIP), with an objective function and one or more constraints that the model should adhere to. If such a model only consists of only linear parts, it is called a Mixed-Integer Linear Programming model (MILP). Models like this can either be solved by commercial solvers, like CPLEX and Gurobi, or by techniques such as column generation and heuristics.

A commercial solver often uses a technique called branch-and-bound to solve the MIP or MILP. In this method first the integrality constraints are relaxed, meaning that some variables do not are not fixed as integers anymore and making it a linear programming (LP) model. By then rounding the value of one of the variables up and down, two branches are created with new MIPs. The solutions for these problems create an upper and lower bound for the solution. The process is then repeated to decrease the gap between these two bounds, until the gap is zero and an optimal solution is found. Branches can also be fathomed, either if the MIP of this branch gives an infeasible solution or if all integrality restrictions are satisfied. If that is the case and the objective value found is lower than any other feasible solution that was found (for a minimisation problem), the value is kept until an even lower value is found. Next to branch-and-bound, solvers use other techniques, like presolving and cutting planes, to quickly arrive at or near the global optimum solution of the problem[2].

In the paper presented by Delgado et al. [9] the model is solved by the Gurobi solver. The result tables show how the solving process behaves for different network sizes. The number of airports is varied, together with the number of original, cancelled and new cargo requests. In the instances with a smaller number of airports and cargo requests the model can find an optimal solution within a few minutes, but this time increases strongly when the size of the network is increased. The number of variables and constraints rapidly grows for an increase in the network, which can result in very high computational times. This is graphically shown by Yan et al. [46] in Figure 4.4, by showing the possible options for a flight between an OD-pair. If no intermediate stops are allowed only four options are possible within this time frame, however the number of options very quickly increases if also stops at intermediate airports are allowed. The same growth in schedule complexity arises if the number of aircraft, airport or cargo requests increases, or if the time discretisation is reduced to a smaller timestep.

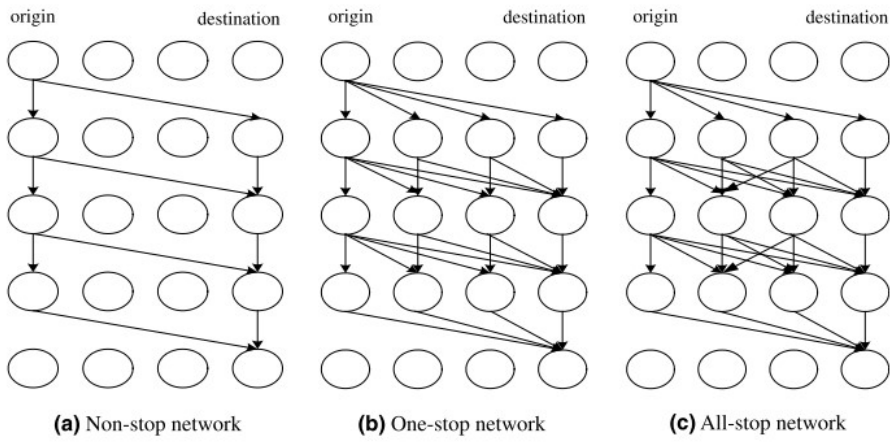

Figure 4.4: Graphical representation of the increase in complexity if multiple stops are allowed in a time-space network [46]

Techniques as heuristics and column generation are developed to be able to solve larger and more complex problems in a faster computational time. Heuristics do not necessarily find the global optimum solution of the problem, but try to make an approximation or find a local optimum. Such an approach is taken by Yan et al. [46] using the different stop strategies shown in Figure 4.4. When the commercial solver does not converge in the all-stop network within a certain computational time, modifications are done by simplifying parts of the network to non- or one-stop restrictions. If this still not converges, a series of network modifications are

[2] https://www.gurobi.com/resource/mip-basics/, accessed on 08-04-2021

proposed in the heuristic to find an approximation to the solution within a reasonable computational time.

Another method that is regularly used in aircraft and cargo routing problems is column generation. This technique is mostly used on problems with a very high number of columns or variables with respect to the other dimensions of the problem. The model is divided into a restricted master problem (RMP) and multiple subproblems. The RMP only contains a subset of the columns and is first solved. Then the model searches in the subproblems for columns that can improve the objective, which continues until no improvement is identified [12]. The column generation approach is mostly used for string-based networks, where the number of feasible strings increases dramatically due to an increase in for example flight legs and airports. Examples can be found in the research of Derigs, Friederichs and Schäfer [12] and Derigs and Friederichs [10]. Finally, a combination between heuristics and column generation is possible, for examle the matheuristic approach as proposed by Delgado and Mora [8]. The RMP is solved using column generation, while a heuristic is created to solve the subproblems. The problem is also solved using Gurobi's MIP solver, which had trouble finding optimality of the larger problems within the given run time.

A trade-off should be made to choose which solution method fits the research goals. If the network is not extremely big, or a larger time discretisation is used, the MILP can be solved within a reasonable computation time using a commercial solver. For larger problems or problems with non-linear terms, column generation or heuristics should be considered, although this will not directly lead to a global optimum solution. This does mean that extra focus and programming time will have to be invested in writing the column generation or heuristic model, while modelling a MILP for a commercial solver is usually easier.

Table 4.1: Summary of the available literature on flight scheduling and cargo modelling, with the key characteristics of each model.

Source	Year	Model Objective	Network Type	Output
Marsten and Muller [33]	1980	Maximising profit	CN Spider Graph	Flight schedule of OD pairs Fleet planning
	Mixed-integer programming model. Cargo transfer only possible at hubs.			
Yan et al. [46]	2006	Minimising system cost	TSN	Airport selection Aircraft fleet routing Fligth timetable
	Integrated scheduling model. Mixed-integer programming model in combination with heuristics. The allowed number of stops for the cargo routing is varied. Projected demand given, applied to a weekly time horizon.			
Yan and Chen [44]	2008	Minimising system cost	TSN	Airport selection Aircraft fleet routing Flight timetable
	Extended model from Yan et al. [46]. Coordinated flight scheduling for freight airline alliances.			
Derigs, Friederichs and Schäfer [12]	2009	Maximising profit	CN and SN	Flight selection Aircraft rotation planning Cargo routing
	Column generation with shortest path algorithms. Methods for both incremental and integrated planning. Modifying an initial flight schedule with mandatory and optional cargo flights. Also bellyhold capacity of passenger aircraft is available.			
Derigs and Friederichs [10]	2013	Maximising profit	TSN and SN	Flight selection Aircraft rotation planning Cargo routing
	Extended model from Derigs, Friederichs and Schäfer [12]. Maintenance modelling added with a string network.			
Derigs and Illing [11]	2013	Maximising profit	TSN and SN	Flight selection Aircraft rotation planning Cargo routing
	Extended model from Derigs and Friederichs [10]. EU ETS costs for CO_2 added to observe schedule changes.			
Delgado et al. [9]	2020	Minimise cost	TSN	Flight schedule Aircraft routing Cargo routing
	Redesign a flight schedule with MILP model to deal with demand disruptions. Three different crew management policies. Random demand generation with cargo request arcs.			
Delgado and Mora [8]	2021	Maximising profit	TSN and SN	Flight schedule Aircraft routing Cargo routing
	Adjusting base flight plan to demand disruptions. Matheuristic solution that solves problem using column generation. Cargo cannot be transferred between aircraft.			

5

Aircraft Emission Modelling

To goal of this research is to combine an air cargo scheduling model with an emission model to investigate how cargo networks would change if a more sustainable schedule would be used. Therefore, this chapter presents an overview of the different aspects of emission modelling that are relevant for this project. First, section 5.1 will briefly discuss the different emissions of air transportation that are relevant for sustainability. This is followed by a closer look on the emissions in two parts of the flight, namely the airport-based emissions from the LTO-cycle (section 5.2) and the full flight emissions which include the cruise phase (section 5.3). In section 5.4 existing emission models are discussed. Finally, in section 5.5 the research that combines air cargo networks with emissions is reviewed. Table 5.1 summarises the used literature in this chapter.

5.1. Aircraft Emissions

When discussing emissions from the aviation industry, mostly CO_2 is mentioned. However, during the combustion of fuel in a jet engine, more (by)products are produced, like H_2O, NO_x, CO, HC, SO_x and soot. These emissions are produced to a different extent in the different phases of the flight and also have other atmospheric and air quality effects. CO_2 and H_2O are products of the combustion and therefore can directly be connected to the amount of fuel consumed by the aircraft [36]. Per kg of kerosene that is burned around 3.15 kg of CO_2 is produced, which is mixed throughout the atmosphere during its long lifetime, meaning that its effect is similar when produced at ground level and at cruise altitude [26]. Because it is a greenhouse gas, it has a warming effect on the atmosphere. 1.26 kg of H_2O is produced per kg of fuel and it is also is defined as a greenhouse gas, but its direct contribution to warming up the atmosphere is fairly small [43]. However, it can lead to the formation of contrails if emitted at high altitudes, which can capture heat within the atmosphere [26]. Trade-offs can be made to fly lower in order to reduce the possibility of contrails, however this does increase the amount of fuel used, as the atmosphere is thicker. This trade-off of the trajectory optimisation is considered outside the scope of this research. The emission modelling of the cruise phase is further explained in section 5.3.

During the idle and taxi phase of the flight, the aircraft engines are running in very low thrust conditions, which is not where the engines are designed for. This inefficient operation results in incomplete combustion and much emissions of CO and HC (hydrocarbons) [36]. The take-off and climb-out phase requires totally different engine characteristics, namely very high thrust. In this complete and high temperature combustion, the CO and HC emissions are very low, however the production of NO_x, SO_x and soot is much larger [36]. These pollutants are mainly affecting the local air quality around airports and thus the health of people living nearby the airport. The modelling of these emissions is further explained in section 5.2. Next to local effects, nitrogen oxides also have an effect on the atmospheric temperature. Due to the emissions of NO_x at altitude, this stimulates the formation of ozone (O_3), which has warming effects on the atmosphere. However, O_3 also helps breaking down methane (CH_4) [26]. This also is a greenhouse gas that warms the atmosphere, and overall effect of NO_x can also slightly decrease the temperature effects of aviation. This cycle is less understood than for example the effect of CO_2 on global warming, therefore it might not have full priority during the emission modelling of the network.

5.2. Airport-based Emissions

The first part of the flight that is looked at is the landing and take-off cycle, also called LTO-cycle. This part combines all emissions of the flight operations that take place around the airport, while the aircraft is below 3000 ft (914 m) above ground level [15]. This contains take-off, climb-out, approach-landing and taxi and ground-idle, also graphically shown in Figure 5.1 [15]. For these phases, the ICAO has defined a standard cycle, which contains the length of each phase and the standard thrust setting of the engine. The emissions for these thrust settings can be found in the ICAO Aircraft Emissions Databank [13]. The high thrust settings during the take-off and climb that were discussed earlier can be found, just as the very low thrust setting during the taxi and idle phase. This phase is by far the longest part of the LTO-cycle and is based on average taxi times. However, more recent and accurate taxi times per airport are defined by Eurocontrol, which can be implemented to create a more realistic estimation of the emissions during the LTO-phase[1]. The European Environment Agency also implemented these data into a LTO emissions calculator[2], which outputs the fuel usage and different emissions for the different phases in the LTO phase for a specific aircraft and airport.

