



Green Flight Trajectories

A REACT4C data analysis

Green Flight Trajectories

A REACT4C data analysis

by

Pieter Verbist

to obtain the degree of Master of Science
at the Delft University of Technology,
to be defended publicly on Friday October 28, 2016 at 14:00.

Student number:	4094697
Thesis registration number:	094#16#MT#FPP
Project duration:	October 15, 2015 – October 28, 2016
Thesis committee:	Prof. Dr. V. Grewe, TU Delft (ANCE), DLR, supervisor Prof. Dr. D. G. Simons, TU Delft (ANCE) Dr. Ir. S. Hartjes, TU Delft (ATO) Dr. Ir. M. Voskuil, TU Delft (FPP)

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

Abstract

REACT4C is a project that received European funding to investigate whether air traffic across the North Atlantic Ocean can be rerouted such that the resulting climate impact is reduced. The project considered 400 daily flights in each direction. Eight frequently occurring weather situations were identified, based on the strength and location of the jet stream between North America and Europe. It was found that the climate impact could be reduced to a large extent, although the impact reduction potential heavily depends on the direction of flight, the weather pattern under consideration, the climate metric used to quantify climate impact, and the relative importance of climate impact and economic cost during the optimization. The way in which climate-optimal trajectories compare to their cost-optimal counterparts, is still largely unknown.

This research examines how the routes are affected by the climate optimization. Two kinds of studies are performed. The first one is a case study, in which the trajectories of one combination of direction, climate metric, weather type and level of climate optimization are compared with the cost-optimal routes. The influence of each of these four case differentiators on the flights is examined as well, using a one-factor-at-a-time approach. A tool is made that can be used to analyze the trajectories of any combination of flight direction, climate metric, weather pattern and level of climate optimization.

The second study is a general analysis of the REACT4C routes, taking into account all combinations of the four case differentiators. This study is conducted to unravel general trends in the way direction, metric, weather and level of climate optimization influence the way in which flights are rerouted.

The following characteristics are used to quantify how routes are altered. First, the percentage of affected flights is determined. Then, the flight duration and flight distance increments with respect to the cost-optimized flights are computed. Furthermore, the shift in latitude and altitude is also considered. Finally, frequently occurring rerouting schemes are identified, and the percentage of flights belonging to each scheme is determined.

It was found in both the case study and the general analysis that the way in which climate-optimized routes compare to their cost-optimized counterparts, highly depends on the direction of flight, the weather pattern and the relative importance of economic cost and climate impact during the optimization. However, the choice of climate metric proved to have hardly any influence on the rerouting strategies. Because of the substantial cross-case differences, the preferred strategy to investigate rerouting characteristics is by making use of the case analysis tool that was created. Nonetheless, the general study unravelled some general trends. The trajectories in general are altered such that the flight duration and flight distance are increased, and the cruise altitude is lowered. Furthermore, generally there are more shifts towards the south than there are towards the north. Increasing the level of climate optimization is shown to result into more extreme route alterations. Flights towards Europe are in general more often north of the original routes than in westbound direction. The climate metric used during the optimization has little to no influence on the way the routes are changed. Finally, no trends in trajectory alterations are distinguishable between the eight weather patterns.

Preface and acknowledgements

This thesis is the final step in obtaining my Master's degree in Aerospace Engineering at the Delft University of Technology. As a Flight Performance & Propulsion student, my knowledge about the climate effects of aviation was limited. When I was looking for a graduation project, my attention was drawn to this unfamiliar field of expertise. The idea of contributing to a more eco-friendly society was the trigger that made me dedicate one year to this research project at the Aircraft Noise and Climate Effects section.

That was a decision I do not regret. Learning more about how to reduce the climate impact of aviation, made me realize that there is so much more that all of us can do to leave a cleaner and cooler planet for future generations, if only we are committed to do so. That is, perhaps, a personal conclusion I draw from this research.

When I was writing the final pages of this text, I was not only finishing one year of hard work and dedication, but I was also turning the final page of my life as a student. In a couple of years, when I look back at this document, not only will my graduation project come back to mind, but also my six happy and eventful years in Delft.

I would like to thank my supervisor, Prof. Dr. Volker Grewe, for giving me the opportunity to contribute to the promising REACT4C project. You always took the time to listen to my questions and answered constructively, even if this meant arranging a web conference. Your recent appointment as professor at TU Delft will give many more students the privilege to be guided towards graduation by someone who is sincerely motivated to arrive at a great end result.

I am grateful to Prof. Dr. Dick G. Simons, who was the person that brought me into contact with Prof. Grewe. Your feedback on the green light presentation was highly valued, as is your presence in my graduation committee. I would also like to extend my appreciation to the two other committee members, Dr. Ir. Sander Hartjes and Dr. Ir. Mark Voskuil, for showing your interest in my thesis.

On a personal level, words cannot express my gratitude towards my girlfriend Lisse, whose perseverance and support was greater than I could have ever imagined. There were times when you were almost more proud than I was myself, and also times when you believed in my capabilities more than I did. Completing this thesis would have been so much harder without you by my side. Thanks again for your unconditional support.

Mom and dad, thank you too, for your support on so many levels. You gave me the opportunity to come to study in Delft in the first place, a gesture that shaped my future and my life. You have done so much for me. I am grateful to my older brother Simon as well, who was always ready to give me advice on how to proceed and who, I know, will be a little bit proud of his younger brother.

*Pieter Verbist
Delft, October 2016*

Contents

Abstract	iii
Preface and acknowledgements	v
List of Figures	ix
List of Tables	xiii
Nomenclature	xv
1 Introduction	1
1.1 Problem statement	1
1.2 Research question and objectives	2
1.3 Document overview	2
2 Literature review: the REACT4C project	3
2.1 Climate Cost Function modeling approach	3
2.1.1 Selection of representative situation	3
2.1.2 Definition of time-regions and computation of climate cost functions	5
2.1.3 Route optimization using the CCFs	5
2.2 Outcome and results	6
3 Methodology	9
3.1 Data description	9
3.1.1 Information quantity and optimization problem	9
3.1.2 Trajectory data	10
3.1.3 Pareto front data	11
3.2 General approach	11
3.3 Cross-case Pareto front examination	13
3.4 Investigated flight characteristics	15
3.4.1 Mean flight duration	15
3.4.2 PDF of flight duration increments	15
3.4.3 Mean flight distance	16
3.4.4 PDF of flight distance increments	16
3.4.5 Percentage of rerouted flights	16
3.4.6 Shift in latitude	16
3.4.7 Trajectory classification based on lateral shift.	18
3.4.8 Altitude shift.	18
4 Case study of trajectory changes	21
4.1 Inter-Pareto trajectory changes	21
4.1.1 Rerouted flights	21
4.1.2 Flight time and distance	22
4.1.3 Shift in latitude and trajectory classification	24
4.1.4 Altitude shift.	25
4.1.5 Conclusion of the inter-Pareto case study.	26
4.2 Influence of flight direction on trajectories	26
4.2.1 Rerouted flights	27
4.2.2 Flight time and distance	27
4.2.3 Shift in latitude and trajectory classification	29
4.2.4 Altitude shift.	30
4.2.5 Conclusion of the case study of the influence of direction	30

4.3	Influence of metric on trajectories	30
4.3.1	Rerouted flights	30
4.3.2	Flight time and distance	30
4.3.3	Shift in latitude and trajectory classification	32
4.3.4	Altitude shift.	32
4.3.5	Conclusion of the case study of the influence of climate metrics.	34
4.4	Influence of weather pattern on trajectories	34
4.4.1	Rerouted flights	34
4.4.2	Flight time and distance	35
4.4.3	Shift in latitude and trajectory classification	37
4.4.4	Altitude shift.	40
4.4.5	Conclusion of the case study of the influence of weather pattern	40
4.5	Total overview	41
5	General climate-optimized trajectory analysis	43
5.1	General inter-Pareto trajectory changes.	43
5.1.1	Rerouted flights	43
5.1.2	Flight time and distance	44
5.1.3	Shift in latitude and trajectory classification	46
5.1.4	Altitude shift.	47
5.1.5	Conclusion of the general influence of Pareto location	48
5.2	General influence of flight direction on trajectories	48
5.2.1	Rerouted flights	48
5.2.2	Flight time and distance	48
5.2.3	Shift in latitude and trajectory classification	50
5.2.4	Altitude shift.	50
5.2.5	Conclusion of the general influence of flight direction	50
5.3	General influence of climate metric on trajectories	52
5.3.1	Rerouted flights	52
5.3.2	Flight time and distance	52
5.3.3	Shift in latitude and trajectory classification	53
5.3.4	Altitude shift.	55
5.3.5	Conclusion of the general influence of climate metric	55
5.4	General influence of weather pattern on trajectories	55
5.4.1	Rerouted flights	56
5.4.2	Flight time and distance	56
5.4.3	Shift in latitude and trajectory classification	56
5.4.4	Altitude shift.	59
5.4.5	Conclusion of the general influence of weather pattern	60
6	Discussion	61
6.1	The big picture	61
6.2	Case study versus general analysis	61
6.3	Relation to previous research	62
6.4	Limitations and recommendations	62
7	Conclusion	65
	Bibliography	67
A	Variability of trajectory classes as a function of Pareto location	69
B	Flight-based box plots	71
B.1	General direction comparison	71
B.2	General climate metric comparison	72
B.3	General weather pattern comparison	73

List of Figures

2.1	Outline of the CCF model chain.	4
2.2	REACT4C's optimal relation between climate impact and economic cost for WP1.	6
2.3	Variation of trajectory characteristics of the REACT4C WP1 case study.	7
3.1	Pareto front of the case westbound/P-AGWP100/WP1.	10
3.2	The Pareto fronts of the cases westbound/P-AGWP100/WP1 and westbound/P-AGWP100/WP2.	14
3.3	Box plots of the normalized inter-Pareto distance between each data set along the Pareto front.	15
3.4	Generic diagram of the 1D interpolation of the same flight at different Pareto locations.	17
3.5	Trajectory scheme of all fully climate-optimized flights of the case westbound/P-AGWP100/WP1.	18
3.6	Example of each trajectory class based on lateral shift.	19
4.1	Normalized mean flight time and distance as a function of Pareto location for the case westbound/P-AGWP100/WP1.	22
4.2	PDFs of flight duration increments as a function of Pareto location for the case westbound/P-AGWP100/WP1.	23
4.3	PDFs of flight distance increments as a function of Pareto location for the case westbound/P-AGWP100/WP1.	23
4.4	Average shift in latitude depending on geographical location for all Pareto points for the case westbound/P-AGWP100/WP1.	24
4.5	Percentage of rerouted flights that fall within each trajectory class as a function of Pareto location for the case westbound/P-AGWP100/WP1.	25
4.6	Average shift in altitude depending on geographical location for all Pareto points for the case westbound/P-AGWP100/WP1.	26
4.7	PDFs of flight duration increments as a function of flight direction for the case P-AGWP100/WP1/ <i>Clim</i> 1.0.	28
4.8	PDFs of flight distance increments as a function of flight direction for the case P-AGWP100/WP1/ <i>Clim</i> 1.0.	28
4.9	Average shift in latitude depending on geographical location for the two directions for the case P-AGWP100/WP1/ <i>Clim</i> 1.0.	29
4.10	Percentage of rerouted flights that fall within each trajectory class as a function of flight direction for the case P-AGWP100/WP1/ <i>Clim</i> 1.0.	29
4.11	Average shift in altitude depending on geographical location for the two directions for the case P-AGWP100/WP1/ <i>Clim</i> 1.0.	30
4.12	Normalized mean flight time and distance as a function of climate metric for the case westbound/WP1/ <i>Clim</i> 1.0.	31
4.13	PDFs of flight duration increments as a function of climate metric for the case westbound/WP1/ <i>Clim</i> 1.0.	31
4.14	PDFs of flight distance increments as a function of climate metric for the case westbound/WP1/ <i>Clim</i> 1.0.	32
4.15	Average shift in latitude depending on geographical location for the three climate metrics for the case westbound/WP1/ <i>Clim</i> 1.0.	33
4.16	Percentage of rerouted flights that fall within each trajectory class as a function of climate metric for the case westbound/WP1/ <i>Clim</i> 1.0.	33
4.17	Average shift in altitude depending on geographical location for the three climate metrics for the case westbound/WP1/ <i>Clim</i> 1.0.	34

4.18	Normalized mean flight time and distance as a function of weather pattern for the case westbound/P-AGWP100/ <i>Clim</i> 1.0.	35
4.19	PDFs of flight duration increments as a function of weather pattern for the case westbound/P-AGWP100/ <i>Clim</i> 1.0.	36
4.20	PDFs of flight distance increments as a function of weather pattern for the case westbound/P-AGWP100/ <i>Clim</i> 1.0.	37
4.21	Average shift in latitude depending on geographical location for the winter weather patterns, for the case westbound/P-AGWP100/ <i>Clim</i> 1.0.	38
4.22	Percentage of rerouted flights that fall within each trajectory class as a function of winter weather pattern for the case westbound/P-AGWP100/ <i>Clim</i> 1.0.	38
4.23	Average shift in latitude depending on geographical location for the summer weather patterns, for the case westbound/P-AGWP100/ <i>Clim</i> 1.0.	39
4.24	Percentage of rerouted flights that fall within each trajectory class as a function of summer weather pattern, for the case westbound/P-AGWP100/ <i>Clim</i> 1.0.	39
4.25	Average shift in altitude depending on geographical location for the winter weather patterns, for the case westbound/P-AGWP100/ <i>Clim</i> 1.0.	40
4.26	Average shift in altitude depending on geographical location for the summer weather patterns, for the case westbound/P-AGWP100/ <i>Clim</i> 1.0.	41
5.1	The percentage of rerouted flights as a function of Pareto location.	44
5.2	Normalized mean flight duration and distance as a function of Pareto location.	45
5.3	Contour plots of the PDFs of flight duration and flight distance increments as a function of Pareto location.	45
5.4	Mean, maximum and minimum lateral shifts as a function of Pareto location.	46
5.5	The percentage of flights within each trajectory class, as a function of Pareto location.	47
5.6	Mean, maximum and minimum altitude shifts as a function of Pareto location.	47
5.7	Box plots of the percentage of flights that are affected by the optimization, as a function of flight direction.	48
5.8	Box plots of the normalized mean flight duration and flight distance per fleet as a function of flight direction.	49
5.9	PDFs of flight duration and distance increments as a function of flight direction.	49
5.10	Box plots of the mean lateral shift per fleet as a function of flight direction.	50
5.11	The percentage of flights within each trajectory class as a function of flight direction.	51
5.12	Box plots of the mean altitude shift per fleet as a function of flight direction.	51
5.13	Box plots of the percentage of flights that are affected by the optimization, as a function of climate metric.	52
5.14	Box plots of the normalized mean flight duration and flight distance per fleet as a function of climate metric.	53
5.15	PDFs of flight duration and flight distance increments as a function of climate metric.	53
5.16	Box plots of the mean lateral shift per fleet as a function of climate metric.	54
5.17	The percentage of flights within each trajectory class as a function of climate metric.	54
5.18	Box plots of the mean altitude shift per fleet as a function of flight direction.	55
5.19	Box plots of the percentage of flights that are affected by the optimization, as a function of weather pattern.	56
5.20	Box plots of the normalized mean flight duration per fleet as a function of weather pattern.	57
5.21	Box plots of the normalized mean flight distance per fleet as a function of weather pattern.	57
5.22	PDFs of flight duration increments as a function of weather pattern.	58
5.23	PDFs of flight distance increments as a function of weather pattern.	58
5.24	Box plots of the mean lateral shift per fleet as a function of weather pattern.	59
5.25	The percentage of flights within each trajectory class as a function of weather pattern.	59
5.26	Box plots of the mean altitude shift per fleet as a function of weather pattern.	60
A.1	Trend line and data points of the percentage of flights within each trajectory class as a function of normalized Pareto location.	69
A.1	Trend line and data points of the percentage of flights within each trajectory class as a function of normalized Pareto location (continued).	70

B.1	Box plots of the normalized flight duration and flight distance per flight as a function of flight direction.	71
B.2	Box plots of the mean latitude and altitude shift per flight as a function of flight direction.	72
B.3	Box plots of the normalized flight duration and flight distance per flight as a function of climate metric.	72
B.4	Box plots of the mean latitude and altitude shift per flight as a function of climate metric.	73
B.5	Box plots of the normalized flight duration per flight as a function of weather pattern.	73
B.6	Box plots of the normalized flight distance per flight as a function of weather pattern.	74
B.7	Box plots of the mean latitude shift per flight as a function of weather pattern.	74
B.8	Box plots of the mean altitude shift per flight as a function of weather pattern.	75

List of Tables

2.1	Characteristics of North Atlantic weather patterns.	3
3.1	Weight factor combinations used in the optimizations.	10
3.2	Explanation of the parameters provided in the trajectory data sets.	11
4.1	Percentage of affected flights as a function of Pareto location of the case westbound/P-AGWP100/WP1.	21
4.2	Probability that a flight is shortened in terms of flight time and flight distance as a function of Pareto location for the case westbound/P-AGWP100/WP1.	23
4.3	Normalized mean flight time and distance as a function of flight direction for the case P-AGWP100/WP1/ <i>Clim</i> 1.0.	27
4.4	Probability that a flight is shortened in terms of flight time and flight distance, as a function of flight direction for the case P-AGWP100/WP1/ <i>Clim</i> 1.0.	27
4.5	Probability that a flight is shortened in terms of flight time and flight distance, as a function of metric for the case westbound/WP1/ <i>Clim</i> 1.0.	32
4.6	Percentage of affected flights as a function of weather pattern for the case westbound/P-AGWP100/ <i>Clim</i> 1.0.	35
4.7	Probability that a flight is shortened in terms of flight time and flight distance, as a function of weather pattern for the case westbound/P-AGWP100/ <i>Clim</i> 1.0.	36
4.8	Overview of the case study results.	42

Nomenclature

List of Acronyms

AEM	Advanced Emission Model
AGWP	Absolute Global Warming Potential
ATR	Average Temperature Response
CATS	Climate Compatible Air Transport System
CCF	Climate Cost Function
ECHAM	European Centre Hamburg general circulation model
EMAC	ECHAM/MESSy Atmospheric Chemistry model
F	(before metric) Climate impact based on Future emissions scenario
ICAO	International Civil Aviation Organization
IQR	Interquartile Range
MESSy	Modular Earth Submodel System
P	(before metric) Climate impact based on Pulse emissions
PDF	Probability Density Function
REACT4C	Reducing Emissions from Aviation by Changing Trajectories for the benefit of Climate
SAAM	System for traffic Assignment and Analysis at a Macroscopic level
SP	Summer weather Pattern
WP	Winter weather Pattern

List of Symbols

α	Weight factor of climate cost in the objective function [-]
C_{clim}	Climate cost [K] or $\left[\frac{W \cdot yr}{m^2}\right]$
C_{eco}	Economic cost [€]
$Clim \alpha$	Fleet optimized with weight factor α in the objective function
Δlat_q	Difference in latitude between cost- and climate-optimized route at query longitude [deg]
f	Objective function of the REACT4C optimization
K	Relation between economic and climate cost [$\text{€}/K$] or [$\text{€}/(W \cdot yr \cdot m^{-2})$]
lat_q	Latitude at query longitude of cost-optimized route [deg]
lat_q^{opt}	Latitude at query longitude of climate-optimized route [deg]
lon_q	Query longitude for interpolation of trajectories [deg]

Introduction

1.1. Problem statement

Climate change has been a controversial issue in the past, but nowadays the scientific world recognizes its existence practically unanimously [1]. The contribution of aviation to the total climate impact caused by human activities, is estimated to be 4.9% with a 90% likelihood range of 2-14% [2].

Scientists are looking for means to reduce this figure for aviation's impact on the climate. A project called "Reducing Emissions from Aviation by Changing Trajectories for the benefit of Climate" (REACT4C) was set up to investigate the potential climate benefits of altering flight trajectories. This project was initiated in 2010 and is part of the European Commission's Seventh Framework Programme. A new modeling approach was used, called the Climate Cost Function approach (CCF), which combined existing well-verified climate models and trajectory optimization tools in a new way [3]. REACT4C rerouted around 400 eastbound flights and 400 westbound flights across the North Atlantic Ocean in 8 different weather situations. In a study of one of these weather types, it was found that the climate impact can be reduced by as much as 25-60% with a cost increase of around 15%, depending on how the climate impact is estimated (the so-called climate metric) and on whether the flights are eastbound or westbound [4].

In total, 288 of these optimizations were performed, one for each combination of flight direction, climate metric, weather pattern, and for six different levels of climate optimization, based on the relative importance of climate impact and economic cost. Each of these optimizations results in its own climate-optimized fleet. The way in which the flights are rerouted, is only concisely discussed for one of these 288 cases in [4], and no general analysis of these flight trajectories was done so far to obtain a general impression of the way in which the REACT4C routes are changed.

Except for the case study performed by Grewe et al. [4], it is unknown how the flight trajectories have been affected in each of the 288 optimizations. In order to explore these optimized fleets, a tool should be built that is able to analyze the route modifications of each of these optimizations. Constructing such a tool allows the user to pick the combination of direction, climate metric, weather pattern and level of climate optimization that is of interest, and examine the way the flights of this case have been rerouted compared to the cost-optimized trajectories.

A case study of this kind is useful for a detailed overview of route modifications in a particular optimization. However, this still does not inform about what climate-optimized routes look like *in general*. Therefore, a second study is required. This is a general study of the REACT4C routes, one that takes into account all available data sets to unravel trends in the way aircraft are diverted. This analysis maintains a high level of generality, and by consequence compromises on the level of detail of the results.

The authors of [4] concisely established trajectory characteristics for one combination of direction, metric and weather pattern, for 5 levels of climate optimization. This means that they investigated the influence of the level of climate optimization on the trajectories for one optimized set of flights. The influence of flight direction, climate metric and weather type on the routes has not been examined before. To expand the knowledge of climate-optimized routes, the effect of these three parameters needs to be investigated, next to the effect of optimization level. This will be done in both the case

study and the general analysis. This research hence needs to establish the influence of level of climate optimization, direction, metric and weather situation on flight trajectories for one case, and the same for all cases combined, examining the trends in the influence of these four parameters.

1.2. Research question and objectives

The research question of this thesis is formulated as follows:

How is the air traffic between North America and Europe restructured when implementing the climate-optimized routes established by the REACT4C project, both on a case-specific level as well as in general?

This research question will only be answered when two studies are conducted: a case study and a general analysis. The case study should primarily be seen as a tool to make a study of trajectory changes that can be used for any combination of direction, metric, weather pattern and level of climate optimization. The general trajectory analysis is meant to provide the reader with a very general idea of how the trajectories change when they are climate-optimized. The research question also implicitly raises the need to determine how the restructuring of air traffic can be quantified.

The research objectives are the following:

1. To establish a set of characteristics that quantify how air traffic is restructured.
2. To create a tool that can be used to analyze air traffic changes for any combination of flight direction, climate metric, weather pattern and level of climate optimization, and to investigate the influence of any of these four parameters on the set of flights under consideration.
3. To find general trends in the way direction, climate metric, weather pattern and level of climate optimization influence the way in which air traffic is altered.

1.3. Document overview

The structure of this document is as follows. Chapter 2 introduces the reader to the REACT4C project by means of a literature review. Chapter 3 then moves on to a description of the methodology that is used to arrive at the envisioned results. The case study in chapter 4 puts the case-specific analysis tool into practice by examining the trajectory changes of the case that was briefly analyzed in [4]. Chapter 5 moves on to the overall analysis, which retains the same format as the case study, but now all data sets are combined and trends are visualized. The results are discussed in chapter 6. Finally, the thesis is concluded in chapter 7.

2

Literature review: the REACT4C project

This chapter aims to familiarize the reader with the REACT4C project. It is necessary to understand how the project members achieved the results that are the basis of this thesis, and to what extent they have analyzed them. This chapter consists of two parts: a part in which the method used for route optimization is explained (section 2.1) and a part in which the results are discussed (section 2.2).

2.1. Climate Cost Function modeling approach

The modeling approach designed and used by REACT4C is described in [3]. REACT4C calls its new flight optimization approach the Climate Cost Function (CCF) approach. This modeling chain makes use of models that have already been applied before. The interaction between these models is new. The structure of the method is given in Figure 2.1. The approach will be explained in this section, by making use of Figure 2.1.

2.1.1. Selection of representative situation

The CCF approach will be explained by applying it directly to the case study considered in [4]. In this case study around 800 daily flights between Europe and North America, across the North Atlantic, are considered. The first step is to obtain representative meteorological data to constitute the weather situations that are to be considered. In total, 8 frequently occurring weather patterns between Europe and North America were identified in [5]: 5 winter patterns, and 3 summer patterns. For each weather scenario, the structure in Figure 2.1 does one iteration. The final climate impact will be the impact of the individual weather patterns, multiplied with their frequency. The division into different patterns is primarily based on the variation of the position and strength of the typical jet stream. The jet stream is a corridor of high-velocity wind speeds in eastbound direction. Table 2.1 shows the characteristics of the jet in each of the weather patterns. In the case study, only one weather pattern is considered. This is the one of the winter scenarios, characterized by a strong, zonal jet stream.

Table 2.1: Characteristics of North Atlantic weather patterns. WP stands for winter pattern, SP for summer pattern. Adapted from [5].

Type	Jet Stream		Frequency (days/season)
	Position	Strength	
WP1	Zonal	Strong	17
WP2	Tilted	Strong	17
WP3	Tilted	Weak	15
WP4	Confined	Strong	15
WP5	Confined	Weak	26
SP1	Zonal	Strong	19
SP2	Weakly tilted	Weak	55
SP3	Strongly tilted	Weak	18

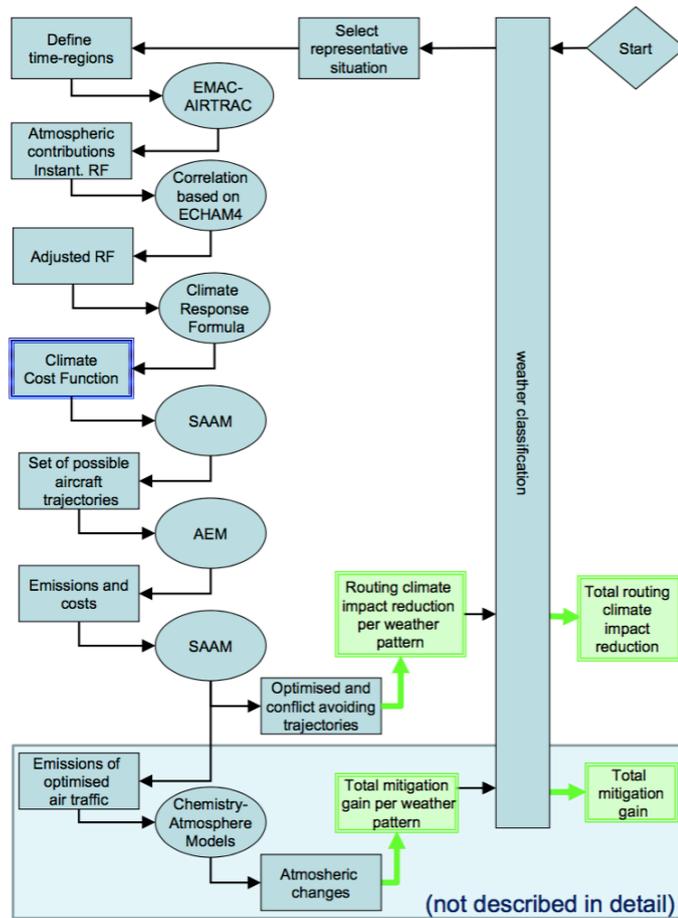


Figure 2.1: Outline of the model chain. Ovals represent models, rectangles contain data and definitions. The green boxes indicate major results. Adopted from [3].

2.1.2. Definition of time-regions and computation of climate cost functions

After establishing the weather patterns, a numerical time-region grid is defined across the North Atlantic region (see Figure 2.1). This is a 4D grid. Each grid point hence represents a 3D location, but also a point in time. This grid is used to compute the CCFs. A CCF contains information on the climate impact of an emission species emitted at each individual grid point. A CCF hence tells the user the impact of emission species A emitted at longitude x , latitude y and altitude z at time t ¹. It can be seen that a climate cost function is five-dimensional. Please note that the climate cost functions will differ if a different weather pattern is used. CCFs represent the climate impact of a unit emission, i.e. they are specific climate metrics. Climate metrics are a vital part of modeling climate impact. They form a common scale that can be used to compare the impact of different emitted substances. There is no such thing as a single perfect metric that suits all purposes. The choice for the metric to be used depends on multiple factors. For instance, what kinds of emissions are considered (pulse emissions, sustained emissions or some particular future scenario)? What is the timeframe of the impact under consideration (the impact integrated over a certain time horizon, or the impact at a certain point in time in the future)? Is it desired to have an evaluation of a parameter, or the rate of change of this parameter? The reader is referred to [6] for an overview of the metrics that are currently used. It is indicated in [7] that the choice of metric is facilitated by posing the right question. REACT4C made use of the following three climate metrics because they are most suited for the question “What would be the short-term and long-term effect on climate, if such a rerouting strategy were applied every day?” [4]: F-ATR20, P-AGWP20 and P-AGWP100. ATR means Average Temperature Response, being a measure for the temperature development over a period of time. In this the ATR is computed for a period of 20 years, based on a future (F) emission scenario. AGWP stands for Absolute Global Warming Potential, which is the time-integrated radiative forcing resulting from a pulse (P) emission, here considered for a time window of both 20 and 100 years. Radiative forcing can be understood as the balance of incoming radiation minus outgoing radiation. For more background of these metrics, the reader is referred to [3] and [6].

The climate cost functions are computed by means of the ECHAM/MESy Atmospheric Chemistry model (EMAC), see Figure 2.1. EMAC simulates atmospheric processes using submodels. Human influences, ocean-air and land-air interactions are included [8]. The fifth European Centre Hamburg general circulation model ECHAM5 is the core model for the atmosphere [9]. The chemical and physical processes that are modeled by the submodels are interlinked by the second version of the MESy2 (Modular Earth Submodel System) [10]. A review of how the EMAC model compares to other climate-chemistry models is made in [11].

2.1.3. Route optimization using the CCFs

Next to a climate model, there are models required that generate possible flight trajectories and corresponding emission profiles. To this end, the models SAAM and AEM are incorporated in the model chain (see Figure 2.1). SAAM stands for System for traffic Assignment and Analysis at a Macroscopic level. SAAM is developed by EUROCONTROL, being a participant of the REACT4C project. It can design air traffic routes taking into account operational constraints such as flight level restrictions. SAAM can optimize these trajectories by minimizing or maximizing an objective function that can be anything. REACT4C was the first to use it to optimize with respect to climate impact. Characteristics of aircraft performance are found using EUROCONTROL’s Base of Aircraft Data. This database makes it possible to estimate parameters like aircraft load and weight. However, to be able to optimize for climate impact, another model called Advanced Emission Model, also developed by EUROCONTROL, was incorporated in the model chain. AEM estimates the fuel burn, the emissions and the locations of the emissions along a certain flight route (4D) for a given aircraft-engine configuration. Estimated emissions are CO_2 , H_2O , NO_x , SO_x , CO and unburnt hydrocarbons. Multiplying the CCF with the outcome of the AEM and feeding this back into SAAM allows to optimize for climate impact.

REACT4C applies the approach to flights across the North Atlantic, as explained before. This is because there is a lot of air traffic in this region. Nevertheless, the CCF approach can be implemented to optimize flights on different routes on the globe as well, e.g. between Europe and Middle East, Middle East and Asia etc. If this is the case, a new aircraft fleet and reference routes should be considered. A new 4D grid and the respective CCFs would have to be generated as well.

¹The symbols A , x , y , z and t are arbitrarily chosen and are only used to illustrate the content of a CCF. Furthermore, they are discrete variables.

2.2. Outcome and results

Figure 2.2 shows what is perhaps the most important graph of the REACT4C project. It shows the optimal relation between climate impact and costs, called the Pareto front, for three metrics: P-AGWP100, P-AGWP20 and F-ATR20. The distinction is made between westbound and eastbound flights, because the presence of a jet stream creates a large difference in achievable climate impact reduction (more about this below). A reference case was defined to which the climate impact and costs of the optimized routes are compared. This case is the situation in which all the flights are optimized for economic costs.

