Reformulating aircraft routing algorithms to reduce fuel burn and thus CO_2 emissions

Cathie A. Wells

Submitted for the degree of Doctor of Philosophy in Mathematics

Mathematics of Planet Earth Centre for Doctoral Training

Department of Mathematics and Statistics

University of Reading

12th September, 2022

Abstract

During the UN Climate Change Conference (COP26), in November 2021, the international aviation community agreed to advance actions to reduce CO_2 emissions. Adopting more fuel efficient routes, now that full global satellite coverage is available, could achieve this quickly and economically. Here flights between New York and London, from 1st December, 2019 to 29th February, 2020 are considered. Trajectories through wind fields from a global atmospheric re-analysis dataset are found using optimal control theory. Initially, time minimal routes are obtained by applying Pontryagin's Minimum Principle. Minimum time air distances are compared with actual Air Traffic Management tracks. Potential air distance savings range from 0.7 to 16.4%, depending on direction and track efficiency. To gauge the potential for longer duration time minimal round trips in the future, due to climate change, trajectories are considered for historic and future time periods, using an ensemble of climate models. Next, fixed-time, fuel-minimal routes are sought. Fuel consumption is modelled with a new physics-driven fuel burn function, which is aircraft model specific. Control variables of position-dependent aircraft headings and airspeeds or just headings are used. The importance of airspeed in finding trajectories is established, by comparing fuel burn found from a global search of optimised results for the discretised approximation of each formulation. Finally, dynamic programming is applied to find free-time, fuel-optimal routes. Results show that significant fuel reductions are possible, compared with estimates of fuel use from actual flights, without significant changes to flight duration. Fuel use for winter 2019–2020 could have been reduced by 4.6% eastbound and 3.9% westbound on flights between Heathrow and John F Kennedy Airports. This equates to a 16.6 million kg reduction in CO_2 emissions. Thus large reductions in fuel consumption and emissions are possible immediately, without waiting decades for incremental improvements in fuel-efficiency through technological advances.

Declaration

I confirm that this is my own work and the use of all material from other sources has been properly and fully acknowledged. Cathie Wells

Acknowledgements

First and foremost I would like to thank my principal supervisor Paul Williams, without whose unfailing enthusiasm, I would not have had the confidence to make this journey. Paul's knowledge of aviation and meteorology as well as his support and encouragement have been invaluable.

I have also been very fortunate in having co-supervisors with a wealth of specialist knowledge. Nancy Nichols' vast experience has been very helpful, not just in tackling Pontryagin's Minimum Principle, but also in applying for travel awards. Dante Kalise at Imperial College, has introduced me to the joys of dynamic programming and none of my research would have been possible without Ian Poll's fuel burn function and his readiness to help out with all questions pertaining to the practicalities of aviation.

I am also thankful for the support of the entire MPE CDT team, past and present. In particular I'd like to thank Jennifer Scott and Helen Dacre for keeping me on track throughout and providing extra opportunities for me beyond my primary research. On the Reading campus Sam Williams, Jan Fillingham and Ruth Harris have provided far more than just administrative support, for which I am very grateful. The previous MPE CDT cohorts have been extremely generous with their time and advice particularly Jemima, Adrian, Jennifer, Ieva and Elena; those reports and posters would never have happened without you. However, it is to the Echoes that I owe my deepest debt of gratitude. You are a great bunch and your camaraderie throughout has kept me moving forwards, even when my trajectories didn't seem to be. I really would not have finished the MRes let alone the thesis without the JJT 211 Office Crew and particularly Oliver and Sam who have helped, inspired and entertained me at many a pause in our work.

Finally I would like to thank my parents, my children and my long suffering husband for putting up with my less cheerful and positive moments, as well as for trying to sound interested when offering to listen to my presentations!

Acknowledgements for published papers:

(Wells et al., 2021)

CAW is supported by the Mathematics of Planet Earth Centre for Doctoral Training at the University of Reading. This work has been made possible by the financial support of the Engineering and Physical Sciences Research Council, grant number EP/L016613/1. DK's involvement in this research benefited from the support of the FMJH 'Program Gaspard Monge for optimization and operations research and their interactions with data science', and from the support from EDF, Thales and Orange. NKN is supported in part by the NERC National Centre for Earth Observation (NCEO).We would like to thank Mr. Ian Jopson from NATS for the provided ATM track usage data.

(Wells et al., 2022)

CAW is supported by the Mathematics of Planet Earth Centre for Doctoral Training at the University of Reading. This work has been made possible by the financial support of the Engineering and Physical Sciences Research Council, grant number EP/L016613/1. DK was supported by the UK Engineering and Physical Sciences Research Council (EPSRC) grants EP/V04771X/1, EP/T024429/1, and EP/V025899/1. NKN is supported in part by the UK Natural Environment Research Council National Centre for Earth Observation (NCEO).

Contents

List of Figures

List of Tables

List of Acronyms

Below are the main acronyms used in this thesis.

ADS	Automatic Dependent Surveillance
ADS-B	Automatic Dependent Surveillance Broadcast
ANSP	Air Navigation Service Provider
ASPIRE	Asia and Pacific Initiative to Reduce Emissions
ATM	Air Traffic Management
BADA	Base of Airline Data
CARATS	Collaborative Actions for Renovation of Air Traffic Systems
CI	Cost Index
CMIP	Climate Model Intercomparison Project
CORSIA	Carbon Offsetting and Reduction Scheme for International Aviation
CPDLC	Controller Pilot Data Link Communications
ERA	Environmentally Responsible Aviation
FL	Flight Level
HJB	Hamilton–Jacobi–Bellman
IATA	International Air Transport Association
ICAO	International Civil Aviation Organisation
ISA	International Standard Atmosphere
JFK	John F Kennedy Airport in New York

KKT	Karush–Kuhn–Tucker
LHR	London Heathrow Airport
MMO	Maximum Operational Mach number
MPC	Model Predictive Control
MTOM	Maximum Take-Off Mass
MZFM	Maximum permitted Zero Fuel Mass
NAO	North Atlantic Oscillation
NCAR	National Center for Atmospheric Research
OCP	Optimal Control Problem
OEM	Operational Empty Mass
OFW	Optimised For Wind
OTS	Organised Track Structure
PMP	Pontryagin's Minimum Principle
PRM	Preferred Route Message
RCP	Representative Concentration Pathway
SAF	Sustainable Aviation Fuels
SESAR	Single European Sky ATM Research
SOCM	Start of Cruise Mass
SQP	Sequential Quadratic Programming
TBO	Trajectory Based Operations
TPBVP	Two Point Boundary Value Problem
UNFCCC	United Nations Framework Convention on Climate Change
UTC	Co-ordinated Universal Time
WMO	World Meteorological Organization

List of Symbols

Below are the main symbols used in this thesis. Any other use of symbols in a particular section is explained in the text.

$\mathcal{A}$	set of all admissible controls
C	target set
d_H	distance around the sphere between $(\lambda_{dest}, \phi_{dest})$ and $(\lambda(t_f), \phi(t_f))$ in m
g	fuel burn in $\mathrm{kg\,s^{-1}}$
h	altitude of aircraft above Earth in m ($h \ll R$)
$H[.]$	Pontryagin Hamiltonian
J	cost functional
L	Lagrangian
M	mass of aircraft in kg
R	radius of Earth ≈6 371 000 m
t	time in s
t_f	final time in s
u	zonal (eastward) wind in $\mathrm{m\,s^{-1}}$
v	meridional (northward) wind in $\mathrm{m\,s^{-1}}$
V	airspeed in $\mathrm{m\,s^{-1}}$
$\mathbf{x}$	vector of state variables

$\boldsymbol{\alpha}$	vector of control variables
β	running cost
δ_{dest}	radius of target around destination airport in m
θ	heading angle in radians
λ	longitude in radians
ϕ	latitude in radians
Ω	closed domain of state space
$.^*$	optimal
$\cdot_{dept}$	quantity at departure
$\cdot_{dest}$	quantity at destination
$\cdot_{ref}$	reference value

Chapter 1

Introduction

Commercial aviation can be held responsible for 3.5% of all anthropogenic climate change (Lee et al., 2020). By 2050 emissions are forecast to be three times those recorded in 2015 (ETS, 2021). We rely on flights, however, to move huge amounts of freight (61.3 million tonnes in 2019) and passengers (4.54 billion in 2019) (Mazareanu, 2020). The United Nations Framework Convention on Climate Change (UNFCCC) has set a goal of commercial aviation being net-zero by 2050 (UKCOP26, 2021), with the European Union also taking the intermediate step of targeting a 55% net reduction in aviation greenhouse gas emissions by 2030 (ETS, 2021). These are worthy ambitions, but there is a heavy reliance in the industry on carbon off-setting, with 80% of the 2021-2035 EU emissions above 2020 levels expected to be offset (ETS, 2021). Not creating the emissions in the first place would obviously be preferable, but this is not so straightforward, with electrification impractical for long haul flights and hydrogen powered aircraft needing further research (Schwab et al., 2021; Airbus, 2021). Clearly we need an immediate answer to limit the impact of aviation emissions. This is offered by improved trajectory planning. By taking advantage of wind conditions and a new satellite communications network, aircraft will be able to fly routes that reduce their fuel use and thus their CO_2 emissions (Aireon, 2020; Girardet et al., 2014).

In this thesis, we look at different ways to plan trajectories across the North Atlantic, between John F Kennedy Airport in New York (JFK) and London Heathrow Airport (LHR). We consider the problem through the framing of Optimal Control Theory and then use three different ways to reach a solution. In the first instance time minimal trajectories at fixed altitude and airspeed are considered, which allows the corresponding Optimal Control Problem (OCP) to be solved using Pontryagin's Minimum Principle

(PMP) (Pontryagin et al., 1962). This method relies on optimisation of the initial system, before discretisation to allow for an approximate numerical solution. When looking at fuel optimisation for a fixed-time flight, we instead discretise the system first and find the control variables which give the minimum value of a cost functional at each time interval, subject to a non-linear constraint. From here a trajectory can be created. Both of these approaches are reliant on values being supplied to initiate the optimisation. In the third method, dynamic programming, this is not necessary, as by solving the Hamilton–Jacobi–Bellman (HJB) equation a value function is obtained which allows all points on a grid to be tested for all combinations of control variables. In this way we recover the optimal feedback control based on discrete time intervals and a spatially discretised state space. From here fuel optimal, free-time trajetories can be found.

Efficient trajectory planning has been the subject of much research in the past, but previously routes have often been optimised without accounting for the wind field, or through simplified wind fields to allow for easier mathematical analysis. In this thesis we use reanalysis data to allow actual wind fields to be considered (Kalnay et al., 1996). This gives a more realistic idea of the complexity of the atmospheric conditions encountered across the North Atlantic and allows for a useful comparison to be drawn between actual flight data and results from modelled, optimised flights. Whilst previous research has tended to concentrate on a selection of a few particular days, we use data describing the atmospheric conditions for all 91 days of an entire winter season. In other trajectory planning research, actual flight data is often not considered or is used for a restricted number of flights. In Chapter 7 we compare simulated fuel and time minimal trajectories with over 3000 actual flight routes. The most important novelty of this research is that a new fuel burn function has been applied to estimate fuel use for both actual flights and simulations (Poll & Schumann, 2021a,b). This aims to give an accurate value for fuel consumption, relying on more transparent methods than Eurocontrol's Base of Airline Data (Eurocontrol, 2021a).

1.1 What questions remain to be answered?

Although this is a well explored area of research, there are still some key questions that await answers:

1. How much difference would horizontal time optimal trajectory planning make to fuel use and thus carbon dioxide emissions, in fixed airspeed transatlantic flights,

compared with the Organised Track Structure (OTS) used currently?

2. To what extent is climate change likely to affect minimum times of flight within the next 35 years?

3. Can fixed-time flights be planned for a whole winter season to ensure that fuel is minimised?

4. In minimising fuel burn of transatlantic flights, can extra benefits result from controlling the airspeed in addition to the heading angle?

5. Would minimising fuel, but allowing time to be free, in a fixed altitude flight produce trajectories that are more fuel efficient than those currently flown and to what extent would the flight duration be changed?

Question 1 will form the basis of the research shown in Chapter 4. The main section of this material has been published in the Environmental Research Letters journal (Wells et al., 2021). Question 2 is answered in Chapter 5 through a project undertaken jointly with the Met Office and forms part of a paper submitted to Meteorological Applications. Results and methods described in this chapter, however, were found independently by the author of this thesis. Questions 3 and 4 are both discussed in Chapter 6, the content of which has been published in the Optimization and Engineering Journal (Wells et al., 2022). The final question is addressed in Chapter 7. A summary of the answers to all of these questions and a brief overview of possible future research directions are given in Chapter 8.

1.2 Outline

This thesis aims to present the answers to these questions as well as providing some background information on motivation, previous research and methodology.

In Chapter 2 the motivation behind this research is explored, looking at why aviation emissions are a priority, how aviation is affecting climate change and also how climate change is affecting aviation. We then look at possible solutions to this problem and show that re-routing is the most immediate way to reduce emissions. There is a brief overview of the origins of aircraft routing, before a thorough review of more recent research in this area is presented. Lastly we ask if any of the ideas discussed could actually be applied

operationally.

Chapter 3 looks in more depth at a variety of optimal control formulations and approaches to solving these. A heuristic proof of PMP is presented as well as a thorough explanation of both non-linear and dynamic programming. The fuel burn rate function used for all fuel minimal optimisation is discussed, along with methods for calculating the take-off mass of an aircraft. A review of the data source for wind and temperature information is presented and an analysis of atmospheric conditions at different altitudes is undertaken, to inform later decisions on model formulation. Finally the OTS, along which transatlantic flights currently fly, is introduced.

The next four chapters all contain original results. Chapter 4 is based on a published paper (Wells et al., 2021), in which PMP is used to find a solution to the time minimal trajectory problem. In this formulation both altitude and airspeed are held constant across each trajectory and heading angle is varied. The initial heading angle, which allows a time minimal trajectory to arrive at the destination airport, is found using a bisection method. The problem is solved approximately, by discretising the dynamical system with respect to time and applying the Euler forward step numerical method. In order to compare the efficiency of results from the simulation and flights along the actual OTS, air distance is used. This metric is the distance flown by an aircraft relative to the air around it. Given that both airspeed and altitude are fixed, fuel burn and thus emissions are directly proportional to this value. Results for the optimised for wind routes (OFW) are calculated for each day between 1st December, 2019 and 29th February, 2020 and similarly air distance estimates along each track are found for each of the daily Air Traffic Management (ATM) tracks in each direction (ATM tracks). It is shown that fuel use could be reduced by up to 16.4% on a single flight. However, it is also conceded that the majority of aircraft will not fly on the least efficient tracks, so data from NATS is used, giving the number of flights along each track each day. In this way a weighted average can be found that allows us to conclude that on average 2.5% of fuel can be saved flying west and 1.7% flying east, amounting to 6.7 million kg of CO_2 emissions being prevented across the whole winter season between these two airports.

In Chapter 5 similar methods are applied in order to compare the flight duration for minimum time trajectories between the same airport pair in a historic time period (1986-

2005) and a future time period (2036-2055). Both summer and winter of each year are considered, but this time the atmospheric conditions are taken from an ensemble of eight different climate models. The trajectory durations based on flying through the daily wind field for each member of this ensemble are averaged to give a final result for each day. Given the number of results to be found, there are certain individual days for which a unique extremal route between the airports does not exist and thus a new method is found to locate the globally time minimal trajectory. Results show that there is very little difference between flight durations for these time periods. These findings are compared with previous research and reasons for the differences between conclusions discussed (Irvine et al., 2016; Williams, 2016).

Fixed-time, fuel-optimal flights are considered in Chapter 6, which is based on a published paper (Wells et al., 2022). Here it is established that time minimal flights are not always convenient for airports and airlines, but that by fixing the flight duration for a whole winter season, fuel could be saved that would otherwise be used in airport stacking patterns. A dynamical system based on change in longitude, latitude and mass forms the basis of the OCP with the cost functional now given as the integral of the fuel burn function introduced in Chapter 3. Two formulations are considered, one in which both the heading angle and the airspeed are controlled and the other in which airspeed is kept constant and heading angle is the only control variable. Once again the a re-analysis dataset is used to provide daily atmospheric data (Kalnay et al., 1996) and the altitude remains constant. The dynamical system and the cost functional are both discretised with respect to time. Then the new discrete approximation to the problem is optimised via Matlab's fmincon function, coded within a global search function. In this way many sets of initial conditions are tried in searching for the minimum fuel trajectory through each daily wind field, in each direction. When results from each of the two formulations are compared, it can be seen that an extra 723 000 kg of CO_2 emissions could be prevented by controlling airspeed in addition to heading angle. A link between average trajectory airspeed and average tailwind along the Great Circle Route (GCR), the shortest path between two points, is also established.

The final results chapter, Chapter 7 addresses the problem of finding both time minimal and free-time, fuel minimal routes for the cruise phase across the North Atlantic, again between LHR and JFK for the same winter period and the same daily atmospheric

conditions. The state space is divided into a discrete grid of longitude, latitude and mass, across which a value function is found by solving the HJB equation numerically, using a Semi-Lagrangian scheme. Thus this method of dynamic programming allows us to retrieve a set of optimal controls for each formulation of the problem, from which time and fuel optimal routes can be found. The same atmospheric data is then applied in estimating the fuel use of actual aircraft routes between this airport pair on each day of the winter season, adjusted to remain at the same altitude. Comparing the data for time minimal, fuel minimal and actual flights, we show that emissions could be reduced by 4.6% flying east and 3.9% flying west, using the fuel minimal formulation and that, although the duration of such flights tends to be longer than for the other two methods, it is not significantly longer, at a 5% level, than routes currently flown. This is achieved by a mixture of lower airspeeds and more efficient routing, accounting for atmospheric conditions.

Finally, the main conclusions are presented in Chapter 8. The research questions discussed in Section 1.1 are answered and we look at how these responses could affect operational transatlantic route planning. Lastly we consider some of the limitations of this research and suggest ways in which it could be enhanced in the future, such as looking at different stretches of airspace, working out how Trajectory Based Operations (TBO) could schedule multiple flights, considering other atmospheric phenomena, such as turbulence and using probabilistic atmospheric data to account for uncertainties in wind fields.

Chapter 2

Motivation and previous research

2.1 Motivation

In this thesis optimal control theory methods will be applied to find new transatlantic flight routes for commercial aircraft in order to reduce emissions. Before previous research is examined and new mathematical models are established, it is important to explain why this research is both essential and timely.