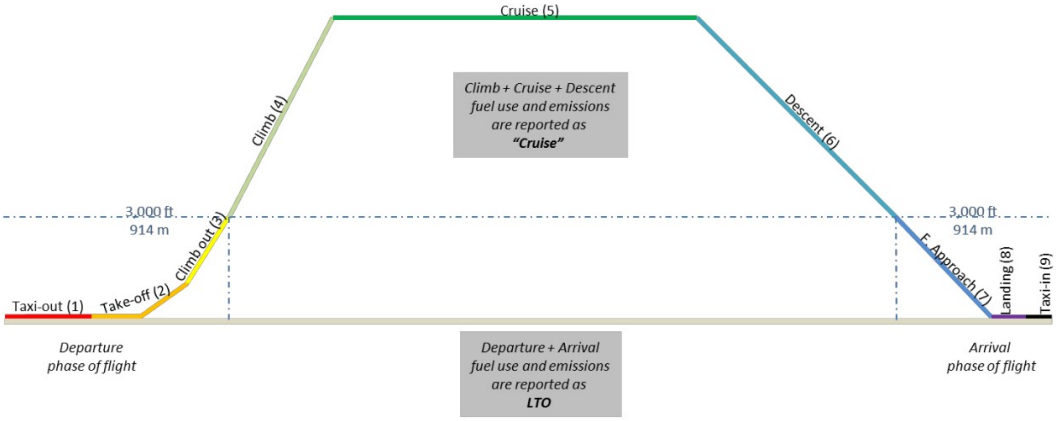

Figure 5.1: Flight phases that are grouped into the LTO-cycle and cruise phase [15]

Operating phase	Time-in-mode (minutes)	Thrust setting (percentage of rated thrust)
Approach	4.0	30
Taxi and ground idle	26	7.0 (in) 19.0 (out)	7
Take-off	0.7	100
Climb	2.2	85

Figure 5.2: Reference times and emissions for the standard ICAO LTO-cycle [23]

The importance of the reduction of the taxi phase is shown in multiple papers that investigate the local emissions at one or more specific airports. Kesgin [27] focuses on Turkish airports and varied the taxi times between 26 and 20 minutes. A decrease of 2 minutes was found to lead to a 6% decrease in LTO emissions, which was defined as the sum of HC, CO, NO_x and SO_2. Further reducing the taxi time to 20 minutes results in a total decrease of 16.5% in LTO emissions. Tokuslu [42] presents similar values for the LTO emissions (HC, CO and NO_x) at a Georgian airport, with a 2 minute shorter taxi time resulting in 5% less emissions. If the taxi times at specific airports in the network are significantly shorter, this might lead to the optimisation model preferring such airports over airports with long taxi times to reduce emissions. Next to the three pollutants

[1] https://www.eurocontrol.int/publication/taxi-times-summer-2019, accessed on 12-04-2021
[2] https://www.eea.europa.eu/publications/emep-eea-guidebook-2019/part-b-sectoral-guidance-chapters/1-energy/1-a-combustion/1-a-3-a-aviation-1-annex5-LTO/view, accessed on 14-04-2021

calculated by Tokuslu, Turgut and Rosen [43] also calculate H_2O and CO_2 emissions for multiple large airports, based on three different databases. For the LTO-cycle it was found that the ICAO database leads to the most accurate results. Large differences were found between the emissions of the different airports, mainly due to the different types of aircraft that operate at these airports. A heavy-wake category aircraft emits way pollutants per LTO-cycle and thus has a larger impact on the total amount of emissions at an airport than a smaller aircraft. Finally, Schürmann et al. [39] compared real measurements of emissions during idling and taxiing at Zurich airport to estimations done using the ICAO emission database. Differences up to a factor of 2 were found, with calculations showing lower values for some aircraft, but higher values for others. This shows that it should be taken into account that the values obtained from emission models do not perfectly match reality.

5.3. Flight Emissions

Next to only looking at the emissions on and close to the airport, also the emissions of the complete flight can be observed. The main pollutant that is researched for the cruise phase and the entire flight is CO_2, due to the less known climate effects of NO_x and H_2O. Because the CO_2 emissions are directly related to the fuel use, this is the main indicator on how much a flight emits. Miyoshi and Mason [34] separately calculate the emissions for cruise and LTO-cycle. The great circle distance is used to determine the flight distance, while the altitude was set to the most used cruise level of the aircraft type. 10 to 15 minutes were added to the cruise phase to compensate for the period from the end of the take-off phase at 3000 ft and the cruise level. Miyoshi and Mason also mention air traffic management inefficiencies, which can cause approximately 10% to the total flight times, however this extra time was not used in the calculations. In the analysis the emissions of complete flights are considered, which show that the carbon efficiency (g CO_2/ passenger km for a fixed load factor) increases for longer flight distances. This is due to the LTO-phase, which represent a large part of shorter flights as the take-off and climb takes a large amount of fuel to complete. Jardine [26] shows that this carbon efficiency slightly decreases for extremely long-haul flights, due to the amount of fuel that needs to be carried for the whole flight. This can be seen for the Boeing 747 and Airbus A340 in Figure 5.3. Also flights shorter than 2000 km become very inefficient for a Boeing 737.

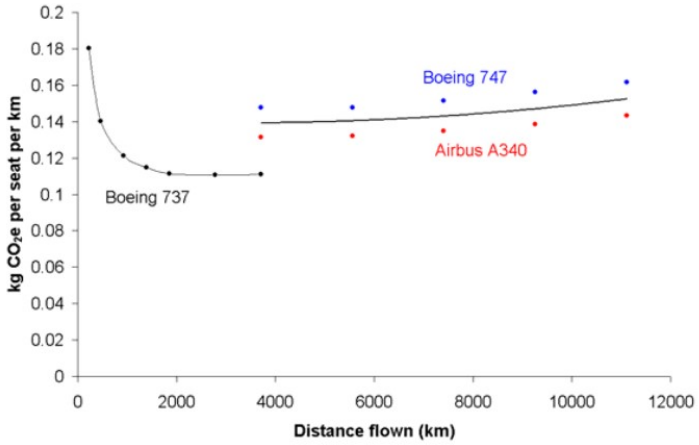

Figure 5.3: Carbon efficiency per passenger as a function of flight distance [26]

Loo et al. [32] use similar methods described above for LTO and cruise emissions, and investigate what effects hub-and-spoke networks have on the total emissions. It is found that on a global level, the introduction of hubs decreases the CO_2 emission per passenger km. However, the large aircraft that are used to connect international hubs leads to a high concentration of emissions at a local level. Also, capacity limits are proposed for hub airports to limit the extra emissions from aircraft flying in holding stacks and or having to taxi longer. Brueckner and Abreu [4] perform a statistical analysis on different factors that influence fuel usage and CO_2 emissions per available seat mile (ASM). For example aircraft fuel efficiency, seats per aircraft, flight length, load factor and average aircraft age. The load factor is interesting, as this is often kept constant in other papers. An increase of 5% in average load factor of an airline would lead to 8.2% more fuel use, due to the larger weight that has to be carried by the aircraft. It has to be noted that this represents a value for an entire airline and might be different for individual flights. Furthermore, also the average age has a large impact and a fleet

that is 3 years 'younger' would reduce the fuel use by 2.2%. This is not something that can be changed in the creation of a flight schedule, therefore aircraft age will not be considered.

5.4. Existing Emission Models

As mentioned in the sections above, almost all papers calculate LTO-emissions using the standard ICAO LTO-cycle and engine emissions database. This method is also implemented in the LTO-emissions calculator from the European Environment Agency[3], which can be used to find emissions for an aircraft type with a specific engine and at a certain airport. This calculator is also expanded to the calculation of the emissions of a complete flight[4] [15]. This model assumes the most used altitude and cruise velocity for a certain aircraft and flight distance in order to calculate the emissions of the flight. A similar model has been created by ICAO[5], which takes the number of passengers and the origin and destination of a flight as an input to calculate an average CO_2 consumption for such a flight. This calculator is designed to determine carbon emissions for emission offset programs and it is not suitable for cargo flights, however the documentation can be used as a reference further in the project [24]. Eurocontrol have developed the Advanced Emissions Model[6] which uses their aircraft performance model Base of Aircraft Data (BADA) [35]. Some researchers also combine two models, like Miyoshi and Mason who determine the total CO_2 emissions by using BADA inputs for the cruise phase and the EEA method for the LTO-cycle. Finally, within the TU Delft the performance model OpenAP is developed, which is an open-source Python toolkit that can be accessed via the GitHub[7]. In this performance model, also fuel use and aircraft emissions can be calculated [41]. This is mostly aimed at the cruise phase, which means the emissions from the LTO-cycle should separately be implemented, for example by using the EEA method.

5.5. Air Cargo Emissions

In the literature, not much research is available that combines air cargo modelling with aircraft emissions. Two papers that do incorporate emissions in their cargo models are Derigs and Illing [11] and Chao [6]. However, both of these papers are not interested in the emissions itself, but in the effects of the emission trading system EU ETS. This system is a so-called 'cap-and-trade', that aims to reduce the EU greenhouse gas emissions in 2030 by 55% and make the EU climate neutral in 2050[8]. Originally this system was only implemented for heavy energy industries, but since 2012 also the aviation industry is implemented. The airlines are capped to a maximum amount of greenhouse gasses that the whole aviation industry can emit, and this cap is reduced over the years to stimulate a decrease in emissions. A part of the allowances to emit greenhouse gasses are given to the airlines for free, based on benchmark ton-kilometers of each airline at a certain reference point, and another part is auctioned to airlines that exceed their allowed emissions. Airlines also can trade these allowances, if a certain airline emits less than it is allowed [14]. The system focuses mainly on CO_2, but the plan is to also cover other greenhouse gasses and, if the understanding is improved, the climate effects of contrails [16]. Only flights between two European airports are considered at this moment, but the goal is to also connect the trading system with other countries and systems, like the world-wide emission trading system CORSIA from ICAO[9].