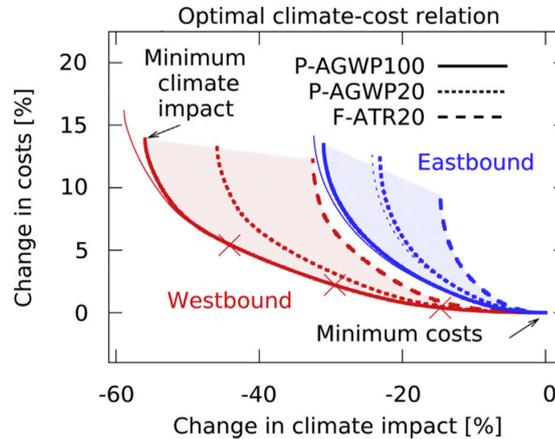


Figure 2.2: Relation between changes in economic costs and changes in climate impact for winter pattern 1. Westbound flights are depicted in red, eastbound flights in blue. Optimization done for three metrics: P-AGWP100, P-AGWP20 and F-ATR20. The crosses indicate 25%, 50% and 75% of the maximum achievable climate impact reduction. Thin lines indicate results when more routing options are considered. Adopted from [4].

Let us denote, for the ease of reading, the change in costs by ΔC (which we define to be positive for an increase in costs) and the change in climate impact by ΔCI (which we define to be positive for a climate impact *reduction*, in contrast to the horizontal axis of the graph in Figure 2.2), both relative to the cost-optimized reference case that is found at the point ($\Delta CI = 0, \Delta C = 0$) in the lower right corner of the graph. At this point, no reroutings have taken place and all aircraft are operated along the cheapest routes.

It can be seen that greatest benefit is achieved when considering long-term impacts (P-AGWP100) of westbound flights. The maximum achievable climate impact reduction, ΔCI_{max} , is around 60%. The associated ΔC is a little less than 15%. However, as large as the benefit for the climate is, airlines are not eager to increase their costs without increase in income. Therefore, it is useful to take a look at how much ΔCI can be achieved with less induced costs. This is indicated by the red crosses in Figure 2.2. They are positioned at the points at which $\Delta CI = 0.25\Delta CI_{max}$, $0.5\Delta CI_{max}$ and $0.75\Delta CI_{max}$. Let us start in the most promising corner of the graph, i.e. at large impact reductions at small costs. At $CI = 0.25\Delta CI_{max}$, ΔC is only about 0.5%. So, for very low incurred costs, the long-term climate impact of westbound trans-Atlantic flights can be reduced by about $0.25 \cdot 0.60 = 15\%$. At $\Delta CI = 0.5\Delta CI_{max} \approx 30\%$, the associated increase in costs is about 2.5%. This ΔC is a little larger but still relatively small. Then, at $\Delta CI = 0.75\Delta CI_{max} \approx 45\%$, ΔC amounts to around 5.5%. The most attractive rerouting option is perhaps one of the cases where the maximum climate impact reduction is not achieved, as still a large benefit can be acquired at smaller costs.

For eastbound flights, there is a smaller potential climate impact reduction. The largest reduction is present when long-term climate impacts are considered. ΔCI_{max} is around 30%, accompanied by a ΔC of around 14%. The reason for the smaller impact reduction potential is as follows. Aircraft flying from North America to Europe make use of the jet stream, which is in the direction of flight (tailwind). Any alternative route that leaves the jet stream would result in a much higher fuel consumption. Hence, there is no possibility to fly outside of the jet stream, reducing the amount of alternative routings. In contrast, westbound flights avoid the jet stream, as the wind direction is opposite to the flight direction (headwind). Any alternative route that enters the jet stream would increase fuel consumption. However, there are still many possible routes around the jet stream to reduce the climate impact.

It strikes the eye that the optimization with respect to long-term climate impact (metric P-AGWP100) shows a greater potential than optimization with respect to short-term impacts (metrics P-AGWP20 and F-ATR20). This is because CO_2 and CH_4 effects are more important for long-term objectives because of their longer lifetimes [2]. Methane shows a great impact reduction potential, and therefore long-term climate optimization leads to better results. As regards the short-term climate impact, the warming effect of ozone, with its smaller lifetime, dominates the longer-term reduced warming (read 'cooling') effect of methane.

The only case of which the trajectory changes are briefly considered is for the westbound flights and P-AGWP100. Figure 2.3 indicates the flight level, latitude and longitude changes when moving along the Pareto front from cost-optimized flights towards climate-optimized flights.

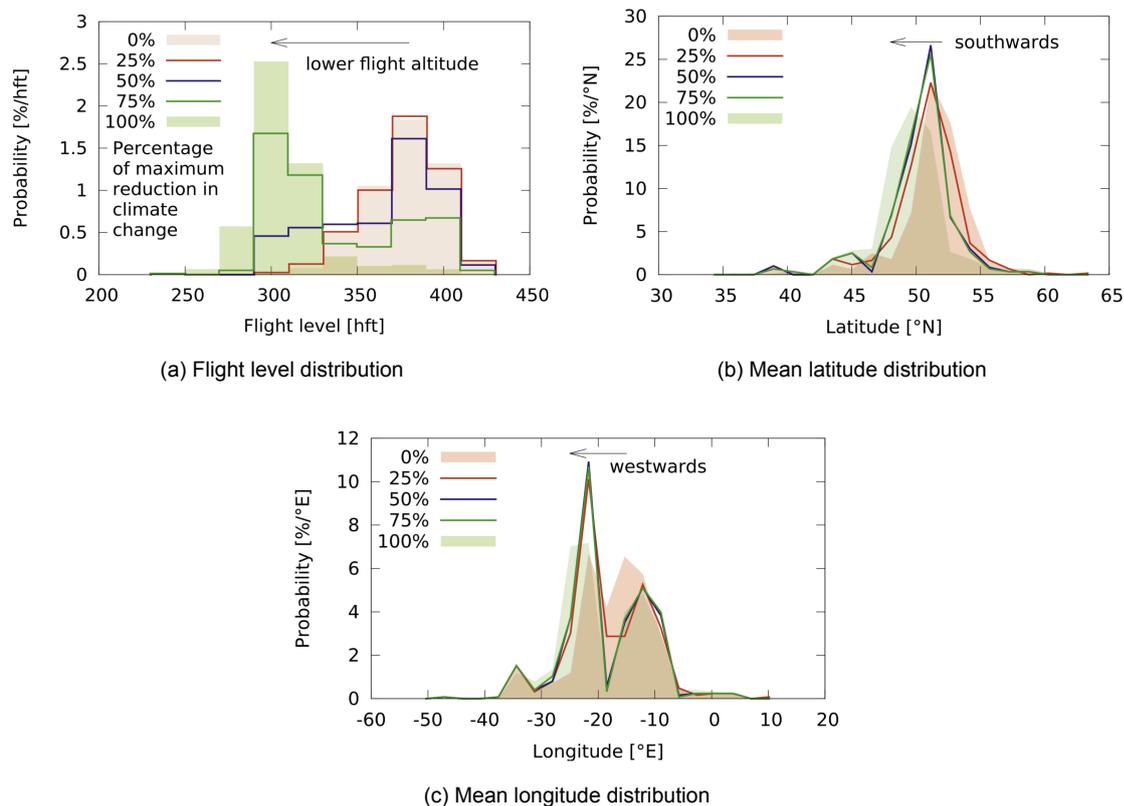


Figure 2.3: Variation of the probability density function of routing characteristics along the Pareto front for westbound traffic, metric P-AGWP100 and winter pattern 1. Adopted from [4].

First the flight altitude is considered. Figure 2.3a displays the flight level changes when moving along the red, solid-lined Pareto front of Figure 2.2. One can see that, as we move from the reference case towards 100% of the ΔCI_{max} , the altitudes are decreasing. The reference case shows a peak at FL380, while the flights under 100% of ΔCI_{max} show a peak at FL300. Furthermore this graph indicates that the altitudes of the flights are hardly altered when comparing the $0.25\Delta CI_{max}$ and the reference case, shown by the red line and red shaded area respectively. Figures 2.3b and 2.3c work in the same way, but show the mean latitude and mean longitude distribution, respectively. The trend is that flights are rerouted southwards and westwards with increasing climate impact reduction. The mean latitude distribution shows that once again there is not a lot of difference between the reference case and the $0.25\Delta CI_{max}$ case. The mean longitude distribution is altered more.

Overall, the results show that there is great potential in this newly developed CCF approach. On top of the results presented here, [4] performed a sensitivity study to investigate the effect of uncertainties in the computation of the CCFs. The hard numbers differ to some extent but the trend remains the same. Grewe et al. are positive about the possibility to implement the CCF modeling approach in air traffic planning, but acknowledge that still a couple of steps have to be completed before being able to do so.

3

Methodology

This chapter presents the methods that are used to arrive at the results that will be provided in chapter 4 and chapter 5. Before being able to do so, one has to be acquainted with the data sets that will be analyzed. This is done in section 3.1. Section 3.2 will explain the general approach that will be applied. After this, section 3.3 examines to what extent Pareto fronts of different combinations of direction, metric and weather type can be compared. This needs to be clarified to be able to do the general analysis of chapter 5. Finally, section 3.4 describes which trajectory characteristics are investigated, and how they are computed.

3.1. Data description

Before any method can be developed, it is required to get to know the data to be analyzed. The data that forms the basis of this research is output that stems from the REACT4C project, and is the result of the optimization performed by Eurocontrol. Two types of data sets are available for the analysis. The first type is the data that contains the information on the trajectories of the flights. The second type are the data sets that hold information on the climate impact and economic cost of the flights, i.e. the data that is used to establish a Pareto front. The two kinds of data sets will be discussed in sections 3.1.2 and 3.1.3, respectively. But first, there will be an elaboration on the amount of information available for analysis in subsection 3.1.1.

3.1.1. Information quantity and optimization problem

The goal of the REACT4C project was to provide very clear instructions on how to reroute flights for the benefit of climate. They performed multiple optimizations on the entire fleet, 288 in total. Each one of these optimizations serves a different purpose and is equally important as another. How the project members arrived at a total of 288 data sets is explained below.

Part of the novelty of the project is that it takes into account weather patterns. A total of eight weather cases were established in [5] (five winter patterns and three summer patterns). Another contributing factor to the number of optimizations that were performed is the fact that multiple climate metrics were considered. In total three of them were used, namely the F-ATR20, P-AGWP20 and P-AGWP100. Furthermore, the distinction was made between the two directions. One is the traffic coming from North America and heading towards Europe, the other is the other way around. Each direction consists of a fleet of approximately 400 flights.

At this point, there are 48 possible combinations of direction, metric and weather pattern. The final contribution to the number of optimizations is the fact that several weight factors were attributed to express the relative importance of operational costs and climate impact in the optimization. The optimizations were done in the form of Equation 3.1.

$$\begin{aligned} \min f &= (1 - \alpha) \cdot C_{eco} + \alpha \cdot K \cdot C_{clim}, \\ \text{with } K &= \frac{C_{eco_{clim-opt}} - C_{eco_{eco-opt}}}{C_{clim_{clim-opt}} - C_{clim_{eco-opt}}}, \\ \text{and } \alpha &= 0, 0.2, 0.4, 0.6, 0.8, 1. \end{aligned} \tag{3.1}$$

In this equation, C_{eco} is the economic cost in euro, and C_{clim} the climate cost in Kelvin when the ATR is considered, or in $\left[\frac{W \cdot yr}{m^2}\right]$ when the AGWP metrics are used. The subscripts *clim-opt* and *eco-opt* stand for fully climate-optimized and fully cost-optimized, respectively. The factor K is a factor that relates economic cost to climate cost, with unit $[\text{€}/K]$ or $[\text{€}/(W \cdot yr \cdot m^{-2})]$, therewith transforming the equation into units of euro. More about the relevance of K in section 3.3. The factor α is a (non-dimensional) scalar, used to attach different levels of importance to economic cost and climate cost in the optimization. It is predefined and varied between zero and one in six steps, as is shown in Table 3.1. These six values of α can be seen as six different levels of climate optimization, starting from *Clim 0.0*, the fully cost-optimized fleet, and ending at *Clim 1.0*, the fully climate-optimized fleet. Each value of α results in an optimal fleet, hence represents one point on the Pareto front. This is visualized in Figure 3.1¹. Throughout this document, the level of climate optimization will be written in the form of *Clim α* .

Table 3.1: Weight factor combinations used in the optimizations.

Weight factor [-]	<i>Clim 0.0</i>	<i>Clim 0.2</i>	<i>Clim 0.4</i>	<i>Clim 0.6</i>	<i>Clim 0.8</i>	<i>Clim 1.0</i>
α (climate cost)	0.0	0.2	0.4	0.6	0.8	1.0
$(1 - \alpha)$ (economic cost)	1.0	0.8	0.6	0.4	0.2	0.0

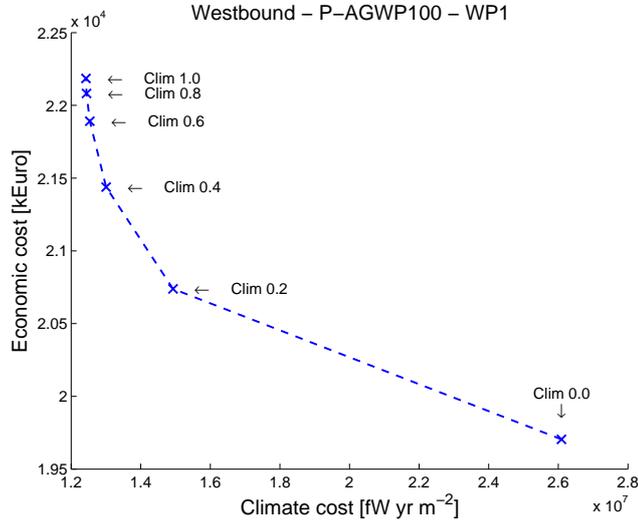


Figure 3.1: Pareto front of the case westbound/P-AGWP100/WP1. Crosses indicate available data sets.

The total amount of optimizations now adds up to 288. When referring to an optimization, it is therefore necessary to state which case is being looked at. The case is unambiguously defined when flight direction, weather pattern, metric and Pareto location are specified. These four variables make the distinction between the optimizations, and they will be called *case differentiators*.

3.1.2. Trajectory data

Each of the 288 cases has its own set of optimized flights. The information of the trajectories was provided in the form of .txt files. Each file presents information of all ≈ 400 flights for 1 combination of direction, climate metric and weather pattern. The flight routes are composed of segments. Each line in these files represents a flight segment. A typical line² looks as follows:

¹The Pareto front in Figure 3.1 is the same front as the one that represents westbound flights in WP1 conditions, optimized using metric P-AGWP100, in Figure 2.2. Note that in Figure 3.1 the climate impact and economic cost are expressed in absolute values, whereas in Figure 2.2 these parameters were expressed relative to the climate impact and economic cost of the cost-optimized fleet. Furthermore, the crosses on the front in Figure 2.2 represent the points at which the climate impact reduction is 25, 50 and 75 percent of the maximum achievable climate impact reduction, whereas in Figure 3.1 the crosses indicate data sets available for this research.

²This is an excerpt of the data set containing flight trajectories of the case eastbound/P-AGWP100/WP1/*Clim 1.0*.

```
$aaav_YQX KEWR EHAM B764 002701 022911 290 290 2 COA70 061230 061230
2486.38 -4370.93 2934 -3272 100910115 2 891.40 3
```

Table 3.2 explains what every parameter in such a line stands for. In this table, the data from the line above is used for illustration and clarification.

Table 3.2: Explanation of the parameters present in the provided trajectory data sets.

#	Example	Parameter	Comment
1	\$aaav_YQX	Segment identifier	Syntax: [first point name] _ [last point name]
2	KEWR	Airport of departure	ICAO code
3	EHAM	Airport of arrival	ICAO code
4	B764	Aircraft type	-
5	002701	Time begin segment	HHMMSS, padded with zeros
6	022911	Time end segment	HHMMSS, padded with zeros
7	290	Flight Level begin segment	-
8	290	Flight Level end segment	-
9	2	Status	0 = climb, 1 = descent, 2 = cruise
10	COA70	Callsign	-
11	061230	Date begin segment	YYMMDD, padded with zeros
12	061230	Date end segment	YYMMDD, padded with zeros
13	2486.38	Latitude begin segment	Expressed in minute decimals
14	-4370.93	Longitude begin segment	Expressed in minute decimals
15	2934	Latitude end segment	Expressed in minute decimals
16	-3272	Longitude end segment	Expressed in minute decimals
17	100910115	Unique flight identifier	Shows to which flight the segment belongs
18	2	Segment number	The number of the segment within the flight (time-sorted)
19	891.40	Segment length	Expressed in nautical miles
20	3	Segment parity/color	Parameter used by Eurocontrol, no further importance

3.1.3. Pareto front data

The information on the climate impact and economic cost of each optimized set of flights was provided in .xls files. There is one such file per combination of direction, metric and weather pattern. Each of these files contains the climate cost, expressed in Kelvin, and the economic cost, expressed in euro, of the six levels of climate optimization. This allows us to make a Pareto front for each case. In fact, one of these data sets was already used to provide the Pareto front in Figure 3.1.

The focus of this document is solely on the trajectories, and not on the climate impact reduction and cost increase they incur. In principle, these Pareto front data sets would not have to be used to be able to do the trajectory analysis. However, they will be needed to check whether the six steps of climate optimization are comparable cross-case. This will be done in section 3.3.

3.2. General approach

The research question to be answered was formulated as:

How is the air traffic between North America and Europe restructured when implementing the climate-optimized routes established by the REACT4C project, both on a case-specific level as well as in general?

This research question consists of two main parts that need to be elaborated further. The first one is the first half of the sentence: "How is the air traffic between North America and Europe restructured?" Note that examining the way, in which air traffic is restructured, is unrelated to the climate impact reduction and cost increase that result from the optimizations. This part of the question makes it clear that only the trajectory changes should be analyzed, and not the corresponding climate impact reduction or cost

increase. On top of that, the word *restructured* emphasizes that we need to investigate the *changes* of the trajectories. Stating what the climate-optimal trajectories look like hence is not satisfactory. It needs to be clear how they are changed with respect to the cost-optimized flights. Furthermore, this first part of the research question also raises the need to determine how air traffic restructuring is measured and quantified. It was determined that the following characteristics combined give a clear understanding of how the traffic is altered. Section 3.4 will give a more detailed explanation of what these characteristics represent, and how they are calculated.

1. The percentage of flights that are affected by the optimization;
2. The mean flight duration of the set of flights, normalized by the mean flight duration of the cost-optimized flights;
3. The probability density function of the flight duration increments with respect to the cost-optimized flights, giving an indication of how many of the flights show a large increase in flight duration, and how many show a small increase, or even a decrease;
4. The mean flight distance of the set of flights, normalized by the mean flight distance of the cost-optimized flights;
5. The probability density function of the flight distance increments with respect to the cost-optimized flights, giving an indication of how many of the flights show a large increase in flight distance, and how many show a small increase, or even a decrease;
6. The mean lateral shift depending on geographical location, i.e. the mean shift in latitude overlaid on a world map, with the shift being the change with respect to cost-optimized flights;
7. A classification of the climate-optimized trajectories, based on the lateral shift with respect their cost-optimized counterparts;
8. The mean altitude shift depending on geographical location, i.e. the mean height shift overlaid on a world map, with the shift being the change with respect to cost-optimized flights.

The second part of the research question is “[...] when implementing the climate-optimized routes established by the REACT4C project, both on a case-specific level as well as in general”. The difficulty here is that there is no such thing as *the* climate-optimized routes established by the REACT4C project. As was discussed in section 3.1, there are 240 data sets of flights that are optimized for climate to some extent (288 minus the *Clim* 0.0 data sets). How will this problem be tackled?

The necessity arises to provide two studies. The first one is a case study. In this case study, there is one base optimization that will be examined. This study hence retains a low level of generality, with results being very specific for that case. A programming script is made so that any case can be examined. Recall that a case was defined by four differentiators: direction, metric, weather pattern and Pareto location. The base case that will be considered in this thesis is westbound/P-AGWP100/WP1/*Clim* 1.0. This case was selected because Figure 2.2 on page 6 showed that this case leads to the largest potential climate impact reduction. The climate-cost relation of this case was discussed in [4]. On top of this, more information on the weather pattern and climate cost functions is provided. Also a sensitivity analysis was performed to investigate the effect of uncertainties in the climate cost functions. It is therefore a good choice to use this optimization for the case study.

The influence of each of the four case differentiators on the trajectories will also be examined in this case study, using a one-factor-at-a-time approach. This means that four comparisons are made: an inter-Pareto comparison, which investigates how each trajectory characteristic listed above varies with the level of climate optimization; a direction comparison, which examines the influence of flight direction on the trajectory changes; a metric comparison, discovering the influence of climate metrics; and finally the weather pattern comparison, again investigating the trajectory parameters listed above, but this time for each weather situation. These four analyses will be done based on the same base case defined above. For instance, the analysis of the influence of direction will make use of two data sets, namely eastbound/P-AGWP100/WP1/*Clim* 1.0 and westbound/P-AGWP100/WP1/*Clim* 1.0. Note that three out of the four case differentiators are kept the same as in the base case, and one case differentiator, in this case direction, is varied. For the research into the influence of the climate metric, the trajectories of three

data sets will be compared: westbound/F-ATR20/WP1/*Clim* 1.0, westbound/P-AGWP20/WP1/*Clim* 1.0 and westbound/P-AGWP100/WP1/*Clim* 1.0. In this case, the only case differentiator that is varied is the metric, the other ones are the same as in the base case.

The second study that is conducted is a general analysis of trajectory changes. In this general analysis, the main focus is on finding trends and making general observations. It retains the four comparisons that are also present in the case study, but now includes all available data sets, therewith increasing the level of generality. For example, the investigation of the *general* influence of flight direction on the trajectories will be done by comparing the average of a certain trajectory characteristic of all data sets in eastbound direction, with the average of the same characteristic of all data sets in westbound direction. Data sets of cost-optimized flights will not be included in these values.

The reason why the distinction is made between a case study and a general analysis is as follows. When performing a data analysis of such a large collection of data sets, the main goal is to discover trends rather than unravelling details. This is useful to have a very general understanding of the way trajectories are altered. However, REACT4C's intention was to be very specific about how flights should be rerouted for the benefit of climate. If the only conducted analysis would be averaging the mean values of mean values, most of the information would be lost and the actual usability of the results in practice would be reduced. The case study is provided to retain a high level of detail and to aid in implementing the routes in practice, if, or when, the aircraft have to be operated on climate-optimal routes. Any case could be studied with the tool that is created, however only one case will be discussed in this thesis. Furthermore, the case study serves as a check to what extent the general trends are also visible in the case-specific analysis, and as a means to obtain an indication of the representativeness of the general analysis.

3.3. Cross-case Pareto front examination

In the previous section, it was explained that a general study will be performed, which examines the influence of the level of climate optimization, the direction, the climate metric and the weather situation on the trajectories. In this study, all data sets are taken into account in the computations. It should therefore be examined whether the four case differentiators can be compared cross-case. Direction, metric and weather pattern can be compared without extra complexities. Direction is either eastbound or westbound, and there is no difference in direction between one eastbound case and another eastbound case. The same holds for the climate metrics: there is no difference in the climate metric between an F-ATR20 case and another F-ATR20 case. Similarly, the weather situation is the same between a WP1 case and another WP1 case. However, there is a difference in level of climate optimization between two fleets, both of which are optimized with e.g. $\alpha = 0.2$ in the objective function (which is repeated below for the ease of reading). Why this statement is true, and what it means for the inter-Pareto comparison of the general analysis, is explained below.

$$\begin{aligned} \min f &= (1 - \alpha) \cdot C_{eco} + \alpha \cdot K \cdot C_{clim} \\ \text{with } K &= \frac{C_{ecoclim-opt} - C_{ecoeco-opt}}{C_{climclim-opt} - C_{climeco-opt}} \\ \text{and } \alpha &= 0, 0.2, 0.4, 0.6, 0.8, 1. \end{aligned}$$

Recall that the parameter K in the objective function is necessary to convert the unit of climate cost to euro. For optimizations based on the ATR, its unit is $[\text{€}/K]$, whereas for AGWP-based optimizations, the unit is $[\text{€}/(W \cdot \text{yr} \cdot \text{m}^{-2})]$. The definition of K shows that this parameter requires estimates of the climate and economic cost of the fully climate-optimized and fully cost-optimized set of flights *before* the optimization is performed. These estimates are different for every combination of direction, metric and weather type, and by consequence, the factor K also changes for each of these combinations (next to the fact that the dimension of K differs between metrics). Therefore, even though α can be kept constant, $\alpha \cdot K$ is different for every direction/metric/weather combination, which has an influence on the location of the data set on the Pareto front.

Consider Figure 3.2, which depicts two Pareto fronts, belonging to different combinations of direction, metric and weather situation. Note that only the weather pattern differs between the two cases. On the Pareto front of the case westbound/P-AGWP100/WP1, the distance between the cost-optimized Pareto point (*Clim* 0.0, i.e. $\alpha = 0.0$) and the first Pareto point that takes climate impact into account

(*Clim* 0.2, i.e. $\alpha = 0.2$) is twice as large as the same distance on the other Pareto front. Note that on both fronts, these points are computed with the same scalars for climate impact, namely $\alpha = 0.0$ and 0.2. The difference in distance between *Clim* 0.0 and *Clim* 0.2 is the influence of K . K hence is responsible for the resolution of the data sets along the Pareto front. The difference in inter-Pareto distance implies that the level of climate optimization of a data set is not unambiguously quantified by the scalar α , but rather by the distance along the Pareto front.

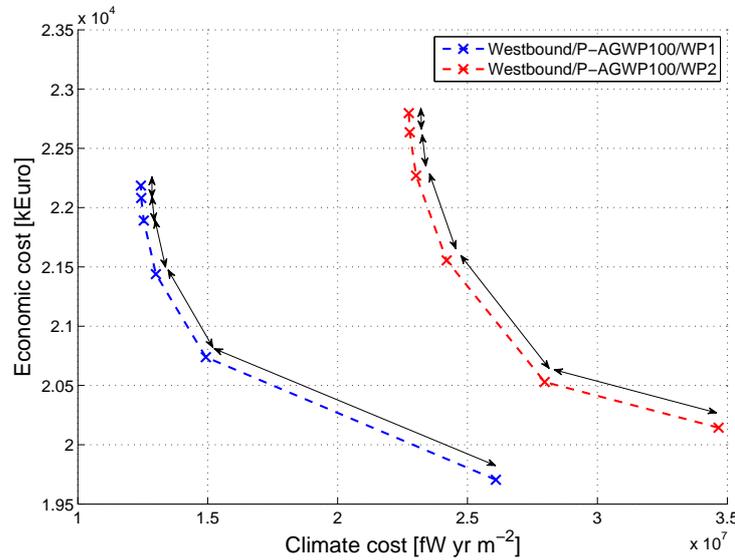


Figure 3.2: The Pareto fronts of the cases westbound/P-AGWP100/WP1 and westbound/P-AGWP100/WP2. Crosses indicate the Pareto locations at which flight trajectory data sets are available.

The two Pareto fronts in Figure 3.2 only represent two out of 48 combinations of direction, metric and weather situation. To examine the variability of the distances between each Pareto point, also called inter-Pareto distances, for all combinations, the distance between two locations is normalized by the total arc length of the Pareto front of the case to which the points belong. This results in the box plots of Figure 3.3. It can be observed that the distance between *Clim* 0.0 and *Clim* 0.2 is not similar between cases. In some cases, *Clim* 0.2 is located at just over 20% of the length of the Pareto front, while in others *Clim* 0.2 is located at over 90 percent of the Pareto arc length. The differences between inter-Pareto distances become smaller when moving towards higher levels of climate optimization. Furthermore, the mean inter-Pareto distance between two data sets becomes increasingly smaller towards *Clim* 1.0. The distances between *Clim* 0.8 and *Clim* 1.0 are very similar cross-case, but also very close to zero. This is also visible in Figure 3.2: towards the end of the Pareto fronts, the data sets become more accumulated.

As long as the trajectory differences are examined within the same Pareto front (like in the case study), it is satisfactory to consider the six data sets along the Pareto front to be six discrete levels of climate optimization (ranging from *Clim* 0.0 to *Clim* 1.0), because K is the same for each of these data sets. However, the box plots confirm that the scalar α is not fit to quantify the level of climate optimization for the inter-Pareto comparison in the general analysis. A Pareto location will at that point be redefined as the arc length along the front, starting from the cost-optimized location (*Clim* 0.0) up until the Pareto point in question, divided by the total arc length of the front³. Hence, this is a value between 0 and 1. All cases will have a data set at Pareto locations 0 and 1. However, with 192 data points in between 0 and 1, Pareto location becomes a quasi-continuous variable.

³The arc length is approximated by summing the straight-line distances between each Pareto location. In Figure 3.2, this would mean adding the double-headed arrows along the Pareto front.

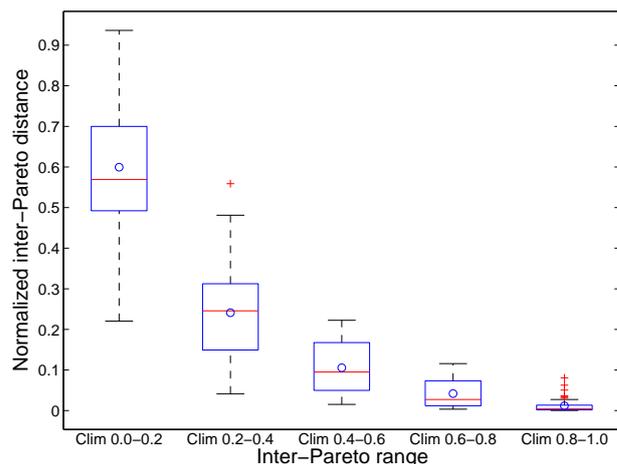


Figure 3.3: Box plots of the normalized inter-Pareto distance between each data set along the Pareto front.

3.4. Investigated flight characteristics

This section explains the way in which the trajectories are analyzed. The flight characteristics that will be examined, were already listed in section 3.2, but now more details will be provided on the way they are computed. Where necessary, the distinction will be made between the case study and the general study.

3.4.1. Mean flight duration

The characteristic that will be explored first are the flight times of the routes. Given are the start and end times of each trajectory segment. These are indicated by numbers 5 and 6 in Table 3.2 on page 11. One can see that the format of this data entry is HHMMSS. To facilitate calculations with these times, they are converted into seconds. This gives the point in time expressed as seconds elapsed since midnight. Now the time spent in a segment can be computed. The next issue is to detect which flight a segment belongs to. This is done by using number 17 in Table 3.2. Then the total flight time is easily computed by adding the durations of each segment. The average flight duration of the whole set of flights in that data set is then computed. To compare between the several Pareto points, this computation is done repeatedly for the six points along the front, after which a plot can be made of average flight time versus Pareto location, normalized by the times of the cost-optimal flights. To compare the influence on travel time of the other case differentiators, the *Clim* 1.0 mean flight time of that direction/metric/weather combination will be normalized with the *Clim* 0.0 of the same case.

For the general analysis, the normalized mean flight duration will be computed for all 288 Pareto locations in the investigation of the influence of Pareto location, and a trend line will be drawn. For the other comparisons, box plots will be made of the normalized mean flight duration for each alternative of the case differentiator under consideration.

3.4.2. PDF of flight duration increments

A probability density function (PDF) will be made for the flight duration increments, such that it is visualized how many of the flights durations are increased to a certain extent. To do this, the difference of the travel time of each of the flights in a data set is computed with respect to the cost-optimized counterpart. There will be multiple PDFs in one plot, depending on the number of alternatives of the case differentiator of which the influence is being examined. For instance, for the case study of the influence of direction, there will be two PDFs in the plot: one for eastbound/P-AGWP100/WP1/*Clim* 1.0 and one for westbound/P-AGWP100/WP1/*Clim* 1.0. These plots will also include the mean time increase, expressed in minutes.