2.1.1 Why are aviation emissions a priority?

Aviation is currently responsible for approximately 2.4% of all anthropogenic sources of CO_2, but this figure is growing (Graver et al., 2019; Grewe et al., 2019; Lee et al., 2020). According to the International Civil Aviation Organisation (ICAO), if action is not taken, aircraft emissions are likely to continue to increase through 2050 (ICAO, 2019). Growth in flight numbers of 80% was recorded between 1990 and 2014 (Hasan et al., 2021) and although the recent pandemic has stalled increases in air travel in the last two years, forecasts post recovery seem to imply a continued growth. The ICAO's 2018 estimate of a compound annual growth rate of 4.3% from 2015 to 2035 still seems likely (Lin et al., 2018; Quadros et al., 2020). Other estimates put the growth in air traffic by 2050 as somewhere between a 47 and a 65% increase on 2019 pre-pandemic levels.

The ICAO has already established a policy of improving the fuel efficiency of international flights by 2% annually, through advances to aircraft technology and ATM and operations (ICAO, 2016, 2019). Failure to comply with this target will not, however, attract any kind of penalty for airlines (Hasan et al., 2021).

During the UN Climate Change Conference, in November, 2021, the international aviation community agreed to work together to accelerate their actions to reduce CO_2 emissions (UKCOP26, 2021). The agreement was for the ICAO to implement short, medium and long term goals in order to keep global average temperatures from exceeding 1.5° C. States are encouraged to produce action plans that will be reviewed by the ICAO, to include deployment of sustainable aviation fuel (SAF) and improvements to the Carbon Offsetting and Reduction Scheme for International Aviation (CORSIA).

SAF, whilst providing an up to 80% reduction in CO_2 emissions, is, however, still too costly and too scarce presently to be economically viable in an industry with historically low profit margins (Uppink, 2021) and a fleet tuned to use kerosene (Hasan et al., 2021). The International Energy Agency through their Sustainable Development Scenario are currently targeting a 10% take-up of biofuel by 2030, rising to 20% within ten years (Hasan et al., 2021; IEA, 2022). When SAF is more readily available, the costly nature of the fuel, which has currently limited it to only 0.05% of the EU market, will mean that airlines will still be interested in the most fuel efficient routes (Uppink, 2021).

In 2016, 192 nations agreed to CORSIA pledging to use offset schemes to maintain net emissions at the 2020 level (Timperley, 2019). CORSIA only provides short term alleviation, as there are difficulties in ensuring that genuine net emissions reduction takes place when investing in a range of carbon reduction projects. CORSIA does have the potential to reduce emissions by 18% by 2039, but in 2021 this figure was just 1.4% and so growth in demand may well negate any emissions reductions made by the scheme, given an annual increase in demand for flights in the Middle East and the Asia Pacific region of 13% and 8.8% respectively (Hasan et al., 2021). The ICAO has an offset target of 2.5 billion tonnes of CO_2 from 2021 to 2035, but as this is a voluntary scheme from 2021 to 2026, benefits may be limited if nations become unable or unwilling to comply (Hasan et al., 2021).

Various alternative solutions to offsetting and the greater use of synthetic bio-fuels include replacing the entire air transport fleet of approximately 31 000 aircraft with updated models (CAPA, 2018; Monbiot, 2007). The current fleet have an average age of 11.3 years, so given that commercial aircraft have an operational lifetime of 25 to 30 years, this will take time (Hasan et al., 2021). Moreover, those technologies that have the potential to

produce significant reductions in fuel use are high-risk, high-cost and have implementation timescales measured in decades (Jensen et al., 2015). By contrast, improvements to current operational procedures, such as routing flights more efficiently, have the potential to provide immediate, low-cost, low-risk and significant reductions, with the largest reductions possible on long-haul flights (Girardet et al., 2014; Németh et al., 2018).

2.1.2 What effect do the emissions have?

Fuel saving through more efficient operations would be a benefit to both the airlines through reduced fuel expenditure and to the environment through reduced emissions. It is difficult to trace all of the problems caused by aviation cruise emissions, as this depends heavily on the state of the winds and other weather effects at the boundary of the upper troposphere and lower stratosphere. However, a number of different effects are worth noting.

Firstly extra CO_2 in the atmosphere accelerates global warming. As a greenhouse gas it absorbs heat and then gradually releases it, which allows the Earth's average temperature to be 16° C rather than below freezing. In the post-industrial age, however, there is an unnaturally high concentration of carbon dioxide in the atmosphere, which means too much heat and so the Earth's average temperature is rising (Buis, 2019). CO_2 has a long term effect as it stays in the atmosphere for between 300 and 1000 years (Lindsey, 2021). The extra heat is projected to melt Arctic sea ice, altering the temperature gradient between the pole and the equator in the northern hemisphere. Such changes will alter weather globally (Eurocontrol, 2021b). The role of CO_2 is also important as it dissolves into water, which has resulted in ocean acidification, posing a threat to marine wildlife (Godbold & Calosi, 2013).

CO_2 is not the only aircraft exhaust gas which contributes to climate change, with three per cent of anthropogenic nitrogen oxides (NO_x) being due to aviation (Schumann, 1997). It is possible to reduce the amount of NO_x emitted by aircraft by reducing the temperature at which fuel is burned, but this leads to an increase in the amount of fuel used and thus an increase in CO_2 emissions. Quantifying the radiative forcing effect of NO_x emissions is far more complicated than for those associated with CO_2, as it lasts for less than a year in the atmosphere and has both a direct warming effect and an indirect cooling effect (Grobler et al., 2019). This happens because the aviation NO_x leads to higher levels of

ozone (O_3) production. Whilst this increases radiative forcing in the short term, with an effect which decays in a few months, some of the extra ozone reacts with methane in the atmosphere, thus reducing the radiative forcing effect of the methane, which has a lifetime of 10 years (Khodayari et al., 2014). It is thought that the net effect of the aviation NO_x is an increase in radiative forcing, but is difficult to estimate just how large this is due to the different timescales and also the fact that the increase depends in a non-linear way on the background NO_x levels present at the point where emissions occur. Khodayari et al. (2014) found a 10-36% difference in estimates of the radiative forcing effect of nitrogen oxides depending on whether both short and long term effects were included. So although nitrogen oxides have an important impact in terms of air quality near airports (Grobler et al., 2019), their net climate effect is less than that of CO_2, which has a lifetime of hundreds of years (Archer et al., 2009).

Contrails (short for condensation trails, line shaped clouds of ice particles left by aircraft), contrail cirrus (contrail clouds that last for at least ten minutes), water vapour and NO_x can all cause warming (Kärcher, 2018). The non-CO_2 impact is said to be double that of CO_2, so why concentrate on reducing fuel use, with its emphasis on CO_2 reduction? Recent work has focused on reducing non-CO_2 emissions via flight routing (Matthes et al., 2020, 2021; Yamashita et al., 2020, 2021), but this has been shown to increase flight costs and times (Grewe et al., 2017; Rosenow et al., 2018; Grewe et al., 2019) and require hourly, finer resolution forecasts (Rosenow et al., 2018; Jackson et al., 2001). Whilst the effect of CO_2 emissions on radiative forcing is known with a less than 1% degree of uncertainty (Lee et al., 2020), the effect of non-CO_2 emissions is less certain (Teoh et al., 2020; Boucher et al., 2013). Teoh et al. (2020) showed that only 2.2% of the flights they considered produced 80% of the contrails observed, so planning contrail avoidance routes for all flights would not be efficient (Teoh et al., 2020). In concluding their research, Grewe et al. (2017) found that the science behind climate routing was not yet robust enough for airlines to implement such routes. An additional problem is the lack of consensus on how to exploit the effects of contrails, as in certain situations these can lead to cooling.

Where flights do intercept supersaturated ice regions that could lead to the formation of contrails, these are often significantly wider than they are deep, so Avila et al. (2019) found that an increase in altitude by 2000 to 4000 feet could ensure a reduction in contrail formation with very little cost in terms of fuel. By optimising just the horizontal trajec-

tory to minimise fuel use, the altitude is left for small tactical changes that may happen en-route. For these reasons, here only CO_2 emissions will be considered from fixed altitude trajectories.

Cruise altitude emissions also have the potential to impact surface air quality. As early as 2010 it was estimated that such emissions were the cause of 8 000 premature deaths a year, particularly through an increase in solid particles less than 2.5 μm in diameter (Barrett et al., 2010). Ground level impacts from cruise level emissions have been shown to be higher than emissions from take-off and landing and to have a higher impact in the winter, due to downward transport from the upper altitudes (Jacobson et al., 2013; Lee et al., 2013; Xing et al., 2016). However, it is hard to quantify surface air quality effects accurately. More recently Vennam et al. (2017) have even shown that previous estimates may have exaggerated some effects due to the coarseness of grid resolution not allowing processes to have been observed in enough detail. By contrast, it has been estimated that civil aviation emissions now contribute 16 000 premature deaths annually (Grobler et al., 2019; Quadros et al., 2020). Such a large effect would normally result in action to reduce this death rate, but regulation is rendered far less straightforward when the cause and the effect are often distant. For example, 88% of the health impacts from aviation emissions over North America actually occur outside of this region (Quadros et al., 2020).

So it seems that re-routing to reduce emissions is the most efficient way at present to counter both environmental and more immediate health impacts. In addition, long-haul flight sustainability, at least until better measurement of the other warming effects is possible, is best tackled by considering just CO_2 effects.

2.1.3 Why are we considering the North Atlantic?

The demand for a quick and safe way to cross the vastness of the Atlantic Ocean dates back to the early twentieth century, well before land planes were capable of taking off with the sheer weight of fuel needed to complete the trip. Early commercial passengers from 1919 to the end of the 1930s relied on airships and flying boats to reduce the five day sea voyage to a 29 hour luxury flight (Airlines, 2019). For the seaplanes this trip required such a weight of fuel to be carried, that they could not actually take off and had to be carried by larger aircraft to the correct altitude and released (Cummins, 2019). However, using technology developed during World War II, such as larger piston engines and longer run-

ways, by the end of the 1950s the transatlantic route was number one in the world in terms of both traffic and revenue (Airlines, 2019; Cummins, 2019). Flights were scheduled to take 17 hours 40 minutes flying west and 15 hours 15 minutes flying east (Cummins, 2019).

Eight years later jet aircraft were introduced, leading to more routes, decreasing flight times and greater flexibility. The next major innovation was the Boeing 747, which, with its greatly increased capacity, made air travel available to all, despite being a huge developmental risk at the time when Boeing designed it in just 29 months (Oakley, 2019). With the exception of supersonic flight, there have been no major technological innovations since this time (Cummins, 2019).

Today transatlantic air routes are the most travelled in the world with about 300 flights in each direction every day (Irvine et al., 2013). During the winter period of 2019 alone, over 3.8 million passengers were flown between New York and London and from April, 2017 to March, 2018 this route was the first to produce more than a billion dollars in revenue for a single airline over a twelve month period (OAG, 2020; Reed, 2018). Globally 92.5% of aviation fuel is burned in the Northern Hemisphere and 74.6% of all fuel burned happens at cruise altitudes (Vennam et al., 2017). By choosing the North Atlantic crossing, we are looking at a region of long-haul flights which spend 92% of the ground distance in cruise (see Chapter 4) and this is where a move to trajectory based operations (TBO) will have most effect (Girardet et al., 2014).

Commercial flying is the most emissions intensive form of transport for short-haul travel, with flights 3 to 1300 times more harmful to the environment than trains across Europe (Schennings et al., 2019). This implies that there is a case to be made for taxing short-haul flights to make them less attractive to consumers, particularly when the time needed for the journeys by air or rail are comparable. However, crossing the North Atlantic is not an easy journey without an aeroplane. According to carbon offsetting companies, a seven night crossing on the Queen Mary 2 ocean liner will produce 0.43 kg per passenger mile of CO_2, compared with their estimate of 0.257 kg per passenger mile for a long-haul flight (CitytoSea, 2022). Not only is the journey longer and less environmentally friendly (both in terms of carbon dioxide emissions and waste produced), but it is also over three times more expensive. Although Ren & Leslie (2019) estimate that four billion dollars of fuel are used annually in flights through the North Atlantic Corridor, their research also

concludes that aircraft travel remains the most energy efficient option for long intercontinental distances.

It would seem, therefore, that a transatlantic crossing is best completed by air, but do we really need to cross the Atlantic? Has the recent Covid-19 enforced use of virtual meetings and "staycations" proved that the cessation of some transatlantic travel is a possibility? According to an article in the Evening Standard on 21st September, 2020, the collapse in air travel between London and New York has led to millions being lost in business deals and tourism (English et al., 2020). At a more intrinsic level, flight has also played an important role in hastening the development of countries and connecting cultures (Singh & Sharma, 2015). Commercial flying has made an: "undeniable contribution to global wealth creation and improved quality of life" (Poll, 2011).

Aside from the $691.3 billion this sector contributed to global GDP in 2019 alone (Hasan et al., 2021), there are also more subtle reasons why commercial aviation is important. What is overlooked by those calling for fewer flights, is our dependence on weather data provided by this sector. In 2008, over 230 000 observations per day came directly from the fleet of commercial aircraft, resulting in a "significantly positive impact on numerical weather prediction and operational forecasting" according to the World Meteorological Organisation (WMO) (Grooters, 2008). This data is used by many scientists and meteorologists worldwide, including those at the National Centers for Environmental Prediction in the United States, due to its excellent quality. For example only 0.8% of recorded wind data was rejected after data cleaning by the US Navy (Moninger et al., 2003). The importance of this method of data collection was further underlined during the COVID-19 pandemic. With far fewer flights being made between March and May, 2020, available data from passenger and freight flights fell by 75% causing forecasts to suffer (Freedman, 2020; Chen, 2020). In tests run in which 80% of aircraft collected data was removed from the global observation system, the reduction in forecast skill was between 30 and 60%, confirming that these data are essential (James et al., 2020).

2.1.4 Why is this research timely?

Re-routing transatlantic flights to save fuel and thus reduce CO_2 emissions sounds like a win-win scenario, with airlines saving money and the environment also benefitting. In which case, why has this not been implemented previously? Historically flight routes

across the North Atlantic have been constrained by the large volume of air traffic and the absence of radar coverage in mid-ocean (Dhief, 2018). However, a new network of 66 low Earth orbit satellites currently being tested is set to revolutionise communications across the Atlantic, with aircraft able to send data continuously during flight. This improved situational awareness, driven by data analytics, machine learning and artificial intelligence allows gate-to-gate tracking of flights, making it possible to consider more flexibility within ATM routing (Aireon, 2020; Poret et al., 2015; NATS, 2019).

Since separation between flights was first introduced in 1955 by the ICAO, it is noticeable that each time an advance in aircraft tracking or communications has occurred, reductions in aircraft spacing have followed. When the twice daily OTS first came into use in the mid 1960s, the time gap between planes on the same track was reduced from 30 minutes (a gap of about 432 km at today's airspeeds) to 10 minutes. When it became illegal to fly without a navigation system, in 1977, the previous lateral safety margin of 120 miles was halved. In 2015 the Future Air Navigation System became mandatory using ADS (Automatic Dependent Surveillance) and CPDLC (Controller Pilot Data Link Communications) to send automatic position updates at agreed times or when an aircraft deviates from its agreed path (NATS, 2007). This surveillance technology, however, requires the aircraft to cooperate in some way with the ground systems, meaning dependence on primary radar.

It is only with the advent of the Aireon satellite system that we can talk not just about bringing tracks closer together, but redesigning the system completely to allow for maximum access to a fuel-optimised route. In fact from 1st March, 2022, the OTS tracks were removed up to a flight level of 33 000 feet and the separation of aircraft changed from the previous 40 to 14 nautical miles. This was as a direct response to the communications upgrade, via the Aireon satellites, that now allows aircraft positions to be updated in less than 8 seconds (Ahlgren, 2022). It should be noted that the paper on which Chapter 4 of this thesis is based, was published before this change to the OTS and was in fact cited in the decision to experiment with such changes (Young, 2021).

My research aims to justify the need for the implementation of fuel and time optimised routes and prove that these will not only reduce fuel use and CO_2 emissions across the Atlantic, but that it is also possible to schedule flights based on each day's weather conditions to ensure that landing slots are not missed. Early flights waiting in holding patterns

for a gate to become available are still burning fuel, with engine efficiency greatly reduced at these lower altitudes. For overnight transatlantic flights there is also the operational issue of arriving in Europe too early to adhere to rules on night time noise curfews (Eurocontrol, 2021b). Delayed flights have a knock-on effect for the rest of the day, causing extra spending on crew, fuel and maintenance. In 2007 costs due to delays amounted to a total of $8.3 billion across the airline industry (Peterson et al., 2013). Improvements to routing, allowing weather to be taken more into consideration in scheduling, could also allow airlines to reduce "padding" of flight times on their timetables, thus enabling them to reduce crew costs. This broad overview justifies my research, but it is vital to explore the large body of material supporting this approach further, in order to focus on all aspects of the problem.

2.2 Emissions and climate change

If the global economy fails to decarbonise sufficiently rapidly, there may be significant consequences for aviation from the ensuing climate change. These consequences include increased turbulence as the jet stream becomes more sheared (Williams & Joshi, 2013; Williams, 2017; Storer et al., 2017; Lee et al., 2019; Kim et al., 2015), modified flight routes and journey times as the prevailing high-altitude winds shift and strengthen (Karnauskas et al., 2015; Irvine et al., 2016; Williams, 2016; Kim et al., 2020) and take-off weight restrictions as warmer air reduces lift and thrust on the runway (Coffel & Horton, 2015; Gratton et al., 2020). Therefore, aviation is not only a contributor to climate change, but may also suffer from its adverse effects increasingly in future.

This is particularly apparent over the North Atlantic due to the prevailing wind pattern, the "jet stream". This is an eddy-driven jet comprising a core of strong winds 8 to 11 km above the surface of the Earth and generally flowing from west to east (McSweeney & Bett, 2020). The jet stream currently appears to have three preferred positions, 35-38° N, 45-47° N and 58-60° N, but there is no annual pattern for how many days it will spend in each position (McSweeney & Bett, 2020).

2.2.1 Climate change impacts affecting flights across the North Atlantic

As the environment through which aircraft fly evolves, so does the need for flight planners to change their approach. Just how certain can we be about the way that the wind fields across the North Atlantic at cruise altitude will be affected by global warming? To predict future wind patterns, we are reliant on climate models and these are evaluated by the Climate Model Intercomparison Project (CMIP) to ensure that predictions are realistic (NCAR, 2022). As research progresses the models being produced need to be checked by making hindcasts for historical periods and comparing these with data from a reanalysis dataset, such as ERA5 to evaluate any systemic biases (Hersbach et al., 2020).

Looking at the North Atlantic region, both the storm track and the jet stream are scrutinised in terms of their spatial patterns and magnitudes. Harvey et al. (2020) compared three generations of weather models (CMIP3, CMIP5 and CMIP6) and found that storm track biases over the North Atlantic were reduced by half in the most recent climate models. Jet stream biases were also reduced in the winter months. Between CMIP3 and CMIP5 there were large changes, but these were less marked between the most recent two generations of models, a pattern also noted by Oudar et al. (2020). This means that the later models give a much better representation of storm tracks and the jet stream over the North Atlantic. This is of interest, as research into aviation routing using CMIP3 data has produced different results to that relying on models from CMIP5 (Eurocontrol, 2021b).