Derigs and Illing [11] investigated the expected effect of EU ETS on the air cargo network that was proposed by Derigs and Friederichs [10]. The goal of the model is to optimise the profits of a cargo airline that operates globally. In the objective function the cost of buying extra allowances is implemented, which dependent on an estimation of the free allowances and on the total CO_2 emissions of the airline. The emissions are calculated based on the fuel used, emitting 2.5 kg CO_2 per litre of kerosene. A fixed fuel consumption per km is assumed for an empty aircraft without cargo, and variable fuel usage is added for each kg of cargo to compensate for the extra fuel that is needed for heavier aircraft. The necessary fuel load for the flight is calculated to find the weight before take-off, with a 10% fuel reserve included. With the allowance costs added, the network

[3] https://www.eea.europa.eu/publications/emep-eea-guidebook-2019/part-b-sectoral-guidance-chapters/1-energy/1-a-combustion/1-a-3-a-aviation-1-annex5-LTO/view, accessed on 14-04-2021
[4] https://www.eea.europa.eu/publications/emep-eea-guidebook-2019/part-b-sectoral-guidance-chapters/1-energy/1-a-combustion/1-a-3-a-aviation-1/view, accessed on 14-04-2021
[5] https://www.icao.int/environmental-protection/Carbonoffset/Pages/default.aspx
[6] https://www.eurocontrol.int/model/advanced-emission-model, accessed on 19-04-2021
[7] https://github.com/junzis/openap, accssed on 19-04-2021
[8] https://ec.europa.eu/clima/policies/ets_en, accessed on 20-04-2021
[9] https://www.icao.int/environmental-protection/CORSIA/Pages/default.aspx, accessed on 23-03-2021

is optimised. It is found that the network is still almost the same, only with some short inner-EU flights being replaced by longer flights to a destination outside the EU, because these flights do not belong to the EU ETS system. Also two more aggressive forms of the ETS are investigated, where no free allowances are given and where the CO_2 emissions are multiplied by 4 to deal with the extra environmental impact of other greenhouse gasses. Due to the system only working in the EU, this mainly results in the cargo network concentrating in a place outside the EU like Asia or America, with only a few connections back to Europe.

In the reseach presented by [6] the main focus lies on how different aircraft types are influenced by ETS regulations for different routes. The fuel use and emissions are calculated separately for the LTO-cycle and cruise phase, using the ICAO standard LTO-cycle and an average fuel consumption for the aircraft in the cruise. This means that unlike Derigs and Friederichs, no variable aircraft weight is used. For shorter routes it is found that the contribution of the LTO emissions is up to 20% of the total flight emissions, while this decreases to around 3% for extremely long flights. The corresponding carbon emission costs are calculated for four different ETS allowance scenarios which showed that costs per ton-kilometer are largest for small cargo aircraft, and smallest for large or new aircraft types. A cargo airline like Cargolux has little variation in their fleet, only using 747's, making this comparison less relevant for this research. However, it does point out that the LTO emissions can take up a large percentage of the flight emissions for shorter flights and thus lowering the carbon efficiency per km.

The only research that is available, which does not look at minimising ETS costs but CO_2 emissions is the research by Yan et al. [47]. They propose a multi-objective optimisation model to minimise the costs and carbon emissions of an air cargo alliance network. The emission weight per LTO cycle is calculated based on the aircraft type and the ICAO LTO-cycle. During the cruise phase a constant cruise speed and fuel usage per unit time is assumed. The cruise time is found by subtracting the time spent in the LTO phase of the complete flight time, which is then multiplied by fuel flow and the emission factor of CO_2, being 3.15 kg per kg of fuel. The model is also able to outsource the service of routes with low demand to airlines in the alliance, which effectively reduces the emissions of the airline itself. This outsourcing of flights is done more if the outsourcing costs are low and this then has a positive effect on the airlines emissions, however the emissions of the alliance airlines are not taken into account. In this master thesis, the airline will be restricted to its own operations and more focus will lay on redesigning the network of the airline itself in order to reduce emissions and still make a profit. The research of Yan et al. also does not take into account the difference in emissions due to different aircraft weights. It can however choose not to serve small demand pairs if an alliance partner can take over, where the model designed in the thesis project will probably have to deal with minimum frequencies per airport.

5.6. Relevant Literature

The research that is used for the review of aircraft emission modelling is summarised in Table 5.1. As said before, the emissions of cargo airlines are only researched in three papers, from which two mainly focus on the influence of the emission trading systems (Derigs and Illing (2013) [11], Chao (2014) [6]). However, these articles can still be used as good references on the effects of dealing with emission costs of emission trading systems on cargo airlines. The methodology of Yan, Zhang and Tang (2020) [47] is the most useful, as this paper directly relates the operational costs with the amount of CO_2 that is emitted. For the research on aircraft emissions, a clear distinction was found between emissions on an airport or local level and the emissions during the entire flight. In this proposed research, the focus mostly lies on the CO_2 emissions throughout the whole flight, which can be more easily compared globally for the complete airline than local emissions at different airports. Therefore, the papers that deal with simulation of the full flights emission are the main source of information. The different studies use similar methodologies that can also be implemented for this research. If the full flight emissions are fully implemented in the model, also the effects on the local emissions at different airports can be considered, however this does not have priority.

Table 5.1: Summary of the available literature on aircraft emission modelling for airport-base, flight-based and cargo airline emissions.

		Airport-based	***Full flight***	***Cargo Airline***
Kesgin (2005)	[27]	✓		
Jardine (2005)	[26]		✓	
Schürmann et al. (2007)	[39]	✓		
Turgut and Rosen (2009)	[43]	✓		
Miyoshi and Mason (2009)	[34]		✓	
Derigs and Illing (2013)	[11]		✓	✓
Chao (2014)	[6]		✓	✓
Loo et al. (2014)	[32]	✓	✓	
Pagoni and Psaraki (2014)	[36]	✓	✓	
Brueckner and Abreu (2017)	[4]		✓	
Tokuslu (2020)	[42]	✓		
Yan, Zhang and Tang (2020)	[47]		✓	✓

6

Conclusion

This literature study has summarised and reviewed the available research in the different fields of interest for this master thesis, namely improving the sustainability of air cargo networks. The aim of the thesis is to create a model that can optimise a cargo and aircraft routing schedule to both maximise the profits of the network, but at the same time minimise the emissions of the flights in the network. With this model it will be investigated if and how a network can be revised, without impacting the financial feasibility of the network.

Three different network setups for airline scheduling modelling are discussed, namely the time-space network (TSN), connection-based network (CN) and the string-based network (SN). The characteristics of each network like the objective function and constraints were analysed. It is found that a TSN is relatively easy to implement for a aircraft and cargo routing model, and that it is very suitable to design a flight schedule from scratch. For CN a preliminary flight schedule should be available, which is then used to connect the flights at an airport. The CN is often used for modelling networks of passenger airlines. The SN models allocate aircraft by choosing between all feasible strings in the network, which are sequences of flights that are linked together. To cope with the large problem size, for this mostly column generation is used. SN also is often focused on maintenance routing problems, which is not a priority for this research.

For integrated cargo and aircraft routing problems, the same three network types can be found, with the most literature using a time-space network set-up. Next to minimising the operational cost of the network, the models often also aim to maximise the profit that can be gained from serving requests. Differences were found in the solution techniques, ranging from commercial MIP solvers, column generation and heuristics.

Next to cargo and aircraft routing, also the literature on aircraft emission models has been reviewed. The emission models can be split up into two types: Local airport emissions and global emissions of the full flight. Methodologies for both emission scales are available, however the focus in this research will mainly lie on the CO_2 emissions of the entire flight, which can be compared the best for the whole airline network. The literature on emissions of air cargo networks are limited and mainly focuses on the effects of emission trading systems, however these papers can still be used as a reference during the thesis.

A research framework and general planning have been created to describe the goals and next steps of the thesis. The knowledge that has been gained from the literature will be used to investigate how the emissions of a cargo network can be reduced by adapting the aircraft and cargo routing problem.

III

Supporting work

A

Input Data

A.1. Available airports and flight data

Table A.1 displays all airports that are used in the experiments, with their associated IATA and ICAO codes. Also the average taxi-out and taxi-in times from the summer of 2019 are retrieved from Eurocontrol[1] and shown in the table.

Table A.1: Available airports, with the corresponding taxi times.

City	Airport	IATA code	ICAO code	Taxi-out [s]	Taxi-in [s]
Luxembourg	Luxembourg Airport	LUX	ELLX	624	258
Amsterdam	Amsterdam Airport Schiphol	AMS	EHAM	852	522
Prestwick	Glasgow Prestwick Airport	PIK	EGPK	540	246
Stansted	London Stansted Airport	STN	EGSS	930	396
Milan	Milan Malpensa Airport	MXP	LIMC	882	348
Vienna	Vienna International Airport	VIE	LOWW	672	384
Budapest	Budapest Ferenc Liszt International Airport	BUD	LHBP	678	325
New York	John F. Kennedy International Airport	JFK	KJFK	2,016	1,104
Miami	Miami International Airport	MIA	KMIA	1,212	504
Chicago	O'Hare International Airport	ORD	KORD	1,530	954
Atlanta	Hartsfield–Jackson Atlanta International Airport	ATL	KATL	1,224	744
Houston	George Bush Intercontinental Airport Houston	IAH	KIAH	1,266	546
Dallas	Dallas/Fort Worth International Airport	DFW	KDFW	1,134	660
Los Angeles	Los Angeles International Airport	LAX	KLAX	1,212	810
Seattle	Seattle-Tacoma International Airport	SEA	KSEA	1,188	480
Calgary	Calgary International Airport	YYC	CYYC	1,038	384
Guadalajara	Guadalajara Airport	GDL	MMGL	1,140	420
Mexico City	Mexico City International Airport	MEX	MMMX	1,740	732

[1] https://www.eurocontrol.int/publication/taxitimessummer2019, accessed on 12-04-2021

In the experiments performed in this study, the airports are grouped into three sub-networks, named EU, EU-NA and NA. The flight distance between each airport pair are assumed to be equal to the great-circle distance. This is calculated using Equation A.1, where the origin and destination airport are denoted by i and j. The earth radius R_e of 6371 km is used. The flight time is calculated using a constant cruise velocity of 900 km/h with the addition of a 30 minute LTO phase t_{LTO}.

$$dist_{ij} = 2 \cdot \arcsin\sqrt{\sin^2\left(\frac{lat(j)-lat(i)}{2}\right) + \cos\left(lat(j)\right)\cdot\cos\left(lat(i)\right)\cdot\sin^2\left(\frac{lon(j)-lon(i)}{2}\right)} \cdot R_e \tag{A.1}$$

$$t_{ij} = \frac{dist_{ij}}{V_{cruise}} + t_{LTO} \tag{A.2}$$

The flight distances and flight times of Network 1: EU are shown in Table A.2 and Table A.3. The flight times that are used in the TSN generation already have a turn-around-time of 1 hour added.

Table A.2: Flight distances for Network 1: EU in kilometres.

	LUX	**AMS**	**PIK**	**STN**	**MXP**	**VIE**	**BUD**
LUX	0	315	978	489	483	776	990
AMS	315	0	708	313	797	960	1169
PIK	978	708	0	513	1441	1662	1867
STN	489	313	513	0	932	1238	1451
MXP	483	797	1441	932	0	657	830
VIE	776	960	1662	1238	657	0	214
BUD	990	1169	1867	1451	830	214	0

Table A.3: Flight times for Network 1: EU in hours.

	LUX	**AMS**	**PIK**	**STN**	**MXP**	**VIE**	**BUD**
LUX	0	0.85	1.59	1.04	1.04	1.36	1.60
AMS	0.85	0	1.29	0.85	1.39	1.57	1.80
PIK	1.59	1.29	0	1.07	2.10	2.35	2.57
STN	1.04	0.85	1.07	0	1.54	1.88	2.11
MXP	1.04	1.39	2.10	1.54	0	1.23	1.42
VIE	1.36	1.57	2.35	1.88	1.23	0	0.74
BUD	1.60	1.80	2.57	2.11	1.42	0.74	0

The flight distances and flight times of Network 2: EU-NA are shown in Table A.4 and Table A.5.