In the general analysis, these PDFs will also be provided. Making use of the same example, i.e. the influence of direction, the PDFs of every case in each direction will be averaged. This way, we still

end up with two PDFs in one plot, representing the mean time increment distribution per direction. For the analysis of inter-Pareto differences, the PDFs of each of the 288 Pareto points will be sorted on Pareto location and stacked behind each other, basically resulting in a 3D plot. A contour plot will then be generated to visualize the spread of the PDFs as a function of Pareto location.

3.4.3. Mean flight distance

The distance of one flight segment is provided in the data sets (number 19 in Table 3.2). Determining the flight distance of each flight is done by summing the segment lengths of segments that belong to the same flight. The mean flight distance for the complete data set is then calculated by averaging the distances of every flight. This value is then normalized by the mean flight distance of the set of flights of the same case, but at the *Clim* 0.0 point. The resulting plots in both the case study and general analysis are of the same kind as for the mean flight duration, but now for distance.

3.4.4. PDF of flight distance increments

The PDFs of the flight distance increments are computed in exactly the same way as the ones of the flight duration increments. They will show how many of the flight distances are increased to what extent.

3.4.5. Percentage of rerouted flights

The percentage of rerouted flights is a good indication of the extent to which the flights are affected by the optimization. At this stage, the flight duration and distance of each individual flight is known. This information can be used to detect whether or not every flight route in the fleet has been altered during the optimization. A flight is considered to be unchanged when it meets these two criteria:

- The flight duration is equal to the flight duration of the same flight in the *Clim* 0.0 situation;
- The flight distance is equal to the flight distance of the same flight in the *Clim* 0.0 situation.

This allows us to determine the percentage of altered flights for every data set. If less flights are affected by the optimization, it would be easier to implement climate-optimal rerouting strategies in practice.

3.4.6. Shift in latitude

To investigate how the air traffic is restructured, one of the most important aspects to examine is the south- or northward shift of the flight trajectories. The goal is to be able to clearly visualize how the trajectory shifts laterally, as a function of Pareto location, direction, metric and weather pattern. This can be achieved by making 2D color-coded plots, showing longitude and latitude on the x- and y-axis respectively. The z-direction (i.e. the color) then depicts the change in latitude relative to the reference case of *Clim* 0.0. Such a graph can then be made for each of the alternatives of the case differentiator under investigation, which enables us to inspect the differences between them.

The method to produce this kind of plots involves a couple of steps, the first one of which is to be able to compare the same flight between different Pareto locations of the same case. A generic diagram of this 1D interpolation is given in Figure 3.4 for a flight between airports *A* and *B*. As the trajectories change when moving along the Pareto front, the number of segments per flight changes. In the illustration, the cost-optimized route consists of 5 segments, whereas the climate-optimized counterpart is composed of 8 segments. At each of the nodes, the longitude and latitude is given. The nodes of the cost- and climate-optimized trajectories differ. Hence, to be able to compute the shift in latitude at a certain longitude, the latitude of each route has to be interpolated at the same query longitudes. Let us consider one query longitude, lon_{q_1} , at which we want to know the lateral shift (see Figure 3.4). This longitude does not coincide with a node, neither on the reference route nor on the climate-optimized route. By consequence, the trajectory has to be interpolated between start and end point of segment 2 of the most economic route, and of segment 4^{opt} of the climate-optimized route. The resulting query latitudes are lat_{q_1} for the reference flight and $lat_{q_1}^{opt}$ for the climate-optimized route. The difference between the two, Δlat_{q_1} , is the lateral shift at query longitude lon_{q_1} . 150 query longitudes are defined, equally spaced between 135°W and 50°E. Hence, we have 150 query longitudes lon_{q_i} , with i ranging from 1 to 150, each resulting in a latitude lat_{q_i} of the reference route, a latitude $lat_{q_i}^{opt}$ of the climate-optimal route, and a lateral shift Δlat_{q_i} . If the route does not exist on a certain query longitude, the query latitude is returned as “not a number”.

A complication arises during the calculation of the lateral shift by means of this 1D interpolation. Near the airports, some trajectories do not follow the direction of flight, but the aircraft fly in a loop

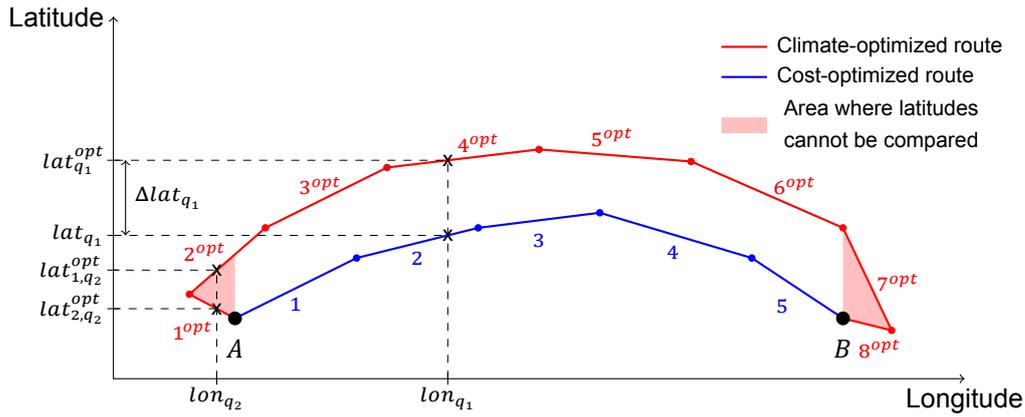


Figure 3.4: Generic diagram of the 1D interpolation of the same flight at different Pareto locations. The two routes belong to the same combination of direction, weather and metric, but to different levels of climate optimization. The blue route is the cost-optimized route, the red one is a route that is optimized for climate (can be any level of climate optimization).

after take-off or before arrival. Segment 1^{opt} in Figure 3.4, for instance, is heading westbound, before continuing in eastbound direction as of segment 2^{opt} . Consider a second query longitude, lon_{q_2} in Figure 3.4. The original route does not pass this longitude, whereas the climate-optimized trajectory does (twice, at latitudes lat_{1,q_2}^{opt} and lat_{z,q_2}^{opt}). Hence, the lateral shift between the two routes cannot be determined at this longitude. The same occurs around airport B . The longitudes at which the lateral shifts between the two routes cannot be computed, are indicated by the shaded areas in Figure 3.4.

The information obtained at this stage can be summarized as: for any flight in any data set, the lateral shift at 150 predefined longitudes between 135°W and 50°E is determined, relative to the cost-optimal trajectory ($Clim\ 0.0$) of the same flight, for the same combination of direction, metric and weather situation. The way in which this information becomes interpretable is when it is visualized in a plot, with longitude in x-direction, latitude in y-direction and the lateral shift in the z-direction. To do this, another interpolation is conducted, this time a 2D interpolation. The two independent variables are the longitude and latitude. The value to be interpolated is the lateral shift. The lateral shift can be seen as a function of the longitude and latitude. The sample longitudes for this interpolation are the same as the query longitudes that were used in the 1D interpolation (lon_{q_i} , with $i = 1, \dots, 150$). The sample latitudes are the query latitudes resulting from the 1D interpolation of the reference Pareto point, $Clim\ 0.0$ (i.e. lat_{q_i} , and not $lat_{q_i}^{opt}$, with $i = 1, \dots, 150$). This is required because the goal is to see what the average lateral shift at a certain longitude-latitude combination is, if the route of the cost-optimized flight passes this longitude-latitude combination. The sample function values are the lateral shifts at these geographic locations (Δlat_{q_i} , with $i = 1, \dots, 150$), the computation of which was made possible because of the 1D interpolation. The sample points consist of every query longitude, latitude and lateral shift of every flight in a data set. This 2D interpolation removes duplicate data points and averages their corresponding function values. This is required, because some flights within one data set (partially) coincide. Hence the average lateral shift of the routes crossing that geographical location should be taken.

A surface plot can now be generated. This is a 3D plot, which becomes easily interpretable when considering the top view of this plot, together with color coding. This view presents a map, i.e. longitude versus latitude, with colors on this map indicating the average lateral shift if a cost-optimized route crosses that location. Making this plot for every Pareto location, direction, metric or weather pattern, depending on the case differentiator that is being examined, allows us to compare the lateral shifts.

In the general analysis, it is impossible to compare such kind of plots between all the cases. The lateral shift will be examined differently and with less details. This is done by generalizing the data in three steps. First, the lateral shift will be averaged for one flight, such that each flight is characterized by one mean shift in latitude. When this is done for every flight in the data set, the mean of these average lateral shifts is computed, such that the result is one mean lateral shift for the entire data set. Finally, these mean values are grouped per alternative of the case differentiator under investigation, after which box plots are made per alternative. These box plots hence depict the variation in the mean latitude change. For the inter-Pareto analysis on the general scale, trend lines will be drawn of the mean, maximum and minimum lateral shift as a function of the Pareto location instead of producing

box plots.

3.4.7. Trajectory classification based on lateral shift

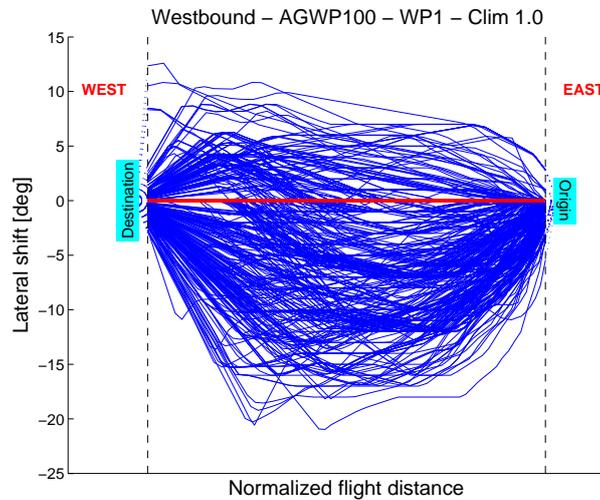


Figure 3.5: Trajectory scheme of all fully climate-optimized flights of the case westbound/P-AGWP100/WP1. The red horizontal line represents the original flight trajectory. The vertical dashed lines represent the borders between which lateral shift can be computed (more in the text).

The lateral shift investigated in the way described in subsection 3.4.6 does not give an impression of how the routes are altered individually. Is there a trend in the way trajectories are altered with respect to the original, cost-optimized trajectories? To investigate this, six trajectory categories are distinguished. Figure 3.5 shows the lateral shift with respect to the cost-optimized trajectory, for every flight in one case. The red line is the reference line: the lateral shift with respect to the original flight is zero along the entire trajectory. The vertical dashed lines delimit the range of longitudes at which the latitudes can be compared (cfr. the shaded areas in Figure 3.4). Recall that this is dependent on whether or not a trajectory makes a loop around the origin or destination airport. The trajectories that reach the destination airport at about 10 degrees of lateral shift, in the top left of the figure, are characterized by a large loop around the airport and are exceptions. The types of trajectory schemes are identified based on visual inspection of this figure.

In total, six different categories were found. Figure 3.6a presents an example of the first category. This is the class in which a flight is rerouted south over the entire trajectory. Part of trajectories may coincide, as long as the lateral shifts that do take place, are towards the south of the original trajectory. A representation of the second class is given in Figure 3.6b. This is similar to the first class, but now the shift should be towards the north. Class three is illustrated in Figure 3.6c. Here, the western part of the flight is rerouted north. Then, the climate-optimized trajectory crosses the cost-optimized trajectory, resulting in a southern shift in the eastern part of the flight. Figure 3.6d depicts a representation of the fourth class. This category is similar to the third one, but here the western part of the flight is rerouted south, and the eastern part north. The fifth category consists of flights that have not been shifted laterally, and the climate-optimal trajectories coincide with the original one. An example of this class is shown in Figure 3.6e. Finally, the sixth class, not provided in Figure 3.6, contains the trajectories that could not be categorized into any of the other classes.

The percentage of flights within each class can be established for every data set and can be compared between the Pareto locations, directions, metrics and weather types.

3.4.8. Altitude shift

The altitude shift is computed in exactly the same fashion as the latitude shift in subsection 3.4.6. However, now the interpolated values in the 2D interpolation are the changes in height instead of the lateral shifts. The result is a color plot of the shift in altitude, with longitude and latitude on x- and y-axis, respectively. For the general analysis, it is not possible to visually compare all cases. Therefore, all

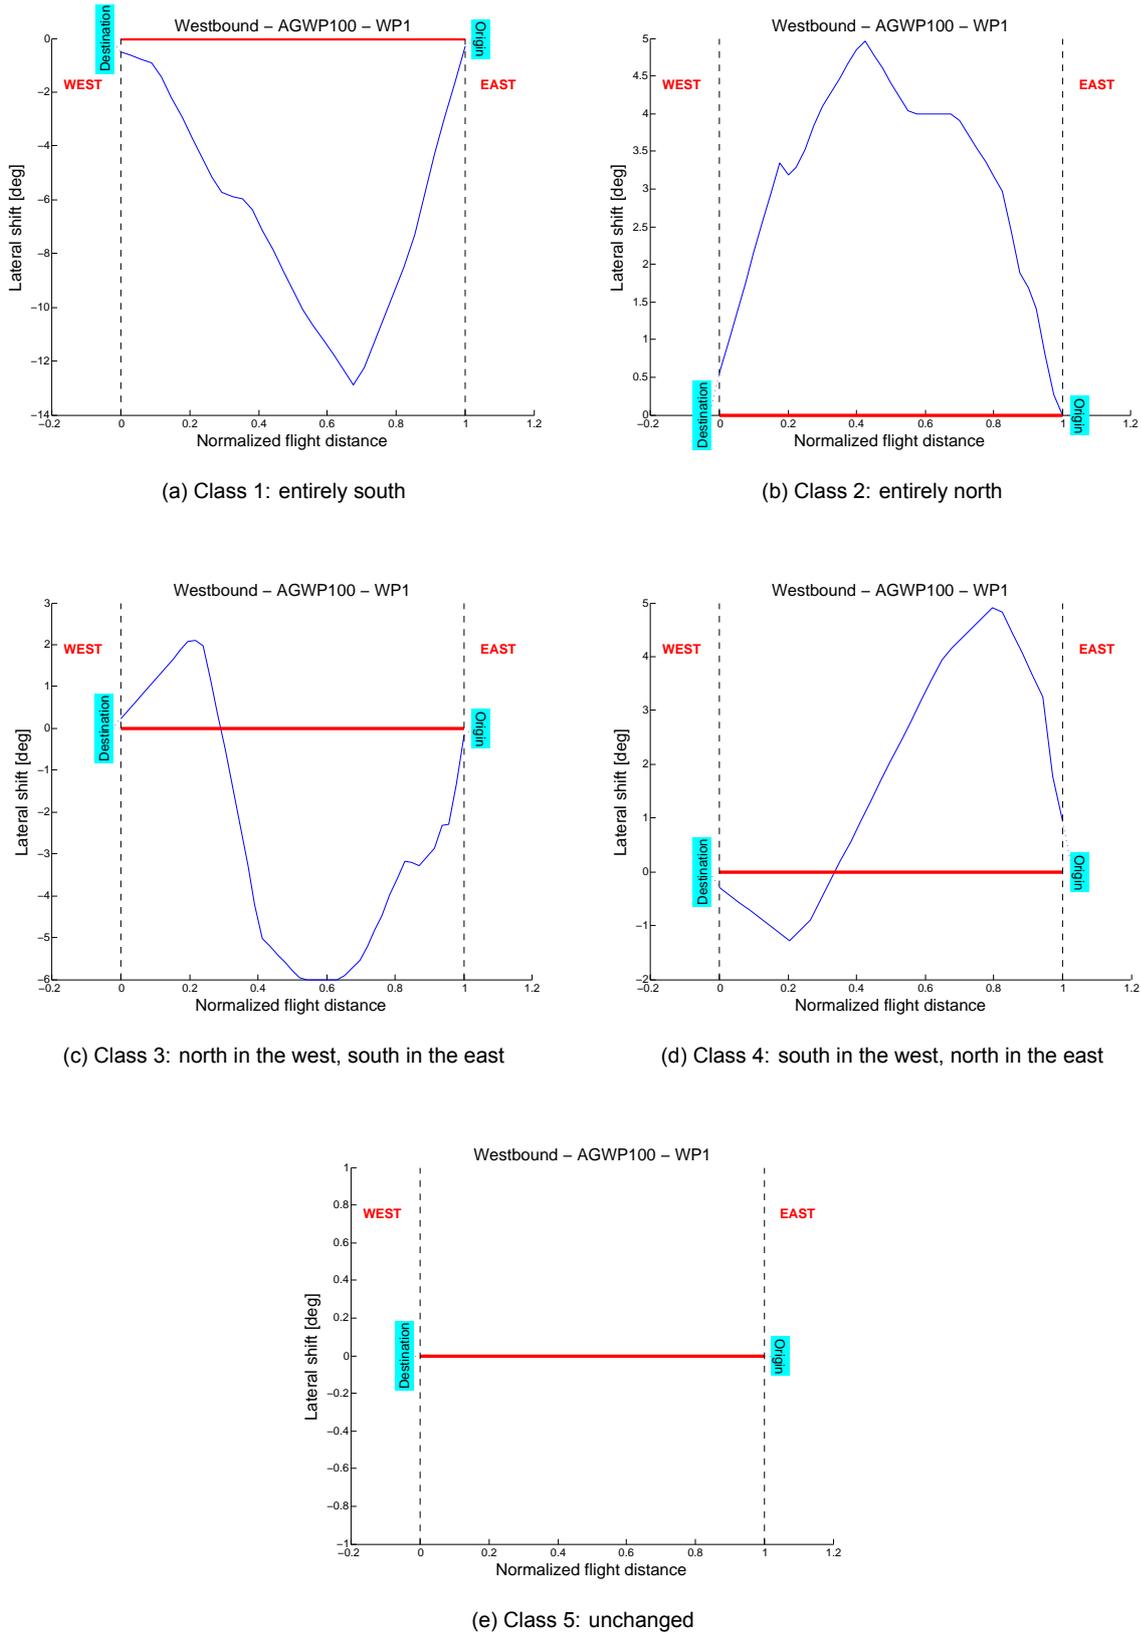


Figure 3.6: Example of each trajectory class based on lateral shift. Sample trajectories are taken from the case westbound/P-AGWP100/WP1/Clim 1.0. The red horizontal line represents the original flight trajectory. The vertical dashed lines represent the borders between which lateral shift can be computed (more in the text).

data sets will be assigned one mean altitude shift, and all these mean values will be presented in box plots.

4

Case study of trajectory changes

This chapter commences the trajectory analysis of REACT4C's climate-optimized flights. In this chapter, one base case is considered. The influence of each of the case differentiators is investigated relative to this base case. The results are very specific to this combination of direction, metric, weather pattern and level of optimization, and the influence of each of these differentiators does not hold for any other case. Chapter 5 is meant to examine whether these influences are general trends.

The base case in the study presented here is the set of flights that stems from the route optimization of westbound traffic, using P-AGWP100 as the metric for climate impact, and the first winter weather pattern that was established by [5]. The choice for this combination of direction, metric and weather pattern was made because of the fact that the climate impact reduction of this case proved to be the most promising of the ones that were examined in [4].

First it will be examined how the trajectory characteristics change when moving along the Pareto front (section 4.1). Section 4.2 investigates the influence of flight direction on the trajectory alteration. Next, it is explored whether, and how, the choice of metric influences the optimization results (section 4.3). Finally, trajectory differences between weather patterns will be discussed in section 4.4. Each of these sections will have a separate conclusion in which the main findings are recapitulated, and at the end of the chapter, a schematic overview of the whole chapter is provided (section 4.5).

4.1. Inter-Pareto trajectory changes

This section will make use of the six data sets belonging to the six different levels of climate optimization of the westbound/P-AGWP100/WP1 flights. Each one relates to a certain point on the Pareto front. Figure 3.1 on page 10 showed the Pareto front of this case, with the crosses indicating the points at which a data set of the flight trajectories is available. Each Pareto location is given a name, ranging from *Clim 0.0* to *Clim 1.0*, expressing the weight factor of climate impact in the objective function. In this section we will show what the effect on the trajectories is when increasing the level of optimization from cost-optimized flights towards fully climate-optimized air traffic, hence moving along the Pareto front.

4.1.1. Rerouted flights

The first parameter that is investigated, is the percentage of flights that is affected by the optimization. The way in which this percentage varies as a function of Pareto location, is shown in Table 4.1.

Table 4.1: Percentage of affected flights as a function of Pareto location. Case: westbound/P-AGWP100/WP1.

<i>Clim 0.0</i>	<i>Clim 0.2</i>	<i>Clim 0.4</i>	<i>Clim 0.6</i>	<i>Clim 0.8</i>	<i>Clim 1.0</i>
0	98.48	99.75	100.0	100.0	100.0

In this case, it is clear that practically all flights are rerouted, no matter which level of climate optimization is being considered. However, there is a slight trend visible that indicates that at lower levels

of climate optimization, less flights are affected. Nevertheless, the differences are very small and can be neglected.

4.1.2. Flight time and distance

Figure 4.1 shows the variation of the mean flight duration and mean flight distance depending on the importance of climate impact in the optimization. The values are normalized with respect to the reference case, being the flights optimized with respect to economic cost.

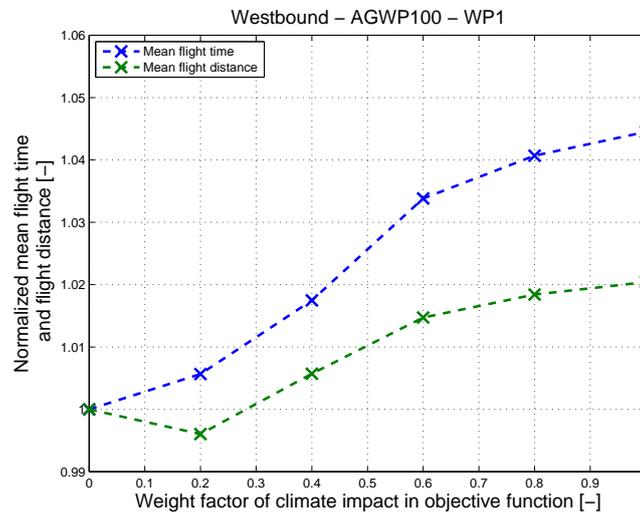


Figure 4.1: Mean flight time and mean flight distance normalized by the mean flight time and distance of the cost-optimized flights, as a function of Pareto location.

The flight duration curve shows a clear trend. As increasing emphasis is put on climate impact during the optimization, the mean flight durations increase. Between *Clim* 0.0 and *Clim* 0.2 the curve shows a moderate positive gradient. Then, between the climate cost weight factor of 0.2 and 0.6, the flight times rise more steeply, after which the curve flattens out again towards *Clim* 1.0.

The curve of the mean flight distances shows almost the same pattern. However, there is a noticeable minimum at the climate cost weight factor 0.2. The normalized mean flight distance here is slightly lower than one. This means that overall, the flights have decreased in length. At first sight, this seems to be in conflict with the fact that at this point the average flight duration is greater than one. However, the reason behind this is that apparently the mean ground speed at which the fleet of aircraft propagates also changes. When the fleet's average velocity decreases, it becomes possible for the mean flight distance to decrease, while simultaneously the average flight time is increased.

The difference between the two curves in Figure 4.1 gives an indication of how the mean velocity of the fleet of aircraft, relative to the ground, has changed with respect to the cost-optimized fleet. When the mean travel time increased more than the mean distance, percentage-wise, it means that the fleet of aircraft was slowed down in general. When moving along the Pareto front, the difference between the two curves becomes larger, hence the average speed of the fleet is more and more reduced.

Figure 4.2 shows the probability density functions of the difference in flight time for the various Pareto points. A vertical dashed line is drawn at the zero on the horizontal axis. This zero represents the point at which the flights are unaltered in terms of flight duration. The odds that a travel time is shortened is given by the area under the curve on the left hand side of the zero, while the probability of an increase in flight time is the area on the right hand side. The probabilities that a flight is shortened in terms of flight distance are also provided in Table 4.2.

The first thing to notice is that flight duration can be changed by anything between -40 minutes and +100 minutes. It can be observed that for *Clim* 0.2 — the first point on the Pareto front where climate impact is taken into account, although to a small extent — there is a peak right before the zero. This means that a lot of flights show a slightly smaller flight time. However, the flights that do fly longer, show a larger duration difference. This is why the *Clim* 0.2 point in Figure 4.1 is greater than one. For all other weight factor combinations, it can be observed that the peak decreases and the curve is more

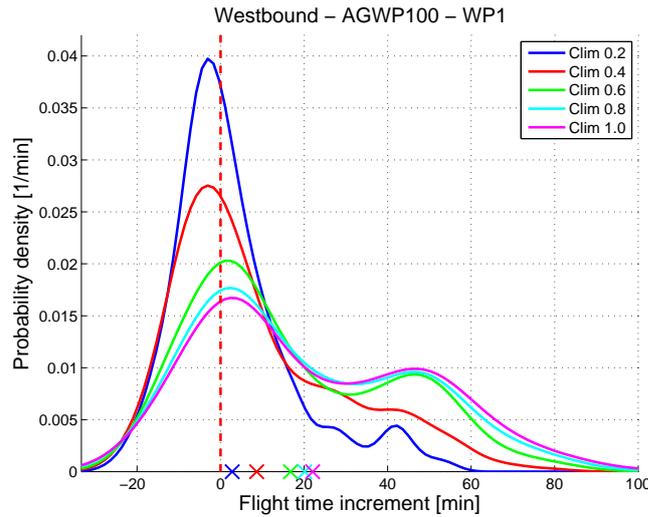


Figure 4.2: Probability density functions of flight duration increments with respect to the cost-optimized flights for the various levels of climate optimization. Zero flight time difference is indicated by the red dashed line. The crosses on the x-axis represent the mean flight duration increase.

Table 4.2: Probability that a flight is shortened in terms of flight time and flight distance as a function of Pareto location. Case: westbound/P-AGWP100/WP1.

	<i>Clim 0.0</i>	<i>Clim 0.2</i>	<i>Clim 0.4</i>	<i>Clim 0.6</i>	<i>Clim 0.8</i>	<i>Clim 1.0</i>
Flight time	0	0.506	0.421	0.296	0.264	0.249
Flight distance	0	0.540	0.414	0.298	0.264	0.247

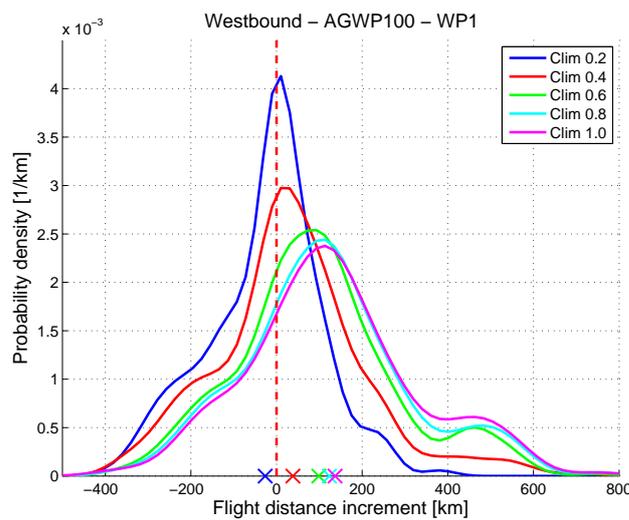


Figure 4.3: Probability density functions of flight distance increments with respect to the cost-optimized flights for the various levels of climate optimization. Zero flight distance difference is indicated by the red dashed line. The crosses on the x-axis represent the mean flight distance increase.

widespread. For all these Pareto locations, the probability that a flight duration is longer, is greater than the probability that it is shorter (see Table 4.2), and the accompanying travel times become greater. Hence, the general trend is that increasing the weight of climate impact in the optimization leads to more flights having a longer flight duration, and the flight duration increments are larger.

The same plot is produced in Figure 4.3, but this time for the change in travelled distance. The respective probabilities that the distance of a flight is decreased are given in Table 4.2. It can be observed that the flight distances are altered by anything between -500 and +800 km. In this plot, the curves are somewhat peakier and less spread than it was the case for time differences. The peak for *Clim* 0.2 arises at zero change in flight distance. This means that a large number of flights stay largely unaffected in terms of travelled distance when optimizing with a small weight factor attributed to climate impact. The probability of a flight being shortened in terms of distance for weight factor 0.2 is 0.540, which is more than half. However, in this case, the distance differences of the shorter flights are greater than the differences of the flights that have become longer. Therefore, the mean distance from Figure 4.1 at *Clim* 0.2 becomes less than one. The peaks gradually decrease and the curves shift to the right when increasing the climate weight factor. The probability of a flight being shortened decreases (see Table 4.2). This implies that more and more flights are rerouted with a deviation rather than a shortcut, and that the distance increments become larger as we move along the Pareto front.

4.1.3. Shift in latitude and trajectory classification

The result of the multi-step interpolation method explained in subsection 3.4.6 is a 3D surface of the average lateral shift of the trajectories depending on the latitude and longitude. Taking the top view of this graph and producing it for all six Pareto points results in the plots depicted in Figure 4.4.

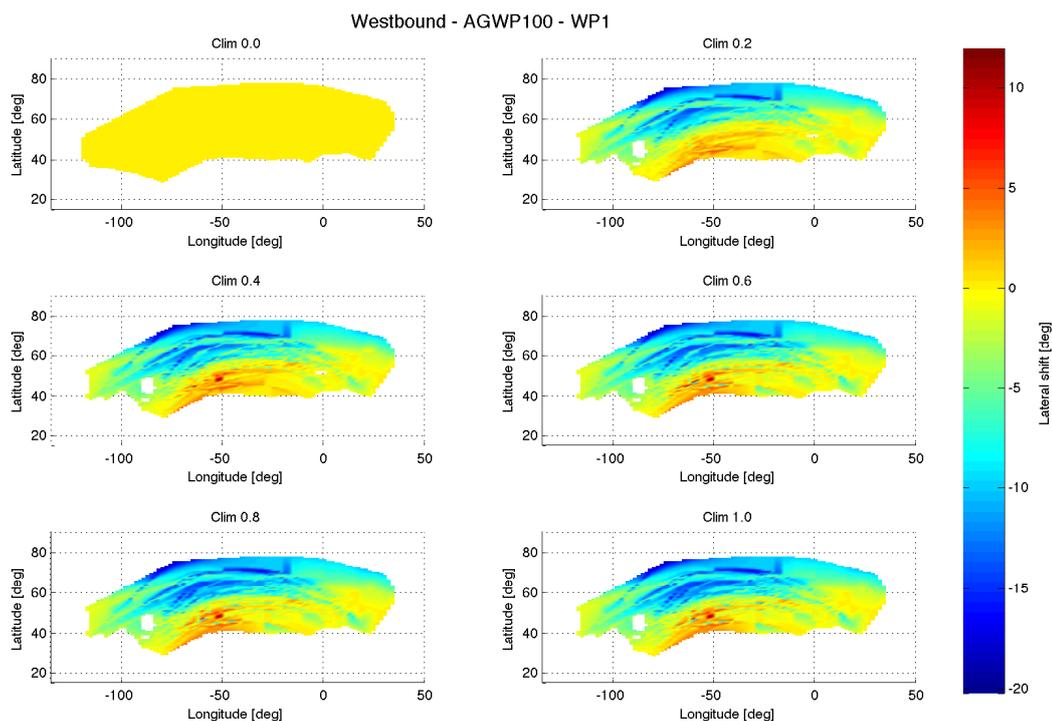


Figure 4.4: Overview of the average shift in latitude depending on geographical location for all Pareto points. A positive lateral shift represents a northward rerouting, a negative one a southward.

The first plot in Figure 4.4 is the reference case, i.e. the flights optimized for economic costs. Hence it is logical that there is no change in latitude. The second Pareto point, *Clim* 0.2, immediately sets the trend for the rest of the figure. The northern half of the set of trajectories is rerouted by between -5 and -20 degrees. This implies that flights are rerouted southwards in this region. The southern part is completely opposite. The shifts in this region are between 0 and 12 degrees, which respectively means that flights that pass through these locations are rerouted in northern direction or not rerouted laterally at all. When progressing to the next Pareto point, there are subtle changes in the intensities of the colors.