Analysing how climate change is likely to affect this region, Harvey et al. (2020) concluded that there would be a poleward shift in the jet stream and that high latitude storminess would be reduced. The MET Office report into "Jet stream position and strength" agreed with these findings predicting a jet stream shift further North in the autumn and summer (McSweeney & Bett, 2020).

Oudar et al. (2020) considered three main reasons why the jet stream might change its position in future. They looked at Arctic amplification (a larger than average increase in net radiation near the poles due to global warming), tropical amplification (stronger warming in the upper than the lower troposphere) and polar vortex strengthening (circulation of high winds in the upper stratosphere). The first of these phenomena had a limited

effect on the jet stream, although there was a slight shift equator-wards. By contrast, both tropical amplification and polar vortex strengthening created a much more marked poleward shift. Research by Ruosteenoja et al. (2019) agrees with these findings, but also suggests that storm activity over the North Atlantic has been extremely variable decade by decade and so shows no clear longterm trend. According to Manzini et al. (2018) there is an added source of uncertainty in the climate model data, due to the non-linear way the jet stream changes under different levels of global warming. For example a higher level of warming may result in a totally different response from a lower level, rather than just a similar change with amplified magnitude.

In Eurocontrol (2021b), which uses an ensemble of eight CMIP5 climate models, the latitude of the jet stream core is shown to vary more each day in the future period considered (2046-2055) than in the historic period (1991-2000). Looking at changes in the jet stream between these two decades there are different patterns in the winter and the summer. The polar jet is shown to be stronger at the core, but narrower, with weaker higher latitude zonal winds in winter. In summer, however, the core spreads and the magnitude of the jet stream winds decreases. Clearly a stronger jet stream with a more northerly position will affect future flight operations over the North Atlantic.

In Storer et al. (2017) the effect that doubling the CO_2 in the atmosphere would have on the clear air turbulence (CAT) encountered by commercial aviation was considered. Results showed that the North Atlantic region would be particularly prone to increases in CAT during the 2050 to 2080 period, with severe incidents at a pressure level of 200 hPa occurring almost three times as often as in the control period, during which atmospheric CO_2 was set to pre-industrial levels. This could affect choice of flight paths in the future. As CAT is not visible to pilots or radar, it is instead forecast depending on a range of different diagnostics, with only a moderate level of skill. This may result in flight planning moving towards lower altitudes to limit the likelihood of severe incidents. At FL340 (34 000 feet), the increase in severe CAT forecast, as an annual average, was 88% compared to 181% just 3000 feet higher (Storer et al., 2017).

Kim et al. (2020) found that winters with a positive North Atlantic Oscillation (+NAO), where differences in sea level pressures across the Atlantic region (between the Subpolar Low and the Subtropical High) lead to a stronger, more northerly jet stream, resulted in

significantly longer round trip flight times between New York and London. These could be from 4.24 to 9.35 minutes longer compared to -NAO seasons. This may not seem like a great difference, but when the 300 flights in each direction each day are taken into account this results in between 66 and 144 million kg of extra CO_2 being pumped into the atmosphere. This research looked at flights at 200, 250 and 300 hPa and calculated wind optimised trajectories as minimum time routes. No distinction was made between results at each altitude and airspeed was fixed, so optimisation relied entirely on the horizontal heading angle. As +NAO years can be forecast 6-9 months in advance, airlines will be able to adjust their seasonal schedules appropriately (Kim et al., 2012; Scaife et al., 2014; Dunstone et al., 2016). Set timings for flight schedules would need to be longer in a strong +NAO year and could be shorter in -NAO years.

It is perhaps counter-intuitive that in periods of stronger winds flights will be delayed in a round trip. This phenomenon can be explained fully as the slower ground speed caused by adverse winds means a larger time discrepancy than can be made up for by a faster ground speed in the opposite direction. This is true as the slower speed is maintained for a longer time period. In Williams (2016) the difference between eastbound time minimal routes between JFK and LHR was found to be 4 minutes shorter when the CMIP3 GFDL CM2 model was used to look at 20 years of doubled CO_2 levels compared with 20 years of pre-industrial CO_2 levels, but westbound routes increased by just over 5 minutes. This work looked at keeping the airspeed fixed to 250 $\mathrm{m\,s^{-1}}$ and the 200 hPa pressure level. In similar research Irvine et al. (2016) found that using an ensemble of five different CMIP5 models (based on the RCP8.5 scenario, which forecasts four degrees of warming by 2100) gave less marked changes when the period from 1979-2005 was compared with the 2073-2099 period. Here eastbound time minimal flights were just under a minute faster and those flying west were about 1 minute slower. Flights were kept to the 250 hPa pressure level and to an airspeed of 250 $\mathrm{m\,s^{-1}}$. Both of these papers found minimal time routes by applying a shooting method to obtain an initial heading angle and then solving a dynamical system, based on the change in latitude, longitude and heading angle of the aircraft, using a numerical method. This approach will be explained in more detail in Chapter 3 and the research considered more comprehensively in Chapter 5.

By contrast, in the 2021 Eurocontrol report, the Met Office's TP algorithm, based on the A star method, was applied to give time minimal routes for the same transatlantic

journey (Eurocontrol, 2021b). This method uses discrete mathematics based on the Dijkstra algorithm to find the time minimal route with reference to a predefined set of grid points (Cheung, 2018). The wind and temperature data used was from an ensemble of 8 CMIP 5 models run under the RCP8.5 emissions scenario, comparing the historic period (1991-2000) with a much closer future period of (2046-2055). In this research the Mach number (ratio of airspeed to speed of sound at the same physical position) rather than the airspeed was held constant, at 0.82 Mach. The pressure level was again restricted to 250 hPa. Results from this research showed that both eastbound and westbound flights were actually faster in the future period. This disagreement with the previous research could be due to the method used for the retrieval of time minimal paths, the differences in models used or the choice of the future time period. This will be examined more closely in Chapter 5.

2.3 What can be done?

If we accept that the aviation industry is producing unacceptable levels of pollution, that are fast accelerating our planet towards a climate catastrophe, then our next logical question must be "What can be done about this?"

2.3.1 Are airlines already making pro-environmental behaviour a priority?

Apparently just 25% of the codes of conduct of the top sixteen airlines in the world (judged on size of fleet and number of scheduled flights) refer to the need to indulge in "climate friendly behaviour", with just three of these actually mentioning the importance of cutting emissions (Ruban, 2020). If consumers start to select tickets based on environmental choices and not just price, then corporate branding will need to embrace an understanding of the need to reduce emissions in the future. However, Schennings et al. (2019) found that only 1.6% of the consumers polled in Sweden chose tickets based on the environmental impact of flights, whilst 81.8% made decisions based on prices. When an algorithm to rank available flights based on emissions was added to Sweden's most popular air travel booking site, only 1.54% of the customers even tried it (Schennings et al., 2019).

Perhaps we need the airlines to act unilaterally, ensuring that all of their flights are more

environmentally friendly. There are several ways this can be done, with the ICAO backing SAFs, operational improvements and aircraft technology (ICAO, 2019).

2.3.2 Changes to aircraft technology and passenger provision

There are currently a number of projects in progress to develop new technology that reduces emissions. According to Economon et al. (2011) improvements in fuel burn performance will need to come from advances to the aerodynamic shaping of an aircraft, its structure or the technology behind its propulsion system. However, working out a priori how fuel burn during operation will be affected by these changes is still very uncertain.

A popular innovation is to use liquid hydrogen as a fuel, thus eradicating all CO_2 emissions. This is not a straightforward fuel swap, as hydrogen liquifies at -252.9°C, so a cryogenic system is needed to keep it well insulated (Grönstedt, 2021). The low density of cryogenic hydrogen necessitates large fuel tanks that could no longer be integrated into the internal wing volume. Rompokos et al. (2020) conclude that for a long-haul flight a blended-wing-body design is most suitable.

Airbus have decided on two research concepts for long-haul aircraft, which it aims to have ready for use by 2035. The first is a turbofan aircraft running on liquid hydrogen. The second is a more revolutionary blended-wing-body design, in which the wide fuselage provides more options for hydrogen storage and passenger seating. Both would have at least a 2000 mile range and carry 200 passengers (Airbus, 2020). Although the aircraft

(a) Turbofan aircraft design

(b) Blended-wing-body design

Figure 2.1: Innovative Airbus designs for hydrogen powered aircraft (Airbus, 2020)

may be available by 2035, many changes will be needed in airport infrastructure to supply the required hydrogen and according to Grönstedt (2021) long-haul flights will not be routinely powered by hydrogen before 2045. Even then there could be issues with producing enough hydrogen. If we were to move to all hydrogen powered long-haul flights, by 2090 making the necessary gas sustainably would use over 30% of all renewable electric power generated (Sethi et al., 2022).

To avoid redesigning the aircraft fuselage, novel engine and propulsive systems have been suggested. One such idea is a water enhanced turbofan engine, in which water is injected into the engine to increase the heat capacity, allowing better heat extraction. This reduces fuel burn and thus carbon dioxide emissions. The extra water is condensed by cold environmental air and reused, so that there is no water in the exhaust, also reducing contrail formation. At the moment this design is at the proof of concept stage as its use may affect the weight and balance of an aircraft (Pouzolz et al., 2021).

Another suggestion is to improve propulsive efficiency by ingesting wake flow and use this extra momentum to increase thrust. In the test stage, this technology has been shown to reduce fuel usage by 3 to 4 % (Seitz, 2012). However, this work has been ongoing for the last ten years and still requires further research before it can be used operationally. Testing of similar boundary layer ingestion fans show that this technology has a detrimental effect on the aerodynamics of an aircraft (Mårtensson, 2021).

Fully electric aircraft are set to be a reality for short-haul commercial flight as early as 2026, with Heart Aerospace developing the ES-19, a 19 passenger airliner with a range of 400 km (HeartAerospace, 2022). The sheer weight of the battery power needed for long-haul flights, due to the inferior energy density of batteries compared to kerosene, makes using just electrical power unfeasible (Doctor et al., 2022). However, another innovative idea would be to print circuits into the fuselage structure of the aircraft, thus allowing the removal of weighty bundles of electric cables and even to replace the cabin floor with a structural battery using carbon reinforced plastic (Linde, 2020). At present, though, these technologies are still a long way from becoming an operational reality.

These upgrades to technology would all be costly to develop and implement, so it can be concluded that operational changes, such as route and airspeed optimisation, are more

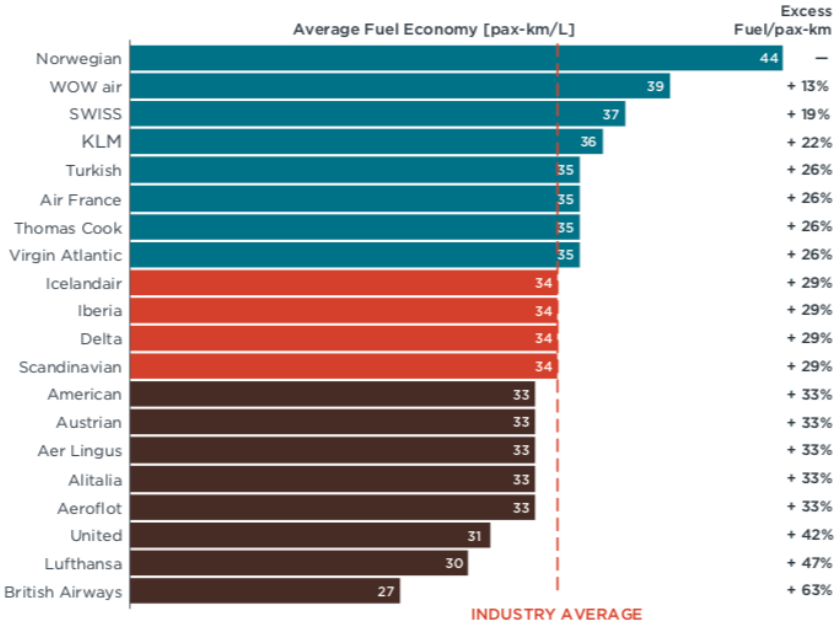

Figure 2.2: Plot from Graver & Rutherford (2018) showing the most and least fuel efficient transatlantic carriers in 2017.

practical in the short term (Economon et al., 2011).

Changing how on-board space is used in the current fleet could be a more immediate way to reduce overall emissions. By allowing less space per passenger, fewer flights could be made. However, this is unlikely to be popular with heritage airlines as although premium passengers only make up 5.2% of air traffic, they generate 30.4% of passenger revenue (Graver & Rutherford, 2018). Increasing flight occupancy reduces emissions per passenger without any need for design changes. Budget airlines currently operate the most fuel efficient services with 84 to 93% occupancy, whilst other groups fly with as few as 76% of the seats filled (Schennings et al., 2019). Both of these factors contribute to the findings displayed in Figure 2.2 in which British Airways is shown to be 63% less fuel efficient than Norwegian Airlines. Although Norwegian's long-haul flights were casualties of the difficulties faced by airlines throughout the COVID-19 pandemic, Norse Atlantic Airlines hope to replace them as the most fuel efficient transatlantic carrier by summer of 2022 (Partridge, 2021).

Another idea to help to reduce emissions lies in flying formations of aircraft. In this way the lead aircraft could provide a slip stream for follower aircraft that would therefore save fuel during the trip. The centralised approach, which requires formations to be

planned in advance of flights, means that follower aircraft would benefit not just from this slip streaming, but also from being able to carry less fuel and thus being lighter, which also reduces fuel burn. Kent & Richards (2015) showed that formations of two aircraft forming at the start of the cruise phase for a transatlantic flight could save 8.7% of fuel compared with a solo flight and formations of three aircraft could lead to savings of 13.1%. However, the method used to generate these savings did not optimise for wind, but instead considered the best way to pair flights up, assuming set routes.

Verhagen et al. (2018) looked at using aircraft communications to set up formations with nearby aircraft during flight. This method would not give the added advantage of reducing the amount of fuel to be carried, but would mean that any airport delays encountered by an aircraft would not have a knock-on effect to other members of a planned formation. The research was limited by a method which assumed a set 10% fuel flow reduction for trailing aircraft and did not take into account issues like turbulent weather causing problems for in flight communications or maintaining the suggested formation.

The airbus "Fello'fly" project tested this idea of wake-retrieval, with two aircraft flying 3 km apart in November, 2021. Results showed a 5% reduction in CO_2 emissions compared with two aircraft flying separately. Whilst the technology for such flights is set to be ready by 2025, the regulatory framework may take longer to adapt to formation flights (Airbus, 2022).

2.3.3 Why is re-routing a good idea?

With airlines, ATM and the ICAO keen to increase the sustainability of commercial flight (ICAO, 2020; Molloy, 2020; AirlinesUK, 2019), it has been shown that TBO provide the key to improving overall efficiency (Wickramasinghe et al., 2012; García-Heras et al., 2014). This approach centres on designing individual flight routes accounting for weather conditions in order to optimize efficiency with regard to a number of different factors. The International Air Transport Association (IATA) states that improvements to aircraft emissions efficiency of up to 12% could be made by improving ATM procedures; a more flexible approach to how airspace is managed will allow use of more of the available capacity (Gardi et al., 2016). For long-haul flights, such as those between LHR and JFK, for which the cruise phase makes up around 92% of the ground distance of the flight, TBO promises the greatest rewards (Girardet et al., 2014).

Currently flight tracks in the North Atlantic's OTS are created on a daily basis by air navigation service providers (ANSPs), NATS in the United Kingdom for the westbound paths and NAV CANADA for those going east. These are based primarily on the need to separate aircraft safely, whilst taking some account of the winds. Airlines request their preferred tracks by submitting Preferred Route Messages (PRMs) in the hours before a flight and the ANSPs create a daily track system that reflects the airlines' wishes as closely as possible. By comparing data from flown routes with those simulated, the amount of fuel savings possible can be quantified.

There are many different methods for defining optimal routes and the efficacy of these is measured in different ways. Time minimal routes aim to take greater advantage of the prevailing eastward winds when flying east and reduce the negative impact of these same air currents when flying west, in order to record the shortest flight time possible. Air distance, the distance flown by an aircraft relative to the surrounding air, is also a good measure of the efficiency of a flight path. As fuel burn is directly proportional to air distance and as emissions, including carbon dioxide, are directly proportional to fuel burn, any saving made in air distance can be translated into a reduction in CO_2 emissions (Henderson et al., 2012; Green, 2009).

Sometimes time optimal routing is not necessarily the most practical option for airlines and airports. Flight capacity in air is far greater than that at each end of a trajectory, with runway availability causing the largest bottleneck in scheduling flights (Soomer & Franx, 2008). For operational and financial reasons airlines need to adhere to their published timetables. Flights arriving early, or late, create additional costs, through extra fuel burned in holding patterns, missed connecting flights and additional crew time. Customer dissatisfaction is also a key issue for airlines when delays occur. If flight trajectories are planned to ensure a fixed flight time, whilst minimising fuel burn, then the costs to both the airlines and the environment associated with early or late arrival can be saved. In this case the fuel burn itself would provide a means of measuring the effectiveness of the new routes. The method of calculating fuel burn in this research, will use an aerodynamic fuel-burn model (Poll & Schumann, 2021a,b), whereas in previous studies there is a standard reliance on EUROCONTROL's BADA (Wickramasinghe et al., 2012; García-Heras et al., 2014; Soler et al., 2020; Yamashita et al., 2020, 2021; Matthes et al., 2021).

2.4 Addressing the routing problem

2.4.1 Origins of aircraft routing

Initial work on the minimum time routing problem was proposed by Ernst Zermelo in his lecture of 1930, which focused on a two-dimensional planar wind field (Zermelo, 1930). By choosing an initial heading angle, the differential equations in zonal (x) and meridional (y) positions can be solved numerically for a particular time step length. Zermelo's equations form the following dynamical system:

$$\frac{dx}{dt} = V\cos\theta + u, \tag{2.1}$$

$$\frac{dy}{dt} = V\sin\theta + v, \tag{2.2}$$

where V is fixed airspeed, θ is heading angle, u is zonal wind and v is meridional wind. Once at the new position a third differential equation concerning the rate of change of the heading angle is solved to give the direction for flight. This new differential equation is derived using PMP (see Appendix A for the derivation and Section 3.3.1 for more information about PMP):

$$\frac{\mathrm{d}\theta}{\mathrm{d}t} = \sin^2\theta\frac{\partial v}{\partial x} - \cos^2\theta\frac{\partial u}{\partial y} + (\frac{\partial u}{\partial x} - \frac{\partial v}{\partial y})\sin\theta\cos\theta. \tag{2.3}$$

In this way progress is planned, timestep by timestep across the trajectory.

Levi-Cività was quick to respond to the initial lecture, pointing out some errors in Zermelo's calculations (Levi-Civita, 1931). This led to Zermelo's publication of a clarified and corrected theory paper the next year (Zermelo, 1931).