Table A.4: Flight distances for Network 2: EU-NA in kilometres.

	LUX	**PIK**	**MXP**	**JFK**	**MIA**	**ORD**	**ATL**	**IAH**
LUX	0	978	483	6,053	7,612	6,854	7,274	8,276
PIK	978	0	1,441	5,160	6,784	5,906	6,375	7,347
MXP	483	1,441	0	6,412	7,919	7,260	7,634	8,659
JFK	6,053	5,160	6,412	0	1,757	1,188	1,222	2,277
MIA	7,612	6,784	7,919	1,757	0	1,930	960	1,550
ORD	6,854	5,906	7,260	1,188	1,930	0	976	1,490
ATL	7,274	6,375	7,634	1,222	960	976	0	1,107
IAH	8,276	7,347	8,659	2,277	1,550	1,490	1,107	0

Table A.5: Flight time for Network 2: EU-NA in hours.

	LUX	**PIK**	**MXP**	**JFK**	**MIA**	**ORD**	**ATL**	**IAH**
LUX	0	1.59	1.04	7.23	8.96	8.12	8.58	9.70
PIK	1.59	0	2.10	6.23	8.04	7.06	7.58	8.66
MXP	1.04	2.10	0	7.62	9.30	8.57	8.98	10.12
JFK	7.23	6.23	7.62	0	2.45	1.82	1.86	3.03
MIA	8.96	8.04	9.30	2.45	0	2.64	1.57	2.22
ORD	8.12	7.06	8.57	1.82	2.64	0	1.58	2.16
ATL	8.58	7.58	8.98	1.86	1.57	1.58	0	1.73
IAH	9.70	8.66	10.12	3.03	2.22	2.16	1.73	0

The flight distances and flight times of Network 3: NA are shown in Table A.6 and Table A.7.

Table A.6: Flight distances for Network 3: NA in kilometres.

	DFW	**LAX**	**SEA**	**YYC**	**GDL**	**MEX**
DFW	0	1,983	2,668	2,451	1,509	1,511
LAX	1,983	0	1,537	1,943	2,106	2,500
SEA	2,668	1,537	0	727	3,448	3,757
YYC	2,451	1,943	727	0	3,528	3,759
GDL	1,509	2,106	3,448	3,528	0	458
MEX	1,511	2,500	3,757	3,759	458	0

Table A.7: Flight times for Network 3: NA in hours.

	DFW	**LAX**	**SEA**	**YYC**	**GDL**	**MEX**
DFW	0	2.70	3.46	3.22	2.18	2.18
LAX	2.70	0	2.21	2.66	2.84	3.28
SEA	3.46	2.21	0	1.31	4.33	4.67
YYC	3.22	2.66	1.31	0	4.42	4.68
GDL	2.18	2.84	4.33	4.42	0	1.01
MEX	2.18	3.28	4.67	4.68	1.01	0

A.2. Demand and request generation

In order to simulate the cargo demand between each airport pair, publicly available flight frequencies are used. An example of the weekly flight frequencies are given for the NA network are given in Table A.8. The weekly cargo demand is calculated using the methods given in the Scientific Paper (Part I) and shown in Table A.9.

Table A.8: Weekly flight frequencies used as input for the demand generation for Network 3: NA.

		Destination					
		DFW	**LAX**	**SEA**	**YYC**	**GDL**	**MEX**
Origin	**DFW**	0	1	0	0	0	1
	LAX	0	0	4	2	0	2
	SEA	0	0	0	2	0	0
	YYC	0	1	2	0	0	0
	GDL	2	0	0	0	0	3
	MEX	1	3	0	0	4	0

Table A.9: Demand matrix with the weekly cargo demand generated for Network 3: NA in tonnes.

		Destination					
		DFW	**LAX**	**SEA**	**YYC**	**GDL**	**MEX**
Origin	**DFW**	0	50	20	16	26	46
	LAX	40	0	240	240	38	80
	SEA	0	48	0	80	0	32
	YYC	1	40	186	0	4	8
	GDL	98	106	32	16	0	148
	MEX	136	208	74	61	160	0

The demand matrix from Table A.9 is used to generate the set of cargo requests. The request generation function loops through each origin-destination pair in the demand matrix and creates the corresponding requests. The demand values are multiplied by $\frac{3}{7}$ to adapt them to the three day simulation period. If the demand between two airports is lower than the minimum request weight (set to 15,000 kg), no cargo request is generated. If the demand is larger, a request is created with a random cargo weight w_r between 15,000 and 30,000 kg. The due time t_r^- is set randomly between the start of day two ($t = 24$) and the end of the time horizon ($t = 72$). The release time t_r^+ has to be at least 24 hours before t_r^-. The lower bound of the release time is the maximum value of either 0 or 48 hours before the due time. All three variables are shown mathematically in Equation A.3, A.4 and A.5. Finally, the strategic weighting factor s_r is added, which is 1 for all requests, except if one of the airports is located in Europe and the other in North-America. In that case s_r is set to 1.5.

$$w_r = \text{random}\big(15000, 30000\big) \tag{A.3}$$

$$t_r^- = \text{random}\big(24, 72\big) \tag{A.4}$$

$$t_r^+ = \text{random}\big(\max(0, t_r^- - 48), t_r^- - 24\big) \tag{A.5}$$

After the request is generated, the sum of all request weights is calculated. As long as this value is smaller than the 3-day cargo demand between the two airports minus the minimum cargo weight of 15,000 kg, new requests are added. If this is not the case, the function moves on to the next origin-destination pair. In order to vary the total number of requests for a certain experiment, the weekly demand matrix can be scaled up or down by multiplying all values by a certain percentage. By doing this, the total cargo volumes can be changed, while keeping the overall demand distribution the same. An example of the corresponding request set that is created with 100% of the cargo demand of Table A.9 is given in Table A.10.

Table A.10: Request set generated using 100% of the cargo demand of Network 3: NA.

Request	Origin	Destination	w_r [kg]	t_r^+ [hour]	t_r^- [hour]	s_r [-]
0	GDL	MEX	15223	24	72	1
1	GDL	MEX	18441	3	30	1
2	GDL	MEX	25435	24	63	1
3	GDL	DFW	29827	3	27	1
4	GDL	LAX	24917	3	27	1
5	GDL	LAX	24086	0	24	1
6	MEX	GDL	25554	15	39	1
7	MEX	GDL	23243	12	39	1
8	MEX	GDL	27903	27	69	1
9	MEX	DFW	15073	0	33	1
10	MEX	DFW	22899	30	57	1
11	MEX	DFW	21897	18	45	1
12	MEX	LAX	21336	15	45	1
13	MEX	LAX	19347	3	45	1
14	MEX	LAX	25106	18	60	1
15	MEX	LAX	15262	3	27	1
16	MEX	SEA	29544	6	36	1
17	MEX	YYC	27611	18	45	1
18	DFW	MEX	21700	0	33	1
19	DFW	LAX	16103	36	66	1
20	LAX	GDL	20329	27	69	1
21	LAX	MEX	23941	6	33	1
22	LAX	DFW	17442	27	72	1
23	LAX	SEA	17409	0	24	1
24	LAX	SEA	19467	6	48	1
25	LAX	SEA	26024	21	54	1
26	LAX	SEA	21382	27	57	1
27	LAX	SEA	21954	48	72	1
28	LAX	YYC	19784	18	45	1
29	LAX	YYC	16976	0	39	1
30	LAX	YYC	27572	0	24	1
31	LAX	YYC	22091	3	27	1
32	LAX	YYC	23258	30	69	1
33	SEA	LAX	17715	6	42	1
34	SEA	YYC	28194	33	63	1
35	YYC	LAX	17589	3	51	1
36	YYC	SEA	27104	15	54	1
37	YYC	SEA	29861	0	24	1
38	YYC	SEA	25466	39	66	1

B

Emission Model

B.1. Calculating aircraft mass and emissions

The pseudo-code of Algorithm 1 shows the steps that are used to determine the fuel flow throughout the flight, iterate to the aircraft take-off weight and calculate the CO_2 emissions of the whole flight.

Algorithm 1 Calculation of fuel mass and emissions for flight ij and aircraft k

Input Flight time t_{ij}, Load factor LF, trajectory data: {airspeed v, vertical velocity vs, altitude h}, operational empty weight OEW, cargo capacity Cap_{max}, OpenAP fuel flow function $FuelFlow$, fuel flow estimate FF_{est}, timestep T
$m_{Fuel\ est} = FF_{est} \cdot t_{ij}$ // Estimated fuel weight
while $Final\ weight = False$ **do**
  $m_{AC}(0) = OEW + LF \cdot Cap_{max} + m_{Fuel\ est}$ // Aircraft weight at take-off ($t = 0$)
  for $t \in \{0, \frac{t_{ij}}{T}\}$ **do**
    $\theta(t) = \arcsin\left(\frac{vs(t)}{v(t)}\right)$ // Flight path angle at time t
    $FF(t) = FuelFlow\Big(m_{AC}(t), v(t), h(t), \theta(t)\Big)$ // Fuel flow per second at time t
    $m_{AC}(t+1) = m_{AC}(t) - FF(t) \cdot T$ // Aircraft weight at next timestep $t+1$
  end for
  $m_{Fuel} = \sum_{t=0}^{t_{ij}/T} FF(t) \cdot T$ // Actual fuel weight
  $gap = \dfrac{m_{Fuel} - m_{Fuel\ est}}{m_{Fuel\ est}}$ // Optimality gap between estimate and actual fuel weight
  if $|gap| < 0.02$ **then**
    $Final\ weight = True$
  else
    $m_{Fuel\ est} = m_{Fuel\ est} + \dfrac{m_{Fuel} - m_{Fuel\ est}}{2}$
  end if
end while
$m_{TO} = OEW + LF \cdot Cap_{max} + m_{Fuel\ est}$ // Take-off weight
$CO_{2\ en-route} = 0$
for $t \in \{0, \frac{t_{ij}}{T}\}$ **do**
  if $h > 3000$ft **then**
    $CO_{2\ en-route} = CO_{2\ en-route} + FF(t) \cdot T \cdot 3.149$ // En-route CO_2 emissions
  end if
end for
return $m_{TO}, m_{Fuel}, CO_{2\ en-route}$
Output Take-off weight, fuel weight and en-route CO_2 emissions

The fuel is first estimated using a constant fuel flow FF_{est} of 3 kg/s, which is multiplied to the flight time to find $m_{fuel\ est}$. With this initial guess, the aircraft take-off weight $m_{AC}(0)$ can be determined by adding the operational empty weight and the payload. The FuelFlow function of OpenAP to find the necessary fuel flow

to fly at the altitude, velocity and vertical velocity that is provided for each time step by the OpenAP TrajGen function. The aircraft mass is updated during each time step to compensate for the burned fuel. The fuel flow for each time step in the flight is summed to find a closer estimate of the fuel that is needed for the flight. If the gap between $m_{Fuel\ est}$ and the newly calculated m_{Fuel} is larger than 2%, the mean of both values is taken and used in the next iteration to find a closer estimate of the fuel needed during the flight. If the difference is smaller, the estimate is assumed good enough and used to determine the take-off weight. This fuel weight is used to determine the CO_2 emissions for the en-route phase, being the part of the flight that takes place above 3000 ft. The emissions are summed to the LTO-emissions to find the emissions of the total flight.