The southern part of the fleet, at around -50° longitude, shows that the lateral shifts towards the north increase a bit in magnitude in that region. The northern part is shifted south somewhat more. Generally, the shifts that were present at *Clim* 0.2 now become a little larger. For the remainder of Pareto points, there are not a lot of differences to be spotted. This is in line with Figure 4.1. The average flight distance curve is flattened towards the more climate-focussed optimizations. This indicates that there is not a lot of route alteration going on in terms of longitudes and latitudes. This is due to the fact these Pareto locations are located very closely together on the Pareto front (see Figure 3.1). In general, the latitude shifts of this case range from 20°S to 12°N .

It can be observed that there are some areas on the plots that are excluded from the interpolation. This is due to the fact that trajectories that make a loop around the airport, no matter how small, could not be interpolated correctly. This was clarified before in subsection 3.4.6, but now the impact on the result can be seen. These hiatuses are small compared to the rest of the plot, and affect the results only in such a way that it cannot be determined what the average shift in latitude is at those locations.

These average lateral shifts are useful to relate geographical location to the route changes. However, they do not provide information about how the trajectories are rerouted individually. Figure 4.5 provides the percentages of flights that fall within the categories identified in subsection 3.4.7, for each level of climate impact importance.

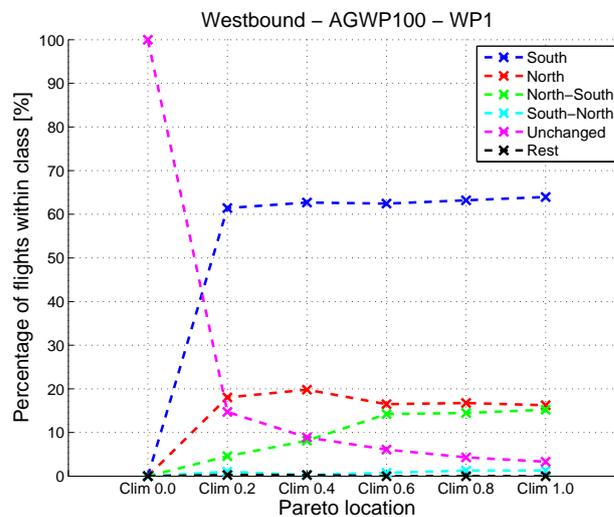


Figure 4.5: Variation with Pareto location of percentage of rerouted flights that fall within each trajectory class.

It is immediately visible that the majority of the flights (around 62% for each optimization level) are rerouted south with respect to their cost-optimized counterparts, independent of which point on the Pareto front is considered. The second largest contribution, with around 17% for each Pareto location, are the flights that are shifted north. The two categories that are significantly changing when moving along the Pareto front are the geographically unaltered flights, and the flights that are shifted north in the western part of the flight and south in the eastern part. The former class is clearly decreasing in frequency, whereas the latter shows an increasing number of occurrences. It can be noticed that the percentages of unaltered flights here are different than the ones presented in Table 4.1. This is because the only criterion here is that the geographical flight path is unaffected, whereas in the table also the flight time had to be the same. For this case, it can be concluded that especially at higher weight factors, there is no significant change in the way the routes are altered individually.

4.1.4. Altitude shift

The same kind of interpolation can be done to obtain the average altitude shift as a function of geographical location. The result can be found in Figure 4.6. As can be observed, *Clim* 0.2 still shows a lot of green areas, which means the flights that originally passed these regions did not shift in altitude. However, there are areas that show a decrease in altitude of approximately 2 kilometers. Moving to the next level of climate optimization, it can be seen that the regions that show a decrease in altitude become more prominent and less flights are unaltered in terms of height. This trend is taken a step

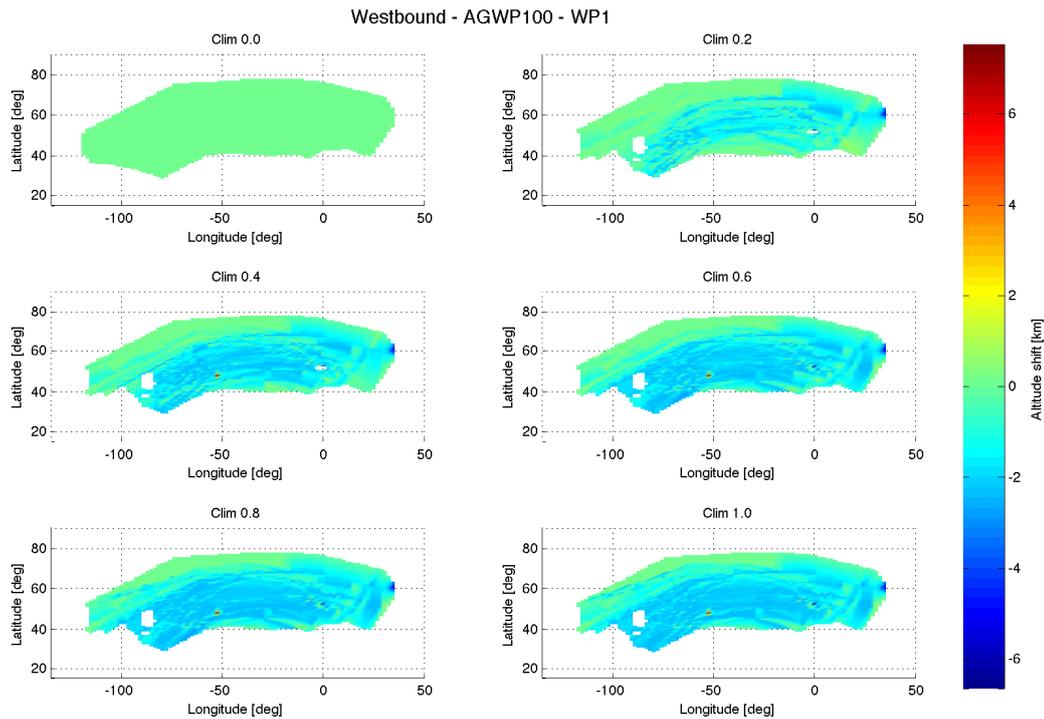


Figure 4.6: Overview of the average shift in altitude depending on geographical location for all Pareto points. A positive shift represents an increase in height, a negative shift a decrease.

further when moving on to *Clim 0.6*. Increasing the weight of climate impact further barely influences the average height shift any further. The extreme altitude shifts that are found in the small areas at approx. $50^{\circ}\text{N } 50^{\circ}\text{W}$, $50^{\circ}\text{N } 0^{\circ}\text{W}$ and $60^{\circ}\text{N } 35^{\circ}\text{E}$ are regions around an airport. This means that the approach was initiated further west, increasing the average height, or that the departure trajectories are stretched in western direction, such that the average height is reduced at these locations.

4.1.5. Conclusion of the inter-Pareto case study

This section compared different levels of climate optimization of westbound transatlantic flights under winter pattern 1 weather conditions and with climate impact quantified by the metric P-AGWP100. This study led to a couple of conclusions. The first and most important one is that increasing the weight of climate impact in the objective function, hence moving along the Pareto front from cost-optimized flights to climate-optimized flights, results in more extreme trajectory changes. However, it was a general observation that at higher climate impact weight factors, the trajectories are only altered marginally and differences can be neglected. This is due to the fact that the higher Pareto points are located much closer to each other, and the inter-Pareto distance is less (see Figure 3.1 on page 10). Next, it was seen that the aircraft are gradually slowed down more when moving from cost-optimized to climate-optimized flights. Another finding was that the southern part of the air traffic between Europe and America is mostly rerouted north, while the northern part of the fleet shifts south, a tendency that also intensifies with Pareto location. Furthermore, it was seen that the majority of the flights are shifted south. Finally, only downward height shifts are present in this case.

4.2. Influence of flight direction on trajectories

This section aims to examine the differences between the optimization between eastbound and westbound route optimizations. To do this, two data sets will be taken a closer look at: the fully climate-optimized (*Clim 1.0*) data sets of west- and eastbound flights with metric P-AGWP100 and winter pattern 1.

4.2.1. Rerouted flights

The number of flights that are rerouted at maximum climate impact reduction does not change with direction. In both cases, every flight has been altered. Recall that this means that either the trajectory has been shifted, and/or the flight time has changed. At maximum climate impact reduction, it is logical that all flights are affected. If not, it would mean that the flights in question were already flying in climate-optimal conditions when they were optimized for economic cost. Hence, this result was to be expected.

4.2.2. Flight time and distance

The mean flight duration and mean flight distance as a function of the flight direction are listed in Table 4.3. The values of eastbound and westbound flights are normalized with respect to the cost-optimized (*Clim* 0.0) flights in eastward and westward direction, respectively.

Table 4.3: Mean flight time and mean flight distance normalized by the mean flight time and distance of the cost-optimized flights, as a function of direction. Case: P-AGWP100/WP1/*Clim* 1.0.

	Eastbound	Westbound
Mean flight time	1.052	1.045
Mean flight distance	1.022	1.020

It can be observed that the increase in mean travel time of the westbound flights is 0.7% less than for its eastbound counterpart. In absolute terms, both directions show an average increase of around 22 minutes per flight. However, the eastbound flights at *Clim* 0.0 are on average already about an hour shorter because of the fact that they make use of the eastbound jet stream, which allows the aircraft to fly faster and consume less fuel. Therefore, the relative increase is greater. Furthermore, in both directions, the flight times are increased more than the distances, percentage-wise. This indicates that the mean ground speed of the fleet reduced, and hence a lower speed proves to be more beneficial for the climate for this weather pattern.

The average travel distance increase is a little smaller in westbound direction as well. Therefore, in general the flights are slightly more elongated in terms of distance. However, the difference is very small and can be considered to be non-existing.

The fact that there is no large difference in mean flight time and duration does not inherently mean that the distribution of the duration differences and distance differences is the same. Indeed, Figures 4.7 and 4.8 and Table 4.4 show that there is a clear distinction between eastbound and westbound reroutings. Figure 4.7 shows that most eastbound travel times are lengthened by between 0 and 50 minutes. Westbound flights are less likely to fall within the same range. However, outside of this range, there are significantly more westbound flights than eastbound flights. This works in two ways: more flights heading towards North America show a shortened travel time (see Table 4.4), and at the same time more routes show a large increase in flight duration. This explains why the averages are so similar.

Table 4.4: Probability that the flights are shortened in terms of flight time and flight distance, as a function of flight direction. Case: P-AGWP100/WP1/*Clim* 1.0.

	Eastbound	Westbound
Flight time	0.113	0.249
Flight distance	0.294	0.247

Figure 4.8 depicts the flight distance increment PDFs. In this case, flights towards Europe are more likely to undergo moderate distance increases of between 0 and 250 km. Outside of this range, the likelihood of flights towards America is greater, especially within the ranges [250, 500] km and [-500, -200] km. This, again, results in the practically equal mean flight distances of Table 4.3, although the probability that a flight is shortened in terms of distance is less for westbound flights (see Table 4.4).

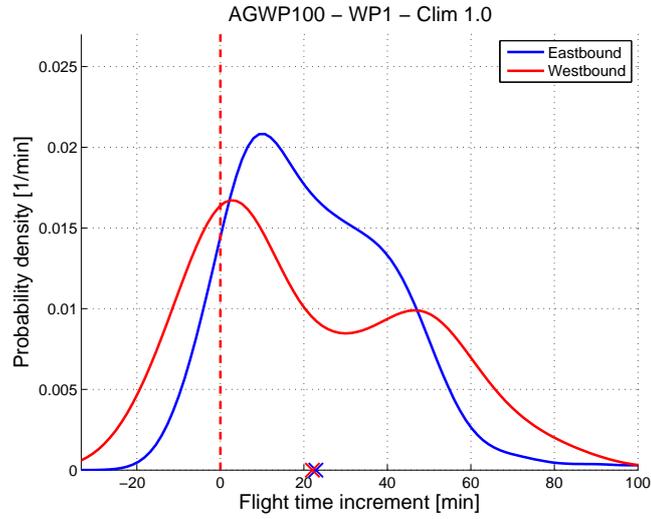


Figure 4.7: Probability density functions of flight duration increments with respect to the cost-optimized flights for the two flight directions. Zero flight time difference is indicated by the red dashed line. The crosses on the x-axis represent the mean flight duration increase.

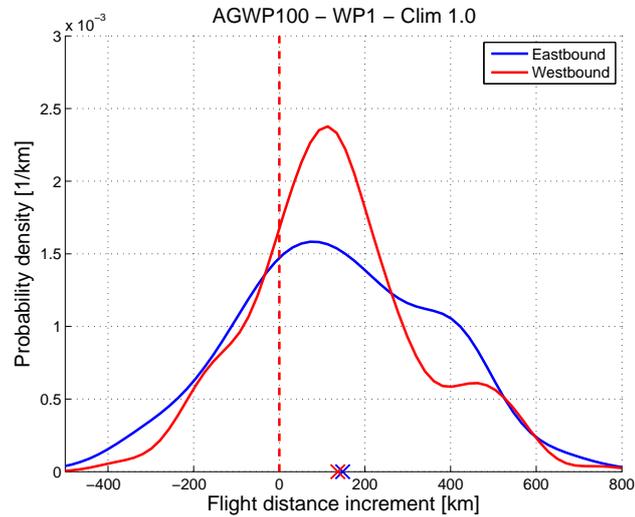


Figure 4.8: Probability density functions of flight distance increments with respect to the cost-optimized flights for the two flight directions. Zero flight distance difference is indicated by the red dashed line. The crosses on the x-axis represent the mean flight distance increase.

4.2.3. Shift in latitude and trajectory classification

Figure 4.9 provides the average lateral shift as a function of geographical location. These are the shifts with respect to different reference cases. The eastbound is constructed by subtracting the fully cost-optimized fleet of the case eastbound/P-AGWP100/WP1 from the fully cost-optimized fleet of the same case. The other plot works in the same fashion, but here the case is westbound/P-AGWP100/WP1.

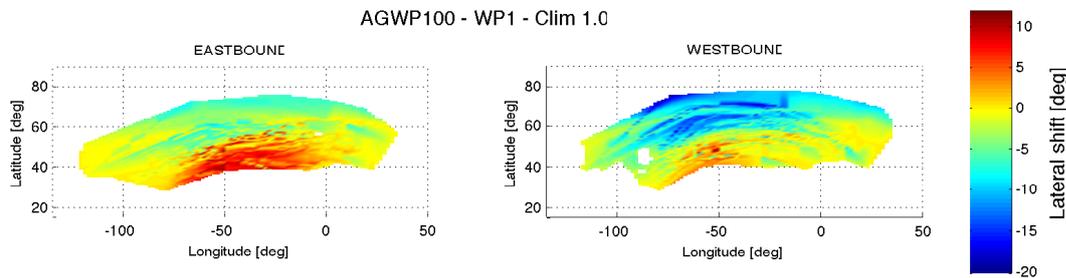


Figure 4.9: Overview of the average shift in latitude depending on geographical location for the two directions. A positive lateral shift represents a northward rerouting, a negative one a southward.

The main trend in the two plots remains the same. The southern part of the set of trajectories is shifted north, while the northern part is shifted south. The big difference here is the intensity of the displacements. The flights operating from west to east are found to be shifted more north. That is, the southward displacements at the higher latitudes are more moderate, with a maximum of around -8 degrees, and the northward shifts more excessive, with a maximum of around 11 degrees. The inverse is true for flights towards America. Here the values range from approximately -20 degrees in the northern part to around 10 degrees at a very small region above the ocean.

Let us now have a closer look at how the routes are categorized into the previously established classes. The percentage of flights that fall within each category can be found in Figure 4.10. One characteristic immediately strikes the eye. The majority of eastbound flights (56.5%) is rerouted north with respect to their original route, in contrast to the flights in the other direction, of which 64.9% are rerouted south. This is in line with the results of Figure 4.9, and has to do with the fact that westbound flights need to avoid the jetstream, whereas eastbound flights eagerly make use of it. Another striking observation is that the North-South class contains almost 12% more flights if the traffic is headed towards America, while in the other direction more flights are shifted south on the western part of the trajectory and north on the eastern part. Finally, it can be observed that there are slightly more flights that are geographically unaffected by the optimization in westbound direction.

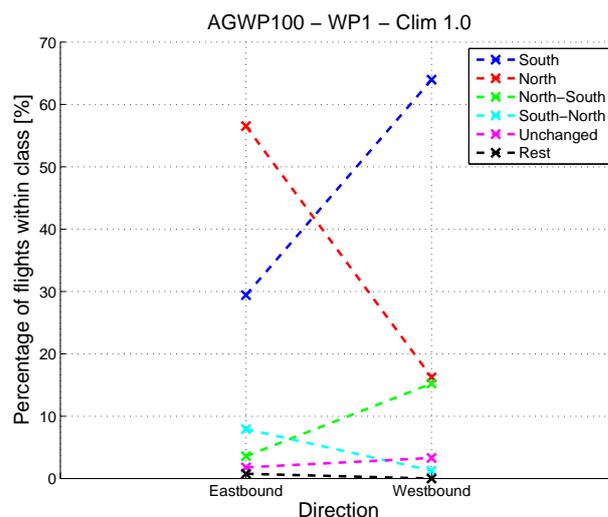


Figure 4.10: Percentage of rerouted flights that fall within each class as a function of flight direction.

4.2.4. Altitude shift

The average height displacements are given in Figure 4.11. The differences are smaller compared to the latitude shift comparison. Generally, the changes in altitude of flights towards the east are somewhat higher than for westbound flights. This is not because of the fact that the maximum altitude decrease is less, but because there are more regions that come close to this maximum (negative) shift of around 2.5 km. Note that some regions near the airports show a slight *increase* in height when the aircraft is flying towards America. This depends on the local approach and departure trajectories, and is not really important for the purpose of understanding how the REACT4C routes are restructured.

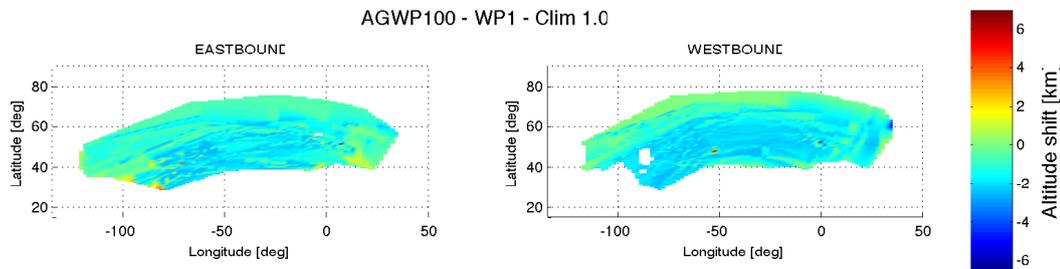


Figure 4.11: Overview of the average shift in altitude depending on geographical location for the two directions. A positive shift represents an increase in height, a negative shift a decrease.

4.2.5. Conclusion of the case study of the influence of direction

To conclude the case study of trajectory differences between the two directions, the main findings are summarized here. First of all, all routes are affected by the optimization in both directions. The mean flight time and duration increases are very similar for east- and westbound flights, percentage-wise. The big difference lies in the distributions of travel time and distance increments and decrements. Westbound flight durations were more often shortened and lengthened by a great amount, whereas more extreme distance differences were more often found in traffic towards Europe. Furthermore, it was found that there are significant differences in the shift in latitude. Eastbound flights tend to be rerouted north by a larger amount of degrees, and the number of trajectories that is shifted completely north is also greater. Westbound trajectories on the other hand are most often shifted south, doing so by larger lateral shifts towards the south. Finally, the altitude proved to be slightly more consistently lowered for flights towards America. The other direction shows the same maximum downward shift, but this shift is less evenly spread geographically.

4.3. Influence of metric on trajectories

The next step in this case study is to compare the impact of the climate metric that is used to quantify climate impact during the optimization. Three data sets are used for this comparison. The combination of direction, weather pattern and level of climate optimization westbound/WP1/Clim1.0 is the case under investigation. The three data sets are the data sets of this case for the three different metrics.

4.3.1. Rerouted flights

As the fully climate-optimized flight trajectories are being considered in this case study, it is to be expected that the percentage of flights that are modified is close to, if not exactly, 100 percent. Indeed, for all three metrics, there is not a single flight that has been left unaltered. Hence, from this perspective, metrics do not have an influence on the number of flights that are rerouted when they are fully optimized with respect to climate impact.

4.3.2. Flight time and distance

Figure 4.12 depicts the variation of the mean travel time and distance, normalized by the respective cost-optimized mean values. It shows that both parameters do not vary significantly when using a different metric in the optimization. The small variation that is visible suggests that the flights are rerouted a little more when using the absolute global warming potential, and still a little more when

targeting long-term climate effects. For all metrics, the flight time increases over 2% more than the travelled distance, indicating that the fleet’s mean velocity is reduced.

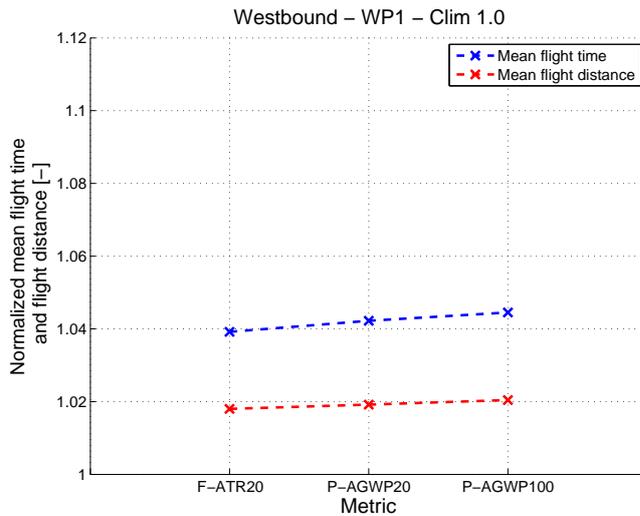


Figure 4.12: Mean flight time and mean flight distance normalized by the mean flight time and distance of the cost-optimized flights, as a function of climate metric.

The absence of substantial differences in mean flight time and distance between metrics makes it interesting to take a look at the probability density functions of the time and distance increments. Figure 4.13 provides the graphs for the increments in flight time. The associated probabilities that a flight time is reduced can be found in Table 4.5. In contrast to the direction comparison, where the mean values of duration and distance were similar but the distributions were not (see subsection 4.2.2), here the PDFs do not show large differences either. The largest difference occurs at small travel time increases of between zero and 20 minutes, where the F-ATR20 peaks and the long-term climate impact metric P-AGWP100 contains the lowest number of flights. At large increments of 50 minutes and more, the roles are reversed and optimization using P-AGWP100 leads to more flights within this range. The total effect of these differences is that the mean flight duration is hardly altered, as was also determined in Figure 4.12.

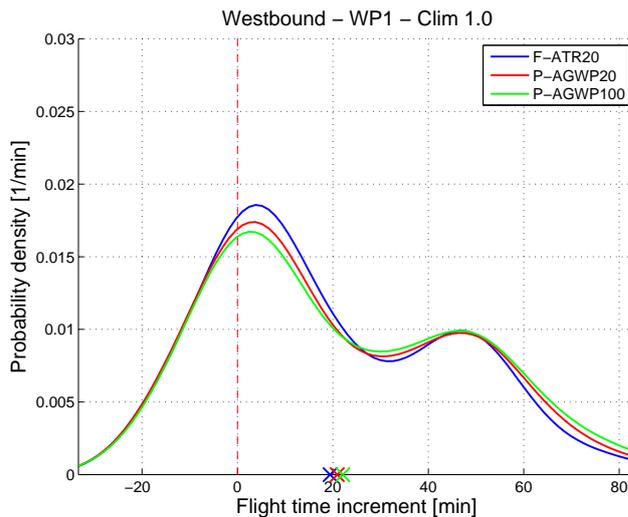


Figure 4.13: Probability density functions of flight duration increments with respect to the cost-optimized flights for the three metrics. Zero flight time difference is indicated by the red dashed line. The crosses on the x-axis represent the mean flight duration increase.

The flight distance increment distributions are depicted in Figure 4.14. Any differences between the three metrics are again very small. It is safe to say that the choice of climate metric does not influence the flight distance difference distribution.

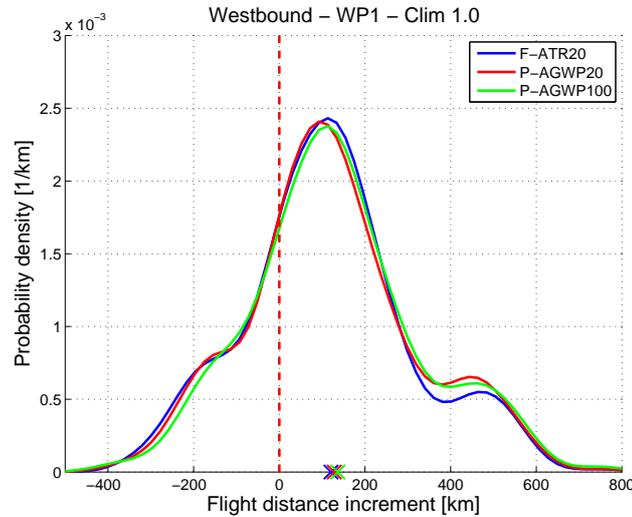


Figure 4.14: Probability density functions of flight distance increments with respect to the cost-optimized flights for the three metrics. Zero flight distance difference is indicated by the red dashed line. The crosses on the x-axis represent the mean flight distance increase.

Table 4.5: Probability that a flight is shortened in terms of flight time and flight distance, as a function of metric. Case: westbound/WP1/Clim 1.0.

	F-ATR20	P-AGWP20	P-AGWP100
Flight time	0.257	0.256	0.249
Flight distance	0.263	0.256	0.247

4.3.3. Shift in latitude and trajectory classification

Figure 4.15 provides the overviews of the average lateral shift depending on geographical location for the various metrics. Once again, there is hardly any difference between the three plots. Only the general trend that northern and southern flights are rerouted differently is visible again. All three plots show the same lateral shift range of around -20 degrees to +12 degrees.

From the previous plot it is to be expected that the categorization of flights into the trajectory classes is very similar from metric to metric as well. Figure 4.16 provides us with the overview of percentages. Indeed, the layout of the trajectories hardly changes. For each metric, the large majority of the flights are rerouted completely south compared to their original trajectories. This category is the one that changes most clearly, although still very subtly, with 3% of difference between F-ATR20 and P-AGWP100. Then, there are almost as many flights that are rerouted completely north as there are trajectories that are shifted north in the West and south in the East. The last remarkable property is that there are still around 4% of routes that are not geographically affected, but only show a change in flight time. This has to do with the fact that the routes are optimized taking into account conflict avoidance, ensuring the safety by maintaining separation between the aircraft. This results in flights flying slower or faster, otherwise they would be separated by less than the minimal required distance. In general, once again not a lot of differences between the metrics are present.

4.3.4. Altitude shift

The last parameter to be investigated is the altitude. The altitude shift with respect to cost-optimized flights for every metric is shown in Figure 4.17. There is a small difference between the metrics. Going

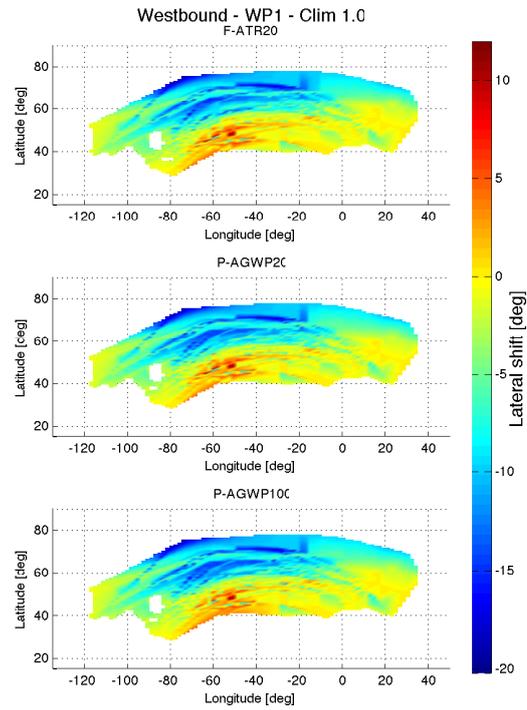


Figure 4.15: Overview of the average shift in latitude depending on geographical location for the three metrics. A positive lateral shift represents a northward rerouting, a negative one a southward.

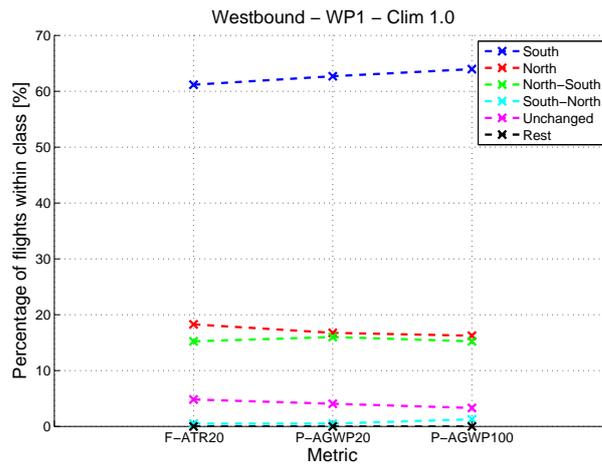


Figure 4.16: Percentage of rerouted flights that fall within each trajectory class as a function of metric.

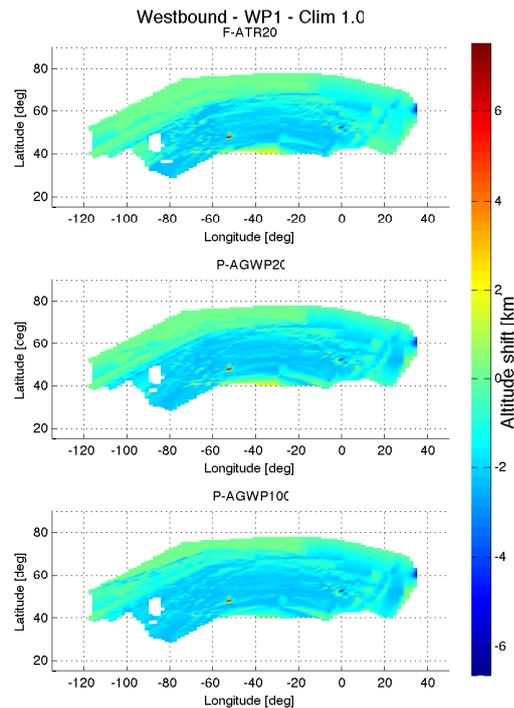


Figure 4.17: Overview of the average shift in altitude depending on geographical location for the two directions. A positive shift represents an increase in height, a negative shift a decrease.

from F-ATR20 to P-AGWP100, the most eastern part of the area spanned by the trajectories shows an increasing shift towards lower altitudes. Apart from that, the graphs look very similar. The average height decrease is around 2 km for all metrics when not taking into account some local extremities near airports. Furthermore, the most northern part of the fleet got shifted by a smaller amount than flights at lower latitudes for all metrics.

4.3.5. Conclusion of the case study of the influence of climate metrics

Comparing the same case for different climate metrics leads to one main conclusion: for this case, the choice of climate metric during the optimization hardly affects the resulting trajectories. All parameters that were investigated, showed very little to no differences. This non-variability suggests that the metrics agree on what the best route solution is for climate. The only difference they seem to make is in the absolute value of the climate impact (and hence climate impact reduction), which is a property that is not treated in this research. Section 5.3 of chapter 5 will examine whether this finding holds in general, i.e. for all cases.

4.4. Influence of weather pattern on trajectories

The last part of this chapter will be a case study on how the weather patterns identified in [5] influence the flight properties. For this study, we will again start with a base case of fully climate-optimized flights in westbound direction and metric P-AGWP100. Recall that there are eight weather patterns in total: five winter patterns and three summer patterns. Hence, there are eight data sets to be compared.

4.4.1. Rerouted flights

First it will be examined whether all flights are rerouted geographically and/or altered in terms of flight time as a function of the weather pattern that was used during the optimization. The percentage of affected flights is provided for each weather situation in Table 4.6. One can observe that in most cases all of the trajectories are changed, which is logical since these are all fully climate-optimized sets of flights. There are three weather patterns in which a couple of flights have not been changed. However, these percentages are so close to 100%, that it can be stated that all traffic is affected, independent of

Table 4.6: Percentage of affected flights as a function of weather pattern. Case: westbound/P-AGWP100/Clim 1.0.