By 1951, Dixon Speas had formed a company to calculate minimum time routes over the Atlantic for airlines. The method used discrete dynamic programming across 15 to 20 flight regions. Paths across these regions were optimised using data about high altitude winds that was gathered from a network of weather balloons (Jardin & Bryson, 2012a).

These Zermelo formulae are extended to a spherical model in Arrow (1949), which concludes that for flights of over 1000 miles, it is worth deviating from the shortest distance route. In the original model, Arrow uses a single heading angle for the whole trajectory, but in the second part of the research a time minimal path with fixed airspeed and altitude is found with a varying heading angle.

After this, the spherical Zermelo system can be found as the basis of many academic works on optimal route planning, such as Warntz (1961), which provides a comprehensive guide to routing through a wind field. In this study Warntz discusses the far greater convenience of using the wind to ensure the minimum flight time is obtained rather than following the shortest distance path, but bemoans the fact that ATM considerations prevent such paths from being followed. It is also mooted that having a centralised routing body for all flights would avoid any issues arising from separate airspace restrictions. This now seems like an early foreshadowing of the European Union's Single European Sky ATM Research (SESAR) (SESAR, 2022).

2.4.2 An overview of current research

Singh and Sharma (2015) reviewed 277 articles published between 1973 and 2014 on every angle of fuel optimisation in aviation. These not only cover the key topics of aircraft design, alternate fuels and aviation operations, but also discuss policies to deal with the excessive fuel consumption of the aviation sector. An increasing trend in pursuing this area of research was highlighted, with 68% of the research being completed after 2000. Another conclusion from the paper was that technical improvements take a long time in development, whilst operational change is quicker to find and implement.

It is also important to remember that there is no single solution to the problem of making air travel more sustainable. Re-routing flights can be undertaken immediately, but this does not mean that finding alternative fuels and improving flight technology cannot also be undertaken to allow fuel optimisation to continue. In the same way, a flight using alternative fuel or incorporating fuel saving technology, will still need to be efficient in its flight path.

Estimation of fuel burn is difficult. Fuel use from actual flights is not available in the public domain, as airlines are not prepared to share this sensitive data. In the past fuel use has been estimated using radar tracks and flight data recorder information giving positions of aircraft at certain times. Next the BADA is applied to polynomial functions, fitted to aircraft performance by manufacturers, to estimate fuel used (Dalmau et al., 2020). It has been found that BADA can underestimate fuel use in some circumstances. In Mouillet et al. (2019) it was found that BADA 4 was not suitable for on-board applications

or economic studies as the estimated airspeed and fuel burn in cruise optimisation could contain errors of up to 5% when compared to actual flights. Both the drag and fuel flow models were found to be in need of improvement. It was, however, deemed suitable for environmental impact assessments and ATM simulations. Up until now, the vast majority of papers discussing the re-routing of aircraft to minimise fuel have made use of the BADA datasets as there has been a shortage of accurate fuel burn data. In this thesis, where fuel burn has been considered, a new physics-driven fuel burn function, which is aircraft model specific, has been used (Poll & Schumann, 2021a,b). This estimates fuel burn at any point during the cruise phase of a flight by considering atmospheric data, as well as the altitude, airspeed and mass of the given aircraft.

Whatever the performance index being minimised or the choice of control and state variables, the basic problem structure for trajectory optimisation is the same. The dynamics of the system are stated, based on a variety of variables describing aircraft performance. Sometimes atmospheric data may be included within this system. Then the cost must be defined, sometimes as a terminal penalty or as a running cost or as a combination of the two. In addition to the boundary conditions on the state and control variables, these elements make up the OCP.

The problem of optimising flight routes has been addressed by aeronautical engineers, mathematicians, computer scientists and meteorologists, for the purposes of designing more fuel efficient flight routes, more environmentally friendly flight routes, more cost efficient flight routes and sometimes the theory is even used to predict what aircraft that are in flight might do next.

Algorithms to use in real time for tactical ATM and the military, tend to look more at the aircraft dynamics and possible routes in terms of an offline data driven exercise, fed into a dynamic programming model. Strategic planning uses a variety of methods to produce optimised routes for the horizontal path, the vertical path or sometimes a combination of the two. Most papers split the problem up into the ascent and descent phases and the cruise phase between these.

Long range cruise trajectories minimising fuel for a fixed mass aircraft have previously been researched by others using control variables of thrust and flight path angle (Schultz

& Zagalsky, 1972; Speyer, 1973) or lift and thrust (Schultz, 1974; Menon, 1989), but these have not accounted for the wind field. In these papers the fuel flow is minimised for a set distance of flight. Wind has been factored into fixed range cruise calculations which compare fixed and free thrust as airspeed varies (Erzberger & Lee, 1980), but here a direct operating cost is minimised, consisting of both time and fuel burn factors.

Fixed-time trajectories minimising direct operating cost have also been generated, but unlike the model formulation for fixed-time flights in Chapter 6, the optimal control model is set up with a free final time (Sorensen & Waters, 1981; Burrows, 1983). The trajectory best fitting the imposed arrival time constraint is then selected. Changes to this arrival time are dealt with by changes to airspeed and altitude in Burrows (1983), in which the use of take-off delays to avoid early arrival and thus increased fuel burn through stacking, are also considered.

Trying to arrive at a fixed-time can be compromised by inflated airspeed during the cruise phase, so in Chakravarty (1985) ways to change the cruise and descent phase tactically, in order to avoid early arrival, are considered. The problem would involve the solution of a ten-order two point boundary value problem (TPBVP), which in real-time would not be feasible, so in this research singular perturbation theory and an energy-state approximation are applied to find sub-optimal solutions of the problem. It is concluded that for fuel efficient delay absorption it is best to reduce both altitude and airspeed in the cruise phase, but to keep the descent as steep as possible. In Chapter 6 fixed-time trajectories are obtained with airspeed as a control variable, so that a piecewise-continuous strategic adjustment of the airspeed should avoid the need for en-route alterations. However, this does depend on the reliability of the atmospheric data available before the flight and the careful and timely scheduling of arrivals at the destination airport.

In Jardin & Bryson (2012a) two different methods are compared for finding time minimal paths through wind fields. The first is the standard Zermelo problem transformed to fit a spherical system and backwards integrated from the destination for a range of initial heading angles, with the minimum time path being chosen from the admissible paths generated. The second method linearises the flight dynamics about a nominal optimal route, which here is the shortest distance path, known as the great circle route (GCR). This does not necessarily guarantee that a path found is a global minimum. It does, however, have

the benefit of greater computational efficiency and so can be applied in a more tactical way. In both cases the airspeed and altitude are kept constant.

2.4.3 Simplified wind fields

Where winds are included in research, a useful simplification lies in modelling these as simple constant headwinds or tailwinds, or creating more complicated functions based on position, which are nonetheless still differentiable.

In Franco & Rivas (2011), a dynamical system is used which describes Mach number as a function of flown distance along a constant heading route, with fixed altitude, but variable mass. The control variable is taken to be the throttle setting, which can vary between 0 (no throttle) and 1 (full throttle). PMP is applied to ensure trajectories found satisfy the necessary conditions for optimality and so are extremal, with the cost functional given as specific fuel consumption multiplied by thrust. Only constant headwinds or tailwinds are considered in this research and constant cruise airspeed is often found to be optimal. In Franco & Rivas (2013) this method is repeated in order to see the fuel cost of meeting a fixed arrival time given that the winds are not correctly modelled. Mismodelled headwinds are found to have a much greater effect than mismodelled tailwinds. This is due to the increase in drag at the high Mach numbers used to arrive at the correct time, given that the flight route throttle setting has been optimised for more favourable en-route winds. Cruising with a constant Mach number is again found to be very often close to the optimal solution, given the constant headwinds and tailwinds applied (between -15 $\mathrm{m\,s^{-1}}$ and 15 $\mathrm{m\,s^{-1}}$). Using accurate atmospheric conditions from re-analysis data, these winds are far more varied, justifying Chapter 6's investigation into the effect of varying airspeed, fixed-time, fuel minimal flights through a data driven wind field.

In Nguyen (2006) the focus is the time optimal climb route of an aircraft. Although in this study it is acknowledged that the inclusion of wind data is important, as this has a large effect, in the research itself, the zonal wind field is modelled as a differentiable function, which changes with altitude, with no meridional effects considered. As real wind fields are not necessarily differentiable, in the research in Chapters 4 and 5 we rely on numerical gradients to approximate the rate of change of the wind speed in each direction and in Chapter 6 we apply a different approach in which these gradients are not required.

2.4.4 Routes avoiding atmospheric phenomena

By using PMP, with a running cost based on both flight time and the amount of turbulence encountered, Kim et al. (2015) create shortest time routes that also avoid areas of turbulence. This is important as two thirds of all severe clear air turbulence is situated near the jet stream (Lester, 1994). However, although factoring convection and turbulence into atmospheric data for route planning is in many ways essential to practical flight planning, prediction accuracy decreases rapidly as the forecast lead time increases (Steiner et al., 2010). Ensemble forecasts used in Kim et al. (2015) over-predicted occurrences of poor conditions by a factor of 2. So in order to make the inclusion of turbulence and convection avoidance more accurate for strategic planning, improvements in long-term, high resolution forecasts are needed. In the meantime, turbulence avoidance can be dealt with tactically by a change of altitude. This means that creating minimum time and fuel burn routes optimised at a single altitude is justified, as planning specifically for altitude will prohibit any altitude changes needed to avoid turbulence. Such climbs will not overly affect the efficiency of the planned horizontal route as there is minimal change in the prevailing air currents encountered at neighbouring altitudes (Krozel et al., 2011).

Sridhar et al. (2011) solve the horizontal optimal trajectory problem for a cost functional involving fuel use, time and penalty areas, which would lead to contrails being formed. The method involves solving a TPBVP, involving aircraft position, but ignoring mass change during flight. The shooting method is applied to find the initial heading angle at a number of different altitudes, so flights can use the most efficient flight level at each stage of the trajectory. Fuel use is based on BADA models, the airspeed is fixed and the terminal time is kept free. This study concludes that avoiding regions which could lead to the formation of contrails requires a 2% increase in fuel use and additional travel time of up to 4.3%. The research is based on deterministic weather data from a single day, but if this were to be used operationally, the extra CO_2 emissions would need to be justified and so more research needs to be completed into the effect of contrails and the certainty with which we can predict their formation hours before a flight.

Soler et al (2020) try to balance the need to arrive promptly, using as little fuel as possible, with the need to avoid areas of convective risk, where storms are likely to propagate. Unlike the prevalent daily wind patterns across the North Atlantic, areas likely to become stormy cannot generally be predicted early. Instead the storm risk must be assessed in the

1 to 3 hour window before a flight. The research shows that reducing exposure to convective regions cost 100 to 200 kg of fuel per flight. Unlike previous papers this research takes into account an ensemble of 51 weather predictions and also includes a penalty for entering areas of high weather uncertainty. The key idea behind this work is to provide airlines with a range of solutions so that the uncertainty in predicting storms can be factored into trajectory choices.

One limitation of Soler et al. (2020), is that flights in the model are kept to a pressure altitude of 200 hPa, which is a flight level of over 38 000 ft. Changes to altitude could in some cases help aircraft to avoid convective regions without more costly horizontal detours. The scope of the findings is also narrowed by the application of a detailed and rigorous optimal control method to just one ensemble forecast made nine hours before the 19th December, 2016. As a principal focus of the research was the variability of the weather across the North Atlantic, discussion of multiple wind fields would have been more appropriate. However, Soler et al. (2020) do suggest that there is now a move towards individual routes rather than "more rigid airspace structures", which reinforces the need for increased research into TBO.

2.4.5 Direct methods for trajectory optimisation

Unlike indirect optimisation methods, in which a dynamical system is first optimised then discretised, direct methods can be applied that reverse this order. The pseudospectral method used in Fahroo & Ross (2002) finds trajectories by fitting Nth order weighted polynomials, comprising N+1 cardinal functions, across a trajectory split into N+1 time intervals. The cost functional is evaluated based on the approximate solutions and derivatives are approximated by polynomial derivatives. By adjusting the weighting of the functions within the polynomial, increasingly cost effective trajectory solutions can be found. In this way the whole problem is first discretised and then optimised. This method forms the basis of the DIDO software which has been used for manoeuvring the International Space Station and a range of other NASA spacecraft (Ross & Karpenko, 2012). While NASA saved an estimated one million dollars in just under three hours using DIDO (Ross, 2015), the software has yet to allow for the use of a position dependent, data-driven wind field and is therefore not suited to searching for fuel optimal real-world aircraft trajectories.

A different direct method, sequential quadratic programming (SQP), is used by Pierson &

Ong (1989) to find fuel minimal aircraft trajectories from a range of different dynamical models. Flight paths are considered in three dimensions with both fixed and free ranges. Four different pairs of control variables are used:

1. the angle of attack and the thrust setting,

2. the flight path angle and the thrust setting,

3. the airspeed and the thrust setting,

4. the angle of attack and the thrust setting including a lag factor.

First control functions are approximated by piece-wise linear or cubic spline functions. The equations of motion are then integrated numerically across the time of flight allowing the cost functional to be evaluated directly. As the OCP is now a non-linear programming problem, the SQP algorithm can be applied until convergence of the control parameters is achieved, resulting in an approximation to an optimal solution. By using this direct approach, changing the dynamic models is made easier, allowing for more straightforward comparisons between results. This showed that engine dynamics were relatively unimportant, but that heading angle was key. Although demonstrating that SQP can solve an OCP directly, this research is a theoretical exploration of solution methods, rather than considering specific flight routes and neglects atmospheric data. In Chapter 6, a similar direct method is used to find fuel minimal, fixed-time trajectories across data driven wind fields. However, in this thesis the numerical procedure used trials different initial conditions to search for a global minimum and thus makes the method more robust.

Four different techniques are applied to the solution of the OCP for a minimum fuel cruise, at constant altitude, with fixed arrival time by García-Heras et al. (2014), including pseudospectral collocation as used by Fahroo & Ross (2002) and described in the first paragraph of Section 2.4.5. The approach is based on flat segments of travel and shows that different results are gained from each method and each has its own advantages and disadvantages in terms of sensitivity to numbers of time steps and computation time. However, the wind field is not considered in this analysis and as is demonstrated in other research, routes through complex wind fields are often very different from their wind-free equivalents (Girardet et al., 2014).

2.4.6 Dynamic programming in route planning

Dynamic programming is most often associated with 4D flight trajectory optimisation, which gives paths as a set of waypoints in space, reached at certain times. A discrete dynamic programming method is used by Ahmed et al. (2021) to examine either fuel reduction or, in a second formulation, a range of different emissions, for short-haul flights. The approach works by moving between three dimensional positions as part of a grid using controls of acceleration, rates of change of flight path pitch angle and heading angle. The conclusions show possible savings of 10% compared with actual flight data. However, the method does have some simplifications which make the results suboptimal: trajectories are forced to move between gridpoints to avoid the problem caused by expanding the grid to allow for every point that could be visited, which can make a numerical solution too computationally expensive. However, the main difficulty in comparing simulations from these calculations with actual flights, is that there is no consideration of wind conditions, the state space does not take into account the curvature of the Earth and the aircraft mass is assumed to remain at its initial value throughout.

Dynamic programming is used by Wickramasinghe et al. (2012) in modelling climb, cruise and descent to optimize fuel flow by controlling flight path angle, heading angle and engine thrust. However, the aircraft paths are limited to 0.5° latitude either side of the GCR, which would be limiting for long-haul, transatlantic flights. This method does not take into account other transatlantic traffic, but it has been shown that aircraft travelling across the Atlantic fly fast enough compared to the evolution of the wind field, to allow optimal trajectories to stay the same despite any small shifts in departure times (Sridhar et al., 2015). This allows for at least some conflict resolution to occur naturally through time planning of flights.

In Girardet et al. (2014) dynamic programming is again applied, but this time over a three dimensional simplex mesh, taking into account wind data, but keeping airspeed constant. At each point the cost of all possible new positions in an aircraft cruise trajectory is calculated and a time minimal front established. During every time interval this front is pushed forwards across the grid. Once the values at each grid point have been computed, the shortest time route from the departure to the destination airport can be retrieved. A method similar to this is used in Chapter 7, but for a grid of cuboids, based on horizontal position and mass and with a cost functional based on fuel use. Girardet et al. (2014)

applies the method to three days of air traffic across Europe and concludes that computing trajectories without using atmospheric data leads to very different routes compared to those taking this into account. It is also found that optimisation is most beneficial for the longest routes flown.

One advantage of using the HJB equation, is that once a map of values across the given grid has been found in advance of a flight, this can be used to obtain a number of flight paths in real time, so that tactical conflict avoidance is possible. Using a grid based method allows certain areas of the airspace, where conflicts would occur to be blocked out. In Parzani & Puechmorel (2017) the airspace is simplified to a 2D planar grid and conflict free minimal time trajectories are found by altering aircraft heading angles, whilst maintaining fixed airspeeds. The HJB equation is solved using the Ultra-Bee numerical scheme, rather than the Semi-Langrangian scheme applied in Chapter 7 of this thesis. In this way the calculation is simplified so that aircraft trajectories can be considered simultaneously, doubling the dimension of the initial problem. Both wind free routes and analytic wind fields are considered, leading to globally optimal solutions without the difficulty of initialising the trajectory search as is inherent in indirect methods.

A combination of both dynamic programming and calculus of variations is developed in Ng et al. (2011, 2014) to find minimum time and fuel burn routes through complex wind fields with step changes to flight altitude. In the first of these papers transpolar routes are considered as a way to reduce emissions on flights between North America and Asia. This initial method involves limiting contrail formation at the expense of the CO_2 emissions. It is found that by changing altitude tactically to avoid contrail formation, the increase in CO_2 can be limited to 1% for the two days trialled (7th August and 10th December, 2010). However, it is noted that the difficulty in applying this methodology operationally lies with the lack of flexibility in the entry and exit points of trans-polar airspace available.

In Ng et al. (2014) a number of step changes is planned to be included in a trajectory strategically. Once this number has been decided, the position of these step changes is found using discrete dynamical programming to minimise fuel use. Then the OCP for the horizontal flight path is solved using calculus of variations. A range of heading angles from the departure point is applied to the dynamical system which is then integrated forward until the first step climb point is reached. Similarly this method is applied with backward

stepping to go from the destination to the first step point. By using the time minimal results for both approaches the heading angles can be triangulated to give the heading angle for the time optimal route where the forwards and backwards generated paths meet. This method is repeated to find the route segments between each pair of step changes until the final step is reached at which point a TPBVP can be solved to find the time optimal route to the final destination. As the airspeed and altitude are kept constant for each step, the time minimal horizontal route will also be fuel minimal. This method was shown to provide savings of between 1 and 3% compared with recorded flight plan trajectories if altitudes were kept constant and up to 10% for trajectories including step climbs. Although these results sound very promising, the airspeed is not varied in the optimisation so its affect on fuel burn rate is kept constant.