Because the aircraft take-off mass is calculated using an iterative process, this can be plotted in a graph. The results of a short range (LUX-PIK) and a long range (LUX-ATL) flight are plotted in Figure B.1. For the short flight, three iterations are needed (Figure B.1(a)). The blue line shows that the initial guess of a constant 3 kg kerosene per second overestimates the actual fuel that is used. The climb phase can be recognised from the aircraft mass quickly decreasing in the beginning of the flight due to the high thrust needed. The final part of the flight shows the descent, which requires much lower fuel flow and thus a slower decrease in aircraft mass. Five iterations are needed for the LUX-ATL flight, with the initial guess underestimating the actual fuel use (Figure B.1(b)). The cruise part of the flight is much larger than in the short-haul flight, with the climb and descent only taking up a small portion of the flight time.

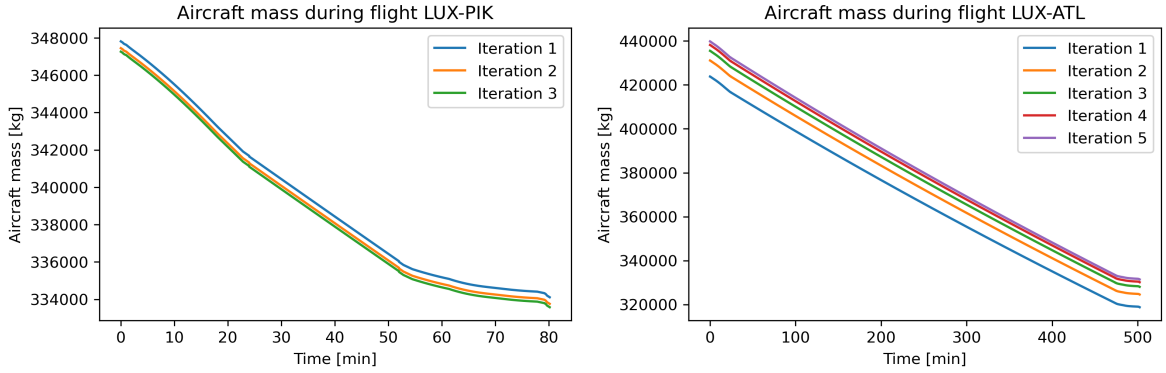

(a) LUX-PIK flight

(b) LUX-ATL flight

Figure B.1: Iteration to find the aircraft take-off mass for two flights, using a B747-8F with a load factor of 1. The figures show how the aircraft mass changes during the flight.

B.2. LTO emissions validation

During the validation of the LTO emissions, some interesting differences appeared between the results from the OpenAP model used in this research and the existing EMEP/EEA LTO-emissions calculator. For the taxi and idling phases, OpenAP found lower emissions than the EEA model, while the other three phases of the ICAO LTO-cycle (take-off, climb and approach/landing) resulted in higher emissions for OpenAP. It was found that these differences are caused by two factors: The approximation of fuel flow in OpenAP and the airport altitude that is used.

The ICAO emissions databank provides the fuel flow for the four thrust settings of the ICAO LTO-cylce, which are plotted in B.2(a). In order to also calculate the fuel flow and emissions at intermediate values, OpenAP approximates the fuel flow by applying a third degree polynomial fit, shown by the blue line in the figure. Next to that, the airport altitude is added in the OpenAP model, which is shown for LUX in the figure. This difference of 350 meters has some impact on the fuel flow needed to fly the same thrust. This causes an over-estimation of the ICAO emission databank values for the take-off (100%), climb (85%) and approach/landing (30%) phases. However, the fuel flow during the taxi phase is expected to be lower than the ICAO emission databank value.

The cumulative CO_2 emissions of a complete LTO-cycle at LUX are shown in B.2(b). The black line indicates

the emissions calculated with the EEA model using the ICAO emission databank values. A gap with the blue and orange lines of the OpenAP CO_2 emissions grows during the taxi-out phase, which correlates to the different fuel flow that is used in the two models. During the take-off, climb and approach/landing phases the gap with the blue line, indicating the OpenAP model at sea level stays relatively constant. This shows that the regular polynomial fit is accurate for these three phases. The gap with the orange line that indicates the OpenAP model at LUX altitude decreases, which could be expected from the higher fuel flows found in B.2(a). This compensates for the lower fuel flow during the taxiing phase. Overall, this analysis shows the origins of the inaccuracies of the emissions during the LTO-cycle.

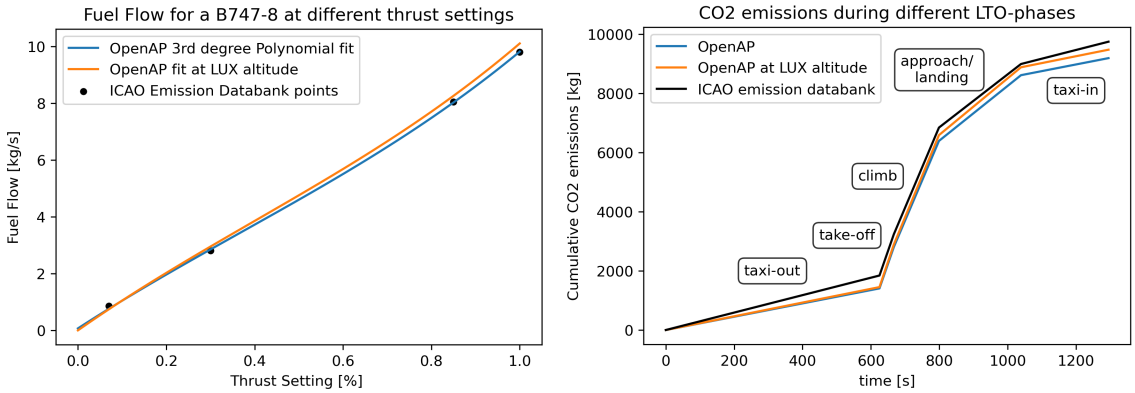

(a) Fuel flow at different thrust settings

(b) CO_2 emissions during a LTO-phase

Figure B.2: Analysis of OpenAP difference in LTO calculations.

B.3. Emission matrix

Finally, Table B.1 shows an example of an emission matrix that is generated using the emission model. The matrix of three flights of a Boeing 747-8F are shown: LUX-PIK, LUX-ATL and MXP-IAH. For each flight the carbon emissions of the full flight (CO_2_tot) are given for 11 different load factors. It can be observed that the CO_2 emissions during the LTO-cycle are constant for the different load factors, as this part of the flight is calculated independently of the aircraft weight. The CO_2 emissions during cruise and the fuel on-board increase for higher load factors. Comparing the different flights, it can be seen for that the two transatlantic flights emit a lot more CO_2 than the short European flight. The CO_2 emitted during the LTO-cycle of the LUX-ATL is higher than the LUX-PIK flight, due to the longer taxi time at the ATL, compared to the relatively small airport at Glasgow Prestwick (PIK). The LTO emissions for the third flight are even larger, due to long taxi-out times at MXP compared to LUX. These taxi times can be found in Table A.1. The flight distance between Milan Malpensa airport (MXP) and IAH is too large for the B747-8F to fly with a full payload capacity. Therefore, the intermediate values are evenly spaced between 0 and the maximum load factor for the flight ($LF = 0.93$).

Table B.1: Example of an emission matrix of a Boeing 747-8F for three origin-destination pairs.

index	orig	dest	LF	CO2_tot [kg]	CO2_ LTO [kg]	CO2_cruise [kg]	fuel [kg]	distance [km]
0	LUX	PIK	0	40,881	8,833	32,048	13,368	978
1	LUX	PIK	0.1	41,506	8,833	32,673	13,707	978
2	LUX	PIK	0.2	42,142	8,833	33,309	13,932	978
3	LUX	PIK	0.3	42,687	8,833	33,854	14,135	978
4	LUX	PIK	0.4	43,289	8,833	34,457	14,354	978
5	LUX	PIK	0.5	43,924	8,833	35,091	14,590	978
6	LUX	PIK	0.6	44,706	8,833	35,873	14,849	978
7	LUX	PIK	0.7	45,721	8,833	36,888	15,269	978
8	LUX	PIK	0.8	46,780	8,833	37,947	15,608	978
9	LUX	PIK	0.9	47,886	8,833	39,053	15,835	978
10	LUX	PIK	1	49,055	8,833	40,222	16,292	978
11	LUX	ATL	0	262,248	9,995	252,253	84,866	7,274
12	LUX	ATL	0.1	268,952	9,995	258,957	86,706	7,274
13	LUX	ATL	0.2	276,333	9,995	266,338	89,356	7,274
14	LUX	ATL	0.3	283,871	9,995	273,877	91,639	7,274
15	LUX	ATL	0.4	291,147	9,995	281,153	92,836	7,274
16	LUX	ATL	0.5	298,638	9,995	288,643	93,848	7,274
17	LUX	ATL	0.6	307,278	9,995	297,284	96,297	7,274
18	LUX	ATL	0.7	316,871	9,995	306,876	99,423	7,274
19	LUX	ATL	0.8	327,088	9,995	317,093	102,955	7,274
20	LUX	ATL	0.9	337,257	9,995	327,262	105,804	7,274
21	LUX	ATL	1	347,899	9,995	337,905	108,841	7,274
22	MXP	IAH	0	315,223	10,089	305,134	101,932	8,659
23	MXP	IAH	0.09	323,543	10,089	313,454	105,054	8,659
24	MXP	IAH	0.19	331,522	10,089	321,434	106,796	8,659
25	MXP	IAH	0.28	340,349	10,089	330,260	109,725	8,659
26	MXP	IAH	0.37	347,800	10,089	337,712	109,725	8,659
27	MXP	IAH	0.47	356,613	10,089	346,524	111,485	8,659
28	MXP	IAH	0.56	366,636	10,089	356,548	114,604	8,659
29	MXP	IAH	0.65	377,626	10,089	367,537	118,353	8,659
30	MXP	IAH	0.74	389,347	10,089	379,258	122,442	8,659
31	MXP	IAH	0.84	400,837	10,089	390,748	125,665	8,659
32	MXP	IAH	0.93	413,755	10,089	403,666	130,276	8,659

C

TSN Pre-processing

The time-space network that is used in this research has been described in the Scientific Paper in Part I. A pre-processing step that is applied after the generation of the TSN network is described in this appendix. A separate set of nodes and arcs is created for each aircraft k in the fleet $\mathcal{K}$. Each aircraft is assigned an origin and final airport, where it has to be at the start and end of the time horizon. Due to its origin and final airport requirements, the aircraft cannot reach some nodes in the beginning and end of the time horizon, making these nodes and adjacent arcs infeasible. By removing these arcs and nodes during pre-processing, the problem space is reduced. A limitation of the time-space network is that the problem size increases very fast for larger networks. Therefore, reducing the problem space in a pre-processing step can help to speed up the computational time of the optimisation.