WP1	WP2	WP3	WP4	WP5	SP1	SP2	SP3
100	100	100	99.24	100	99.75	100	99.75

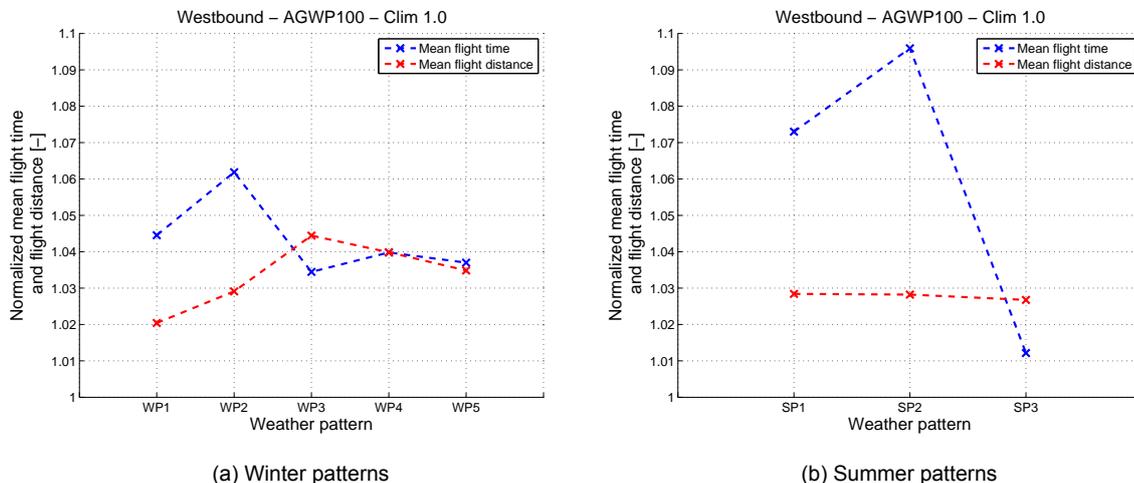


Figure 4.18: Mean flight time and mean flight distance normalized by the mean flight time and distance of the cost-optimized flights, as a function of (a) winter weather pattern and (b) summer weather pattern.

weather pattern.

4.4.2. Flight time and distance

It was determined that in every weather scenario, practically all trajectories are modified. How does this translate into changes in mean travel time and distance? Figure 4.18 provides the answer to that question.

First the winter patterns in Figure 4.18a will be looked at. The first thing that strikes the eye is that there is quite a variation, in contrast to the comparison of flight direction and climate metrics. Previously, the increase in duration was always larger than the increase in distance, percentage-wise. However, this comparison shows that this is not always the case. Winter pattern 3 displays an average distance increase that is 1% greater than the flight time increase. This indicates that the average velocity of the fleet is augmented under winter pattern 3 conditions for these westbound flights, compared to the cost-optimized fleet. WP1 and WP2 show the largest duration increases and at the same time the smallest distance increases. This translates into a significant mean ground speed reduction of the fleet. The increments of travel time and distance under WP4 and WP5 are roughly the same, which indicates that the fleet mean speed is not affected as much.

The summer patterns in Figure 4.18b also show variability in mean travel duration. For SP1 and SP2, there is a big difference between the travel time and travel distance increase, again indicating the reduction of aircraft mean velocity relative to the cost-optimized flights. The increase in flight time is also very large, and larger than any winter pattern with almost 10%. For summer pattern 3, the distance increases more than the time the aircraft fly at higher speeds, although less extremely than patterns 1 and 2. In fact, the average flight duration is only increased by little over 1%, which is the lowest amount of any weather pattern. Remarkable is that the flight distance is augmented by an approximately constant percentage for each summer pattern.

Figure 4.19a provides the time increment distributions for every winter pattern, with the accompanying probabilities of a flight being shortened listed in Table 4.7. As already suggested by the mean time and distance increase, there is also a lot of variation visible in these distributions. Some of the peaks occur near the zero time increment (WP1, 4 and 5), one at a moderate time increase (WP3) and one at a large difference (WP2). However, this does not reflected in the mean time increase indicated by the crosses. Because of the large differences in the spread, WP3 shows the lowest average time elon-

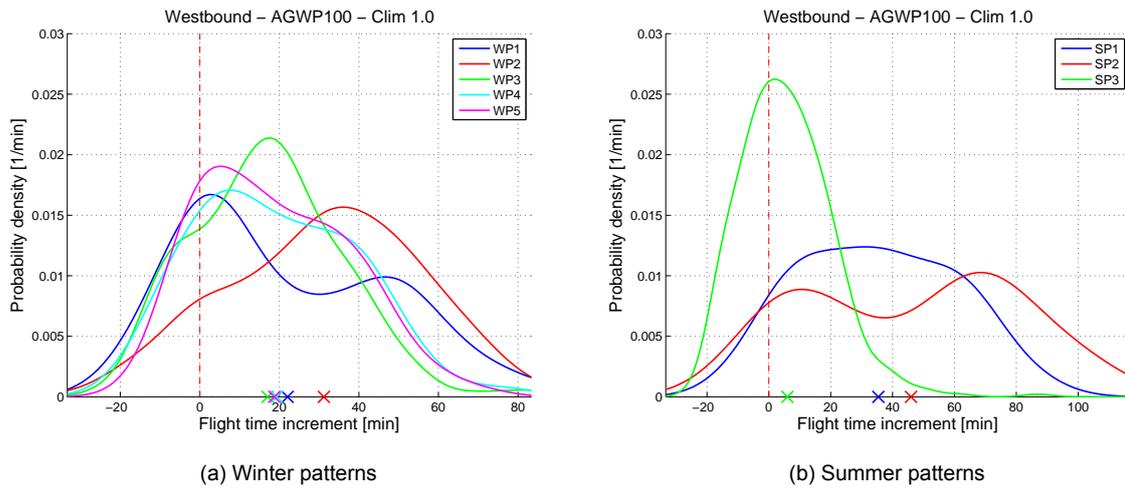


Figure 4.19: Probability density functions of flight duration increments with respect to the cost-optimized flights for (a) winter weather patterns and (b) summer weather patterns. Zero flight time difference is indicated by the red dashed line. The crosses on the x-axis represent the mean flight duration increase.

Table 4.7: Probability that a flight is shortened in terms of flight time and flight distance, as a function of weather pattern. Case: westbound/P-AGWP100/Clim 1.0.

	WP1	WP2	WP3	WP4	WP5	SP1	SP2	SP3
Flight time	0.249	0.129	0.209	0.209	0.192	0.102	0.127	0.381
Flight distance	0.247	0.184	0.0521	0.0921	0.0877	0.190	0.194	0.190

gation, even though its peak is located second furthest down the x-axis. WP2 shows the largest trend towards increasing flight times, with the probability of a flight being shortened of just 0.129 (Table 4.7). Of the winter patterns that peak at a little over zero, WP1 is the one that is most broadly spread, hence shows the greatest average time increment of the three. In short, because there is a large variability of the location of the peaks and the spread of the curves, the winter patterns are very distinct in how their travel times are affected.

The travel time distribution of the summer patterns, provided in Figure 4.19b, are very distinct patterns. The travel times of SP3 have been altered relatively little, with the mean flight increment being only a couple of minutes and a high probability that the flight is shortened (0.381, see Table 4.7). The curve shows a high peak around zero and is not widespread. SP1 displays a very wide peak that is lower than the one of SP3. Finally, the travel time difference distribution of SP2 contains two local maxima, one at around 10 minutes and one at a very large increment of around 70 minutes. It is also the most widely spread PDF, resulting in the highest mean duration increment of all three summer patterns.

Figure 4.20 then presents the probability density functions of distance increments with respect to the cost-optimized sets of flights. The winter patterns in Figure 4.20a again do not show many commonalities. WP1 shows the smallest average distance increment, in agreement with Figure 4.19a, because its maximum is located at around 110 kilometers increment, where the maximum peaks of the other winter distributions are at larger increments. From Table 4.7 it is also clear that the probability of a flight being trimmed in terms of distance is greatest for this winter pattern. The other flights show a significant greater likelihood that the distance is lengthened. WP3 shows the largest increase in distance. The distribution reaches its maximum at around 300 extra kilometers and only a small portion of the area under this curve finds itself at the left hand side of the dashed line at zero increment, resulting in no more than about one in 20 flights travelling a shorter distance (see Table 4.7). The PDFs of the other three winter patterns are between those of WP1 and WP3, and so are the corresponding average increments.

The PDFs of the summer patterns, depicted in Figure 4.20b, tell another story. They confirm that the

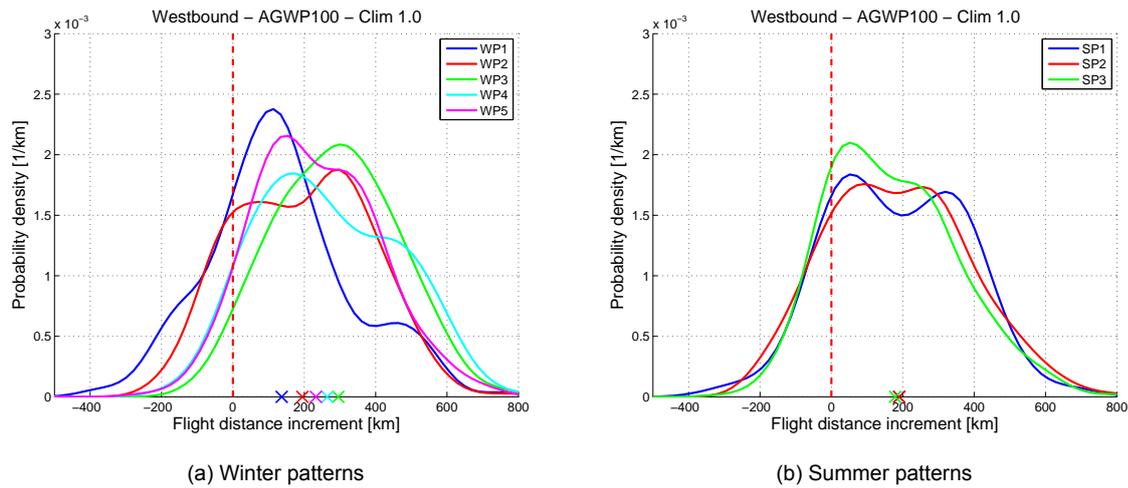


Figure 4.20: Probability density functions of flight distance increments with respect to the cost-optimized flights for (a) winter weather patterns and (b) summer weather patterns. Zero flight distance difference is indicated by the red dashed line. The crosses on the x-axis represent the mean flight distance increase.

average distance increase almost does not change depending on the weather pattern. This average is approximately 180 kilometers. The differences between the increment distributions are fairly small. The patterns that show more flights with small distance increments show less flights with larger increments, which apparently cancels out in the mean. Thus far in this weather comparison, travel distance increase and the according increment distributions are the only characteristics that show similarity.

4.4.3. Shift in latitude and trajectory classification

With the travel duration and distance increments discussed, this section focusses on the lateral shift and the categorization of the trajectories based on these lateral shifts. First the focus will be on winter patterns, after which summer patterns will be treated. To start with, the average shift in latitude depending on geographical location is shown in Figure 4.21 for every winter weather pattern. The interpretation of these plots is facilitated by immediately introducing the trajectory classes as well. The percentage of flights within each class as a function of winter weather pattern is given in Figure 4.22.

The trend that we have seen thus far in every comparison is still present. The southern part of the fleet has the tendency to be shifted north, while the northern part is mainly rerouted south. Furthermore, we see that most of these westbound flights are shifted south throughout the complete trajectory, irrespective of weather pattern. There are two extreme winter patterns that can be identified from Figure 4.22.

The first extreme pattern is WP2. This weather situation shows that more than 82% of the flights are rerouted completely south of their original cost-optimized routes. The number of flights that are shifted north is the lowest of all winter patterns, with just over 10%. Does this agree with the overview of the lateral shift in Figure 4.21? At first sight, one might think it does not, because of the large, red colored area at the most southern part of the trajectories, indicating a lateral shift towards the north. This region, however, is somewhat misleading. The fact that the transition from red to yellow happens so smoothly and gradually, indicates that there is only one flight, originally located at the bottom of the area spanned by the set of flights, that was shifted north by approximately 8 degrees (because the lateral shifts are interpolated between flights). This is confirmed when comparing the bottom of the plot of WP2 to the other winter patterns. In the other cases, one can observe that the southern boundary of the area spanned by the original set of flights shows more of an arch, whereas for WP2 this is practically a straight line. This indicates that there was indeed an exceptional flight in the cost-optimized data set of WP2. When recognizing the fact that the region of large northward lateral shifts in the south is caused by only one flight, it is feasible that less flights are rerouted south than other winter patterns, as the number of regions indicating northward shifts, aside from the region in the south, is small.

The second extreme winter pattern is WP5. From Figure 4.22 it can be observed that this case still has most trajectories shifted south. However, the difference with the number of flights that are rerouted completely north is much smaller. This checks with Figure 4.21. In contrast to WP2, the region of

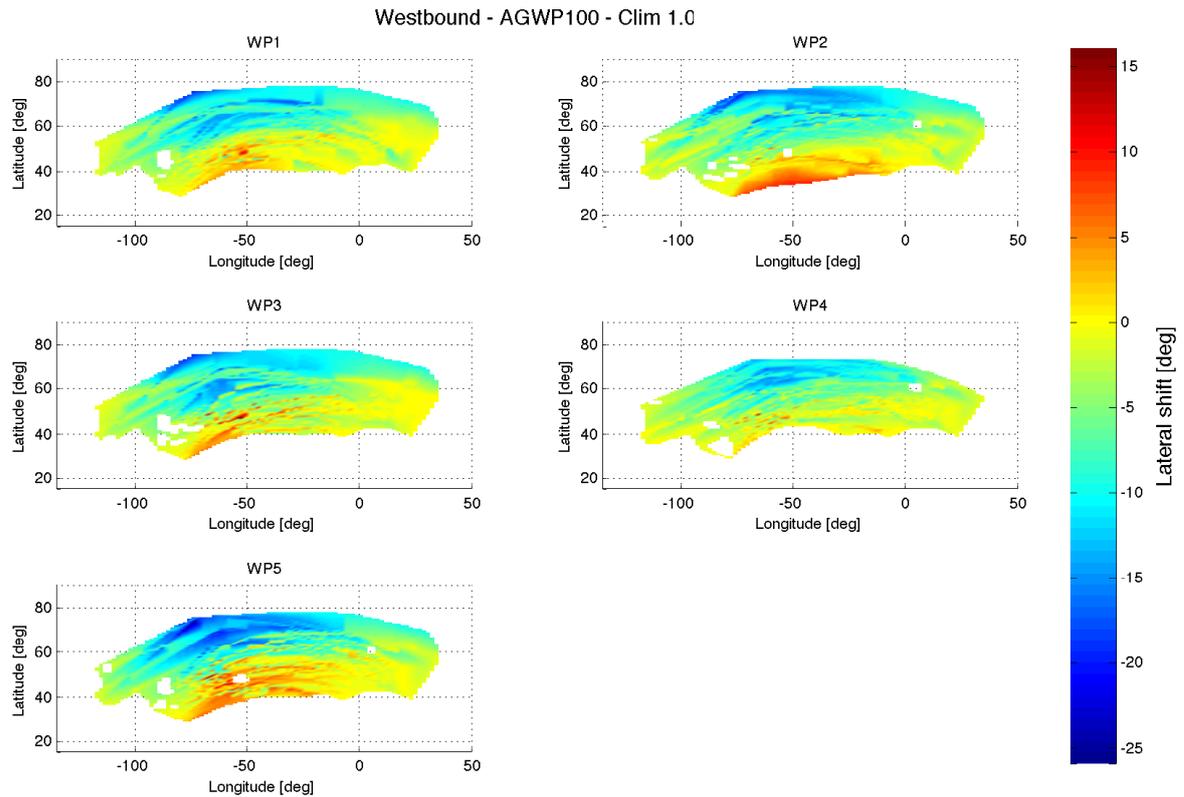


Figure 4.21: Overview of the average shift in latitude depending on geographical location for the winter weather patterns. A positive lateral shift represents a northward rerouting, a negative one a southward.

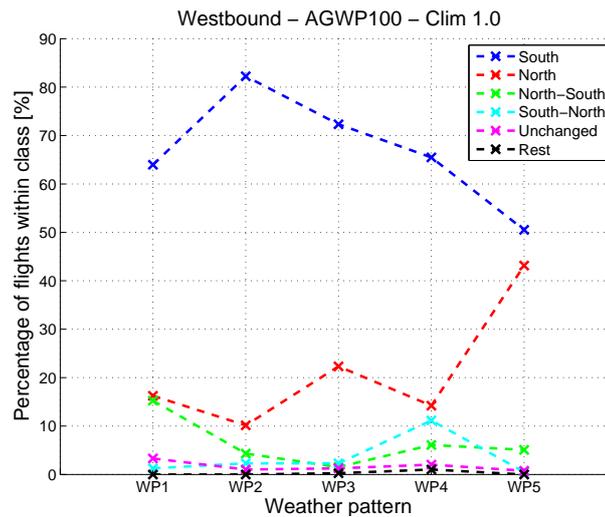


Figure 4.22: Percentage of rerouted flights that fall within each trajectory class as a function of winter weather pattern.

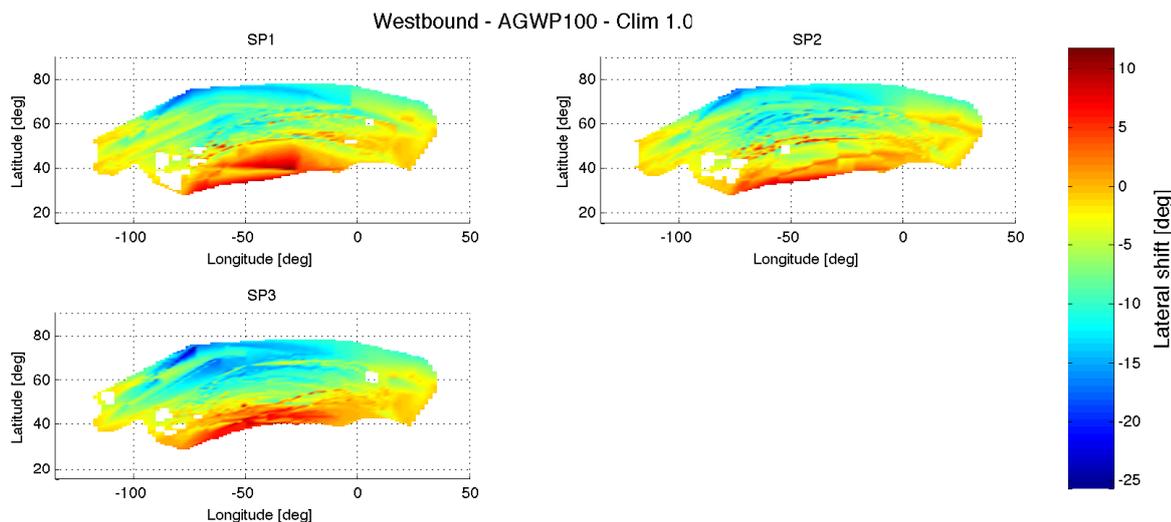


Figure 4.23: Overview of the average shift in latitude depending on geographical location for the summer weather patterns. A positive lateral shift represents a northward rerouting, a negative one a southward.

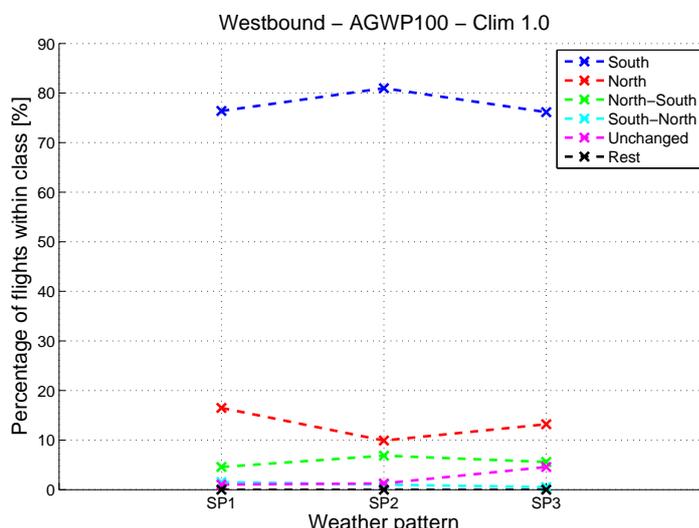


Figure 4.24: Percentage of rerouted flights that fall within each class as a function of summer weather pattern.

northward reroutings in the south under WP5 is far from smoothly and gradually transitioning into lower or southward lateral shifts. This means that there are a lot of flights in that region, each one rerouted north to a different extent.

The remaining winter patterns have more moderate characteristics. Two more remarks can be made. First, WP1 contains almost as many flights that are relocated completely in northern direction as flights that show a north-south type of lateral shift. Second, WP4 displays a more-than-average percentage of flights that are rerouted south on the western part of the trajectory and north on the eastern part. Over 11% of the routes are affected this way, which is the highest number that we have seen so far throughout this case study.

The average lateral shifts and route classification for summer weather patterns can be found in Figures 4.23 and 4.24, respectively. The route categorization is remarkably similar for all summer patterns. Again southward reroutings are the most prominent with up to around 80% of the flights. The second largest category is the one that contains the northward relocations. The average lateral shift depending on the geographical location is similar too. The regions of the mean shift in northern direction somewhat change shapes, but the intensities are comparable.

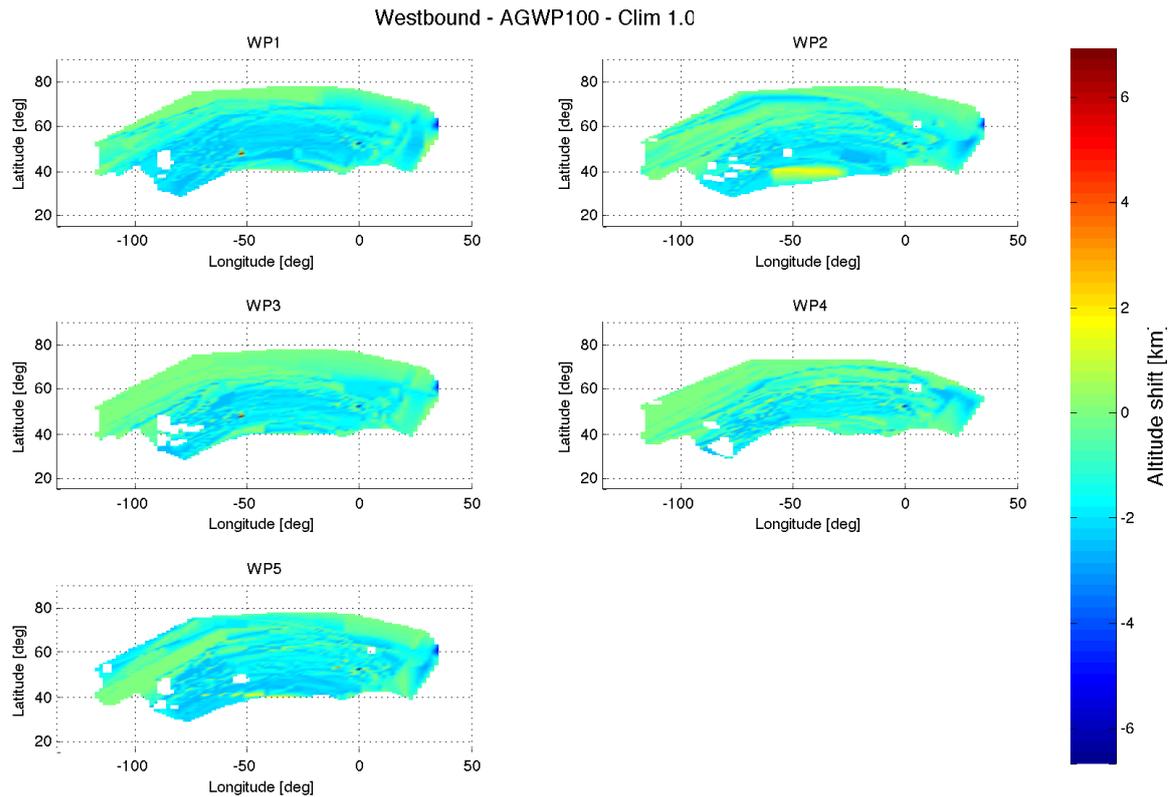


Figure 4.25: Overview of the average shift in altitude depending on geographical location for the winter weather patterns. A positive shift represents an increase in height, a negative shift a decrease.

4.4.4. Altitude shift

Finally, the changes in height will be discussed. The overviews of winter pattern cases are given in Figure 4.25. Overall, the shifts in height are similar for every winter pattern. Every pattern shows that the northern part of the fleet experiences less of a decrease in height than the flights over lower latitudes. The maximum height decrease is around 2.5 kilometers in all plots. The height shift plot of WP2 shows a somewhat unusual height *increase* in the southern part of the fleet. The gradual transition into a downward shift indicates that there is just one flight in this region, and it happens to be shifted upward by about 1.5 kilometers. The other plots show no irregularities of that kind (apart from the ones near airports, where the local approach or departure trajectories are altered).

Figure 4.26 presents the same plots for summer weather patterns. In all plots, the maximum downward route modification is again around 2.5km, but mostly height decreases of between 1 and 2 kilometers are seen. There are some noticeable differences between the summer patterns. The northern parts of the climate-optimized flights of SP1 and SP3 generally show a smaller decrease in height than the southern parts. SP2, on the contrary, also shows significant altitude decrements in the northern regions. Furthermore, the altitudes of SP1 seem to be least affected by the optimization¹. The plots show that the climate-optimized routes in SP3 are most generally shifted downward, i.e. there are not many regions that show small average height decreases, apart from the northern part. The altitude changes under SP2 do depend more on the latitude-longitude combination.

4.4.5. Conclusion of the case study of the influence of weather pattern

This section about the influence of weather patterns on the trajectory changes studied the fully climate-optimized trajectory data sets of westbound flights for the metric P-AGWP100 and eight different weather patterns. Here is a recapitulation of the main findings.

To start with, it was found that as good as all flights are affected by the optimization, independent

¹Note that the small extremity of a 5 km height increase in the SP1 altitude shift plot at around 47°N 52°W is a modification of the local approach trajectory around St. John's airport in the most western part of Canada (IATA code: YYT).

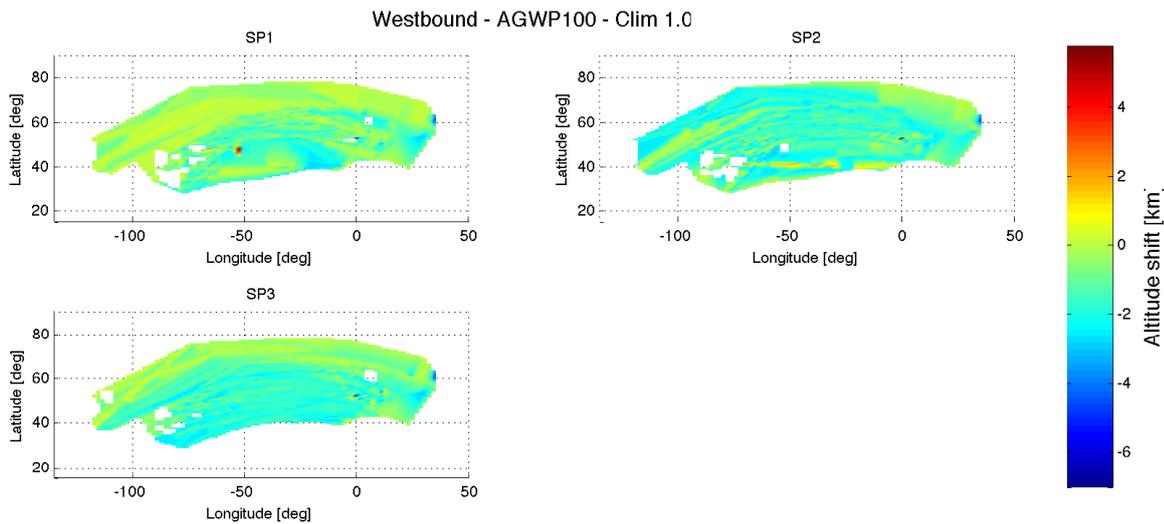


Figure 4.26: Overview of the average shift in altitude depending on geographical location for the summer weather patterns. A positive shift represents an increase in height, a negative shift a decrease.

of weather pattern. Then it was observed that mean flight time and distance increments for winter patterns is heavily dependent on the weather situation. For summer patterns the mean flight distance surprisingly remains almost the same. The flight duration and flight distance increment *distributions* also showed a large variation with weather pattern, except for distance increment distributions of the summer patterns, which were very similar. Furthermore, it was found that for all weather patterns, the majority of the flights is rerouted completely south of the cost-optimized trajectories, and the second largest category are flights that are shifted completely north. The exact numbers differ quite a lot for winter patterns, but are very similar between summer patterns. The plots of average lateral shift as a function of longitude and latitude differ from case to case, but it still holds in general that the most northern part of the fleet is rerouted south and the most southern part north. Finally, it was determined that the shift towards lower altitudes was mostly between 1 and 2.5 kilometers for every weather pattern. The plots of altitude shift versus longitude and latitude vary from case to case, but the main trend was that flights in higher longitudes showed a smaller decrease in height than more southern flights.

4.5. Total overview

Table 4.8 provides an overview of the results obtained in this chapter. It indicates the way in which each of the four case differentiators influences the flight time, flight distance, the lateral shift, the trajectory classification and the altitude shifts.

Table 4.8: Overview of the influence of level of climate optimization, flight direction, climate metric and and weather pattern on trajectory changes.

	Time	Distance	Latitude	Trajectory classifica- tion	Altitude
Optimization	Increases with in-creasing impact reduction	Increases with in-creasing impact reduction	Shifts are slightly larger with increasing impact reduction	Changes slightly, mostly south	Decreases with in-creasing impact reduction
Direction	Mean flight time incre-ment the same, PDFs different	Mean flight distance increment the same, PDFs different	Eastbound flights more north, and less south	Eastbound flights more often north, and less often south	Altitude shifts of west-bound routes slightly more geographically consistent
Metric	No considerable differ-ences	No considerable differ-ences	No considerable differ-ences	No considerable differ-ences	No considerable differ-ences
Weather	Both mean flight times and PDFs of time increments highly dependent on weather pattern	WP: Both mean flight distances and PDFs of distance increments highly variable. SP: Mean flight distances similar, PDFs different.	Highly dependent on weather pattern	Highly dependent on weather pattern	Always towards lower alitudes, geographical spread differs.

5

General climate-optimized trajectory analysis

The previous chapter studied the influence of the level of climate optimization, direction, climate metric and weather pattern on one base case. It did so by taking one combination of these case-defining parameters, namely *westbound/P-AGWP100/WP1/Clim 1.0*, and then varying one of these differentiators depending on which comparison was being made. This resulted in a very detailed analysis, specific to that case. The current chapter will also examine the influence of each of the case differentiators, but then at a high level of generality. Instead of a case, all available data sets will be considered. For example, if the influence of flight direction is examined, then all possible combinations of metric, weather pattern and level of climate optimization in eastbound direction will be taken into account in the calculations, and the same for the westbound direction. In other words, this chapter aims to extend the reach of the results provided in chapter 4 and examines whether trends that arose during the case analysis also hold in general. Because of this generalization, the level of detail will not be as high anymore, and the graphs will depict trends rather than detailed results.

The current chapter will retain the same format as the previous one. First general trends of differences in trajectories between Pareto locations are discussed in section 5.1. Section 5.2 treats the general trajectory changes between the two directions. Then we move on to section 5.3, which investigates the overall influence of metrics on the aircraft routes. Finally, section 5.4 examines the general differences in how the routes are affected between weather patterns. Once again, each section will contain its own conclusion to highlight the main findings.