Direct and indirect methods to solve the control problem of optimising an approach trajectory for both fuel use and noise level are compared to using a dynamic programming method in Khardi (2012). Results show that the direct and indirect methods give the same results both in terms of position and throttle settings. However, by using HJB to simplify the discretisation, a different solution is obtained, which is easier to implement for an in-flight management system, given the reduction in computer time needed for its calculation. Therefore, all three methods are of use in trajectory planning, but each is best fitted to a particular scenario. This is exploited in my research, with an indirect method used for time minimal fixed airspeed trajectories with an assumed constant mass, a direct method applied to reduce fuel burn for fixed-time flights with varying airspeed and mass and dynamical programming via the HJB equation employed to allow for changes to airspeed when fuel is minimised for free-time flights.

2.4.7 Dynamic programming used with indirect methods

Two of the three main methods for finding flight trajectories are combined in both Villarroel & Rodrigues (2016) and Fan et al. (2020). In these papers the cruise trajectory of a flight is optimised for a cost index (CI) calculated as the ratio of time-related costs to fuel-related costs. First the OCP is formulated to find the associated HJB equation, with airspeed as the control variable, then Pontryagin's Maximum Principle is applied to the Hamiltonian of the HJB equation to reduce it to a simpler state. In so doing a more time efficient, though sub-optimal solution can be obtained. In neither case are atmospheric conditions considered.

The cruise flights in Villarroel & Rodrigues (2016) are restricted to a longitudinal track between the top of ascent and the top of descent waypoints. Altitude is considered as fixed, with airspeed controlled to minimise CI. An analytical formula for optimal velocity based on a co-state of the weight is proved and then this is further simplified by assuming that this co-state is zero, resulting in suboptimal analytic solutions. These solutions are shown to be very close to results gained from a more computationally expensive shooting method. By applying the approximate method, the calculation becomes possible in real time and so could be used as part of the Flight Management System on board a commercial aircraft. A sensitivity analysis of the CI, weight and specific fuel consumption (S_{FC}) is performed, showing that when CI is very small (so fuel cost is the main component) the sensitivity to S_{FC} is zero and sensitivities to weight and CI are very small. However, for a large CI, which means the time related costs are dominant, only the change in weight has a small effect on the optimal airspeed.

In Fan et al. (2020) the simplified HJB equation can be solved by a MATLAB function to give the best airspeed for each altitude trialled. The aim of the system is to find the vertical profile that is most economical for a flight, based on keeping thrust and pitch angle within certain operational constraints. In both of these studies the state space is viewed as a planar slice through the atmosphere. The calculation of best airspeed depends on the particular model of the aircraft, its mass at each waypoint on the journey and the thrust applied. It does not, however, account for wind, so airspeed adjustments due to changing winds at different altitudes are not considered.

2.4.8 Non-analytical trajectory search methods

Not all solutions to flight trajectory problems involve continuous optimal control methods. In Ho-Huu et al. (2019) departure flight routes are found from Schiphol Airport in the Netherlands that minimise both the fuel use and the noise annoyance caused by the aircraft. This is done by approximating a take-off flight path by straight and fixed radius curved sections in the horizontal and vertical planes. Possible routes are compared by using a multi-objective evolutionary algorithm to make the problem into a set of sub problems, which can in turn be solved simultaneously using a genetic algorithm. Once the routes are set, specific flights can be assigned to different routes using an integer optimisation model to allow for most efficient use of the available airspace. This research

was limited by not taking atmospheric conditions into account, but shows that discrete optimisation can be applied to flight trajectories if legs are limited to certain shapes. The researchers also found that the process had to be completed in two stages, as combining the routing and allocation algorithms led to a prohibitively large computational cost.

Another example of a non-analytical trajectory search approach is the AirTraf Simulation which is used as part of an atmospheric chemistry model (EMAC) together with BADA and ICAO engine and fuel use data. This again finds aircraft cruise trajectories by way of a genetic algorithm. The benefit of such a system is that computation of derivatives is not required. First the GCR between the two airports is found. The longitude between the airports is split into four equal parts using three points, which are placed at the latitudes of the GCR in each case. These points serve as the centres of three user defined rectangles. Control points are chosen randomly inside each rectangle (which are used to bound all subsequent control point choices) and these can move independently of each other at each iteration of the system. A B-spline curve is then plotted in the horizontal plane, using the control points to locate it. Five vertical control points are now added along this route, equally spaced vertically and another B spline curve is generated to fit through these to make a 3D path. At each iteration the cost of a flight is evaluated by splitting the trajectory into a user defined number of waypoints so that every section uses relevant atmospheric data from a 2.8° resolution grid. Next selection (deciding which individual control points will be used to generate others), crossover (deriving new control points from two previous points) and mutation (the introduction of random noise to control point position) are used to create the next set of control points within their bounding rectangles horizontally and within the cruise range vertically. From these the new trajectory is found and so the process continues until the difference in cost between subsequent iterations is below a user set tolerance. This method does not guarantee a global minimum, but it is hoped that it converges to a cost minimal route. It is limited by the use of a constant Mach number for each trajectory. One major advantage of this method is that many different factors can be optimised by just changing the cost functional to be evaluated, so although in the first instance only wind optimal routes are plotted (Yamashita et al., 2015), in the second version of the model time optimal, cost optimal, fuel optimal and climate optimal paths are all considered (Yamashita et al., 2020). The drawback to such a method is that each route takes about 20 hours to compute (Yamashita et al., 2015) and that this reduces the number of days of atmospheric

conditions that can be trialled (Yamashita et al., 2021). To try to limit this effect, the position and strength of the jet stream on different days is used to characterise weather conditions and then five examples of each weather type are chosen for running the simulations. In Yamashita et al. (2021) climate optimised routing, trading off between economic cost and climate impact is investigated. The conclusion shows that the disadvantage of the increased costs involved would only be reduced if environmental taxes were levied globally.

Both Murrieta Mendoza et al. (2020) and Cheung (2018) use methods derived from graph theory to find aircraft trajectories. Murrieta Mendoza et al. (2020) apply the Floyd-Warshall Algorithm to minimise a cost functional based on fuel burn, flight time and an airline CI. In effect they turn the space between the departure and destination airports into a directional weighted graph across a three dimensional mesh. The flight model used takes into account deterministic weather forecasts from Environment Canada and uses BADA fuel burn data. Flight routes obtained have a constant Mach number and varying mass, but do not include a required time of arrival. Compared with actual flight data simulated trajectories were more fuel and time efficient when both heading angle and altitude were varied, rather than altitude alone. It was noted that climbing to a new altitude is very expensive in terms of fuel burn, a finding echoed in Lovegren & Hansman (2011) in which fixed horizontal flight path trajectories showed a larger potential to improve speed performance, with nearly 2.4% savings possible from speed optimisation compared to 1.5% for altitude optimisation. From these findings and analysis of the differences between atmospheric data at different altitudes, it was decided to use only fixed altitude flights throughout this thesis.

Cheung (2018) uses a different discrete method for finding optimal routes, based on an enhanced Dijkstra's algorithm. The area between the airports is split into a square mesh with 25 points zonally and the corresponding number of points meridionally to keep the grid square. Wind and temperature data are interpolated to get values at the chosen grid points. In the transatlantic examples completed, the 48 grid points closest to the current position are accepted as possible destinations. Then the great circle distance to each is divided by the ground speed of the aircraft to give the cost functional for the Dijkstra algorithm, for that particular stage of the journey. This would be too computationally expensive as is, especially as points that backtrack on the previous route would be tried. To avoid this an extra heuristic function is added to the sum of previous stages resulting in

the A* cost functional. The heuristic function works out the great circle distance from the current point to the destination and divides this by the airspeed of the aircraft added to a nominal 80 $\mathrm{m\,s^{-1}}$ tailwind. In this way, any parts of the journey that appear to go back towards the departure airport are identified and can be ruled out. This allows a shortest path across the network to be chosen, which has the advantage of being completely robust, (there will always be a path that goes from the departure to the destination airport). Another advantage of this method is that areas of the grid can easily be blocked out, to simulate regions of storms or no-fly zones. However, the trajectories are limited by the available grid points and the method is computationally expensive.

To avoid the need for complicated optimisation techniques during time sensitive manoeuvres, Dancila et al. (2019) take an already optimised horizontal trajectory and create a vertical profile which constrains airspeeds and pitch angles in order to reduce wear and tear on the aircraft, as well as fuel burn. This is formed using basic geometry and aircraft specific data allowing for fast calculations with limited computing resources, but it does not account for mass changes or atmospheric conditions and so it is less useful in strategic flight trajectory planning.

2.4.9 Machine learning for route prediction

A newer tactic for trajectory prediction in the case of an approaching collision is to use machine learning. In Hashemi et al. (2020) machine learning is applied to find a more time efficient solution to this problem than traditional methods. Model Predictive Control (MPC), where new trajectories depend on creating dynamic and kinematic models can be slow and costly. Hashemi et al. (2020) showed that data driven solutions outperformed MPC if enough training data is used. Unfortunately, the drawback to this method was found to be its susceptibility to "adversarial attack". If data was added to the training data that might lead to a collision, then 80 to 90% of the time the result would be corrupted by this. In the future it would be interesting to see if fuel optimal routes could be planned more quickly by a machine learning approach, trained on data from fuel optimal flights through deterministic wind fields.

Similarly in Jiang et al. (2021) the trajectory of an aircraft already in flight is predicted by using a combination of a database of aircraft performance models (drawn from BADA), gridpoint wind and temperature data and a large set of historical data. The data is mined

offline to formulate a horizontal flight path to which altitude and airspeed profiles are then added. As a flight progresses, this initial prediction can be tweaked using current information about the aircraft's actual route, in the same way that data assimilation is used to adjust numerical weather predictions. This method is time efficient and shows a better level of accuracy than a prediction model based purely on aircraft performance. Landing times predicted by this system are shown to be within 5% of actual times.

Lin et al. (2020) apply neural networks to solve the HJB equation approximately, in order to control automated land vehicles. By minimising a cost functional dependent on steer angle error and distance error, the vehicle is guided on a path. The neural networks increase the speed of mapping of the system states onto value functions and control inputs by 500 times, compared with a standard MPC solver. So in the future this method could perhaps be applied in order to create a real time aircraft trajectory map that could respond almost instantaneously to weather updates.

2.4.10 Route planning algorithms based on nature

Some authors have chosen to approach trajectory planning by borrowing techniques from the natural world. For example Murrieta Mendoza et al. (2016) use an algorithm based on ants finding food and Zhang et al. (2018) mirror the predatory behaviour of the antlion larva.

Ants when following paths to find food, leave behind a trace of pheromone. Other ants know from the strength of the pheromone how popular a path has been and therefore, where the food is most likely to be. In Murrieta Mendoza et al. (2016) lateral airspace is marked out as a grid in which the nodes are waypoints that aircraft may pass. The Ant Colony Optimization algorithm is a swarm intelligence algorithm. In the first sweep, many random paths are travelled, with weighting given to each to mirror the way pheromone is deposited. Then the 5 paths with the lowest performance index (based on time and fuel use) are selected and more weighting is given to these. With time, some of the weighting "evaporates", so when an "ant" is now placed at the departure node it will select a route to the destination node based on the strength of "pheromone" on each route. This is repeated a few times until convergence to a single route occurs. Improvements were shown between the cost of the routes obtained in this way and the shortest distance routes, taking into account the winds encountered across the grid. This approach, though innovative,

gives no guarantee that routes found are minimal. The grid system is predetermined and so fixes a selection of waypoints that may not be the best positions to fly to. The search technique should converge to a minimum cost path for this grid, but again, this is not guaranteed.

In Zhang et al. (2018) the behaviour of the antlion larva when hunting prey is exploited to create an optimisation process that finds weights and thresholds for a Long Short-Term Memory neural network. This network is used to predict the movements of one fighter plane relative to another based on nine different input variables, such as current 3D positions and relative speeds of the fighters. This system is trained in advance to allow for real time results in the air and provides a glimpse at where whole flight trajectory optimisation may go in the future. Obviously the need to take weather data into account will make the training more complicated and increase the number of inputs dramatically, but this is an interesting potential area for future research.

2.4.11 Innovations applied to drone flight planning

Trajectory optimisation is not just restricted to manned aircraft. Recent research has used a range of optimisation techniques to reduce power requirements of drone flights.

In Dobrokhodov et al. (2020) energy optimal trajectories for unmanned hybrid aircraft powered by both hydrogen fuel cells and solar photovoltaic cells are found by solving a TPBVP for aircraft travelling in two dimensions through a time-varying atmospheric field. It is assumed that both the air speed and heading angle of the craft are continuously adjustable and that air speed will remain within practical bounds. We make similar assumptions in Chapters 6 and 7. PMP is applied to derive the TPBVP, but due to the convoluted non-linearities of the fuel burn determination and wind fields, rather than trying to find an analytical evaluation of the second order necessary conditions to confirm that a minimum has been found, a numerical test is performed on the results instead. The optimal controls are perturbed at every instance of time by a small amount. This perturbation leads to convex energy and fuel costs, as the Hessian matrix is confirmed to be positive definite and so the retrieved solution is a minimum.

Henninger et al. (2020) minimise the energy for a journey with fixed waypoints by controlling drone airspeed. Here PMP is applied to the trajectory dynamics to minimise a

quadratic cost based on kinetic energy throughout the flight. The continuous system is discretised and then the discrete problem solved to give a good approximation of the solution to the original problem. As the initial co-state value needed to solve the system is not known, a simplified structure is used to allow a reasonable initial guess and then a shooting method based on Matlab's fminunc, which uses unconstrained non-linear optimisation, is applied to find the closest path to the waypoints.

2.4.12 What is the best way to plan a flight?

It is clear from this review of the research completed into aircraft trajectory planning that many methods have been applied to find efficient routes in a theoretical setting. No one method is the best, but rather each has its advantages and disadvantages, making it most applicable for a particular scenario.

Data from airlines is not usually shared for reasons of confidentiality, but there are now online platforms that do give information about flight tracks, if not about fuel use (Flightradar24, 2020). By plotting tracks of recent flights, some of the operational strategies can be deduced. For example in Chapter 6, Figure 6.12 shows that actual flights tend to be spread either side of the time optimal route. Furthermore, the parallel nature of the tracks flown shows that these paths have been dictated largely by the OTS. However, with the new focus on TBO this is likely to change.

2.4.13 Can any of these ideas be applied operationally?

Having considered the large body of literature available that details theoretical approaches to trajectory routing it is worth questioning if using TBO is actually practical. Can a theoretically conceived trajectory be used operationally?

There have been a number of initiatives to try to reduce flight emissions over the past two decades such as SESAR in Europe, the Environmentally Responsible Aviation project (ERA) in the US, the Asia and Pacific Initiative to Reduce Emissions (ASPIRE) in the Asia-Pacific region and the Collaborative Actions for Renovation of Air Traffic Systems (CARATS) in Japan (Gardi et al., 2016). These projects aim to bring theoretical methods for improving trajectories into a more practical setting, so that after modelling and simulation, actual trial flights can be used to test emissions reduction predictions fully. If

aviation is to become sustainable, then it will happen because Multi-Objective Trajectory Optimisation algorithms are applied in conjunction with Dynamic Airspace techniques, ensuring that the benefits predicted by mathematical research are not compromised by inflexible airspace rules and limited communications technology.

2.5 Chapter summary

Due to the competitive nature of commercial aviation, it is not possible to gain information about how airlines plan the routes they send to the ANSPs in their PRMs. However, we do know that the ANSPs do try to reflect these in their route planning for the OTS. The Met Office in the United Kingdom does provide time minimal trajectories based on the A-star method (Cheung, 2018), but many of the other research papers summarised here look instead at minimising a CI based on airline expenses, which combine time and fuel costs.

A large number of research papers has been written about finding optimal flight routes, of which this chapter has provided a summary. In general the key methods are indirect optimisation, direct optimisation and dynamic programming and in Table 2.1 advantages and disadvantages noted about these three approaches from a range of research papers are listed.

Where time minimisation is required, then using an indirect method such as PMP allows extremal solutions to be found, without the need for adjoint variables to be estimated. This method requires a system for finding the initial values of the control variables, so can be computationally expensive compared to direct methods and dynamical programming. Although it allows inclusion of general atmospheric conditions, like the wind field, within the working, trying to avoid areas of convection or turbulence would be more problematic. Equally, trying to apply this method when minimising fuel burn would require an approximate analytic formula for fuel burn, making the results less applicable to actual flights.

Direct methods are often used where a CI must be optimised, but simplified atmospheric conditions are acceptable. In a fixed time problem, these are often the best methods to use, as they provide more flexibility in choice of discretisation and optimisation methods. However, they are not guaranteed to even give an extremal like the indirect methods,

although a unique route will be obtained every time, making the method robust, if suboptimal. Methods that apply enough initial conditions to allow more certainty that a global minimum has been found can be slow, as the optimisation must be completed from each starting set and the lowest local minimum located amongst the candidate solutions.

Dynamic programming has the advantage of trying all points across a mesh and as long as representative atmospheric conditions are gained via interpolation, then the resulting trajectory is guaranteed to be optimal, for the chosen discretisation of the variables. The problem with this method is that all states must be tested with each possible set of controls and so it is easy for code to become overloaded with data in infeasibly large arrays. Although this is not a time constraint where parallelisation is possible, it can cause computer memory overload. However, the method of dynamic programming does allow sections of state space to be blocked out to represent no-fly areas, which could be used for storms or super-saturated icy regions.

Research into tactical flight planning often uses information from offline data mining that can then be applied in real time. In future, it might be possible for a method of machine learning to match previously planned routes to the encountered atmospheric conditions so that observed relationships can be used to generate likely optimal routes far faster.

In Chapters 4 and 5 of this thesis, we use time minimal route planning so that extremal routes are found and can be compared to OTS tracks used or to similarly simulated trajectories from different time periods. In Chapter 6 fuel minimal optimisation is favoured using a direct method for fixed time flights to allow comparison between formulations with different numbers of controls. Finally dynamic programming is used in Chapter 7 so that globally optimal fuel minimal routes can be compared with actual flights encountering the same atmospheric conditions.

Table 2.1: Summary of main trajectory optimisation methods from literature: indirect (I), direct (D) and dynamical programming (P).