A similar preprocessing technique is applied for each cargo request. First, all nodes and adjacent nodes with a timestamp before the release time or after the due time of the request are removed. Because the requests are only available for a period between 24 and 48 hours, this quickly reduces a large part of the problem size. An example is shown in Figure C.2, where the infeasible nodes and arcs of the network in Figure C.1 are removed. Request r_1 is available from timestamp 1 at Airport 2, which makes the nodes associated with the other two airports at time-stamp 1 impossible to reach. Therefore, these two nodes and all flight arcs and ground arcs starting from these nodes can be discarded from the TSN. At the end of the time period even more nodes and arcs can be removed, due to the cargo having to be delivered at time-stamp 4. Immediately all nodes and flights ending at time-stamp 5 can be taken out of the network. Finally, all nodes from which it is not possible to reach airport 3 in time are removed, together with all arcs that start and end in these nodes. This only leaves the relevant nodes and arcs for cargo request r_1.

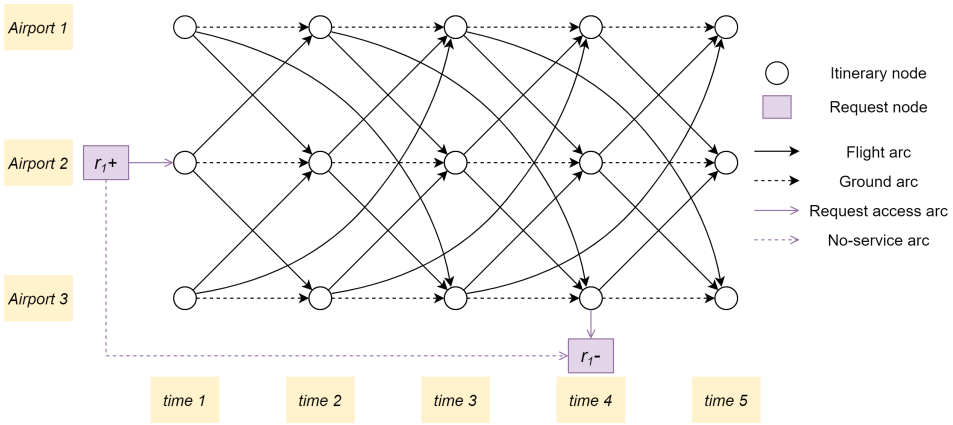

Figure C.1: Example of a time-space network consisting of node set $\mathcal{N}$ and arc set $\mathcal{A}$

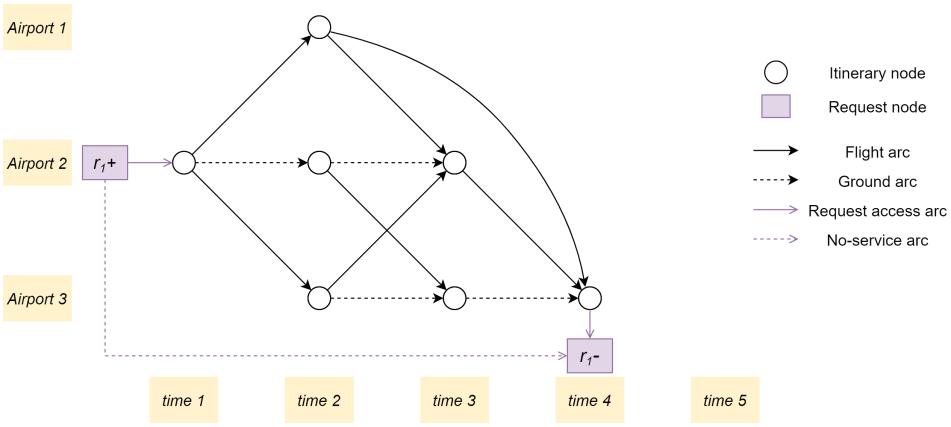

Figure C.2: A time-space network pre-processed for example request r_1

In order to show the results of this step, the amount of arcs and nodes in the TSN is compared to the problem size after pre-processing. Table C.1 gives the values for the nodes and Table C.2 shows the reduction of the number of arcs. These results clearly show that the pre-processing step discards a large portion of the arcs and nodes in the network. The largest reduction can be found for the transatlantic Network 2, with a 70.6% decrease in the amount of nodes and 77.5% in the amount of arcs. This is mainly caused by the long flight times between the European and American airports. Because of this, a part of the airports can only be reached after 7 or more hours into the simulation horizon, making a lot of nodes and arcs infeasible. The same is valid for time steps closer to the due time of a cargo request or the end of the schedule for an aircraft. Multiple nodes and arcs can be discarded because it is no longer feasible to reach the final destination of the cargo or aircraft in time. The problem sizes before the pre-processing steps can directly be related to the size of the network, leading to a smaller problem size for the 6-airport NA network and the largest amount of arcs and nodes for the EU-NA network with 8 available airports.

Table C.1: Amount of nodes before and after pre-processing for an instance with 3 aircraft and 25 cargo requests for all three networks.

	Nodes	Nodes after pre-processing	Difference [%]
Network 1: EU	4,375	1,743	-60.2
Network 2: EU-NA	5,200	1,528	-70.6
Network 3: NA	3,750	1,288	-65.7

Table C.2: Amount of arcs before and after pre-processing for an instance with 3 aircraft and 25 cargo requests for all three networks.

	Arcs	Arcs after pre-processing	Difference [%]
Network 1: EU	29,200	10,980	-62.4
Network 2: EU-NA	37,544	8,447	-77.5
Network 3: NA	20,950	6,385	-69.5

D

Pareto front flight schedules

D.1. Network 1: EU

Tables D.1 and D.2 show the flight schedules which are graphically shown in the Pareto front search in Experiment 2 in Part I. These schedule are the first two schedules of the Pareto front, with an emission weighting factor of 0 and 0.1 respectively. Table D.3 gives the request set that is available for these instances, where also it is noted if the specific cargo request is served.

Table D.1: Flight schedule for aircraft 0 for Network 1: EU with 30 available cargo requests and $w_{CO_2} = 0$

	aircraft	arc	orig	dest	t_{dep}	t_{arr}	requests	payload	LF	CO2	dist	t_{ij}
0	0	101	LUX	STN	6	9	[3]	21456	0.160	26715	489	1.04
1	0	171	STN	PIK	9	12	[3, 22]	38992	0.291	28257	513	1.07
2	0	212	PIK	AMS	12	15	[17, 19, 21, 22]	74749	0.558	35603	708	1.29
3	0	253	AMS	LUX	15	18	[14, 15, 17, 19, 21, 22]	124530	0.929	21499	315	0.85
4	0	298	LUX	MXP	18	21	[7, 9, 10, 15, 21]	113993	0.851	27996	483	1.04
5	0	377	MXP	BUD	21	24	[9, 10, 15, 25]	98892	0.738	42217	830	1.42
6	0	440	BUD	VIE	24	27	[25, 28, 29]	64251	0.479	20490	214	0.74
7	0	477	VIE	LUX	27	30	[25, 26, 27, 28, 29]	106964	0.798	39739	776	1.36
8	0	492	LUX	PIK	30	33	[2, 4]	46289	0.345	44056	978	1.59
9	0	554	PIK	LUX	33	36	[16, 18, 20]	52590	0.392	44381	978	1.59
10	0	643	LUX	BUD	39	42	[0, 1, 5, 6, 8]	115362	0.861	48990	990	1.60
11	0	730	BUD	AMS	42	45	[0, 1, 5, 6]	87149	0.650	54743	1169	1.80
12	0	745	AMS	STN	45	48	[5, 6, 11, 12, 13]	117806	0.879	21633	313	0.85
13	0	806	STN	LUX	48	51	[11, 12, 13, 23, 24]	96912	0.723	28530	489	1.04

Table D.2: Flight schedule for aircraft 0 for Network 1: EU with 30 available cargo requests and $w_{CO_2} = 0.1$

	aircraft	arc	orig	dest	t_{dep}	t_{arr}	requests	payload	LF	CO2	dist	t_{ij}
0	0	149	LUX	PIK	9	12	[2, 3, 4]	67745	0.506	45434	978	1.59
1	0	212	PIK	AMS	12	15	[16, 17, 19, 20, 21]	94754	0.707	36433	708	1.29
2	0	253	AMS	LUX	15	18	[15, 16, 17, 19, 20, 21]	118913	0.887	21426	315	0.85
3	0	298	LUX	MXP	18	21	[7, 9, 10, 15, 21]	113993	0.851	27996	483	1.04
4	0	377	MXP	BUD	21	24	[9, 10, 15, 25]	98892	0.738	42217	830	1.42
5	0	440	BUD	VIE	24	27	[25, 28, 29]	64251	0.479	20490	214	0.74
6	0	477	VIE	LUX	27	30	[25, 26, 27, 28, 29]	106964	0.798	39739	776	1.36
7	0	689	LUX	STN	42	45	[0, 1, 5, 6]	87149	0.650	27962	489	1.04
8	0	758	STN	AMS	45	48	[0, 1, 23, 24]	66255	0.494	21497	313	0.85
9	0	792	AMS	LUX	48	51	[11, 12, 13, 14, 23, 24]	122534	0.914	21473	315	0.85

Table D.3: Request set available for Network 1: EU.

	orig	dest	weight	t_{rel}	t_{due}	s_r	Served in $w_{CO_2} = 0$	Served in $w_{CO_2} = 0.1$
0	LUX	AMS	17081	39	63	1	yes	yes
1	LUX	AMS	15854	0	48	1	yes	yes
2	LUX	PIK	16418	9	33	1	yes	yes
3	LUX	PIK	21456	0	24	1	yes	yes
4	LUX	PIK	29871	9	33	1	yes	yes
5	LUX	STN	26888	27	60	1	yes	yes
6	LUX	STN	27326	30	66	1	yes	yes
7	LUX	MXP	28013	3	48	1	yes	yes
8	LUX	BUD	28213	33	69	1	yes	no
9	LUX	BUD	19211	0	24	1	yes	yes
10	LUX	BUD	26657	3	30	1	yes	yes
11	AMS	LUX	18005	30	69	1	yes	yes
12	AMS	LUX	25755	45	72	1	yes	yes
13	AMS	LUX	19832	21	69	1	yes	yes
14	AMS	LUX	25622	12	51	1	yes	yes
15	AMS	BUD	24159	3	30	1	yes	yes
16	PIK	LUX	15260	0	42	1	yes	yes
17	PIK	LUX	17930	0	24	1	yes	yes
18	PIK	LUX	15049	27	51	1	yes	no
19	PIK	LUX	23330	0	24	1	yes	yes
20	PIK	LUX	22281	0	39	1	yes	yes
21	PIK	MXP	15953	3	30	1	yes	yes
22	STN	LUX	17536	0	24	1	yes	no
23	STN	LUX	17846	15	57	1	yes	yes
24	STN	LUX	15474	45	69	1	yes	yes
25	MXP	LUX	28865	15	45	1	yes	yes
26	VIE	LUX	24757	0	45	1	yes	yes
27	VIE	LUX	17956	18	57	1	yes	yes
28	BUD	LUX	15958	21	45	1	yes	yes
29	BUD	LUX	19428	24	57	1	yes	yes

D.2. Network 2: EU-NA

Tables D.4 and D.5 show the flight schedules which are graphically shown in the Pareto front search in Experiment 2 in Part I. These schedule are the first two schedules of the Pareto front, with an emission weighting factor of 0 and 0.55 respectively. Table D.6 gives the request set that is available for these instances, where also it is noted if the specific cargo request is served.