5.1. General inter-Pareto trajectory changes

When aiming to compare Pareto locations on the general scale, an extra complexity arises compared to the case study. In the case study only six Pareto points, all belonging to the same case, were examined. In this overall analysis all available data sets, and thereby also all Pareto points, are included. The inability to compare Pareto points cross-case, which was uncovered in section 3.3, necessitates considering Pareto location to be a continuous variable with values between zero and one, rather than 6 discrete levels of climate impact importance, depending on the scalar used in the objective. A Pareto location is then defined as the arc length along the front, starting from the cost-optimized point up until the Pareto location in question, divided by the total arc length of the Pareto front of the case to which the Pareto point belongs¹. Notice that every case has a Pareto point at zero and at one, representing the fully cost-optimized and fully climate-optimized fleet, respectively. All data sets are included in the graphs that will follow, i.e. a total of 288.

5.1.1. Rerouted flights

The first parameter that is examined is the percentage of flights that experience a change in travelled distance and/or a change in flight time, as a function of Pareto location (Figure 5.1). In the figure, the

¹The arc length is approximated by the sum of the straight-line distances between the Pareto points.

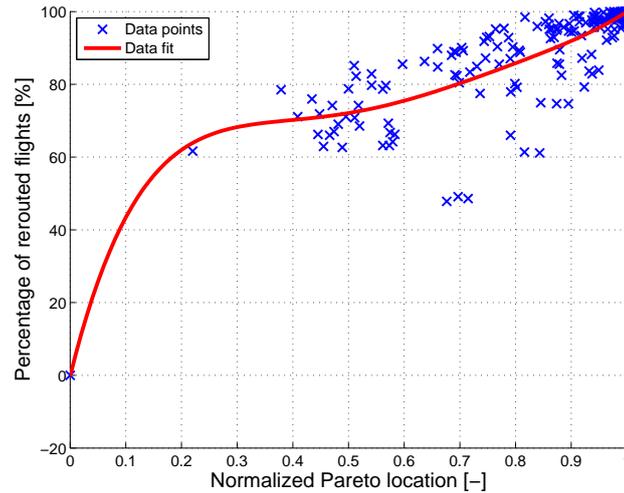


Figure 5.1: The percentage of rerouted flights as a function of Pareto location. Crosses indicate the points at which the percentages are computed. The trend line is a 5th degree polynomial fit of the computed data points.

crosses indicate the computed percentage of affected flights at each Pareto location. A 5th order polynomial fit is added. It can be observed that the bigger the role of climate impact during the optimization, the more flights are affected. Because this is a general analysis, this holds as a general observation for all cases. At Pareto position 1, the trend shows that every flight is affected in terms of flight distance and/or flight duration.

5.1.2. Flight time and distance

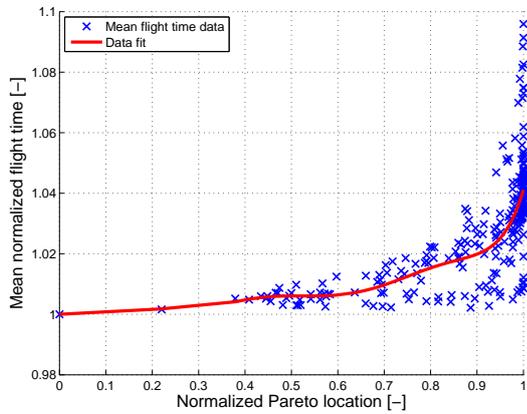
Figure 5.2a provides the normalized mean flight times as a function of Pareto location. Immediately it is clear that moving from cost-optimized to climate-optimized routes increases the mean flight duration. In fact, there is not a single data set that undergoes a decrease in mean flight duration. The mean travel duration at the fully climate-optimized Pareto location is increased by approximately 4.1% on average. The range of time increment values at Pareto location 1 is 0.9%–9.6%.

Figure 5.2b represents the normalized mean flight distance for all Pareto points. Again, increasing the climate impact weight during the optimization augments the mean flight distances. In contrast to the duration, the mean flight distance is in some cases shortened. These cases are in the minority and do not affect the trend line in such a way that it becomes negative at some point. The average flight time increase at Pareto location 1 is 3%. This is more than one percent lower than the mean value of the duration increase, indicating that overall, fully climate-optimized flights are also slowed down with respect to the cost-optimized flights. The range of distance increments at Pareto location 1 is 1.8%–4.6%.

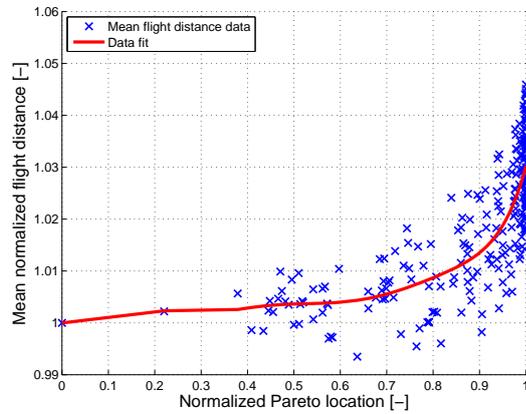
The plots in Figure 5.3 are variants of the PDFs of the duration and distance increments from subsection 4.1.2. These plots are constructed as follows. The PDFs of every case (i.e. Pareto location) are sorted based on their Pareto location. Then they are interpolated to form a 3D surface: the x-axis is the Pareto location, the y-axis contains the time (or distance) increments and the z-axis is the probability density, which is a function of Pareto location and the time (or distance) increments. The plots depicted here are contour plots of this surface. The dashed white line is the line that denotes an increment of zero minutes or kilometers.

Figure 5.3a shows peaks around the zero increment line at lower Pareto locations. This means that many flight times have only been altered by a small amount at these points. When moving towards the fully climate-optimized Pareto point, these peaks are decreased. Next to this, larger time increments are found when moving towards maximum climate impact reduction. Hence the PDFs of the higher Pareto points are more widespread, and the probability that the flight duration is increased by a large amount is greater.

The same can be said about Figure 5.3b, which is the same contour plot but now for the distance increments. The probability that a route is altered heavily is larger near Pareto location 1.

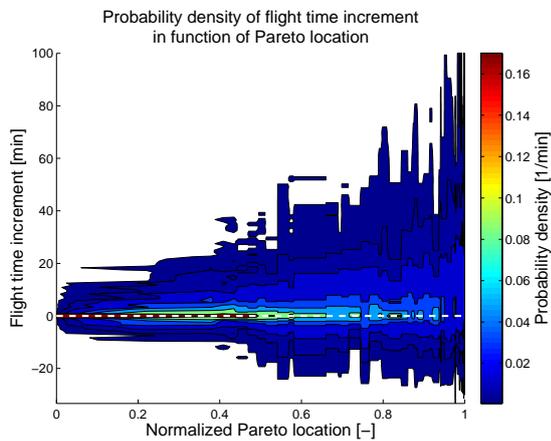


(a) Duration

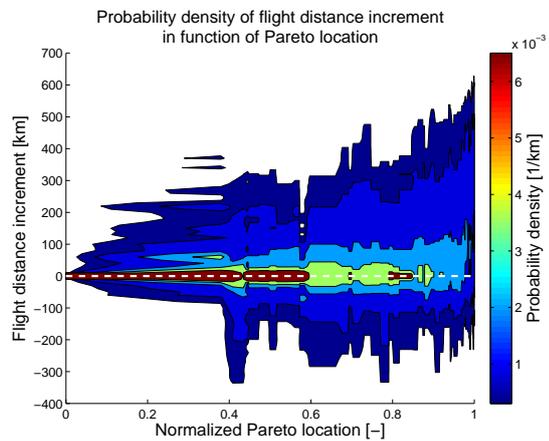


(b) Distance

Figure 5.2: The normalized mean flight duration (a) and normalized mean flight distance (b) as a function of Pareto location. Crosses indicate the points at which the mean flight durations and distances are computed. The trend lines are 5th order polynomial fits of the computed data points.



(a) Duration



(b) Distance

Figure 5.3: Contour plots of the probability density functions of flight duration and flight distance increments with respect to the cost-optimized flights, as a function of Pareto location. Zero flight time and distance difference is indicated by the white dashed line.

5.1.3. Shift in latitude and trajectory classification

The lateral shifts are analyzed a bit differently than it was done in the case study (see subsection 4.1.3). In the case study, the average lateral shift was plotted as a function of longitude and latitude for the six Pareto points of that combination of direction, metric and weather pattern. In this general analysis, it is impossible to do this, because such kind of plot would have to be made for every Pareto location. No trend of the influence of Pareto location would become visible. The way a clear trend can be identified is by computing the mean, maximum and minimum lateral shift for each Pareto point and then making trend lines. The result is depicted in Figure 5.4. Before interpreting this figure, it is necessary to note that each of these Pareto points could belong to any combination of direction, metric and weather pattern.

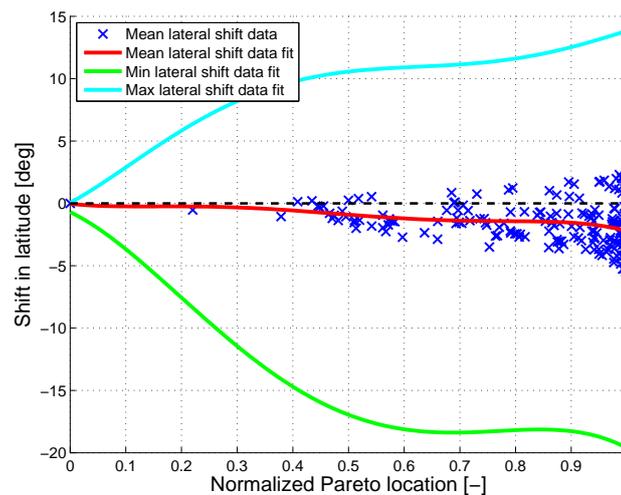


Figure 5.4: The mean, maximum and minimum lateral shifts as a function of Pareto location. Positive values represent shifts to the north, negative values shifts to the south. Crosses indicate the computed mean lateral shift data points. The computed maximum and minimum lateral shift data points are not shown for clarity. 5th order polynomial fits of the mean (red), minimum (green) and maximum (cyan) lateral shifts are added.

Firstly, Figure 5.4 shows that the trend of the mean latitude shifts is southward. Nevertheless, there are some cases that show an average lateral shift towards the north. The average shift at Pareto location 1 is 2.2 degrees south. The range of values at this location is -6.2 to 2.6 degrees. Next, it can be seen that the values of the minimum and maximum shifts increase (in absolute sense) towards the climate-optimal Pareto point, which indicates more extreme traffic displacements. Finally, the shifts south are larger than the ones in northern direction. At Pareto location 1, the average maximum shift is 13.8 degrees, whereas the mean minimum shift is as high as 19.8 degrees.

So how does the route classification vary along the Pareto front? Figure 5.5 shows the trend lines, and the computed data points are omitted for clarity. In Appendix A, the lines are depicted separately, together with the computed data points. Figure 5.5 shows that at the cost-optimized Pareto point, all flights are unchanged in terms of lateral shift. That is logical, as the cost-optimized fleet is the reference. Then there is a steady decrease in the number of unchanged trajectories when moving along the Pareto distance, down to less than 3% at the Pareto location 1. Furthermore, at any point along the Pareto front, the trend shows that there are more flights rerouted south of their original trajectories than north. In fact, the climate-optimized sets of flights on average show over 60% southward route alterations, whereas only around 27% of the flight trajectories are relocated north. These findings are in line with the variation of lateral shift that was shown in Figure 5.4. Then, there are the more special classes, into which in general not many flights are categorized. Little over 6% of the routes are shifted north in the western part of the trajectory, and south in the eastern part. The other two categories in general contain an insignificant number of flights.

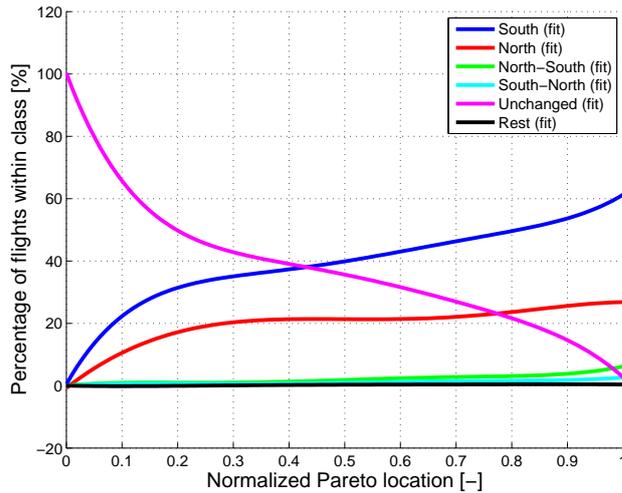


Figure 5.5: Trend lines of the percentage of flights within each trajectory class, as a function of Pareto location. The trend lines are 5th degree polynomial fits of the computed data points. The variability of the data points to which the lines are fitted, can be found in Appendix A.

5.1.4. Altitude shift

Figure 5.6 presents the mean, maximum and minimum altitude shifts as a function of Pareto location. The trend lines of minimum and maximum height shift are not as representative in this case, because these mostly occur in the vicinity of airports. These high and low height modifications are the result of local approach or departure trajectories that are altered. They will therefore not be discussed, but are included for completeness. The mean altitude shift is shown to increase in absolute sense when nearing climate-optimal air traffic. The emphasis is definitely on downward route relocations. In fact, there is not a single data set that shows an increase in mean height. At the fully climate-optimized Pareto location, the mean height decrease is 1.43 kilometer, and the range at this location is 0.82–1.94 kilometer.

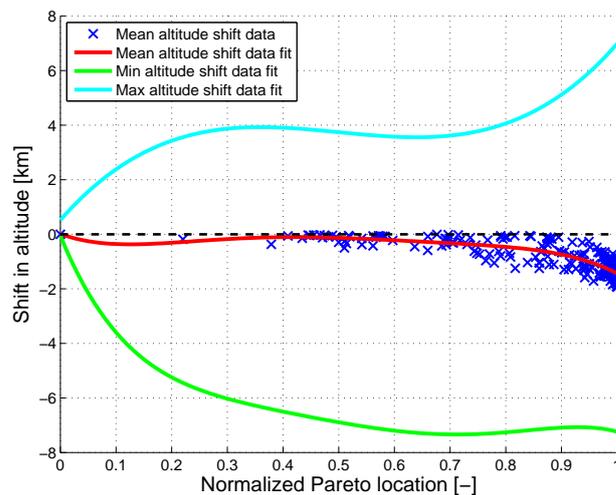


Figure 5.6: The mean, maximum and minimum altitude shifts as a function of Pareto location. Positive values represent shifts to higher altitudes, negative values shifts to lower altitudes. Crosses indicate the computed mean altitude shift data points. The computed maximum and minimum altitude shift data points are not depicted. 5th order polynomial fits of the mean (red), minimum (green) and maximum (cyan) altitude shifts are added.

5.1.5. Conclusion of the general influence of Pareto location

This general analysis of inter-Pareto trajectory differences showed that in general, route alterations get more extreme when progressing from cost-optimized to fully climate-optimized air traffic. The percentage of affected flights increases towards Pareto location 1, with on average 100% of the flights rerouted at this point. Furthermore, the mean flight duration and distance are also augmented. Climate-optimal routes on average show a 4.1% increase in flight time and a 3% increase in travelled distance. The mean lateral shift is southward throughout the complete Pareto range, with the mean value at the climate-optimized location being a 2.2 degree shift towards lower latitudes. Flights are primarily rerouted completely south of their cost-optimized counterparts. At Pareto location 1, on average around 61% of the flights are rerouted south, 27% north, 6% north in the western part of the trajectory and south in the eastern part, 2.7% is geographically unaltered, 2.6% of the reroutings show the south-north pattern and 0.4% do not belong to any of the classes. Finally, also the mean altitude decrease proved to become larger with Pareto location. Not a single case showed an increase in mean flight altitude, possibly because this would result in a higher probability of contrail formation. Fully climate-optimized flights are lowered by 1.43 kilometer on average.

5.2. General influence of flight direction on trajectories

Instead of Pareto location, in this section the directions will be compared. This means that except for the sets optimized for economic cost, all available data sets in eastbound direction on one side and all available data sets in westbound direction on the other will be compared to each other. Therefore, there are 120 data sets per direction.

5.2.1. Rerouted flights

First it is investigated whether there is a difference in the number of flights that have been altered geographically and/or velocity-wise between the two directions. Figure 5.7 provides the box plots of this parameter for each direction. It is clear that westbound traffic is altered to a larger extend. 50% of the data points are greater than or equal to 98.9%, whereas the median of eastbound flights is only at 96.5% and the lower whisker extends much further down. This is because of the jet stream that complicates climate optimization for westbound flights.

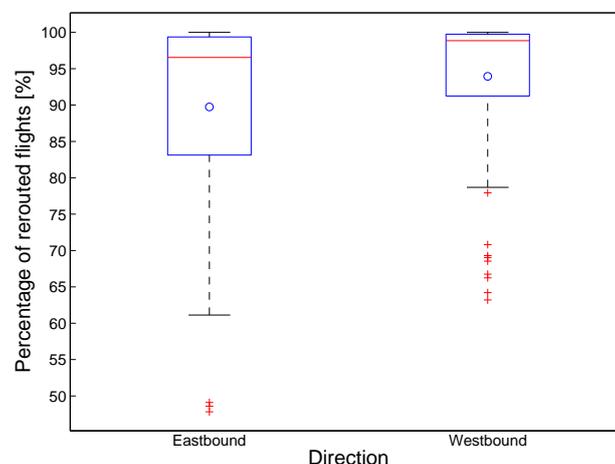
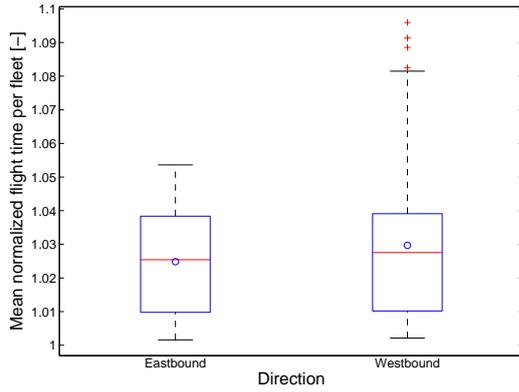


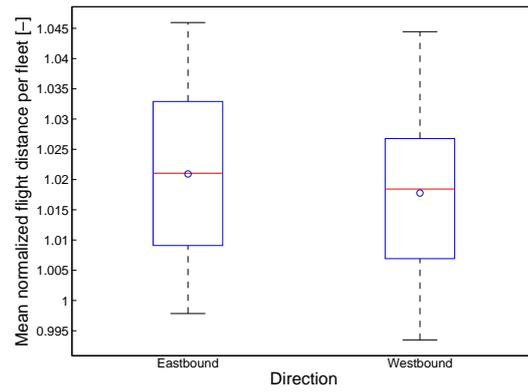
Figure 5.7: Box plots of the percentage of flights per fleet that are affected by the optimization, as a function of flight direction. The box depicts the interquartile range (IQR). Whiskers extend to 1.5·IQR or minimum or maximum if there is no data point outside 1.5·IQR. Outliers are data points greater or less than 1.5·IQR and are indicated by plus signs. Mean values are indicated by circles.

5.2.2. Flight time and distance

The normalized mean flight times and distances are assembled in box plots as a function of direction in Figure 5.8. The box plots are constructed based on fleet-averaged flight times and distances to



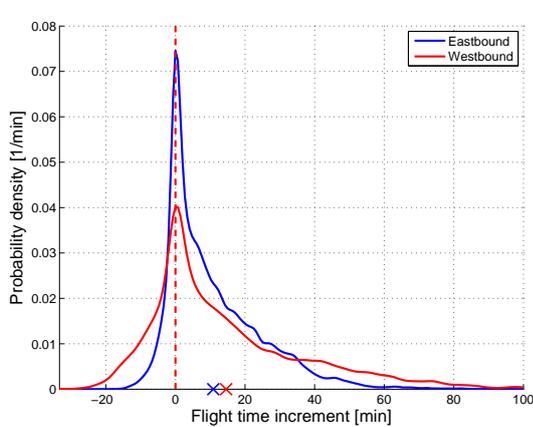
(a) Duration



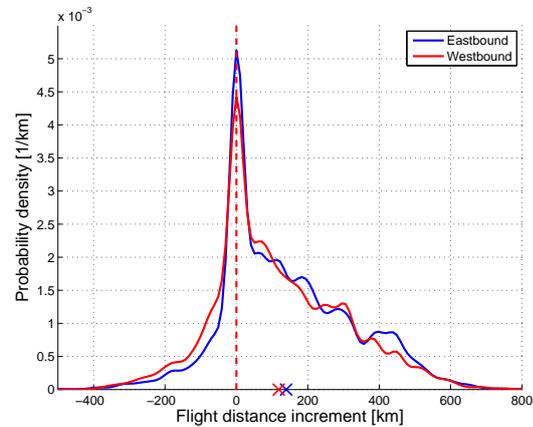
(b) Distance

Figure 5.8: Box plots of the normalized mean flight duration and flight distance per fleet as a function of flight direction. The box depicts the interquartile range (IQR). Whiskers extend to $1.5 \cdot \text{IQR}$ or minimum or maximum if there is no data point outside $1.5 \cdot \text{IQR}$. Outliers are data points greater or less than $1.5 \cdot \text{IQR}$ and are indicated by plus signs. Mean values are indicated by circles.

decrease the range of the plots and by doing so, magnify the differences between the boxes. For the box plots that are constructed using data points per flight, the reader is referred to Appendix B. It can be observed in Figure 5.8a that the medians of the mean flight durations of east- and westbound flights are very similar, with around a 2.5% time increase. The biggest difference is the fact that westbound direction shows a larger whisker towards higher mean flight duration increases. This also causes the average to be 0.5% higher. The mean distances are presented in Figure 5.8b. The two box plots are very similar. Westbound flights show a somewhat lower mean flight distance increase. The averages and medians of both plots are around the 2% increment. Noticeable here is that in both directions, there are data sets that show a mean distance reduction.



(a) Duration



(b) Distance

Figure 5.9: Probability density functions of (a) flight duration increments and (b) flight distance increments with respect to the cost-optimized flights, as a function of flight direction. Zero difference is indicated by the red dashed line. The crosses on the x-axis represent the mean increments.

Figure 5.9 displays the mean probability density functions of time and distance increments. These PDFs are the result of averaging the PDFs of every westbound and eastbound data set, respectively. Figure 5.9a shows that the time increment PDF peaks at 0 for both directions. Westbound flights in general show a higher likelihood of a large increase in flight duration. They also show a greater probability that a flight is being shortened, but in combination with the lower peak and the higher probability of very high duration increments, the average increment is greater for westbound than for eastbound

flights. This is in line with Figure 5.8a. The mean PDFs of distance increments in Figure 5.9b also peak at zero. The two curves are very similar, and so are the respective mean flight distance increments. This is also in line with the box plots in Figure 5.8b.

5.2.3. Shift in latitude and trajectory classification

The box plot of the mean lateral shift as a function of direction is given in Figure 5.10. The figure shows that the mean shift in latitude is quite different for the two directions. Flights towards Europe are in general shifted less south than westbound flights. In fact, more than 25% of the eastbound cases show an average shift towards the north. The value of the median is -0.65 degrees. All data sets containing flights towards North America, on the other hand, exhibit a negative lateral shift. The median is -3.0 degrees.

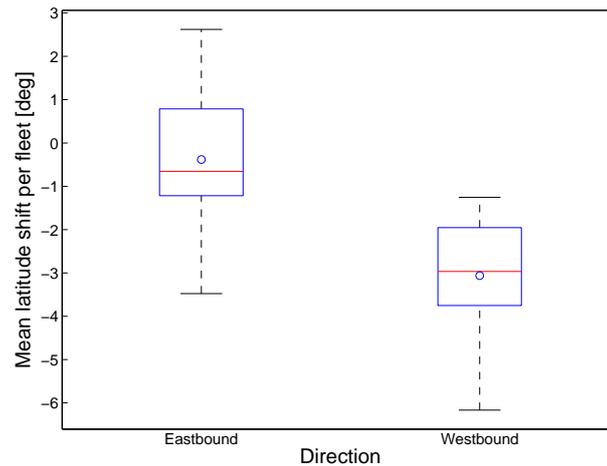


Figure 5.10: Box plots of the mean lateral shift per fleet as a function of flight direction. Positive shifts are shifts in northern direction, negative ones shifts in southern direction. The box depicts the interquartile range (IQR). Whiskers extend to 1.5-IQR or minimum or maximum if there is no data point outside 1.5-IQR. Outliers are data points greater or less than 1.5-IQR and are indicated by plus signs. Mean values are indicated by circles.

The fact that the mean trajectory shifts for both directions is towards the south is also visible from the route classification in Figure 5.11. The majority of the trajectories is rerouted south with respect to cost-optimized flights. Air traffic towards America is rerouted south around 16% more often. The percentage of flights that is shifted north on the other hand is 12% less than in eastbound direction. Also the number of flights that are unaltered in terms of latitude reduces. Another remarkable property is that flights that have not been geographically altered are relatively well represented, with about 15% for eastbound traffic and 11% for westbound traffic. This is because this direction comparison does not make a difference between the levels of climate optimization. The lower Pareto points show a larger amount of unaltered flights, which was established in Figure 5.5 on page 47. The comparison of trajectory classes are in line with the box plots of the mean lateral shift, which also indicated that flights heading for Europe in general undergo a shift that is less south compared to traffic towards America.

5.2.4. Altitude shift

The final examination of this general analysis of the influence of flight direction is of the average altitude shift. The box plots for both directions are provided in Figure 5.12. The plots indicate that generally, the traffic towards America is lowered more than in the other direction. The difference between the two medians is 414 meters. Notice that each of the directions contain sets of flights in which the mean height is practically unaffected after optimization.

5.2.5. Conclusion of the general influence of flight direction

This section did research on what the climate-optimized trajectory differences are between eastbound and westbound flights taking into account all available data sets. In general, westbound flights:

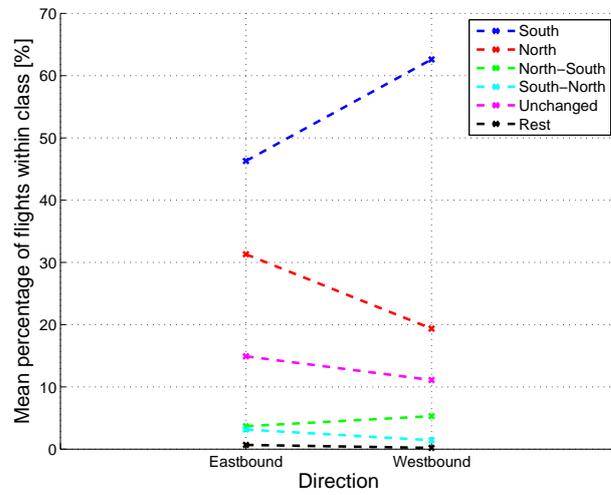


Figure 5.11: The percentage of flights within each trajectory class as a function of flight direction.

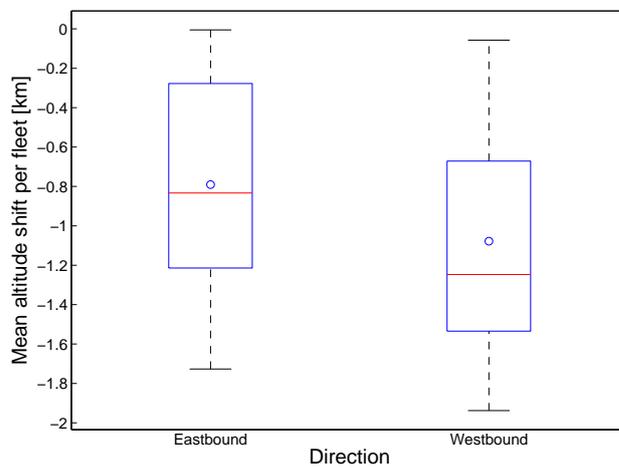


Figure 5.12: Box plots of the mean altitude shift per fleet as a function of flight direction. Positive shifts are shifts towards higher altitudes, negative ones towards lower altitudes. The box depicts the interquartile range (IQR). Whiskers extend to 1.5-IQR or minimum or maximum if there is no data point outside 1.5-IQR. Outliers are data points greater or less than 1.5-IQR and are indicated by plus signs. Mean values are indicated by circles.

- are affected a bit more often by the optimization than eastbound flights, either geographically or velocity-wise;
- show an average flight duration that is 0.5% *greater* and an average flight distance that is 0.3% *less* than eastbound flights, resulting in a small decrease in mean velocity;
- are generally shifted 2.7 degrees more south than eastbound flights;
- are 16% *more* often shifted south over the entire trajectory w.r.t. the cost-optimized flights than eastbound flights, and 12% *less* often shifted entirely north;
- are on average lowered 290 meters more than eastbound flights.

5.3. General influence of climate metric on trajectories

The next step in this chapter on the general analysis of REACT4C routes is examining what happens to the trajectories when different climate metrics are used during the optimization. Recall that a climate metric quantifies climate impact by computing a measure, in this case the average temperature response or the absolute global warming potential, over a defined time window, here 20 or 100 years, using a defined emission scenario, in this case either a future increasing emission scenario or pulse emissions. This section will again mainly make use of box plots to compare the metrics. Per box plot, all available data sets for one metric are included, except for the ones that are optimized for economic cost. That means that a box plot contains data from 80 data sets.

5.3.1. Rerouted flights

First it is investigated whether the percentage of flights that are affected in terms of latitude and/or in terms of flight time, changes with metric. Figure 5.13 provides the box plots. One can observe that the boxes become increasingly compact from F-ATR20 to P-AGWP100. The medians are already very high, but also increase from left to right. The biggest differences are to be found between F-ATR20 and P-AGWP20, which both cover short-term climate impact.

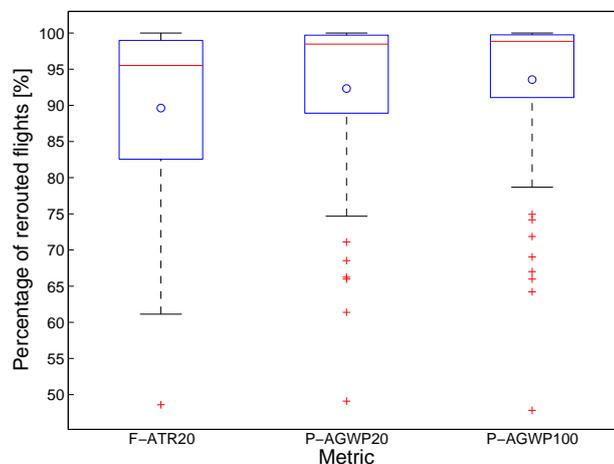


Figure 5.13: Box plots of the percentage of flights that are affected by the optimization, as a function of climate metric. The box depicts the interquartile range (IQR). Whiskers extend to 1.5-IQR or minimum or maximum if there is no data point outside 1.5-IQR. Outliers are data points greater or less than 1.5-IQR and are indicated by plus signs. Mean values are indicated by circles.

5.3.2. Flight time and distance

The plots of average flight times and distances are displayed in Figure 5.14. In Appendix B these plots are made based on flight values instead of fleet averages, allowing for a more detailed overview of the variability. The box plots of the mean flight durations in Figure 5.14a are quite similar. The difference between the greatest and lowest box plot average is 0.6%. Note that for an 8-hour flight, this is equivalent to a difference of just 2.9 minutes. Nevertheless, this means that the mean travel

time increases a little bit when moving from F-ATR20 to P-AGWP100. The same determination can be made for the mean distances in Figure 5.14b. The plots are very similar, and the difference between the lowest and highest mean is again 0.6%. For a flight of 6,000 kilometers, this is only 36 kilometers. The mean flight distance hence increases from F-ATR20 to P-AGWP100, although to a small extent.