Paper	Method	Advantages	Disadvantages
Schultz & Zagalsky (1972) Schultz (1974) Speyer (1973) Menon (1989)	I	Optimised for fuel.	Set distance. No wind.
Erzberger & Lee (1980)	I	Use wind data.	Fixed range cruise.
Sorensen & Waters (1981) Burrows (1983)	I	Take-off delays included.	Final time from many runs.
Chakravarty (1985)	I	Delay absorption.	Approximates fuelburn.
Jardin & Bryson (2012b)1 Jardin & Bryson (2012b)2	I I	Finds extremal. Faster as approximate.	Tests many angles. Finds suboptimal.
Franco & Rivas (2011) Nguyen (2006)	I	Finds extremal.	Simplified winds.
Kim et al. (2015)	I	Avoids convection.	Over predicts.
Sridhar et al. (2011)	I	Avoids contrails.	No mass change,slow.
Soler et al. (2020)	I	Avoids convection.	One day tested.
Fahroo & Ross (2002)	D	Robust, finds route.	Suboptimal, no wind.
Pierson & Ong (1989) García-Heras et al. (2014)	D	Compares methods. Minimises fuel.	Suboptimal. No wind data.
Ahmed et al. (2021)	P	Minimises fuel. 3 controls.	No wind data. Planar. Fixed mass.
Wickramasinghe et al. (2012)	P	Minimises fuel.	Limited path area.
Parzani & Puechmorel (2017)	P	Blocks areas.	Simplified wind.
Ng et al. (2011) Ng et al. (2014)	P	Minimises fuel. Minimises time.	Fixed airspeed. Fixed mass.
Khardi (2012)	P	Fuel/noise optimsed	Only descent phase.

Chapter 3

Background theory, models and data

3.1 Introduction

This chapter describes in detail the methods that can be used to solve optimal control theory problems, such as those arising from the need to find optimal routes for crossing the North Atlantic between LHR and JFK.

In each case deterministic wind fields and temperatures from re-analysis data or predictions of these wind fields and temperatures from an ensemble of climate models, is used to give realistic atmospheric conditions for such a flight.

Flight dynamics for rate of change of latitude and longitude around a spherical shell at a fixed altitude above the Earth's surface are used as the dynamical system, with an extra line describing mass change for the examples with varying mass.

Optimal control requires the solution of these systems of non-linear differential equations subject to certain constraints, costs and conditions. These vary depending on how the OCP is formulated.

3.2 OCP specific to transatlantic flights

3.2.1 Optimal control formulation

In general an optimisation problem will involve changing certain variables (the control variables) in order to navigate from an initial to a final state, whilst minimising or maximising a particular cost associated with the journey. Here we assume that a minimum is required in each case. The state variables describe the trajectory which gives the minimum cost, once the corresponding control variables have been found.

Problems are stated as:

- cost functional: cost to minimise,
- dynamical system: set of differential equations describing evolution of states,
- boundary conditions on each state: normally not all of these are known,
- constraints on states/control variables: values not to exceed en route.

We define the state equations as:

$$\dot{\boldsymbol{x}} = \boldsymbol{f}(\boldsymbol{x}, \boldsymbol{\alpha}, t), \tag{3.1}$$

where $\boldsymbol{x}$ is a vector of n state variables, which lie in a closed domain $\Omega \subset \mathbb{R}^n$, $\boldsymbol{\alpha}$ is a vector of m piecewise continuous controls, which lie in the admissible set of all controls $\mathcal{A} \subset \mathbb{R}^m$ and t is time, which runs from t_0 to t_f.
The associated cost functional, or performance index, is composed of the terminal condition and the integral of the running cost, which describes what is to be minimised:

$$J = \beta[\mathbf{x}(t_f), t_f] + \int_{t_0}^{t_f} L[\mathbf{x}(t), \alpha(t), t]\mathrm{d}t. \tag{3.2}$$

However, in the following research, rather than including a terminal cost, which requires finding β, a non-linear constraint is used. This means that only routes that reach the specified target set, C, will be accepted as admissible trajectories. This involves measuring the distance between the final destination and the end of each trajectory. If the distance between the two lies within a certain range, it is considered to have reached the target set, where C is a compact subset of $\mathbb{R}^n$. In some contexts this cost functional is known as the objective function (Cartis et al., 2021).

3.2.2 Parameters, states and controls

The OCP itself can be defined in a general way, as a set of parameters, states and control variables which are all part of the dynamical system under consideration.

In each part of this thesis, the parameter is time, t. This is measured in seconds. In the infinite horizon case, where fuel burn is to be minimised:

$$t_0 = 0,$$
$$t_f = T,$$

where t_f is final time and t_0 is initial time. Here T is a given, set final time. However, in the time minimal formulation and the HJB formulation in Chapter 7, t_f is unknown a priori and will be found as part of the optimisation.

The state variables describe the position of the aircraft at any time t and in the case where mass is not considered as constant, mass is an additional state variable. For example in the constant altitude, constant mass case, the state variables ($\mathbf{x}$) are:

$$x_1 = \lambda(t) \in [-\pi, \pi],$$
$$x_2 = \phi(t) \in [-\pi/2, \pi/2],$$

where λ and ϕ are positions of longitude and latitude, in radians, respectively, but in later problems three different states are considered so we extend this list to include:

$$x_3 = M(t) \in [M_{min}, M_{max}],$$

where M is the mass of the aircraft in kg. The maximum and minimum masses of an aircraft are governed by the model of aircraft used and its expected time of flight.
Control variables ($\boldsymbol{\alpha}$) are the values which are varied to ensure an optimal result. In the simplest problem only heading angle is controlled, but later airspeed is added to the list of controls:

$$\alpha_1 = \theta(t) \in [0, 2\pi],$$
$$\alpha_2 = V(t) \in [199, 252],$$

where θ is the heading angle in radians and V is the airspeed in $\mathrm{m\,s^{-1}}$. The airspeed constraints are the boundary values for the fuel burn calculation to be valid at cruise altitude given the maximum and minimum temperatures across the North Atlantic.

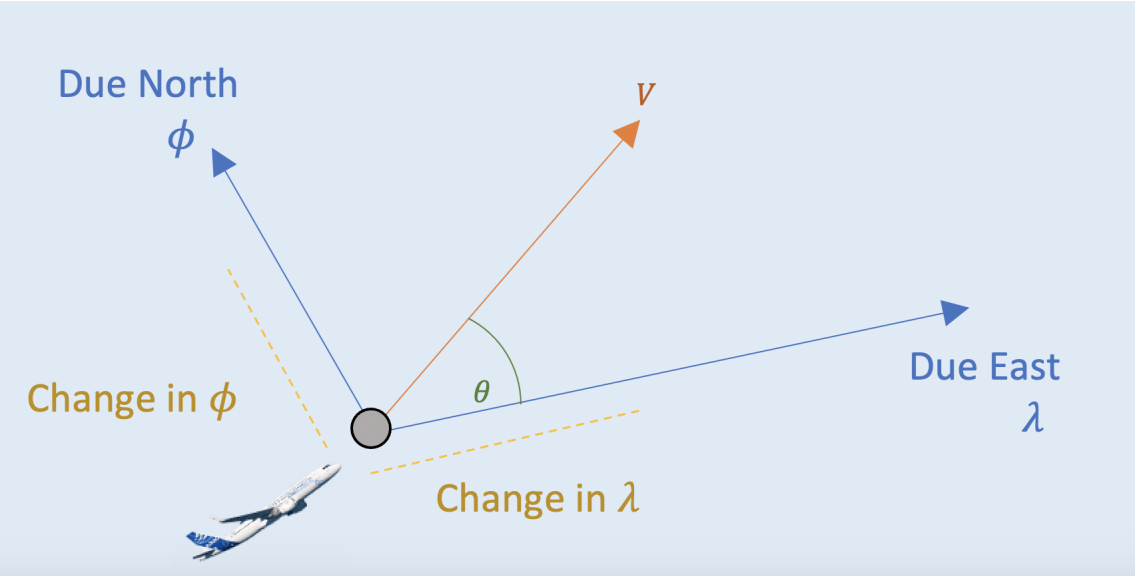

Figure 3.1: Diagram showing the velocity of an aircraft in flight, showing the state variable axes and control variables.

A diagram showing the motion of the aircraft is shown in Figure 3.1.

3.2.3 Dynamical system and cost functional for the time minimal problem

The dynamical system is derived from a planar longitude/latitude grid system, but this is then scaled conformally to fit the spherical model (as shown in Appendix B). As the radius of the Earth, R is very much greater than the altitude of the aircraft, h, we assume that h is negligible in this mapping and use just R as the radius of the spherical shell. In the time minimal case in Chapters 4 and 5, the mass of the aircraft is considered to be constant and taken as an average across the whole trajectory. This gives the following dynamical system to describe the movement of the aircraft:

$$\dot{\lambda} = \frac{1}{R\cos\phi}(V\cos\theta + u(\lambda,\phi)), \tag{3.3}$$

$$\lambda(0) = \lambda_{\text{dept}}, \tag{3.4}$$

$$\dot{\phi} = \frac{1}{R}(V\sin\theta + v(\lambda,\phi)), \tag{3.5}$$

$$\phi(0) = \phi_{\text{dept}}, \tag{3.6}$$

where V is airspeed in $\mathrm{m\,s^{-1}}$, R is the radius of the Earth (taken to be 6 371 km), u is zonal wind velocity and v is meridional wind velocity (in $\mathrm{m\,s^{-1}}$). Both wind components are position dependent, but not time dependent as wind speeds are taken across a position grid, but averaged over each day. The dynamical system is referred to as f.

In Chapter 4 the airspeed of the aircraft is kept constant throughout each trajectory. For the system used in Chapter 5 the Mach number of the aircraft is kept constant, so at each time step the airspeed is slightly altered to allow for any changes to the temperature of the air. Mach number is the ratio of the airspeed to the speed of sound at any particular point, which depends on air temperature. It should be noted that across the North Atlantic ambient temperature changes with location are very small. In the time minimal calculations in Chapter 7 airspeed V is included as a second control variable, although it is always chosen to be at its maximum by the algorithm in order to minimise time.

The cost functional is the value to be minimised by varying the control variables. The basic performance index for an OCP is composed of the terminal cost and the integral of the Lagrangian, L, which is the running cost, subject to the state vector, $\boldsymbol{x}$, the control vector,$\boldsymbol{\alpha}$, and time, the parameter t. In the current model the terminal cost is removed and applied as a non-linear inequality constraint leaving only the running cost:

$$J = \int_{t_0}^{t_f} L[\mathbf{x}(t), \alpha(t), t]\mathrm{dt}. \tag{3.7}$$

In the case where time is being minimised, the cost functional is reduced to:

$$J = \int_{t_0}^{t_f} L[\mathbf{x}(t), \alpha(t), t]\mathrm{dt} = \int_{t_0}^{t_f} 1dt = t_f. \tag{3.8}$$

The non-linear inequality constraint ensures that the spherical distance between the co-ordinates of the end of a trajectory and the co-ordinates of the destination airport is less than or equal to a given value. The spherical distance is calculated using the Haversine formula, which is a standard method in this field (Veness, 2019):

$$\begin{aligned} a_H &= \sin^2(\Delta lat/2) + \cos(lat1)\cos(lat2)\sin^2(\Delta lon/2), \\ c_H &= 2\mathrm{atan2}(\sqrt{a_H}, \sqrt{1-a_H}), \\ d_H &= Rc_H, \end{aligned} \tag{3.9}$$

where $lat1$ is the final latitude in the state vector, $lat2$ is the latitude of the destination airport, Δlat is the difference in latitudes between these two points, Δlon is the difference

in their longitudes and R is the radius of the Earth. The atan2 function is the four-quadrant inverse tangent of the two real values in the function bracket. This computes arctan $\frac{\sqrt{a_H}}{\sqrt{(1-a_H)}}$, but includes an extra term of $\frac{\pi}{2}\text{sign}(\sqrt{(}1-a_H))(1-\text{sign}(\sqrt{a_H}))$ to extend the range of results to cover $[-\pi, \pi]$. It also accepts the case where $a_H = 1$, returning $\frac{\pi}{2}$.

3.2.4 Dynamical system and cost functional for the fuel minimal problem

The fuel minimal problem requires a new dynamical system, where the changing mass is included. This is essential when minimising fuel, as the fuel used during each time step causes the mass to change by this amount, with a knock-on effect to the calculation for the next fuel burn rate. This dynamical system is stated as:

$$\dot{\lambda} = \frac{1}{R\cos\phi}(V\cos\theta + u(\lambda, \phi)), \tag{3.10}$$

$$\lambda(0) = \lambda_{\text{dept}}, \tag{3.11}$$

$$\dot{\phi} = \frac{1}{R}(V\sin\theta + v(\lambda, \phi)), \tag{3.12}$$

$$\phi(0) = \phi_{\text{dept}}, \tag{3.13}$$

$$\dot{M} = -g(\lambda, \phi, M, V), \tag{3.14}$$

$$M(0) = M_{\text{dept}}, \tag{3.15}$$

where V is the control variable airspeed in $\text{m}\,\text{s}^{-1}$, R is the radius of the Earth, u is zonal wind velocity and v is meridional wind velocity (in $\text{m}\,\text{s}^{-1}$). The fuel burn function, g is described in more detail in Section 3.4 of this chapter.

In Chapter 6 this system is solved by a direct method, for both fixed airspeed flights and those with airspeed as a second control variable.

When dynamic programming is applied to solve this problem in Chapter 7, both heading angle and airspeed are used as control variables. In this case fuel is being minimised at fixed altitude, so the cost functional is based on the fuel burn function with altitude kept constant:

$$J = \int_{t_0}^{t_f} L[\mathbf{x}(t), \alpha(t), t]\text{dt} = \int_{t_0}^{t_f} g(\lambda, \phi, M, V)dt. \tag{3.16}$$

3.3 A summary of optimal control approaches

Finding a minimum fuel burn route is an optimisation problem with a large number of potential state and control variables. There are three main ways to resolve optimisation problems: direct methods, indirect methods and via dynamic programming. Each of these methods has advantages and disadvantages and normally the statement of the problem will give an indication of which method to apply.

3.3.1 Indirect methods

Indirect methods involve first optimising and then discretising. The method makes use of PMP to derive necessary first order optimality conditions, by taking derivatives of the Pontryagin Hamiltonian. If problem boundary conditions are known this reduces to a TPBVP, with adjoint values at the final time given by the transversality conditions. Indirect methods are then needed to solve the resulting system and the solution is not necessarily a minimum. By working in this way the number of states and controls is not limited. However, for every starting value the system of differential equations should have a unique solution. We then look for the initial conditions that give the solution we are trying to find. For example in the case of minimum time flights, we find the initial conditions that allow the aircraft to arrive within a given distance of the airport in the shortest time possible. The minimum time methods used in Chapters 4 and 5 of my thesis rely on PMP.

Mathematical theory underpinning PMP

PMP gives the necessary conditions for a first order optimal control to exist based on the Pontryagin Hamiltonian defined as:

$$H(\boldsymbol{x}, \boldsymbol{p}, \boldsymbol{\alpha}) = L[\mathbf{x}(t), \boldsymbol{p}, \boldsymbol{\alpha}(t)] + \langle \boldsymbol{p}, f[\mathbf{x}(t), \boldsymbol{\alpha}(t)] \rangle, \tag{3.17}$$

where L is the running cost, $\boldsymbol{p}$ is a vector of co-states and $\langle \cdot, \cdot \rangle$ is the scalar product (Kirk, 1970; Lee & Markus, 1967; Macki & Strauss, 1982; Athans & Falb, 1966).

For the preparation of the formulation of PMP, we follow the method of Athans & Falb (1966).

Assuming that partial derivatives of $\boldsymbol{f}$ and L with respect to both the state variables and time, are continuous on $\mathbb{R}^n \times \mathcal{A} \times (t_0, t_f)$, then functions $H(\boldsymbol{x}, \boldsymbol{p}, \boldsymbol{\alpha})$ and $\frac{\partial H}{\partial \boldsymbol{x}}(\boldsymbol{x}, \boldsymbol{p}, \boldsymbol{\alpha})$ are

continuous on $\mathbb{R}^n \times \mathbb{R}^n \times \Omega$. This also means that $\frac{\partial H}{\partial \boldsymbol{\alpha}}(\boldsymbol{x}, \boldsymbol{p}, \boldsymbol{\alpha})$ is well defined and given by:

$$\frac{\partial H}{\partial \boldsymbol{\alpha}}(\boldsymbol{x}, \boldsymbol{p}, \boldsymbol{\alpha}) = f[\mathbf{x}(t), \boldsymbol{\alpha}(t)]. \tag{3.18}$$

Now let $\hat{\boldsymbol{x}}$ be the trajectory starting from state vector $\boldsymbol{x}$ at t_0 generated by the control vector $\hat{\boldsymbol{\alpha}}$ and $\boldsymbol{\pi}$ be any n vector. For any function $\boldsymbol{p}(t)$ we get:

$$\dot{\hat{\boldsymbol{x}}} = \frac{\partial H}{\partial \boldsymbol{p}}[\hat{\boldsymbol{x}}(t), \boldsymbol{p}(t), \hat{\boldsymbol{\alpha}}] = \boldsymbol{f}[\hat{\boldsymbol{x}}(t), \hat{\boldsymbol{\alpha}}(t)], \tag{3.19}$$

$$\dot{\boldsymbol{p}}(t) = -\frac{\partial H}{\partial \boldsymbol{x}}[\hat{\boldsymbol{x}}(t), \boldsymbol{p}(t), \hat{\boldsymbol{\alpha}}], \tag{3.20}$$

where Equation (3.20) has a unique solution $\boldsymbol{p}(t, \boldsymbol{\pi})$ satisfying $\boldsymbol{p}(t_0, \boldsymbol{\pi}) = \boldsymbol{\pi}$.

Pontryagin's Minimum Principle

From these definitions, PMP for the case where:

- final time is free,
- final position lies within a target set,
- system of state equations $\dot{\mathbf{x}}$ do not depend explicitly on t,
- cost functional $\mathbf{J}$ does not depend explicitly on t,

can now be quoted:

Theorem 1: The Minimum Principle of Pontryagin (Athans & Falb, 1966)
Let $\boldsymbol{\alpha}^*(t)$ be an admissible control which transfers $(\mathbf{x}_0, t_0)$ to target C. Let $\boldsymbol{x}^*(t)$ be the trajectory corresponding to $\boldsymbol{\alpha}^*(t)$ starting at $(\mathbf{x}_0, t_0)$ and meeting C for the first time at t_f, so that $\boldsymbol{x}^*(t_f) \in C$.
In order for $\boldsymbol{\alpha}^*(t)$ to be optimal, it is necessary that there exists a function $\boldsymbol{p}^*(t)$ such that:

1. $\boldsymbol{p}^*(t)$ corresponds to $\boldsymbol{\alpha}^*(t)$ and $\boldsymbol{x}^*(t)$ so that $\boldsymbol{p}^*(t)$ and $\boldsymbol{x}^*(t)$ are a solution of the canonical system:

$$\dot{\boldsymbol{x}^*}(t) = \frac{\partial H}{\partial \boldsymbol{p}}[\boldsymbol{x}^*(t), \boldsymbol{p}^*(t), \boldsymbol{\alpha}^*(t)],, \tag{3.21}$$

$$\dot{\boldsymbol{p}^*}(t) = -\frac{\partial H}{\partial \boldsymbol{x}}[\boldsymbol{x}^*(t), \boldsymbol{p}^*(t), \boldsymbol{\alpha}^*(t)], \tag{3.22}$$

satisfying the boundary conditions:

$$\boldsymbol{x}^*(t_0) = \boldsymbol{x}_* \qquad \boldsymbol{x}^*(t_f) \in C, \tag{3.23}$$

2. $H[\boldsymbol{x}^*(t), \boldsymbol{p}^*(t), \boldsymbol{\alpha}(t)]$ has an absolute minimum as a function of $\boldsymbol{\alpha}$ over $\mathcal{A}$ at $\boldsymbol{\alpha} = \boldsymbol{\alpha}^*$ for $t \in [t_0, t_f]$, which can be expressed as:

$$\min_{\boldsymbol{\alpha} \in \mathcal{A}} H[\boldsymbol{x}^*(t), \boldsymbol{p}^*(t), \boldsymbol{\alpha}(t)] = H[\boldsymbol{x}^*(t), \boldsymbol{p}^*(t), \boldsymbol{\alpha}^*(t)], \tag{3.24}$$

3. $H[\boldsymbol{x}^*(t), \boldsymbol{p}^*(t), \boldsymbol{\alpha}^*(t)] = 0 \quad t \in [t_0, t_f]$,

4. Vector $\boldsymbol{p}^*(t_f)$ is normal to target C at $\boldsymbol{x}^*(t_f)$.

Towards a heuristic proof of PMP

The proof of PMP is stated in full in Pontryagin's original text: Pontryagin et al. (1962). This is described in the introduction to Barron & Jensen (1986) as "very long and technical", so in their paper the case with just running cost is proved more succinctly by deriving PMP from the viscosity solution of the HJB equation, which will be discussed further in Section 3.3.3. More heuristic forms of the proof for particular cases have appeared in a range of texts including in Athans & Falb (1966), which is summarised here. The material splits into nine steps, which gradually lead to the proof of PMP.