Table D.4: Flight schedule for aircraft 0 for Network 2: EU-NA with 30 available cargo requests and $w_{CO_2} = 0$

	aircraft	arc	orig	dest	t_{dep}	t_{arr}	requests	payload	LF	CO2	dist	t_{ij}
0	0	193	LUX	PIK	9	12	[0, 5]	39832	0.297	43642	978	1.59
1	0	269	PIK	ORD	12	21	[5]	22648	0.169	223317	5906	7.06
2	0	494	ORD	ATL	21	24	[22, 23]	55363	0.413	47974	976	1.58
3	0	566	ATL	ORD	24	27	[25, 26]	44069	0.329	46822	976	1.58
4	0	620	ORD	JFK	27	30	[16, 17, 21, 24, 25]	121231	0.905	61537	1188	1.82
5	0	665	JFK	LUX	30	39	[10, 11, 16, 17, 25]	132111	0.986	288218	6053	7.23
6	0	835	LUX	JFK	39	48	[2, 3, 4, 6]	88977	0.664	263736	6053	7.23
7	0	1054	JFK	ATL	48	51	[4, 6, 24]	64133	0.479	58840	1222	1.86
8	0	1205	ATL	MIA	54	57	[4, 22]	51078	0.381	45730	960	1.57
9	0	1404	MIA	ORD	66	72	[]	0	0.000	75417	1930	2.64

Table D.5: Flight schedule for aircraft 0 for Network 2: EU-NA with 30 available cargo requests and $w_{CO_2} = 0.55$

	aircraft	arc	orig	dest	t_{dep}	t_{arr}	requests	payload	LF	CO2	dist	t_{ij}
0	0	385	LUX	PIK	18	21	[0, 2, 3, 4, 5]	105426	0.787	47854	978	1.59
1	0	459	PIK	JFK	21	30	[2, 3, 4, 5]	88242	0.659	222677	5160	6.23
2	0	669	JFK	ORD	30	33	[5]	22648	0.169	54694	1188	1.82
3	0	748	ORD	JFK	33	36	[21, 22, 23, 24]	96262	0.718	59581	1188	1.82
4	0	796	JFK	MIA	36	42	[4, 22, 23, 24]	96113	0.717	82399	1757	2.45
5	0	998	MIA	ATL	45	48	[23, 24]	45035	0.336	45749	960	1.57
6	0	1270	ATL	ORD	57	60	[]	0	0.000	44109	976	1.58

Table D.6: Request set available for Network 2: EU-NA.

	orig	dest	weight	t_{rel}	t_{due}	s_r	Served in $w_{CO_2} = 0$	Served in $w_{CO_2} = 0.55$
0	LUX	PIK	17184	3	27	1	yes	yes
1	LUX	MXP	24794	27	51	1	no	no
2	LUX	JFK	19420	18	60	1.5	yes	yes
3	LUX	JFK	23438	12	48	1.5	yes	yes
4	LUX	MIA	22736	9	57	1.5	yes	yes
5	LUX	ORD	22648	6	33	1.5	yes	yes
6	LUX	ATL	23383	24	72	1.5	yes	no
7	PIK	LUX	16565	33	57	1	no	no
8	PIK	LUX	25221	36	66	1	no	no
9	MXP	LUX	15466	0	24	1	no	no
10	JFK	LUX	23777	3	51	1.5	yes	no
11	JFK	LUX	28002	0	42	1.5	yes	no
12	MIA	LUX	28462	3	36	1.5	no	no
13	MIA	IAH	19970	18	54	1	no	no
14	MIA	IAH	20088	15	42	1	no	no
15	ORD	LUX	23391	0	39	1.5	no	no
16	ORD	LUX	27816	24	66	1.5	yes	no
17	ORD	LUX	27070	12	51	1.5	yes	no
18	ORD	LUX	23179	24	57	1.5	no	no
19	ORD	PIK	23666	27	60	1.5	no	no
20	ORD	MXP	18674	30	66	1.5	no	no
21	ORD	JFK	22885	24	51	1	yes	yes
22	ORD	MIA	28342	21	60	1	yes	yes
23	ORD	ATL	27021	9	48	1	yes	yes
24	ORD	ATL	18014	27	63	1	yes	yes
25	ATL	LUX	25446	15	51	1.5	yes	no
26	ATL	ORD	18623	9	36	1	yes	no
27	IAH	LUX	24041	3	33	1.5	no	no
28	IAH	LUX	24593	3	48	1.5	no	no
29	IAH	PIK	21686	30	57	1.5	no	no

D.3. Network 3: NA

Tables D.7 and D.8 show the flight schedules which are graphically shown in the Pareto front search in Experiment 2 in Part I. These schedule are the first two schedules of the Pareto front, with an emission weighting factor of 0 and 0.2 respectively. Table D.9 gives the request set that is available for these instances, where also it is noted if the specific cargo request is served.

Table D.7: Flight schedule for aircraft 0 for Network 3: NA with 30 available cargo requests and $w_{CO_2} = 0$

	aircraft	arc	orig	dest	t_{dep}	t_{arr}	requests	payload	LF	CO2	dist	t_{ij}
0	0	140	MEX	LAX	9	15	[4, 5, 10]	68179	0.509	109607	2500	3.28
1	0	190	LAX	GDL	15	21	[4, 5, 15, 16]	86612	0.646	93678	2106	2.84
2	0	281	GDL	MEX	21	24	[1, 2, 3, 16]	89632	0.669	30067	458	1.01
3	0	319	MEX	DFW	24	30	[2, 3, 6, 8, 9, 12]	132874	0.992	77648	1511	2.18
4	0	361	DFW	LAX	30	36	[3, 8, 9, 12]	91365	0.682	89648	1983	2.70
5	0	441	LAX	YYC	36	42	[12, 17, 18, 21, 22, 24]	129700	0.968	93022	1943	2.66
6	0	525	YYC	SEA	42	45	[17, 18, 27, 28, 29]	108853	0.812	39659	727	1.31
7	0	555	SEA	YYC	45	48	[25]	25690	0.192	35983	727	1.31
8	0	596	YYC	LAX	48	54	[]	0	0.000	75765	1943	2.66
9	0	656	LAX	SEA	54	60	[19, 20]	55771	0.416	67271	1537	2.21
10	0	770	SEA	LAX	63	69	[]	0	0.000	62070	1537	2.21

Table D.8: Flight schedule for aircraft 0 for Network 3: NA with 30 available cargo requests and $w_{CO_2} = 0.2$

	aircraft	arc	orig	dest	t_{dep}	t_{arr}	requests	payload	LF	CO2	dist	t_{ij}
0	0	107	MEX	GDL	6	9	[4, 5]	46712	0.349	30405	458	1.01
1	0	137	GDL	MEX	9	12	[1, 3]	44130	0.329	29287	458	1.01
2	0	320	MEX	LAX	24	30	[3, 8, 9, 10, 12]	112832	0.842	117868	2500	3.28
3	0	404	LAX	SEA	33	39	[12, 17, 18, 21, 22, 24]	129700	0.968	75094	1537	2.21
4	0	483	SEA	YYC	39	42	[12, 21, 22, 24, 25]	117662	0.878	39886	727	1.31
5	0	524	YYC	LAX	42	48	[27, 28, 29]	71125	0.531	85627	1943	2.66
6	0	584	LAX	SEA	48	54	[19, 20, 27, 28, 29]	126896	0.947	74798	1537	2.21
7	0	662	SEA	LAX	54	60	[]	0	0.000	62070	1537	2.21

Table D.9: Request set available for Network 3: NA.

	orig	dest	weight	t_{rel}	t_{due}	s_r	Served in $w_{CO_2} = 0$	Served in $w_{CO_2} = 0.55$
0	GDL	MEX	27566	48	72	1	no	no
1	GDL	MEX	19762	0	30	1	yes	yes
2	GDL	DFW	25989	6	51	1	yes	no
3	GDL	LAX	24368	0	48	1	yes	yes
4	MEX	GDL	27662	0	39	1	yes	yes
5	MEX	GDL	19050	6	33	1	yes	yes
6	MEX	DFW	15520	18	42	1	yes	no
7	MEX	DFW	22669	48	72	1	no	no
8	MEX	LAX	23955	24	57	1	yes	yes
9	MEX	LAX	20224	21	57	1	yes	yes
10	MEX	LAX	21467	9	39	1	yes	yes
11	MEX	SEA	16102	6	42	1	no	no
12	MEX	YYC	22818	18	57	1	yes	yes
13	DFW	MEX	28474	36	60	1	no	no
14	DFW	LAX	19980	0	24	1	no	no
15	LAX	GDL	20387	0	24	1	yes	no
16	LAX	MEX	19513	0	24	1	yes	no
17	LAX	SEA	19309	6	45	1	yes	yes
18	LAX	SEA	18419	9	45	1	yes	yes
19	LAX	SEA	25912	45	72	1	yes	yes
20	LAX	SEA	29859	21	60	1	yes	yes
21	LAX	YYC	28615	18	48	1	yes	yes
22	LAX	YYC	18183	12	42	1	yes	yes
23	LAX	YYC	15593	3	30	1	no	no
24	LAX	YYC	22356	30	72	1	yes	yes
25	SEA	YYC	25690	30	57	1	yes	yes
26	YYC	LAX	26541	3	30	1	no	no
27	YYC	SEA	24140	30	54	1	yes	yes
28	YYC	SEA	19599	33	57	1	yes	yes
29	YYC	SEA	27386	27	54	1	yes	yes

Bibliography

[1] Jeph Abara. Applying Integer Linear Programming to the Fleet Assignment Problem. Interfaces, 19(4): 20–28, 1989. ISSN 0092-2102. doi: 10.1287/inte.19.4.20.

[2] Cynthia Barnhart, Natashia L. Boland, Lloyd W. Clarke, Elles L. Johnson, George L. Nemhauser, and Rajesh G. Shenoi. Flight string models for aircraft fleeting and routing. Transportation Science, 32(3): 208–220, 1998. ISSN 00411655. doi: 10.1287/trsc.32.3.208.

[3] Boeing. Boeing World Air Cargo Forecast (WACF). Technical report, 2020. URL http://www.boeing.com/resources/boeingdotcom/market/assets/downloads/2020_WACF_PDF_Download.pdf.

[4] Jan K. Brueckner and Chrystyane Abreu. Airline fuel usage and carbon emissions: Determining factors. Journal of Air Transport Management, 62:10–17, 2017. ISSN 09696997. doi: 10.1016/j.jairtraman.2017.01.004.