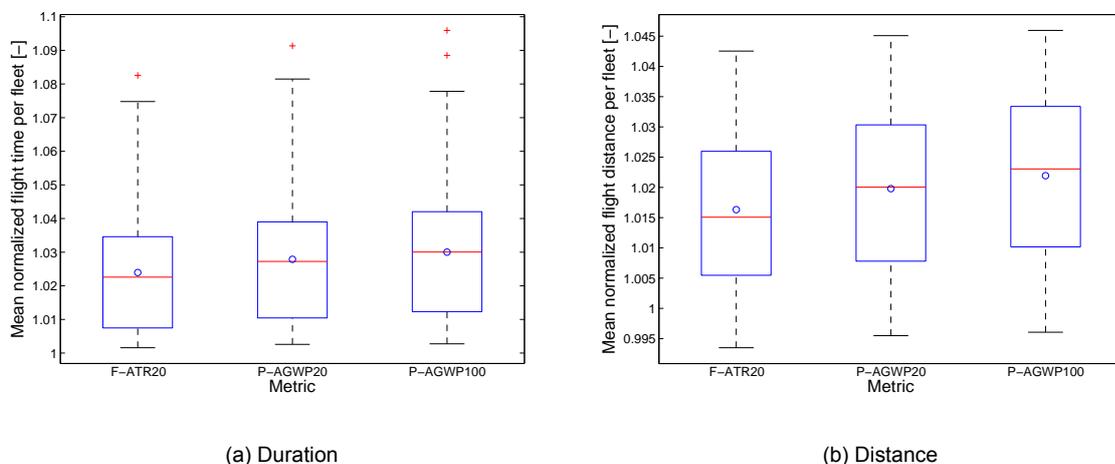


Figure 5.14: Box plots of (a) the normalized mean flight duration per fleet and (b) the normalized mean flight distance per fleet as a function of climate metric. The box depicts the interquartile range (IQR). Whiskers extend to 1.5-IQR or minimum or maximum if there is no data point outside 1.5-IQR. Outliers are data points greater or less than 1.5-IQR and are indicated by plus signs. Mean values are indicated by circles.

The time and distance increment distributions that accompany the box plots are provided in Figure 5.15. These were established by averaging the probability functions of all 80 data sets for the metric under consideration. From these plots it is exceptionally clear that the choice of metric does not affect the flight times and distances to a large extent. The average time increases are between 11 and 14 minutes, whereas the mean distance increments range from 109 to 146 kilometers. The peaks take place around the zero increment point, indicating that in general, many flights are affected by a very small amount.

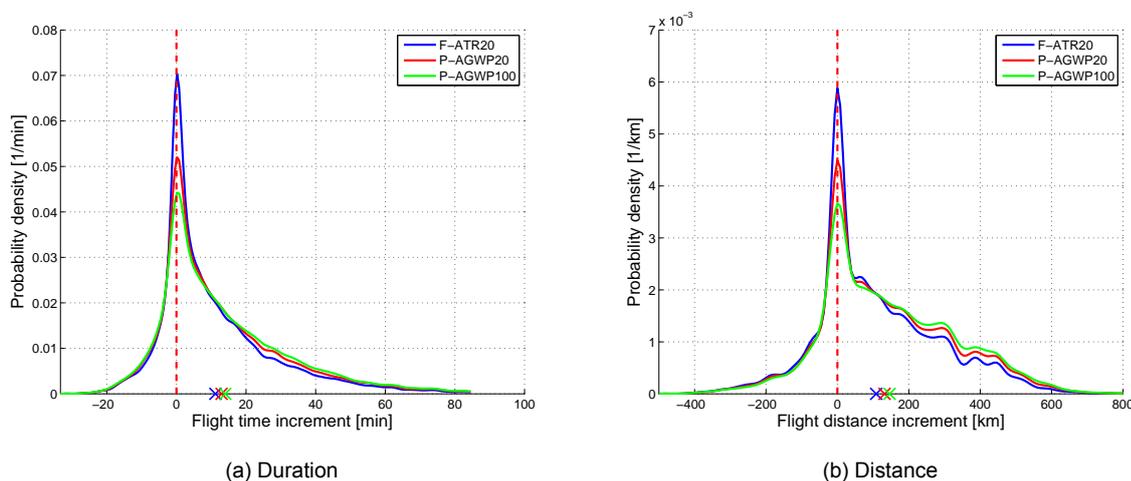


Figure 5.15: Probability density functions of flight duration and flight distance increments with respect to the cost-optimized flights as a function of climate metric. Zero difference is indicated by the red dashed line. The crosses on the x-axis represent the mean increments.

5.3.3. Shift in latitude and trajectory classification

Let us now take a look at how the mean lateral shift behaves for every climate metric. Figure 5.16 displays the box plots. The plots indicate that the lateral shift varies very little with metric. The average

of the F-ATR20 box plot and the one of the P-AGWP100 plot differ only by 0.32 degrees latitude. The range of mean latitude shifts does not vary significantly cross-metric either.

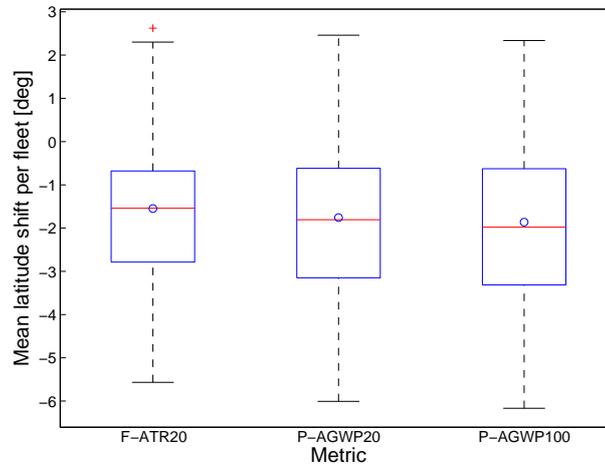


Figure 5.16: Box plots of the mean lateral shift per fleet as a function of climate metric. Positive shifts are shifts in northern direction, negative ones shifts in southern direction. The box depicts the interquartile range (IQR). Whiskers extend to 1.5-IQR or minimum or maximum if there is no data point outside 1.5-IQR. Outliers are data points greater or less than 1.5-IQR and are indicated by plus signs. Mean values are indicated by circles.

The mean lateral shift box plots give reason to believe that the way in which flights are rerouted will not be varying much either. Figure 5.17 provides the graph that confirms this. In line with any general analysis done so far, it can be observed that most flights are shifted south of their original, cost-optimized counterpart. This category gains flights when moving towards the long-term climate impact metric. The second largest category are reroutings towards the north. This category is not dependent on climate metric. Unchanged flights, being the ones which are not rerouted or rerouted on the exact same trajectory but with different speed, are on average well presented. This trajectory class is reduced in size towards long-term climate optimization, and it seems that most of these flights are rerouted south when leaving this class. The remaining categories do not vary either, and all represent less than 5% of the flights.

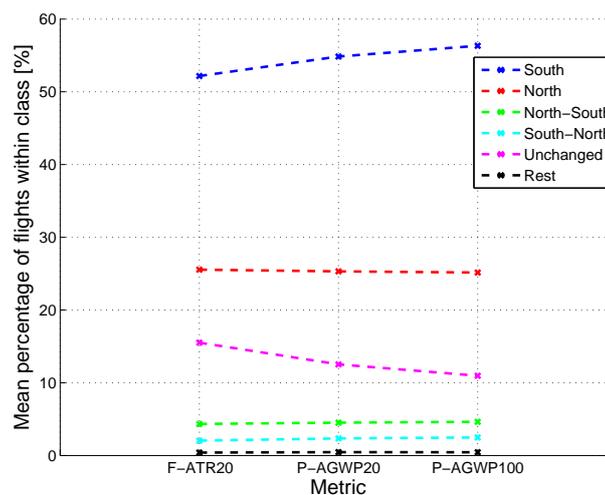


Figure 5.17: The percentage of flights within each trajectory class as a function of climate metric.

5.3.4. Altitude shift

The mean altitude shift box plots are depicted in Figure 5.18. Firstly, there is a very wide range of mean latitudes. At P-AGWP100 for example, there are data sets that show practically no height alteration on average, and at the same time there are also data sets in which the mean altitude decrease is almost 2 kilometers. Next to this, altitude shift seems to be the only trajectory characteristic that depends considerably on the choice of metric. The averages and medians of the box plots vary with the metric. The medians are altered quite considerably, with the difference between the first and the third metric being 470 meters. The average values of the box plots on the other hand vary less, with a difference between the first and last metric of 330 meters. It seems that in general, long-term climate impact is reduced by lowering the altitude further down than for short-term climate impact.

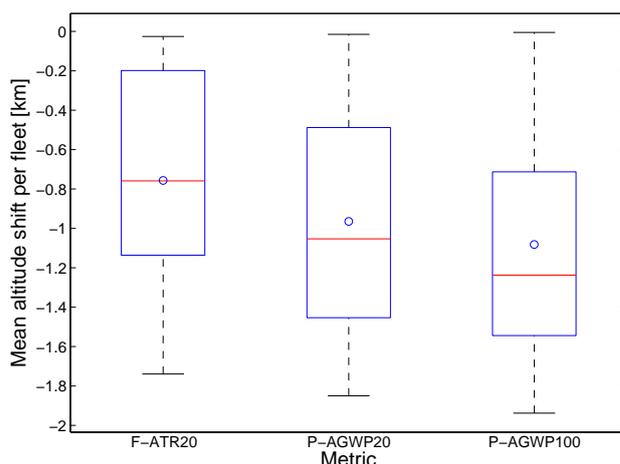


Figure 5.18: Box plots of the mean altitude shift per fleet as a function of climate metric. Positive shifts are shifts in northern direction, negative ones shifts in southern direction. The box depicts the interquartile range (IQR). Whiskers extend to 1.5-IQR or minimum or maximum if there is no data point outside 1.5-IQR. Outliers are data points greater or less than 1.5-IQR and are indicated by plus signs. Mean values are indicated by circles.

5.3.5. Conclusion of the general influence of climate metric

This section compared several route characteristics between the three climate metrics used in the REACT4C project. In general, the choice of climate metric does not impact the way in which trajectories are rerouted to a large extent. The percentage of rerouted flights, mean flight duration, mean flight distance, latitude shift and trajectory classification are affected marginally. From the differences that do exist, it can be observed that the way in which flights are rerouted becomes a little more extreme when moving from F-ATR20 to P-AGWP20 and then again when progressing towards P-AGWP100. The mean altitude shift depends on the climate metric a little more. The difference of the mean value of the mean lateral shift between F-ATR20 and P-AGWP100 is 330 meters.

5.4. General influence of weather pattern on trajectories

After having provided the influence of Pareto location, direction, and climate metric on flight routes, the final step in this general analysis is to present route characteristics for each weather pattern. Again, the main trajectory characteristics will be compared by means of box plots. The plot of one weather pattern includes all available data sets of that weather pattern, which means that no distinction is made between direction, metric and level of climate optimization. Note that the goal for this section is not to discuss all of the information in the plots. It is not desirable to discuss every route characteristic for every weather pattern in depth, because that would not contribute to a general understanding.

5.4.1. Rerouted flights

The box plots of the percentage of flights that are in some way affected by the optimization are given in Figure 5.19. Right from the start, this figure suggests that the position and strength of the jet stream has a large influence on the way the routes are affected. It is clear that in general many flights are altered. However, WP4 (a confined, strong jet) and SP1 (a zonal, strong jet) show box plots that are much wider, and SP1 shows a much lower median, indicating that less flights are influenced under these weather conditions.

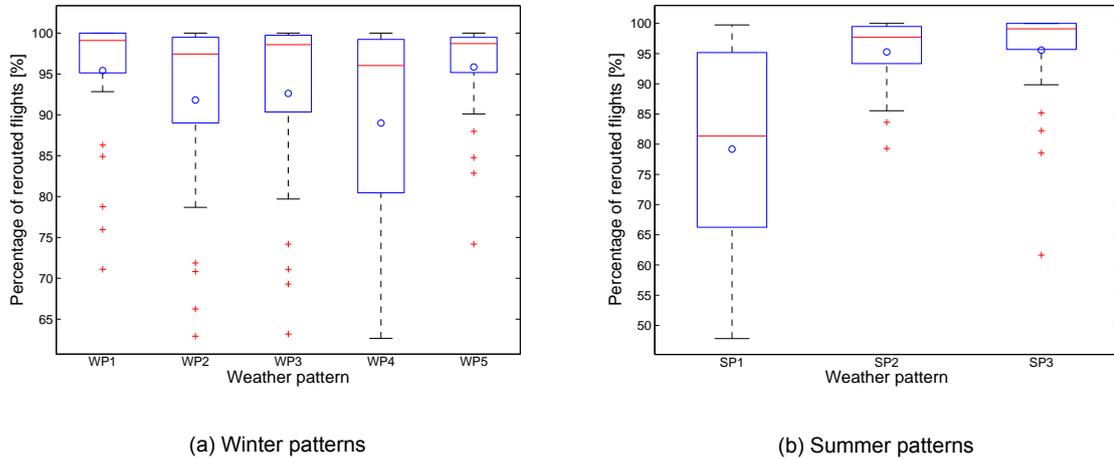


Figure 5.19: Box plots of the percentage of flights per fleet that are affected by the optimization, as a function of (a) winter weather patterns and (b) summer weather patterns. The box depicts the interquartile range (IQR). Whiskers extend to 1.5-IQR or minimum or maximum if there is no data point outside 1.5-IQR. Outliers are data points greater or less than 1.5-IQR and are indicated by plus signs. Mean values are indicated by circles.

5.4.2. Flight time and distance

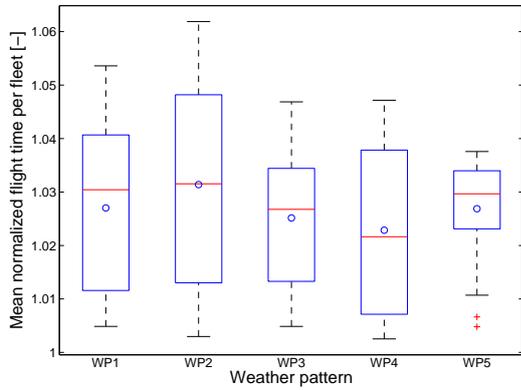
The box plots of mean flying times and durations for every weather pattern are given in Figure 5.20 and Figure 5.21, respectively. The flight-based boxes are presented in Appendix B. In Figure 5.20, it is remarkable that SP2 (a weekly tilted, week jet) shows a much higher increase in mean flight time than the other summer patterns. For winter patterns, the difference is less. The averages of the winter patterns lie within 1 percent from each other. The mean distances of winter patterns depicted in Figure 5.21 show a greater variation. WP1 is characterized by an exceptionally low mean flight distance increase of less than 1 percent, and 25% of the cases in this weather pattern show a mean distance reduction. SP1 shows the least distance increase of all summer patterns, the mean increase being just over 1 percent.

Figures 5.22 and 5.23 provide the mean probability density functions of time and distance increments, respectively. In all these plots, WP4 and SP1 represent the highest peaks at the zero increment point. This means that the probability that a flight is not altered in terms of duration or distance is on average highest for these weather patterns. This completely in line with the percentages of unaltered flights in Figure 5.19.

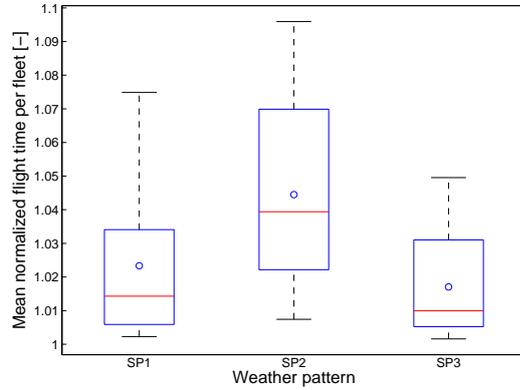
5.4.3. Shift in latitude and trajectory classification

The mean lateral shift box plots are presented in Figure 5.24 for each weather type. For the winter patterns in Figure 5.24a it can be observed that WP 1 (zonal, strong jet) and 5 (confined, weak jet) are rerouted least south (read: most north) on average. Note also that pattern 4 indicates the greatest similarity between its cases, because the range of the box plot is relatively confined. The summer patterns in Figure 5.24b show that flights under these weather situations are in general shifted more south than winter types, because the averages of these box plots are at most -1.85 degrees, whereas for winter patterns they are *at least* -1.85 degrees. Furthermore, SP3 depicts that the mean and median of the mean latitude shifts are very south-heavy, both values being more than 3 degrees south.

The way in which routes are on average categorized into the classes that were defined in subsection 3.4.7, is given in Figure 5.25. It is visible that the weather situations with the least amount of

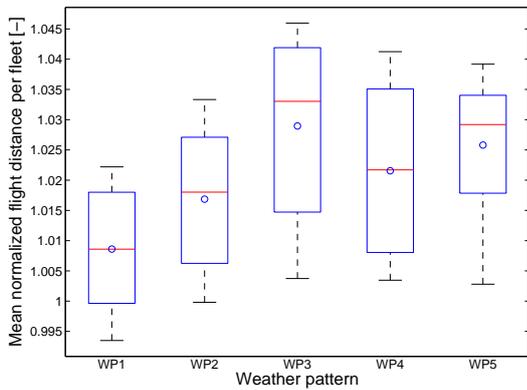


(a) Winter patterns

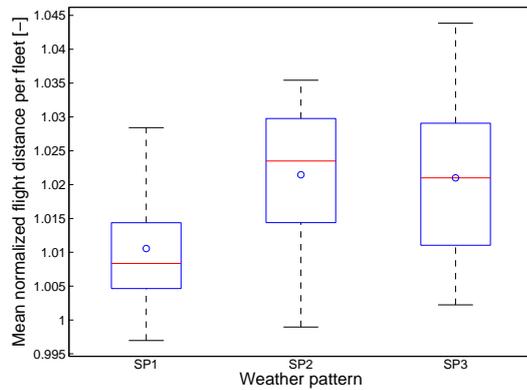


(b) Summer patterns

Figure 5.20: Box plots of mean flight time per fleet, normalized by the mean flight time of the cost-optimized fleets, for (a) winter weather patterns and (b) summer weather patterns. The box depicts the interquartile range (IQR). Whiskers extend to 1.5·IQR or minimum or maximum if there is no data point outside 1.5·IQR. Outliers are data points greater or less than 1.5·IQR and are indicated by plus signs. Mean values are indicated by circles.

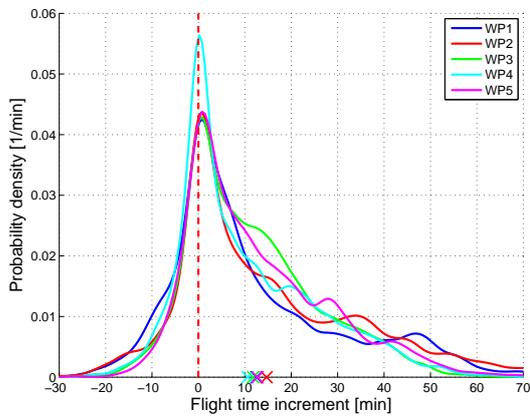


(a) Winter patterns

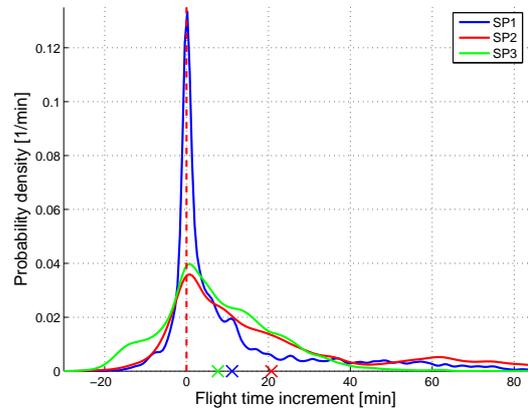


(b) Summer patterns

Figure 5.21: Box plots of the mean flight distance per fleet, normalized by the mean flight distance of the cost-optimized fleets, for (a) winter weather patterns and (b) summer weather patterns. The box depicts the interquartile range (IQR). Whiskers extend to 1.5·IQR or minimum or maximum if there is no data point outside 1.5·IQR. Outliers are data points greater or less than 1.5·IQR and are indicated by plus signs. Mean values are indicated by circles.

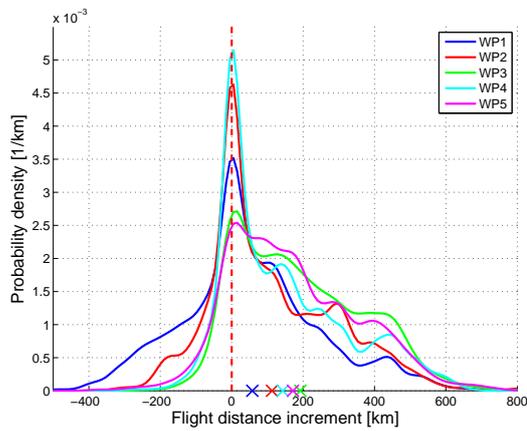


(a) Winter patterns

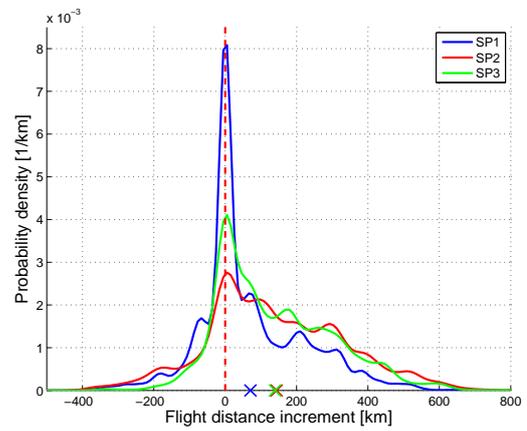


(b) Summer patterns

Figure 5.22: Probability density functions of flight duration increments with respect to the cost-optimized flights for (a) winter weather patterns and (b) summer weather patterns. Zero difference is indicated by the red dashed line. The crosses on the x-axis represent the mean increments.

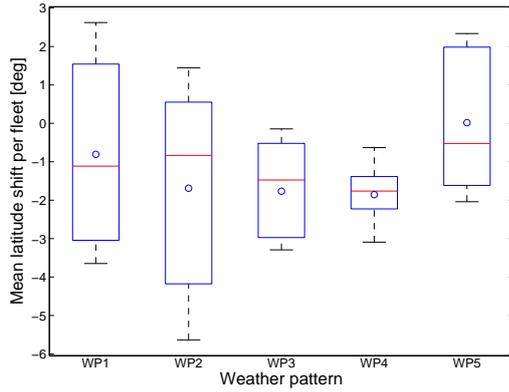


(a) Winter patterns

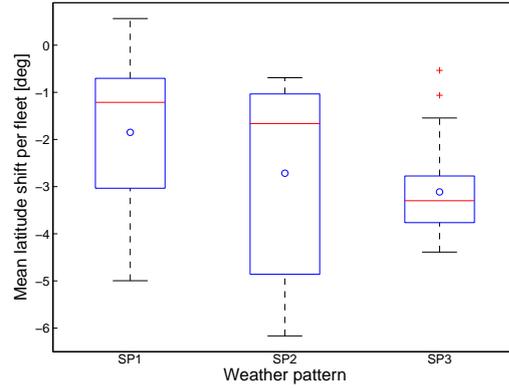


(b) Summer patterns

Figure 5.23: Probability density functions of flight distance increments with respect to the cost-optimized flights for (a) winter weather patterns and (b) summer weather patterns. Zero difference is indicated by the red dashed line. The crosses on the x-axis represent the mean increments.



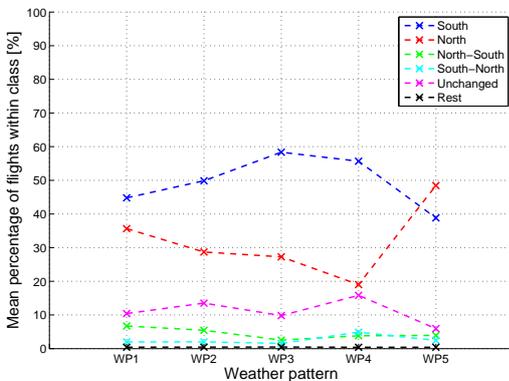
(a) Winter patterns



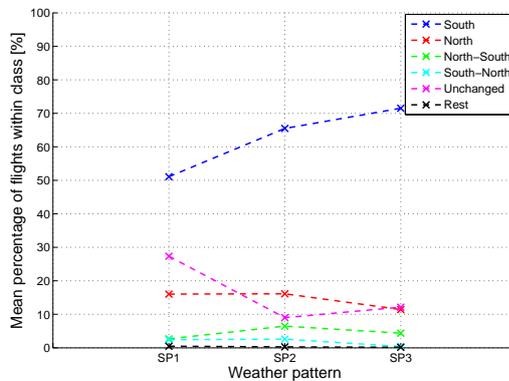
(b) Summer patterns

Figure 5.24: Box plots of the mean lateral shift per fleet for (a) winter weather patterns and (b) summer weather patterns. Positive shifts are shifts in northern direction, negative ones shifts in southern direction. The box depicts the interquartile range (IQR). Whiskers extend to 1.5-IQR or minimum or maximum if there is no data point outside 1.5-IQR. Outliers are data points greater or less than 1.5-IQR and are indicated by plus signs. Mean values are indicated by circles.

geographically rerouted flights are WP4 and SP1. This matches the findings from Figure 5.19, which showed that these weather patterns had the least amount of affected flights in terms of flight distance and/or flight time. Furthermore, WP5 is the only weather pattern in which flights are on average rerouted more north than south. This is in line with the fact that this winter pattern showed a mean lateral shift that was more north than all the others in Figure 5.24. Finally, in summer type 3 the mean percentage of flights that are rerouted south over the entire trajectory is with over 70% the highest of all weather types. This finding is also in line with Figure 5.24, which showed the most southern mean shift of all weather patterns.



(a) Winter patterns



(b) Summer patterns

Figure 5.25: The percentage of flights within each trajectory class as a function of (a) winter weather pattern and (b) summer weather pattern.

5.4.4. Altitude shift

The final trajectory parameter that is computed for each weather pattern is the mean altitude shift. The corresponding box plots are provided in Figure 5.26. Once again, there is a large variability between the weather types. Especially the summer types show very large differences. The median of SP1 shows only a slight reduction in mean altitude of 130 meters. The median of SP3 on the other hand, indicates that 50% of the flights are lowered in height by more than 1.5 kilometers. This is also the lowest median of all weather situations. Furthermore, all winter patterns contain cases that undergo a

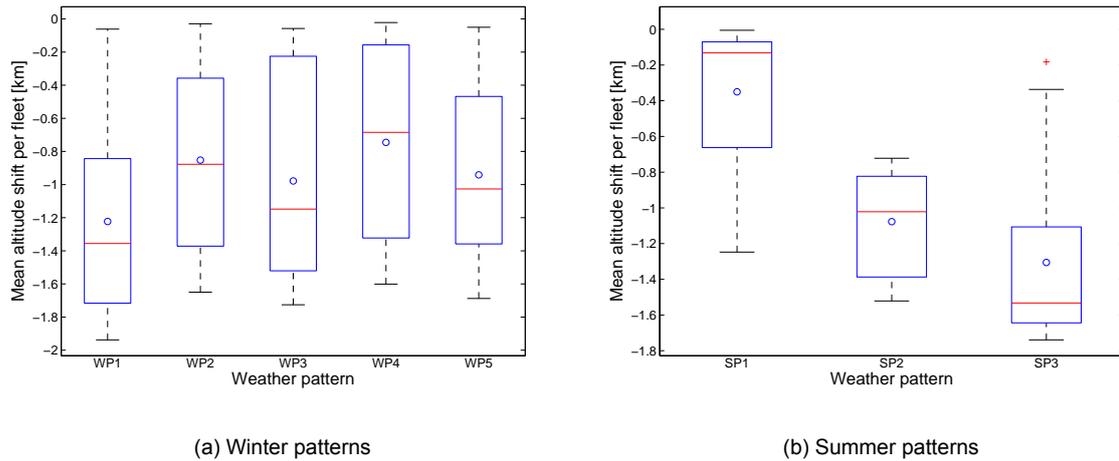


Figure 5.26: Box plots of the mean altitude per fleet shift as a function of (a) winter weather pattern and (b) summer weather pattern. Positive shifts are shifts towards higher altitudes, negative ones towards lower altitudes. The box depicts the interquartile range (IQR). Whiskers extend to 1.5·IQR or minimum or maximum if there is no data point outside 1.5·IQR. Outliers are data points greater or less than 1.5·IQR and are indicated by plus signs. Mean values are indicated by circles.

very small decrease in altitude, visualized by the extent of the upper whiskers.

5.4.5. Conclusion of the general influence of weather pattern

This section divided all available data sets based on their weather patterns. It is only possible to provide a few general conclusions here, as the weather patterns just differ too much for trends to become visible. First of all, the medians of the percentage of affected flights are all greater than 95%, except for summer type 1. Then, the mean flight times and distances are increased after optimization, for all weather patterns. Thirdly, the average values of the mean latitude shift indicate a clear southward shift for all patterns, except for winter type 5. Summer patterns are rerouted south by a larger mean shift than winter patterns. Furthermore, in every weather situation, the trajectory class that is characterized by a shift towards the south over the entire trajectory is the most dominant, except for WP5. For SP3, this percentage even increases to over 70%. Finally, every weather situation is characterized by a mean height shift towards lower altitudes. The median of the mean altitude shifts of SP1 is only 130 meters. SP1 is therefore by far the least affected weather type in terms of altitude shift.

6

Discussion

The previous two chapters provided the necessary material to get an idea of how air traffic is restructured when flying climate-optimal routes. What follows is a discussion on the results. The discussion is divided into four sections. The first one will discuss where to look for the most general observations. The second section examines differences and similarities between the case study and the general analysis. After this, the relation of this study to previous research will be highlighted. The final section concludes the discussion by pointing out the limitations of this study and presenting recommendations.

6.1. The big picture

It is important to be able to see the wood for the trees. The most general trends in trajectory changes are to be found in the general inter-Pareto analysis, at Pareto location 1. This point constitutes the combined trajectory changes of all fully climate-optimized data sets. The general Pareto location 1 is the only place in this study where the results are independent of flight direction, climate metric and weather pattern, and furthermore the only place in the Pareto range where every combination of direction, metric and weather type has got a data point, except for Pareto location 0. As stated in the conclusion of the general inter-Pareto analysis, the climate-optimal REACT4C trajectories on average experience a 4.1% increase in flight time, a 3% increase in flight distance, a 2.2 degree shift towards the south, and a 1.43 kilometer shift towards lower altitudes. Furthermore, the majority of flights (61%) is rerouted completely south of their cost-optimized counterparts, whereas 27% is relocated entirely to the north of the original trajectory.

6.2. Case study versus general analysis

The case study and the general analysis serve two different purposes. The former is a very specific investigation, focussed on providing detailed insight into the trajectory changes for one case. A tool was made that could be used to analyze any case. The case study presented in this research considers the case that proved to have a great potential to reduce the climate impact in [4]. The general study is conducted to unravel trends in the trajectory changes that hold for all optimizations performed by the REACT4C project. To get a grasp of the representativeness of the general results, let us look for similarities and/or differences between the two studies.

The inter-Pareto trends are the same for both the case and general analysis. Each kind of trajectory change becomes more distinct when progressing towards higher Pareto locations, whether it is the duration or distance increase, the lateral shift or the altitude shift.

The influence of direction is generally also very similar between the two studies. The mean flight time and distance increase is almost the same in both studies. Eastbound flights are in both cases rerouted less south than westbound flights, and are lowered less. There is, however a considerable difference in the trajectory classification. Where the general analysis shows that the majority of flights in eastbound direction is shifted entirely south, the case study tells that most of the flights are rerouted entirely north.

The metric comparison from both the case study and the general analysis proves that the choice of metric hardly influences the way in which trajectories are altered. In both analyses, the altitude shift is

the trajectory characteristic that is most influenced by the choice of metric, albeit still to a small extent. It seems to be the case that the climate metrics “agree” on what the best rerouting solution is for the climate. The main difference between them is the absolute value of the climate impact.

Finally, trends between weather patterns are as indistinguishable in the case study as they are in the general analysis. In that sense, there is no considerable difference between the two studies.

The few differences between the case study and the general analysis imply that the general results should not be used to get an overview of the trajectories of a specific case. They *can* be used to obtain a good impression of what the effect on the trajectories in general is when operating climate-optimal routes, as the most important trends remain largely the same.