Firstly an auxiliary variable is introduced into the state vector, so that the the running cost is expressed as the rate of change of x_0 and the system is now of order $n + 1$:

$$\dot{x}_0(t) = L[\boldsymbol{x}(t), \boldsymbol{\alpha}(t)],$$
$$\dot{\boldsymbol{x}}(t) = \boldsymbol{f}[\boldsymbol{x}(t), \boldsymbol{\alpha}(t)].$$

The optimal trajectory is $\boldsymbol{y}^*(t)$, which can be projected onto the x-axis as $\boldsymbol{x}^*(t)$, allowing the cost to be plotted on the vertical axis and the state space to be compressed to a single horizontal axis. At t^*, the final time for the optimised trajectory starting at t_0, the optimal trajectory is given as:

$$\boldsymbol{y}^*(t^*) = \begin{bmatrix} J^* \\ \boldsymbol{x}_1 \end{bmatrix} = \begin{bmatrix} x_0^*(t^*) \\ \boldsymbol{x}^*(t^*) \end{bmatrix}. \tag{3.25}$$

If we now change the start time to $t_0 + r$, the control vector, $\boldsymbol{\alpha}^r(s + \tau)$, is seen to be the same as the original control vector for $s \in (t_0, t^*)$, $\boldsymbol{\alpha}^*(s)$. This means that a translation in time results in the same optimal trajectory, with no change to the cost.

Next the principle of optimality can be demonstrated by showing that the optimal control $\boldsymbol{\alpha}^*$ is optimal for any point on the optimal trajectory. By considering smaller segments of the optimal trajectory it is shown that any trajectory sections plotted below the optimal, on a cost against state graph, will prevent $\boldsymbol{\alpha}^*$ from being optimal, so a proof by contradiction means that geometrically any $\boldsymbol{y}$ trajectories meeting the target must lie above the trajectory generated by $\boldsymbol{\alpha}^*$. Therefore any portion of the optimal trajectory must also be optimal. This also allows us to generate a tube-like region around the optimal trajectory within which all control vectors are approximately optimal, thus broadening the target, which will help in the next stages in the heuristic proof, as the trajectory is perturbed temporally and spatially.

The third step looks at the effect of a small change to the initial conditions, so that:

$$\boldsymbol{y}^*(t) + \epsilon\boldsymbol{\psi}(t) + \boldsymbol{O}(\epsilon), \tag{3.26}$$

represents the perturbed trajectory, where a small change in starting position $\boldsymbol{\xi}$ generates a change to the state $\boldsymbol{\psi}(t)$ at each time t with:

$$\lim_{\epsilon \to 0} \frac{\| \boldsymbol{O}(\epsilon) \|}{\epsilon} = 0. \tag{3.27}$$

In Figure 3.2 the effect of a shift in initial conditions of $\boldsymbol{\xi}$ is shown. The perturbed trajectories always lie in a narrow region around the optimal trajectory, which has approximately constant width as time changes.

In the next step the adjoint variables are considered and these are shown to describe hyperplanes, using the fact that the adjoint vector at any time t is perpendicular to the corresponding hyperplane. The first of these, P_0, passes through y_0 and is perpendicular to $\boldsymbol{\xi}$, which is the vector showing the perturbation of the initial values (shown in Figure 3.2). We can then describe P_t, which is the hyperplane at any time t and will be perpendicular to $\boldsymbol{\psi}(t)$. In effect P_0 is moved along the optimal trajectory to give P_t.

We now consider a set of convex cones, C_t, whose vertices are at $\boldsymbol{y}^*(t)$. The rays of these cones result from perturbing the optimal control vector, $\boldsymbol{\alpha}^*$ and are called the cones of attainability. These will always lie on the same side of the hyperplane defined by the perpendicular vector of adjoint variables.

The fifth step involves considering the effect to the optimal control, $\boldsymbol{\alpha}^*$, of a small change,

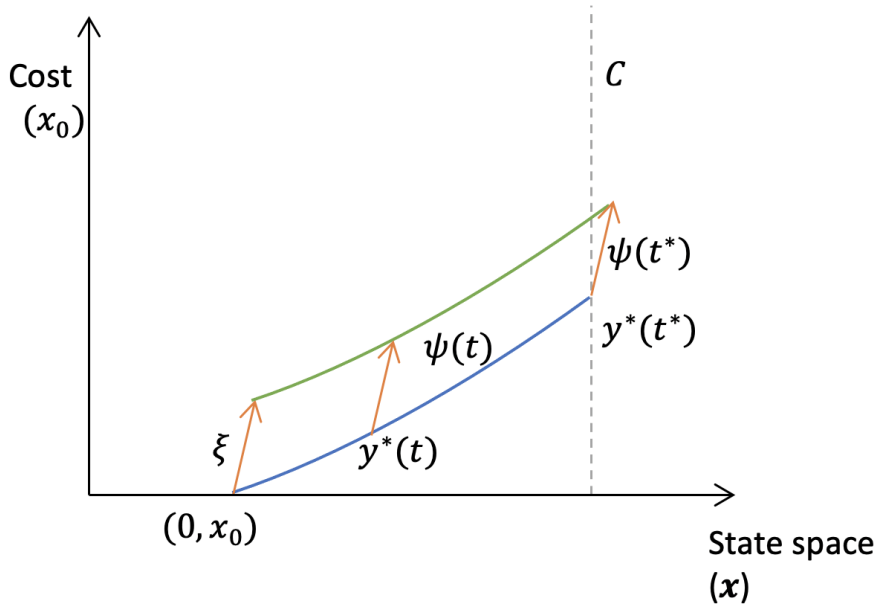

Figure 3.2: Plot of cost against position in state space, demonstrating the effect of a small change to the initial position, $\boldsymbol{\xi}$. This has been exaggerated for clarity. This graph has been adapted from Figure 5-21 on page 317 of Athans & Falb (1966),

τ, to the terminal time, t^*:

$$\boldsymbol{\alpha}[\tau](t) = \begin{cases} \boldsymbol{\alpha}^*(t) & t_0 \leq t \leq t^* + \epsilon\tau \qquad \tau < 0, \\ \boldsymbol{\alpha}^*(t) & t_0 \leq t \leq t^*, \\ \boldsymbol{\alpha}^*(t) & t_0 \leq t \leq t^* + \epsilon\tau \qquad \tau > 0. \end{cases}$$

The new end points of these longer or shorter duration versions of the optimal trajectory are shown to lie along the same ray, $\overrightarrow{p}$, which goes through $\boldsymbol{y}^*(t^*)$ and is tangent to the slope of $\boldsymbol{y}^*(t)$ at this point. This is shown in Figure 3.3.

Step six involves perturbing the optimal control spatially, just across one very short time interval of length ϵa finishing at time b, during which the control is taken to have a constant value, ω. The linear sums of such perturbations are also considered. Examples of spatial perturbations are shown in Figure 3.4.

All of the combinations of the rays $\overrightarrow{p}$ form a cone $\overrightarrow{P}$, whose convex cover co($\overrightarrow{P}$) is a convex cone with a vertex at $\boldsymbol{y}^*(t^*)$. (We take the cover to ensure convexity. It is the smallest convex set containing $\overrightarrow{P}$).

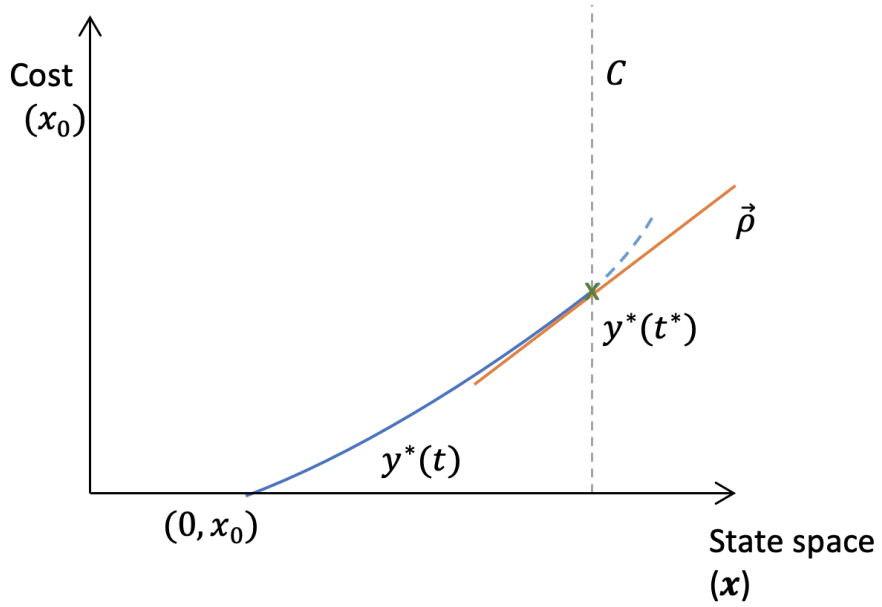

Figure 3.3: Plot of cost against position in state space, showing the optimal trajectory $\boldsymbol{y}^*$ and the ray $\overrightarrow{\rho}$, which is generated by perturbing the duration of the control vector. This graph is based on part of Figure 5-26 on page 324 of Athans & Falb (1966).

In step seven, the two perturbations whose effects were established in the previous two stages, are now combined to give $C_{t^*} = \overrightarrow{\rho} + \text{co}(\overrightarrow{P})$ which is also a convex cone with a vertex at $\boldsymbol{y}^*(t^*)$. All points included in this cone are linear combinations of parts of $\overrightarrow{\rho}$ and parts of co$(\overrightarrow{P})$ and all cones of attainability will therefore lie somewhere within this cone. As C_{t^*} can be shown to lie between two half hyperplanes, which coincide along the ray $\overrightarrow{\rho}$, we can see that all cones of attainability must lie the same side of the moving hyperplane P_{t^*}, which is defined as perpendicular to the vector of adjoint variables.

By considering all of these findings it is possible to show that a ray in the direction of decreasing cost from $\boldsymbol{y}^*(t^*)$, $\overrightarrow{\mu}$, does not meet the interior of C_{t^*}, as the hyperplane P_{t^*} separates them. The existence of this hyperplane means that there is a non-zero vector of adjoints, normal to the hyperplane:

$$\begin{bmatrix} p_0^* \\ \boldsymbol{p}^*(t^*) \end{bmatrix}. \tag{3.28}$$

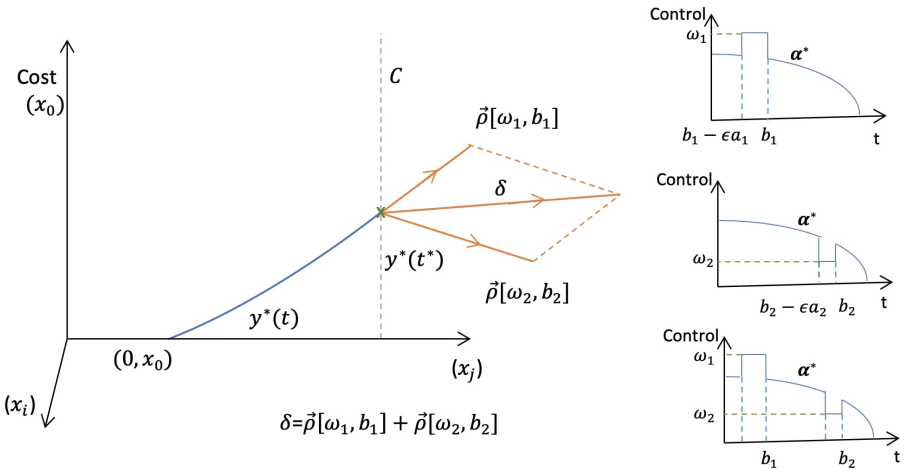

Figure 3.4: Plot of cost against position in state space, showing the optimal trajectory $\boldsymbol{y}^*$ and the rays $\overrightarrow{\rho}$, which are generated by perturbing the optimal control vector for short duration time intervals at different times during a trajectory. This effect is shown to be additive, with vector δ being the sum of $\overrightarrow{\rho}[\omega_1, b_1]$ and $\overrightarrow{\rho}[\omega_2, b_2]$, resulting from the two corresponding perturbations of the control vector as shown in the graphs on the right. These graphs are based on those in Figure 5-31 on page 330 of Athans & Falb (1966).

This means that:

$$P_{t^*} = \left\langle \begin{bmatrix} p_0^* \\ \boldsymbol{p}^*(t^*) \end{bmatrix}, \boldsymbol{y} - \boldsymbol{y}^*(t^*) \right\rangle = 0. \tag{3.29}$$

Motion along the optimal trajectory can be described by:

$$\dot{\boldsymbol{\psi}}(t) = \nu_*(t)\boldsymbol{\psi}(t), \tag{3.30}$$

where ν_* is the $n+1 \times n+1$ matrix:

$$\nu_* = \begin{bmatrix} 0 & \frac{\partial L}{\partial x_1}|_* & \frac{\partial L}{\partial x_2}|_* \cdots \frac{\partial L}{\partial x_n}|_* \\ \mathbf{0} & \frac{\partial \boldsymbol{f}}{\partial \boldsymbol{x}}|_* & \end{bmatrix},$$

which gives a linear and homogeneous system adjoint:

$$\begin{bmatrix} \dot{p}_0 \\ \dot{\boldsymbol{p}}(t) \end{bmatrix} = -\nu_*'(t) \begin{bmatrix} p_0 \\ \boldsymbol{p}(t) \end{bmatrix}. \tag{3.31}$$

This means that:

$$\begin{aligned} p_0^*(t) &= p_0^*, \\ \dot{\boldsymbol{p}}(t) &= -\frac{\partial L}{\partial \boldsymbol{x}}\bigg|_* p_0(t) - \left(\frac{\partial f}{\partial \boldsymbol{x}}\bigg|_*\right)' p(t), \end{aligned}$$

and so on the optimal path:

$$\dot{\boldsymbol{p}}^*(t) = -\frac{\partial L}{\partial \boldsymbol{x}}\bigg|_* p_0^*(t) - \left(\frac{\partial f}{\partial \boldsymbol{x}}\bigg|_*\right)' p^*(t), \tag{3.32}$$

This can then be expressed using the Hamiltonian:

$$H[\boldsymbol{x}, \boldsymbol{p}, \boldsymbol{\alpha}, p_0] = p_0 L(\boldsymbol{x}, \boldsymbol{\alpha}) + \langle \boldsymbol{p}, \boldsymbol{f}(\boldsymbol{x}, \boldsymbol{\alpha}) \rangle, \tag{3.33}$$

which gives Equation (3.22).

To prove the second part of PMP we show that $\boldsymbol{\alpha}^*$ minimises the Hamiltonian. Take a point $t_0 \leq b \leq t^*$ and $\boldsymbol{\omega} \in \mathcal{A}$. Then we consider the difference in the extended dynamical system for $\boldsymbol{y}^*(b)$ when the trajectory depends on control $\boldsymbol{\omega}$ compared to the optimal control at b, $\boldsymbol{\alpha}^*(b)$:

$$\boldsymbol{\xi}_b(\boldsymbol{\omega}) = \begin{bmatrix} L[\boldsymbol{x}^*(b), \boldsymbol{\omega}] - L[\boldsymbol{x}^*(b), \boldsymbol{\alpha}^*(b)] \\ \boldsymbol{f}[\boldsymbol{x}^*(b), \boldsymbol{\omega}] - \boldsymbol{f}[\boldsymbol{x}^*(b), \boldsymbol{\alpha}^*(b)] \end{bmatrix}. \tag{3.34}$$

We now define $\Phi_*(t, b)$ as the fundamental matrix of $\dot{\boldsymbol{\psi}}(t) = \nu_*(t)\boldsymbol{\psi}(t)$, so that for small ϵ:

$$\epsilon \Phi_*(t^*, b)\xi_b(\boldsymbol{\omega}) \in C_{t^*}. \tag{3.35}$$

As this is an element of the cone of attainability on the optimal trajectory, the inner product of the vector of adjoints and $\boldsymbol{\xi}_b(\boldsymbol{\omega})$ must be non-negative:

$$\left\langle \begin{bmatrix} p_0^* \\ \boldsymbol{p}^*(b) \end{bmatrix}, \xi_b(\boldsymbol{\omega}) \right\rangle \geq 0, \tag{3.36}$$

which means that:

$$H[\boldsymbol{x}^*(b), \boldsymbol{p}^*(b), \boldsymbol{\omega}, p_0^*] \geq H[\boldsymbol{x}^*(b), \boldsymbol{p}^*(b), \boldsymbol{\alpha}^*, p_0^*], \tag{3.37}$$

which establishes Equation (3.24).

The next part of PMP can be proved by showing first that the Hamiltonian, H, is zero at the terminal time, t^* and then that it is constant on the time interval $[t_0, t^*]$.