[5] Edmund K. Burke, Patrick De Causmaecker, Geert De Maere, Jeroen Mulder, Marc Paelinck, and Greet Vanden Berghe. A multi-objective approach for robust airline scheduling. Computers and Operations Research, 37(5):822–832, 2010. ISSN 03050548. doi: 10.1016/j.cor.2009.03.026.

[6] Ching Cheng Chao. Assessment of carbon emission costs for air cargo transportation. Transportation Research Part D: Transport and Environment, 33:186–195, 2014. ISSN 13619209. doi: 10.1016/j.trd.2014.06.004.

[7] Lloyd Clarke, Ellis Johnson, George Nemhauser, and Zhongxi Zhu. The aircraft rotation problem. Annals of Operations Research, 69:33–46, 1997. ISSN 15729338. doi: 10.1023/a:1018945415148.

[8] Felipe Delgado and Julio Mora. A matheuristic approach to the air-cargo recovery problem under demand disruption. Journal of Air Transport Management, 90:Article 101939, jan 2021. ISSN 09696997. doi: 10.1016/j.jairtraman.2020.101939.

[9] Felipe Delgado, Cristóbal Sirhan, Mathias Katscher, and Homero Larrain. Recovering from demand disruptions on an air cargo network. Journal of Air Transport Management, 85:Article 101799, jun 2020. ISSN 09696997. doi: 10.1016/j.jairtraman.2020.101799.

[10] Ulrich Derigs and Stefan Friederichs. Air cargo scheduling: Integrated models and solution procedures. OR Spectrum, 35(2):325–362, 2013. ISSN 14366304. doi: 10.1007/s00291-012-0299-y.

[11] Ulrich Derigs and Stefan Illing. Does EU ETS instigate Air Cargo network reconfiguration? A model-based analysis. European Journal of Operational Research, 225(3):518–527, 2013. ISSN 03772217. doi: 10.1016/j.ejor.2012.10.016.

[12] Ulrich Derigs, Stefan Friederichs, and Simon Schäfer. A New Approach for Air Cargo Network Planning. Source: Transportation Science, 43(3):370–380, 2009. doi: 10.1287/trsc.l090.0282.

[13] EASA. Introduction to the ICAO Engine Emissions Databank. Technical report, Cologne, 2021.

[14] European Commission. EU ETS Handbook. Technical report, 2015. URL https://ec.europa.eu/clima/sites/default/files/docs/ets_handbook_en.pdf.

[15] European Environment Agency. EMEP/EEA Air Pollutant Emission Inventory Guidebook, Part B: Aviation. Technical report, European Environment Agency, Copenhagen, 2019. URL https://www.eea.europa.eu/publications/emep-eea-guidebook-2019/part-b-sectoral-guidance-chapters/1-energy/1-a-combustion/1-a-3-a-aviation/at_download/file.

[16] European Federation for Transport and Environment. Including Aviation in the EU's Emissions Trading Scheme (EU ETS) Background Briefing. Technical report, 2008. URL http://www.transportenvironment.org/Article201.html.

[17] Bo Feng, Yanzhi Li, and Zuo Jun Max Shen. Air cargo operations: Literature review and comparison with practices. Transportation Research Part C: Emerging Technologies, 56:263–280, 2015. ISSN 0968090X. doi: 10.1016/j.trc.2015.03.028.

[18] Gary Froyland, Stephen J. Maher, and Cheng Lung Wu. The recoverable robust tail assignment problem. Transportation Science, 48(3):351–372, 2014. ISSN 15265447. doi: 10.1287/trsc.2013.0463.

[19] Brandon Graver, Kevin Zhang, and Dan Rutherford. CO2 emissions from commercial aviation, 2018. Technical report, International Council on Clean Transportation, 2019. URL https://theicct.org/sites/default/files/publications/ICCT_CO2-commercl-aviation-2018_20190918.pdf.

[20] Christopher A. Hane, Cynthia Barnhart, Ellis L. Johnson, Roy E. Marsten, George L. Nemhauser, and Gabriele Sigismondi. The fleet assignment problem: Solving a large-scale integer program. Mathematical Programming, 70(1-3):211–232, 1995. ISSN 00255610. doi: 10.1007/BF01585938.

[21] Mohamed Haouari, Shengzhi Shao, and Hanif D. Sherali. A lifted compact formulation for the daily aircraft maintenance routing problem. Transportation Science, 47(4):508–525, 2013. ISSN 15265447. doi: 10.1287/trsc.1120.0433.

[22] ICAO. Environmental Trends in Aviation to 2050. Technical report, 2016. URL https://www.icao.int/environmental-protection/Documents/EnvironmentalReports/2016/ENVReport2016_pg16-22.pdf.

[23] ICAO. Airport Air Quality Manual. Technical report, 2016.

[24] ICAO. ICAO Carbon Emissions Calculator Methodology. Technical report, 2018. URL https://www.icao.int/environmental-protection/CarbonOffset/Documents/MethodologyICAOCarbonCalculator_v11-2018.pdf.

[25] International Air Transport Association. Aviation & Climate Change. Technical report, 2020. URL https://www.iata.org/en/iata-repository/pressroom/fact-sheets/fact-sheet--climate-change/.

[26] C.N. Jardine. Calculating the Environmental Impact of Aviation Emissions. Technical report, Oxford University Centre for the Environment, Oxford, 2005. URL http://www.hwa.uk.com/site/wp-content/uploads/2020/10/CD-16.6-Calculating-the-Environmental-Impact-of-Aviation-Emissions-June-2005-SSE04.pdf.

[27] Ugur Kesgin. Aircraft emissions at Turkish airports. Energy, 31(2-3):372–384, 2006. ISSN 03605442. doi: 10.1016/j.energy.2005.01.012.

[28] Oumaima Khaled, Michel Minoux, Vincent Mousseau, Stéphane Michel, and Xavier Ceugniet. A compact optimization model for the tail assignment problem. European Journal of Operational Research, 264(2):548–557, 2018. ISSN 03772217. doi: 10.1016/j.ejor.2017.06.045.

[29] Zhe Liang and Wanpracha Art Chaovalitwongse. The aircraft maintenance routing problem. In Wanpracha Art Chaovalitwongse, Kevin C. Furman, and Panos M. Pardalo, editors, Optimization and Logistics Challenges in the Enterprise, pages 327–348. Springer, New York, 2009.

[30] Zhe Liang and Wanpracha Art Chaovalitwongse. A network-based model for the integrated weekly aircraft maintenance routing and fleet assignment problem. Transportation Science, 47(4):493–507, 2013. ISSN 15265447. doi: 10.1287/trsc.1120.0434.

[31] Zhe Liang, Wanpracha Art Chaovalitwongse, Huei Chuen Huang, and Ellis L. Johnson. On a new rotation tour network model for aircraft maintenance routing problem. Transportation Science, 45(1):109–120, 2011. ISSN 15265447. doi: 10.1287/trsc.1100.0338.

[32] Becky P.Y. Loo, Linna Li, Voula Psaraki, and Ioanna Pagoni. CO2 emissions associated with hubbing activities in air transport: An international comparison. Journal of Transport Geography, 34:185–193, 2014. ISSN 09666923. doi: 10.1016/j.jtrangeo.2013.12.006.

[33] Roy E Marsten and Michael R Muller. A Mixed-Integer Programming Approach to Air Cargo Fleet Planning. Technical Report 11, 1980.

[34] C. Miyoshi and K. J. Mason. The carbon emissions of selected airlines and aircraft types in three geographic markets. Journal of Air Transport Management, 15(3):138–147, 2009. ISSN 09696997. doi: 10.1016/j.jairtraman.2008.11.009.

[35] Angela Nuic, Damir Poles, and Vincent Mouillet. BADA: An advanced aircraft performance model for present and future ATMsystems. International Journal of Adaptive Control and Signal Processing, 24 (10):850–866, 2010. ISSN 08906327. doi: 10.1002/acs.1176.

[36] Ioanna Pagoni and Voula Psaraki. A tool for calculating aircraft emissions and its application to Greek airspace. Transportation Planning and Technology, 37(2):138–153, 2014. ISSN 03081060. doi: 10.1080/03081060.2013.851510.

[37] Maurice Pollack. Some aspects of the aircraft scheduling problem. Transportation Research, 8(3):233–243, 1974. ISSN 00411647. doi: 10.1016/0041-1647(74)90010-0.

[38] Nima Safaei and Andrew K.S. Jardine. Aircraft routing with generalized maintenance constraints. Omega (United Kingdom), 80:111–122, oct 2018. ISSN 03050483. doi: 10.1016/j.omega.2017.08.013.

[39] Gregor Schürmann, Klaus Schäfer, Carsten Jahn, Herbert Hoffmann, Martina Bauerfeind, Emanuel Fleuti, and Bernhard Rappenglück. The impact of NOx, CO and VOC emissions on the air quality of Zurich airport. Atmospheric Environment, 41(1):103–118, 2007. ISSN 13522310. doi: 10.1016/j.atmosenv.2006.07.030.

[40] Hanif D. Sherali, Ebru K. Bish, and Xiaomei Zhu. Airline fleet assignment concepts, models, and algorithms. European Journal of Operational Research, 172(1):1–30, 2006. ISSN 03772217. doi: 10.1016/j.ejor.2005.01.056.

[41] Junzi Sun, Jacco M Hoekstra, and Joost Ellerbroek. OpenAP: An Open-Source Aircraft Performance Model for Air Transportation Studies and Simulations. Aerospace, 7(104), 2020. doi: 10.3390/aerospace7080104.

[42] Aydin Tokuslu. Estimation of aircraft emissions at Georgian international airport. Energy, 206:Article 118219, 2020. ISSN 03605442. doi: 10.1016/j.energy.2020.118219.

[43] E. T. Turgut and M. A. Rosen. Assessment of emissions at busy airports. International Journal of Energy Research, 34(9):800–814, 2009. ISSN 0363907X. doi: 10.1002/er.1601.

[44] Shangyao Yan and Chia Hung Chen. Optimal flight scheduling models for cargo airlines under alliances. Journal of Scheduling, 11(3):175–186, 2008. ISSN 10946136. doi: 10.1007/s10951-007-0020-1.

[45] Shangyao Yan and Hwei-Fwa Young. A Decision Support Framework for Multi-Fleet Routing and Multi-Stop Scheduling. Transportation Research Part A: Policy and Practice, 30(5):379–398, 1996.

[46] Shangyao Yan, Shin Chin Chen, and Chia Hung Chen. Air cargo fleet routing and timetable setting with multiple on-time demands. Transportation Research Part E: Logistics and Transportation Review, 42(5):409–430, 2006. ISSN 13665545. doi: 10.1016/j.tre.2005.02.002.

[47] Yan Yan, Jin Zhang, and Qiuyu Tang. Aircraft fleet route optimization based on cost and low carbon emission in aviation line alliance network. Journal of Intelligent and Fuzzy Systems, 39(1):1163–1182, 2020. ISSN 18758967. doi: 10.3233/JIFS-192041.

[48] Lei Zhou, Zhe Liang, Chun-An Chou, and Wanpracha Art Chaovalitwongse. Airline planning and scheduling: Models and solution methodologies. Fronteirs of Engineering Management, 7(1):1–26, 2020. doi: 10.1007/s42524-020-0093-5.