6.3. Relation to previous research

Research that has been done before in the field of REACT4C’s climate-optimal trajectories is limited. A very basic analysis of the case that was also treated in this case study presented in this research, was done in [4]. Figures 5 and 6 and Table 2 of that reference can serve as a verification of this code. These figures indicate that the altitude is decreased by approximately 80 flight levels, which is equivalent to 1.4 kilometer. The same kind of displacement was found in the plot of the altitude shift depending on the geographic location. Furthermore, the lateral shift in [4] was found to be 2.2 degrees south on average. The case study presented in this thesis did not present a mean latitude shift, as this was considered to be not detailed enough for a case study. However, the provided plots of the lateral shift as a function of longitude and latitude indicated that the emphasis was on southward reroutings. Hence, the mean shift is likely to be similar in both studies. The increase in mean flight duration according to Grewe et al. is 7.2%. Our analysis showed that this was only 4.5%. Furthermore, the mean distance in [4] increased by 3.8%, whereas in this study the increment was only 2.0%. The origin of these differences could not be determined, as Grewe et al. provide no explanation of how they computed the mean duration and distance from the available data sets. Although the percentages somewhat differ, the trend that flight distance is increased less than the flight duration is the same.

This research went much further in the analysis of REACT4C trajectories. Next to the fact that this study provides PDFs of flight duration and distance increments and examines the correlation between latitude shifts and altitude shifts on one hand and geographic location on the other, also the influence of direction, metric and weather pattern on the flight trajectory changes are investigated for the first time. This leads to a more complete and well-founded case study.

Furthermore, it is the first time that REACT4C routes were analyzed on a general scale. Incorporating all of the data sets that stem from the REACT4C project allows for a very general impression of the effect of climate optimization on flight trajectories.

The way in which REACT4C routes are changed, is similar to previous studies on climate-optimal trajectories. For instance, the TRADEOFF project investigated the possibility to adapt the cruise altitude in order to reduce the climate impact [12]. The height was altered between 2000 feet towards higher altitudes and 6000 feet towards lower altitudes. It was found that the climate impact is lowest when flying at lower altitudes. That finding matches the result from this research, which showed that the trajectories tend to be shifted downward when optimizing for climate impact.

The CATS project took the optimization one step further, and allowed the Mach number to be varied on top of the cruise altitude [13]. It was found once again that climate impact was reduced when lowering the cruise altitude, which is similar to the TRADEOFF findings, and to the results established in this thesis on REACT4C routes. Furthermore, CATS found that there is an optimal Mach number at which climate impact is minimal. For a specific route between Detroit and Frankfurt, the optimal Mach number was found to be 0.625, which is lower than the cruise velocities generally adopted in today’s jet propelled air transport. The authors indicate that this is a general trend. The analysis of REACT4C routes provided in this document showed that the travel times are increased more than the travel distance, percentage-wise. This means that the mean velocity of the fleet decreases. The mean Mach number is not provided in this research, and even though this number is dependent on height, a clear trend of speed reduction is present in both the REACT4C routes and the CATS routes.

6.4. Limitations and recommendations

There are three limitations to this study. First, this research focussed on the trajectory changes, and did not investigate why the routes are altered the way they are. If the reasons behind the differences

are known, one would have a complete understanding of the results of the REACT4C project. Especially the reasons behind the trajectory differences between weather patterns would benefit general understanding, as no clear trends between them are present. The driving parameters behind the route alterations are the main topic of another thesis that is currently being made.

The second issue is based on the fact that the climate cost functions, which play a prominent role in REACT4C's optimization approach, are five-dimensional quantities. They contain the climate impact per unit emission, depending on 1) the climate agent, 2-4) the longitude, latitude and altitude at which the emission takes place, and 5) the point in time at which the emission takes place. The fact that emission time is incorporated, allowed the flights to be shifted within the period of a day. The question whether, and if so to what extent, flights are rescheduled throughout the day, was not included in this research. This could be done in a further investigation.

A third limitation to this research is that no data was available of the geographical location of the jet stream in each weather pattern. This made it impossible to examine how the rerouted trajectories relate to the jet stream. Only the trajectory changes, relative to the cost-optimized routes, were considered. A further study of the way in which climate-optimal routes avoid or make use of the jet stream, would add to the scientific knowledge of REACT4C's climate-optimal routes.



Conclusion

The REACT4C project aimed to reduce aviation-induced climate warming by rerouting air traffic between North America and Europe. The project members made a distinction between east- and west-bound flights, between three climate metrics to quantify climate impact, and between eight weather situations. For each combination of these three variables, they optimized the fleet for six different levels of climate impact importance. This resulted in a collection of 288 optimized fleet data sets.

The lack of insight in how the climate-optimized trajectories compare to their cost-optimized counterparts led to the establishment of this research. The goal was to do a trajectory analysis on two levels: one at a low level of generality, in which the route alterations of one case are examined, the other at high level of generality, in which all cases are combined to identify existing trends in the way trajectories are changed. A tool was created to do the case study, such that any case can be examined.

In both studies, the influence of the level of climate optimization, direction, metric and weather type on several trajectory characteristics was investigated. The following characteristics were established to get a good impression of how a fleet is rerouted: the mean flight duration and distance increase, distributions of the flight duration and distance increments, and latitude and altitude shift, depending on geographic location. Furthermore, reroutings were categorized into six different classes, based on the lateral shift along the trajectory, relative to the cost-optimized reference flight.

It was found that generally, fully climate-optimized flights are in the air for a 4.1% longer period of time. The distance they travel is increased by 3% on average, which indicates that the mean velocity of the flights decreases. It was shown that not all routes undergo an increase in time and distance, and that some are shortened. Furthermore, it was observed that the emphasis of lateral relocation is on southward shifts. Not only is the mean latitude shift of the fully climate-optimized flights 2.2 degrees south, but also 61% of flights are shifted south along the entire trajectory with respect to the cost-optimized counterparts, on average. The second most prominent route category is the northern relocation along the complete flight path, displayed by 27% of the flights on average. The remaining 12% consists of routes shifted north in the west and south in the east, routes that show the opposite pattern, namely a southward relocation in the west and a northward relocation in the east, and routes that show no lateral difference with the original ones. Routes that can not be categorized, constitute only 0.4% of the total number of flights, on average. Finally, it was shown that the fully climate-optimized flights in general undergo a 1.43 kilometer shift towards lower altitudes.

Both the case study and the general analysis agreed on certain trends that became visible in the influence of level of climate optimization, flight direction, climate metric and weather pattern. The statistics outlined above hold for the data sets that result in maximum climate impact reduction. Reducing the level of climate optimization was shown to result in the same trajectory changes, but then milder and to a lesser extent.

The flight direction analyses showed that the percentage of increase in mean travel times and distances does not differ much between the two directions. However, traffic towards America is characterized by a larger spread in terms of travel time increments; more flights show decrease in time, and more show a large increase. The lateral shift of eastbound flights is very distinctively less south than westbound flights. Flights towards Europe are less often rerouted entirely south and more often entirely

north. Lastly, the downward shift in altitude is less extreme for eastbound flights than for westbound flights.

Climate metrics did not seem to have a lot of impact on the way the traffic is restructured. The only characteristic that showed a noteworthy cross-metric difference is the mean altitude shift. The difference between the highest and lowest mean height shift is 290 meters, with the metric P-AGWP100 characterized by the greatest mean altitude decrease.

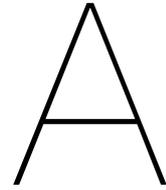
Finally, weather patterns were shown to largely affect the way in which the trajectories are rerouted. Because the weather situations are very distinct in terms of position and strength of the jet stream, no trends in the differences between the trajectory changes were found.

On top of a scientific understanding of the REACT4C routes, let this research serve as a means to let air traffic management entities get a grasp of what climate-optimal trajectories look like, if or when the day comes that climate-optimal routes are put into practice. Let it also be useful for an analysis or prediction of the large-scale economic impact that would result when implementing the REACT4C optimal routes, due to the typical flight duration and distance alterations, and the extra fuel consumption and rescheduling problems that come with it. But foremost, let it be an inspiration to other scientists and non-scientists to pursue climate-friendly solutions to any kind of problem, and to question whether the status quo of technological advancement should really be this harmful for future generations.

Bibliography

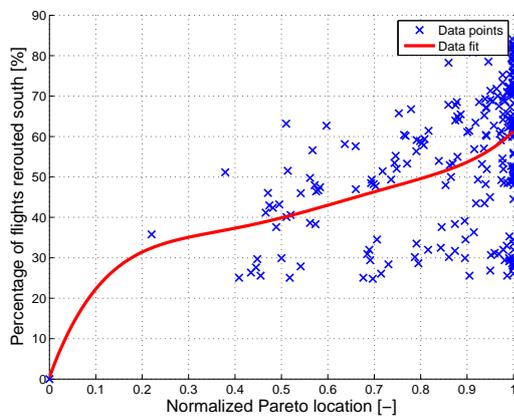
- [1] L. Whitmarsh, "Scepticism and uncertainty about climate change: Dimensions, determinants and change over time," *Global Environmental Change*, vol. 21, pp. 690–700, may 2011.
- [2] D. Lee, G. Pitari, V. Grewe, K. Gierens, J. Penner, A. Petzold, M. Prather, U. Schumann, A. Bais, and T. Berntsen, "Transport impacts on atmosphere and climate: Aviation," *Atmospheric Environment*, vol. 44, pp. 4678–4734, dec 2010.
- [3] V. Grewe, C. Frömming, S. Matthes, S. Brinkop, M. Ponater, S. Dietmüller, P. Jöckel, H. Garny, E. Tsati, K. Dahlmann, O. A. Søvde, J. Fuglestedt, T. K. Berntsen, K. P. Shine, E. A. Irvine, T. Champougny, and P. Hullah, "Aircraft routing with minimal climate impact: The REACT4C climate cost function modelling approach (V1.0)," *Geoscientific Model Development*, vol. 7, no. 1, pp. 175–201, 2014.
- [4] V. Grewe, T. Champougny, S. Matthes, C. Frömming, S. Brinkop, O. A. Søvde, E. A. Irvine, and L. Halscheidt, "Reduction of the air traffic's contribution to climate change: A REACT4C case study," *Atmospheric Environment*, vol. 94, pp. 616–625, 2014.
- [5] E. A. Irvine, B. J. Hoskins, K. P. Shine, R. W. Lunnon, and C. Froemming, "Characterizing North Atlantic weather patterns for climate-optimal aircraft routing," *Meteorological Applications*, vol. 20, pp. 80–93, mar 2013.
- [6] J. S. Fuglestedt, K. P. Shine, T. Berntsen, J. Cook, D. S. Lee, A. Stenke, R. B. Skeie, G. J. M. Velders, and I. a. Waitz, "Transport impacts on atmosphere and climate: Metrics," *Atmospheric Environment*, vol. 44, no. 37, pp. 4648–4677, 2010.
- [7] V. Grewe and K. Dahlmann, "How ambiguous are climate metrics? And are we prepared to assess and compare the climate impact of new air traffic technologies?," *Atmospheric Environment*, vol. 106, pp. 373–374, apr 2015.
- [8] P. Jöckel, H. Tost, A. Pozzer, C. Brühl, J. Buchholz, L. Ganzeveld, P. Hoor, A. Kerkweg, M. G. Lawrence, R. Sander, B. Steil, G. Stiller, M. Tanarhte, D. Taraborrelli, J. van Aardenne, and J. Lelieveld, "The atmospheric chemistry general circulation model ECHAM5/MESy1: consistent simulation of ozone from the surface to the mesosphere," *Atmospheric Chemistry and Physics*, vol. 6, no. 12, pp. 5067–5104, 2006.
- [9] E. Roeckner, R. Brokopf, M. Esch, M. Giorgetta, S. Hagemann, L. Kornblueh, E. Manzini, U. Schlese, and U. Schulzweida, "Sensitivity of Simulated Climate to Horizontal and Vertical Resolution in the ECHAM5 Atmosphere Model," *Journal of Climate*, vol. 19, pp. 3771–3791, jan 2006.
- [10] P. Jöckel, A. Kerkweg, A. Pozzer, R. Sander, H. Tost, H. Riede, A. Baumgaertner, S. Gromov, and B. Kern, "Development cycle 2 of the Modular Earth Submodel System (MESy2)," *Geoscientific Model Development*, vol. 3, no. 2, pp. 717–752, 2010.
- [11] O. Morgenstern, M. A. Giorgetta, K. Shibata, V. Eyring, D. W. Waugh, T. G. Shepherd, H. Akiyoshi, J. Austin, A. J. G. Baumgaertner, S. Bekki, P. Braesicke, C. Brühl, M. P. Chipperfield, D. Cugnet, M. Dameris, S. Dhomse, S. M. Friith, H. Garny, A. Gettelman, S. C. Hardiman, M. I. Hegglin, P. Jöckel, D. E. Kinnison, J.-F. Lamarque, E. Mancini, E. Manzini, M. Marchand, M. Michou, T. Nakamura, J. E. Nielsen, D. Olivié, G. Pitari, D. A. Plummer, E. Rozanov, J. F. Scinocca, D. Smale, H. Teyssède, M. Toohey, W. Tian, and Y. Yamashita, "Review of the formulation of present-generation stratospheric chemistry-climate models and associated external forcings," *Journal of Geophysical Research: Atmospheres*, vol. 115, no. D3, 2010.

- [12] C. Frömming, M. Ponater, K. Dahlmann, V. Grewe, D. S. Lee, and R. Sausen, "Aviation-induced radiative forcing and surface temperature change in dependency of the emission altitude," *Journal of Geophysical Research: Atmospheres*, vol. 117, oct 2012.
- [13] A. Koch, B. Lührs, K. Dahlmann, F. Linke, V. Grewe, M. Litz, M. Plohr, B. Nagel, V. Gollnick, and U. Schumann, "Climate impact assessment of varying cruise flight altitudes applying the CATS simulation approach," in *Third International Conference of the European Aerospace Societies, Venice, Italy, 2011*.

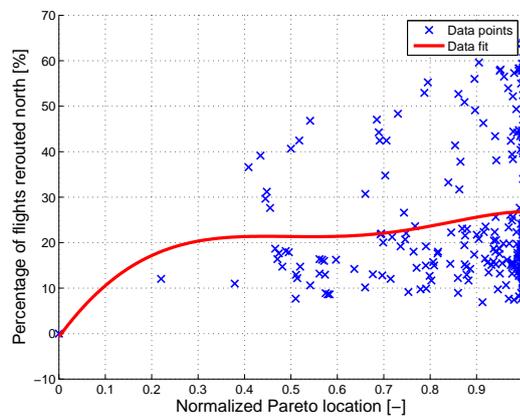


Variability of trajectory classes as a function of Pareto location

In Figure 5.5 on page 47, data fits were presented to indicate trends in how the trajectories are shifted with respect to their cost-optimal counterparts as a function of normalized Pareto location. The data points were excluded to retain the clarity of the figure. In Figure A.1, each trend line is plotted separately, together with the data points that are being fitted. This way, the extent of the variability becomes clear.

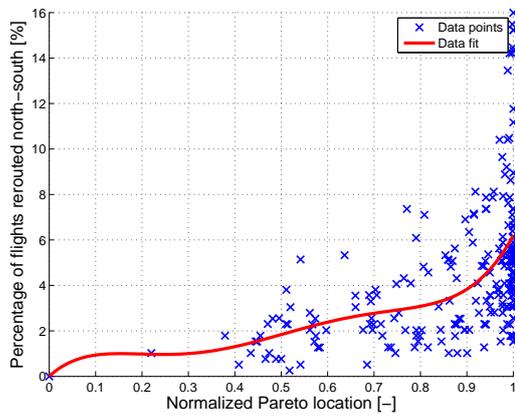


(a) Entirely south

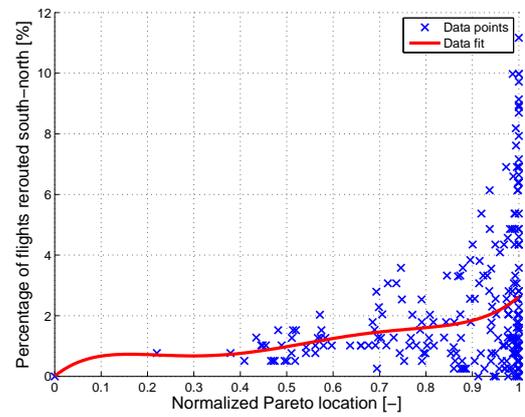


(b) Entirely north

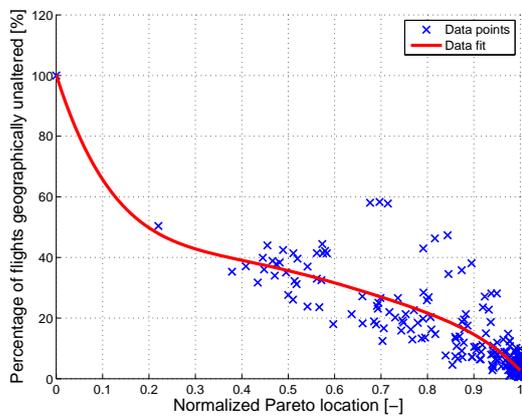
Figure A.1: Trend line and data points of the percentage of flights within each trajectory class as a function of normalized Pareto location.



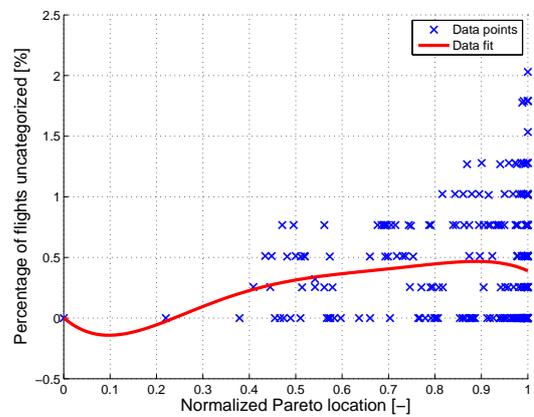
(c) North in the west, south in the east



(d) South in the west, north in the east



(e) Geographically unaltered



(f) Uncategorized

Figure A.1: Trend line and data points of the percentage of flights within each trajectory class as a function of normalized Pareto location (continued).

B

Flight-based box plots

In the general study of climate-optimal routes, flight characteristics were compared by making use of fleet averages. This was done to magnify the differences between the alternatives of the case differentiator of which the influence was examined. This appendix provides the reader with box plots of flight characteristics that are constructed by means of flight values instead of fleet averages. The box plots presented in Figures B.1 to B.8 allow the reader to have a more detailed impression of the variability of flight characteristics within fleets.

B.1. General direction comparison

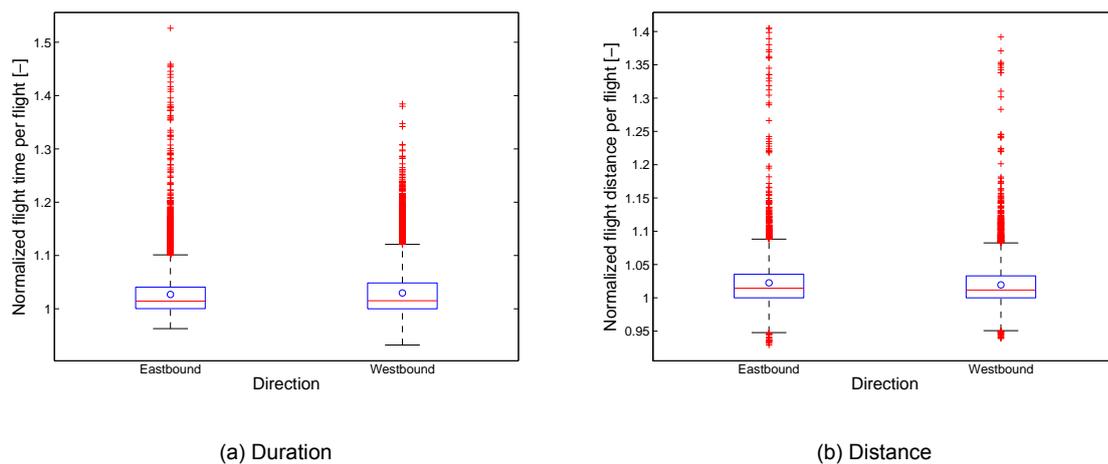
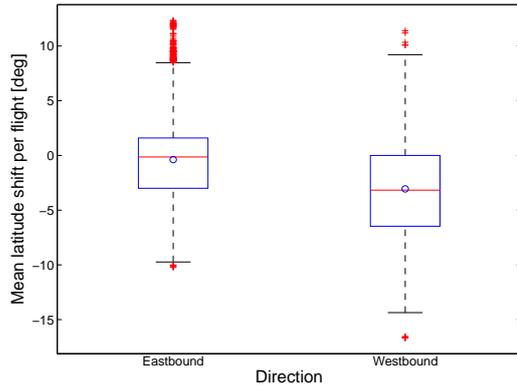
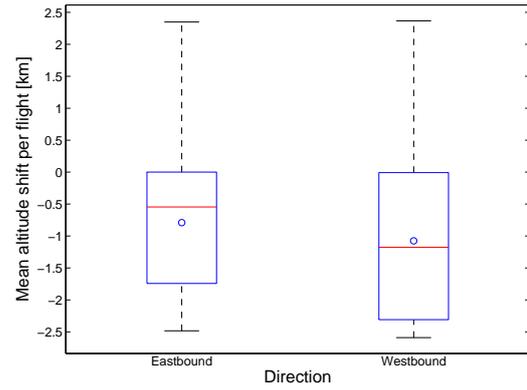


Figure B.1: Box plots of (a) the normalized flight duration (b) the normalized flight duration as a function of flight direction. The box depicts the interquartile range (IQR). Whiskers extend to 1.5·IQR or minimum or maximum if there is no data point outside 1.5·IQR. Outliers are data points greater or less than 1.5·IQR and are indicated by plus signs. Mean values are indicated by circles.



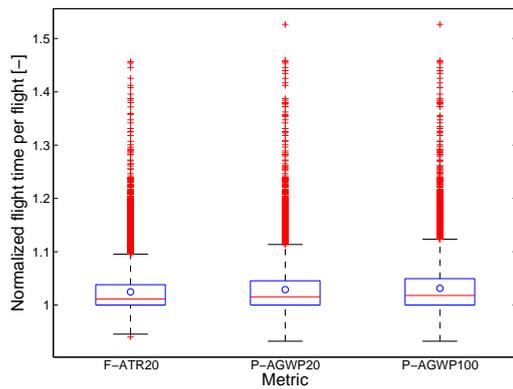
(a) Duration



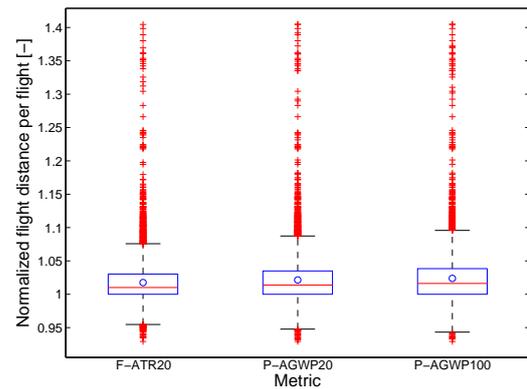
(b) Distance

Figure B.2: Box plots of (a) the mean lateral shift per flight and (b) the mean altitude shift per flight, as a function of flight direction. The box depicts the interquartile range (IQR). Whiskers extend to 1.5-IQR or minimum or maximum if there is no data point outside 1.5-IQR. Outliers are data points greater or less than 1.5-IQR and are indicated by plus signs. Mean values are indicated by circles.

B.2. General climate metric comparison

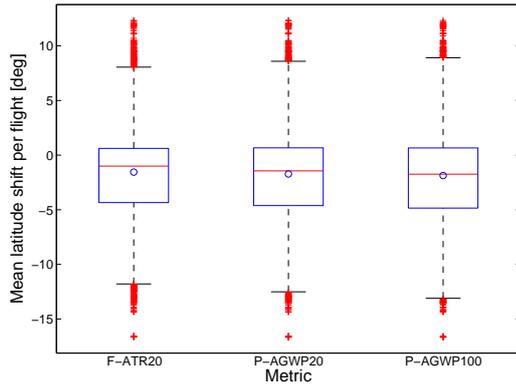


(a) Duration

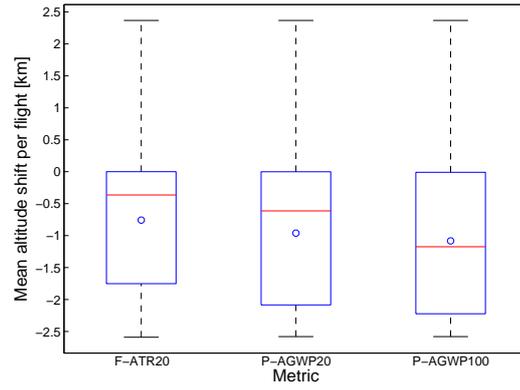


(b) Distance

Figure B.3: Box plots of the normalized flight duration (a) and distance (b) as a function of climate metric. The box depicts the interquartile range (IQR). Whiskers extend to 1.5-IQR or minimum or maximum if there is no data point outside 1.5-IQR. Outliers are data points greater or less than 1.5-IQR and are indicated by plus signs. Mean values are indicated by circles.



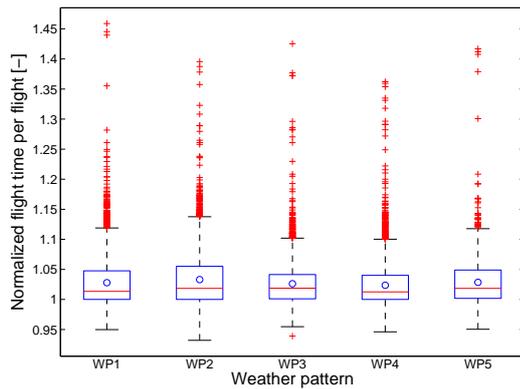
(a) Duration



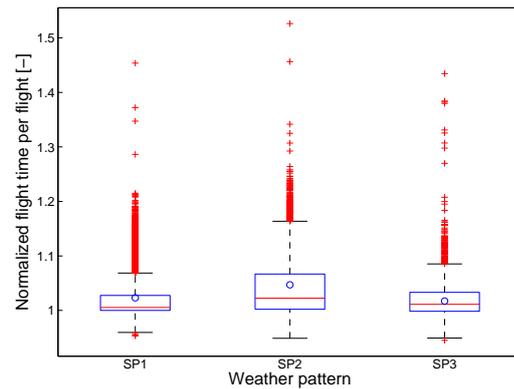
(b) Distance

Figure B.4: Box plots of (a) the mean lateral shift per flight and (b) the mean altitude shift per flight, as a function of climate metric. The box depicts the interquartile range (IQR). Whiskers extend to 1.5·IQR or minimum or maximum if there is no data point outside 1.5·IQR. Outliers are data points greater or less than 1.5·IQR and are indicated by plus signs. Mean values are indicated by circles.

B.3. General weather pattern comparison

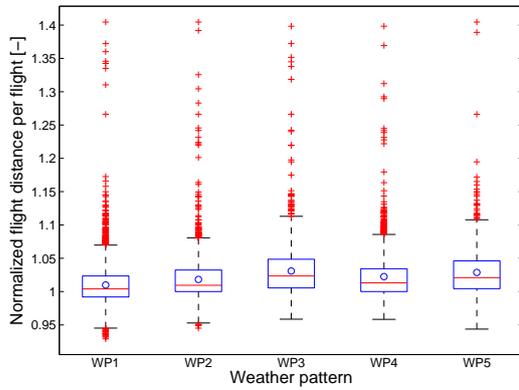


(a) Winter patterns

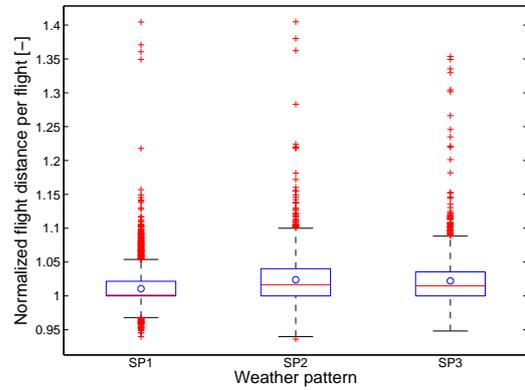


(b) Summer patterns

Figure B.5: Box plots of the normalized flight duration for (a) winter weather patterns and (b) summer weather patterns. The box depicts the interquartile range (IQR). Whiskers extend to 1.5·IQR or minimum or maximum if there is no data point outside 1.5·IQR. Outliers are data points greater or less than 1.5·IQR and are indicated by plus signs. Mean values are indicated by circles.

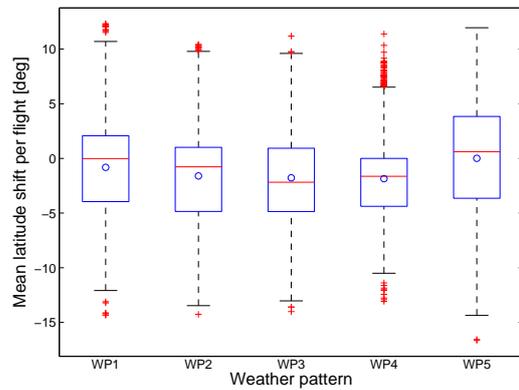


(a) Winter patterns

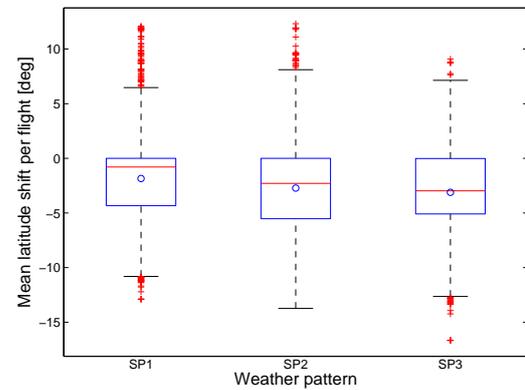


(b) Summer patterns

Figure B.6: Box plots of the normalized flight distance for (a) winter weather patterns and (b) summer weather patterns. The box depicts the interquartile range (IQR). Whiskers extend to 1.5·IQR or minimum or maximum if there is no data point outside 1.5·IQR. Outliers are data points greater or less than 1.5·IQR and are indicated by plus signs. Mean values are indicated by circles.

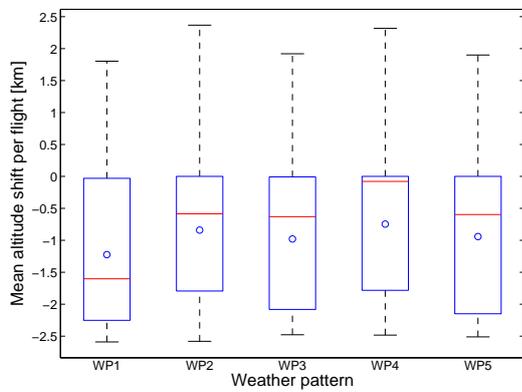


(a) Winter patterns

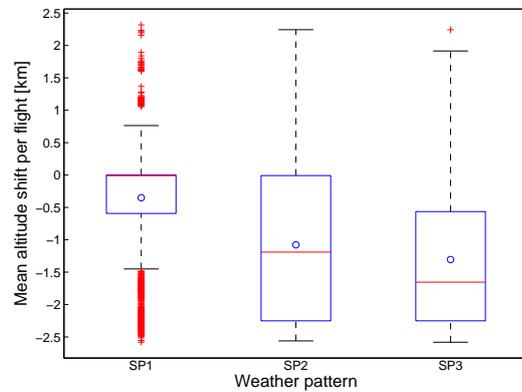


(b) Summer patterns

Figure B.7: Box plots of the mean lateral shift per flight for (a) winter weather patterns and (b) summer weather patterns. The box depicts the interquartile range (IQR). Whiskers extend to 1.5·IQR or minimum or maximum if there is no data point outside 1.5·IQR. Outliers are data points greater or less than 1.5·IQR and are indicated by plus signs. Mean values are indicated by circles.



(a) Winter patterns



(b) Summer patterns

Figure B.8: Box plots of the mean altitude shift per flight for (a) winter weather patterns and (b) summer weather patterns. The box depicts the interquartile range (IQR). Whiskers extend to 1.5·IQR or minimum or maximum if there is no data point outside 1.5·IQR. Outliers are data points greater or less than 1.5·IQR and are indicated by plus signs. Mean values are indicated by circles.