Consider $\boldsymbol{\delta}(\tau)$ which is a vector attached to $\boldsymbol{y}^*(t^*)$, and part of ray $\overrightarrow{\rho}$:

$$\boldsymbol{\delta}(\tau) = \begin{bmatrix} L[\boldsymbol{x}^*(t^*), \boldsymbol{\alpha}^*(t^*)] \\ \boldsymbol{f}[\boldsymbol{x}^*(t^*), \boldsymbol{\alpha}^*(t^*)] \end{bmatrix} \tau, \tag{3.38}$$

where $\tau \in \mathbb{R}$. For this to be true:

$$\left\langle \begin{bmatrix} p_0^* \\ \boldsymbol{p}^*(t^*) \end{bmatrix}, \begin{bmatrix} L[\boldsymbol{x}^*(t^*), \boldsymbol{\alpha}^*(t^*)] \\ \boldsymbol{f}[\boldsymbol{x}^*(t^*), \boldsymbol{\alpha}^*(t^*)] \end{bmatrix} \tau \right\rangle \geq 0, \tag{3.39}$$

which means that:

$$H[\boldsymbol{x}^*(t^*), \boldsymbol{p}^*(t^*), \boldsymbol{\alpha}^*(t^*), p_0^*]\tau \geq 0. \tag{3.40}$$

However, τ could be positive or negative, so it follows that $H = 0$ at t^*.

We now assume that $\boldsymbol{\alpha}^*$ has only as finite number of discontinuities and that in the continuous interval between two given times $[t_1, t_2]$, $\boldsymbol{\alpha}^*$ is continuous. We consider closed and bounded sets of states, adjoints and controls on this interval, X_1 , P_1 and $\mathcal{A}_1$. The Hamiltonian is therefore continuous on $[X_1 \times P_1 \times \mathcal{A}_1]$ and on this set it is also assumed to have continuous partial derivatives with respect to both $\boldsymbol{x}$ and $\boldsymbol{p}$.

We now define a function:

$$m(\boldsymbol{x}, \boldsymbol{p}) = \inf_{\boldsymbol{\alpha} \in \mathcal{A}_1} H[\boldsymbol{x}, \boldsymbol{p}, \boldsymbol{\alpha}, p_0^*]. \tag{3.41}$$

As $\boldsymbol{\alpha}^*$ is continuous in the given time interval we know that m is a continuous function when applied to the optimal trajectory. Taking two distinct values of t, t_a and t_b, such

that $t_1 \leq t_a < t_b \leq t_2$:

$$m[\boldsymbol{x}^*(t_a), \boldsymbol{p}^*(t_a)] \leq H[\boldsymbol{x}^*(t_a), \boldsymbol{p}^*(t_a), \boldsymbol{\alpha}^*(t_a), p_0^*], \tag{3.42}$$

which means that:

$$\frac{m[\boldsymbol{x}^*(t_b), \boldsymbol{p}^*(t_b)] - m[\boldsymbol{x}^*(t_a), \boldsymbol{p}^*(t_a)]}{t_b - t_a} \leq \frac{H[\boldsymbol{x}^*(t_b), \boldsymbol{p}^*(t_b), \boldsymbol{\alpha}^*(t_b), p_0^*] - H[\boldsymbol{x}^*(t_a), \boldsymbol{p}^*(t_a), \boldsymbol{\alpha}^*(t_a), p_0^*]}{t_b - t_a}. \tag{3.43}$$

If we narrow the time between t_a and t_b, as t_b tends to t_a from the right, this inequality becomes (where s is a general time variable):

$$\frac{d}{ds} m[\boldsymbol{x}^*(s), \boldsymbol{p}^*(s)]\Big|_{t^+} \leq \left\langle \frac{\partial H}{\partial \boldsymbol{x}}\Big|_*, \dot{\boldsymbol{x}}^*(t_a) \right\rangle + \left\langle \frac{\partial H}{\partial \boldsymbol{p}}\Big|_*, \dot{\boldsymbol{p}}^*(t_a) \right\rangle = 0, \tag{3.44}$$

which is true as $\dot{\boldsymbol{x}}^*(t_a)$ and $\dot{\boldsymbol{p}}^*(t_a)$ satisfy Equations (3.21) and (3.22) of the canonical system respectively.

By reversing the narrowing of the time interval we get the opposite result with:

$$\frac{d}{ds} m[\boldsymbol{x}^*(s), \boldsymbol{p}^*(s)]\Big|_{t^-} \geq 0. \tag{3.45}$$

This shows that:

$$\frac{d}{ds} m[\boldsymbol{x}^*(s), \boldsymbol{p}^*(s)]\Big|_{t} = 0, \tag{3.46}$$

and therefore:

$$H[\boldsymbol{x}^*(t), \boldsymbol{p}^*(t), \boldsymbol{\alpha}^*(t), p_0^*] = 0 \ \ \forall\ t \in [t_0, t^*]. \tag{3.47}$$

The last part of PMP is called the transversal condition and holds for when there is a destination target to reach that is smaller than the whole set $\mathbb{R}^n$ and made up of more than just one point. We assume that this target C is k-dimensional and includes the terminal optimal trajectory point, $\boldsymbol{x}^*(t^*)$. Next we define two new planes, $M[x^*(t^*)]$ and $N[y^*(t^*)]$ which can be seen in the representation given in Figure 3.5. $M[x^*(t^*)]$ is the tangent plane to the target and $N[y^*(t^*)]$ is the projection of $M[x^*(t^*)]$ to the hyperplane where cost is kept constant at the optimal value.

We also define $\overrightarrow{\mu}$ as the set of all vectors $\mathscr{B}\boldsymbol{\varrho}$, where:

$$\boldsymbol{\varrho} = \begin{bmatrix} -1 \\ 0 \\ 0 \\ \vdots \\ 0 \end{bmatrix}, \tag{3.48}$$

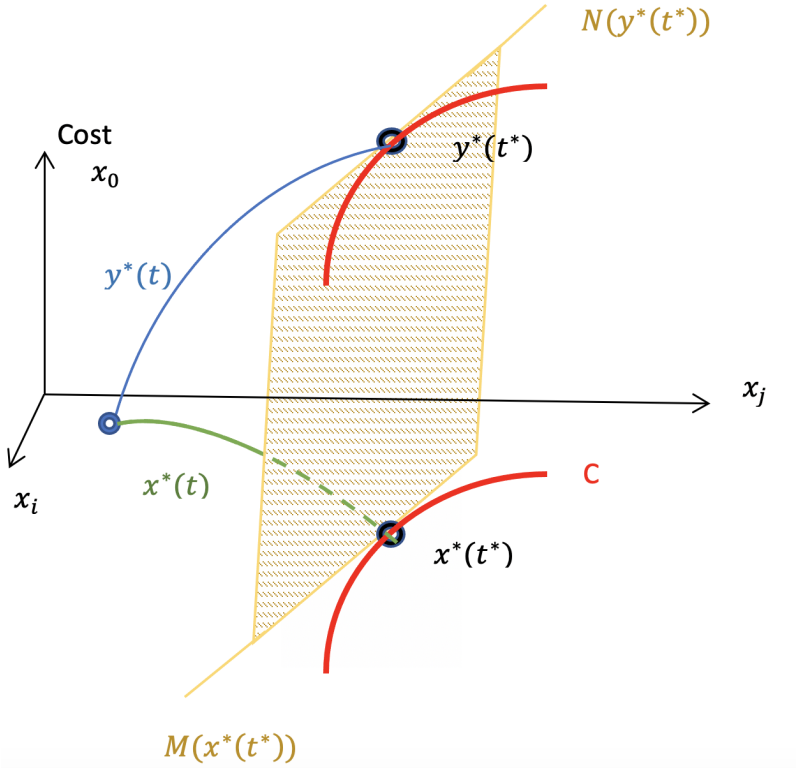

Figure 3.5: Plot of cost against position in state space, showing the planes $M[x^*(t^*)]$ and $N[y^*(t^*)]$, where the first is the tangent plane to target C and the second is the projection of this plane to the hyperplane where cost is kept constant at the optimal value.This graph has been adapted from Figure 5-35 on page 341 of Athans & Falb (1966). The two planes and the target appear as lines and curves in this image as the state space has been illustrated as 2D.

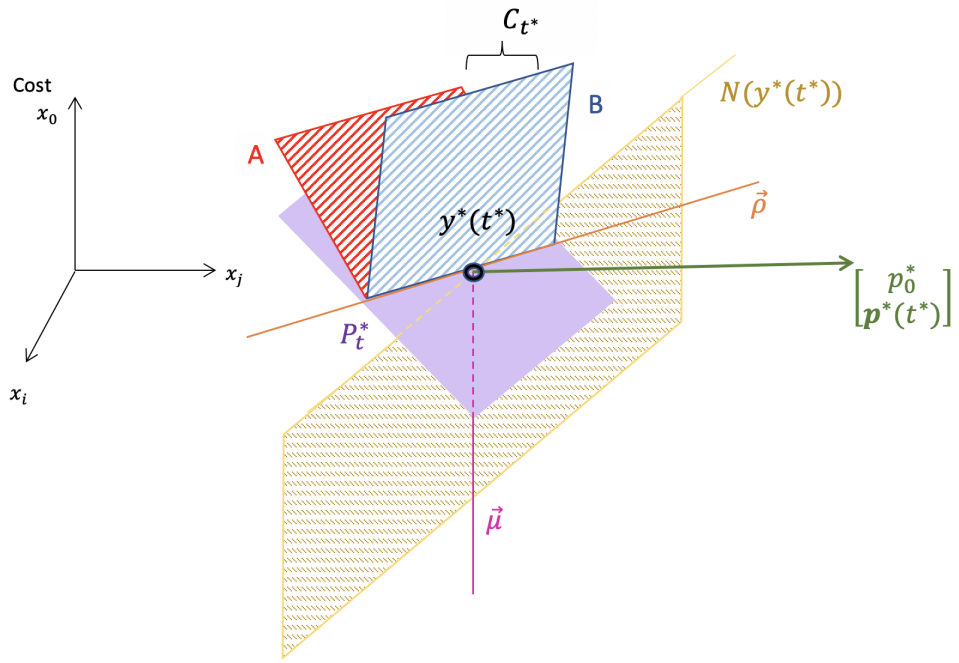

Figure 3.6: Plot of cost against position in state space, showing the plane $N[y^*(t^*)]$ (depicted as a line in this representation in 2D state space), where cost is kept constant at the optimal value, the convex cone C_{t^*}, pictured as the valley between two half hyperplanes A and B and its support P_{t^*}. The ray of decreasing cost, $\overrightarrow{\mu}$ and the ray along which A and B coincide, $\overrightarrow{\rho}$ are also shown, along with the vector of adjoints. This graph has been adapted from Figure 5-36 on page 342 of Athans & Falb (1966).

and $\mathscr{B} \geq 0$. This is a ray pointing in the direction of decreasing cost. By adding this vector to $N[y^*(t^*)]$ we create a new convex cone with a vertex at $y^*(t^*)$:

$$N[\boldsymbol{y}^*(t^*)] + \overrightarrow{\mu} = \left\{ \boldsymbol{y} : \boldsymbol{y} = \boldsymbol{y}^*(t^*) + \hat{\boldsymbol{y}} + \mathscr{B}\boldsymbol{\varrho}, \hat{\boldsymbol{y}} \in \hat{N}, \mathscr{B} \geq 0 \right\}, \tag{3.49}$$

where $\hat{N}$ is a subspace of $\mathbb{R}^{n+1}$ that is k dimensional and:

$$\hat{N} = N[\boldsymbol{y}^*(t^*)] - \boldsymbol{y}^*(t^*). \tag{3.50}$$

Planes M and N are shown in Figure 3.6, with ray $\overrightarrow{\mu}$ also pictured. The target convex cone C_{t^*} is represented by the valley between two half hyperplanes, A and B, which coincide along $\overrightarrow{\rho}$. The support to C_{t^*} is the hyperplane P_{t^*}. It is clear that the hyperplane P_{t^*}, has C_{t^*} on one side and $N[\boldsymbol{y}^*(t^*)] + \overrightarrow{\mu}$ on the other, given the proof of the first part of

PMP, Equation (3.22). This means that there must be an $n+1$ adjoint vector, such that:

$$\left\langle \begin{bmatrix} p_0^* \\ \boldsymbol{p}^*(t^*) \end{bmatrix}, \hat{\boldsymbol{y}} + \mathscr{B}\varrho \right\rangle \leq 0 \quad \forall \hat{\boldsymbol{y}} \in \hat{N}, \mathscr{B} \geq 0, \tag{3.51}$$

$$\left\langle \begin{bmatrix} p_0^* \\ \boldsymbol{p}^*(t^*) \end{bmatrix}, \delta \right\rangle \geq 0 \quad \forall \delta \in C_{t^*}. \tag{3.52}$$

When $\mathscr{B} = 0$:

$$\left\langle \begin{bmatrix} p_0^* \\ \boldsymbol{p}^*(t^*) \end{bmatrix}, \hat{\boldsymbol{y}} \right\rangle \leq 0, \tag{3.53}$$

but if $\hat{\boldsymbol{y}} \in \hat{N}$ then $-\hat{\boldsymbol{y}} \in \hat{N}$ as $\hat{N}$ is a subspace of $\mathbb{R}^{n+1}$. This means that:

$$\left\langle \begin{bmatrix} p_0^* \\ \boldsymbol{p}^*(t^*) \end{bmatrix}, \hat{\boldsymbol{y}} \right\rangle = \mathbf{0} \quad \forall \hat{\boldsymbol{y}} \in \hat{N}, \tag{3.54}$$

and thus the transversality conditions, the fourth point of the statement of PMP, is proved.

3.3.2 Direct methods

Direct methods work in the opposite order to indirect methods, with the system discretised on a mesh before it is optimised. Normally polynomial approximation is used to represent the vectors of state and control variables that define a trajectory. These variables are retrieved as the unknowns of a non-linear programming problem. This method is easy to apply, but, although solutions are optimal for the discrete problem, they have the disadvantage of not guaranteeing a minimum of the continuous problem. However, optimal trajectories obtained from software are shown to be very close to exact solutions in a range of test cases. In Chapter 6, the control trajectory is discretised into linear time steps, leading to a piecewise continuous approximation to the control function.

Mathematical theory underpinning Non-Linear Programming

In the current research the direct method is used by replacing the continuous OCP by one split into time steps that uses the conditions at each state reached to define the dynamical system for the subsequent step. Once a route is obtained, the cost of this is compared with that of an initial trajectory, which is chosen by reviewing the physical system being optimised. Improvements are made by searching for a local minimum of the cost functional. As the running cost in the current work is not analytical, derivatives are approximated by finite differences. Matlab's fmincon function takes a vector of values, $\boldsymbol{\alpha}$, and returns a vector $\boldsymbol{\alpha_{opt}}$. This second vector of values results in a local minimum for a scalar function,

$J(\boldsymbol{\alpha})$, subject to linear constraints, non-linear constraints and bounds on $\boldsymbol{\alpha}$.

An interior point algorithm is applied to the input data. This algorithm is a combination of two methods. The first uses direct linear algebra to solve the Karush–Kuhn–Tucker (KKT) conditions followed by a line search to compute step sizes. These conditions are first derivative tests to see if a solution is optimal, given that it stays within certain constraints. Further details can be found in Bertsekas (1999). If this is unable to produce an acceptable step towards a reduced cost functional, a trust region method using conjugate gradient iterations is applied. This technique is computationally expensive, but does guarantee progress towards a stationary value. By combining these two methods, an approach is created which balances robustness and efficiency. A detailed description of each part of the algorithm is given below based on Matlab function documentation and associated papers (Waltz et al., 2006; Byrd et al., 2000, 1999).

Direct step

The primary step starts with the construction of a sequence of barrier functions to approximate the original cost functional. The original problem:

$$\begin{aligned} &\min_{\boldsymbol{\alpha}} J(\boldsymbol{\alpha}) \\ &\text{subject to } h(\boldsymbol{\alpha}) = 0, g(\boldsymbol{\alpha}) \leq 0, \end{aligned} \tag{3.55}$$

is replaced by the first order optimality conditions for the barrier problem:

$$\begin{aligned} &\min_{\boldsymbol{\alpha},\mathbf{s}} J_\mu(\boldsymbol{\alpha},\mathbf{s}) = \min_{\boldsymbol{\alpha},\mathbf{s}} J(\boldsymbol{\alpha}) - \mu \sum_i \ln s_i \\ &\text{subject to } h(\boldsymbol{\alpha}) = 0,\ g(\boldsymbol{\alpha}) + \mathbf{s} = 0, \end{aligned} \tag{3.56}$$

where $\mathbf{s}$ is a vector of positive slack variables, with length equal to the number of inequality constraints, g and μ is the barrier parameter. Here μ is chosen so that as it tends to 0, J_μ tends to the minimum value of the cost functional. By removing the inequality constraints and replacing them with a series of equality constraints, the system becomes easier to solve. The Lagrangian, according to Byrd et al. (1999, 2000) is now:

$$L(\boldsymbol{\alpha}, \boldsymbol{s}, \boldsymbol{\lambda}; \mu) = J_\mu(\boldsymbol{\alpha}, \mathbf{s}) + \boldsymbol{\lambda}_g^T (g(\boldsymbol{\alpha}) + \mathbf{s}) + \boldsymbol{\lambda}_h^T h(\boldsymbol{\alpha}). \tag{3.57}$$

To take a step, we need to know what the change is in both the control vector, $\boldsymbol{\alpha}$ and in the slack variables vector, $\mathbf{s}$. To do this the KKT equations are solved using the Lagrangian

in a linearised form. This is best illustrated by using the matrix equation:

$$\begin{bmatrix} Hn & 0 & Jn_h & Jn_g \\ 0 & S\Lambda & 0 & -S \\ Jn_h & 0 & I & 0 \\ Jn_g & -S & 0 & I \end{bmatrix} \begin{bmatrix} \Delta\boldsymbol{\alpha} \\ \Delta\mathbf{s} \\ \Delta\boldsymbol{\lambda_h} \\ \Delta\boldsymbol{\lambda_g} \end{bmatrix} = - \begin{bmatrix} \frac{\partial \mathbf{J}}{\partial \boldsymbol{\alpha}} - Jn_h^T\boldsymbol{\lambda}_h - Jn_g^T\boldsymbol{\lambda}_g \\ S\boldsymbol{\lambda}_g - \mu\epsilon \\ h \\ g + s \end{bmatrix}, \tag{3.58}$$

where Hn is the Hessian of the Lagrangian of J_μ (or if no exact second order derivatives of the objective function can be found, then a quasi-Newton approximation is applied), Jn_h is the Jacobian of the linear equality constraints, Jn_g is the Jacobian of the linear inequality constraints, S is a diagonal matrix with the elements of $\mathbf{s}$ along its diagonal, $\boldsymbol{\lambda_h}$ and $\boldsymbol{\lambda_g}$ are the vectors of Lagrange multipliers associated with h and g respectively, Λ is a diagonal matrix with the $\boldsymbol{\lambda_g}$ elements along its diagonal, I is an identity matrix of the same size as Jn_h and ϵ is a vector of the same length as g. If the Hessian is positive definite, then the problem is solved to find $(\Delta\boldsymbol{\alpha}, \Delta\mathbf{s})$. The step is now tested using the merit function:

$$J_\mu(\boldsymbol{\alpha}, \mathbf{s}) + v\|(h(\boldsymbol{\alpha}, g(\boldsymbol{\alpha}) + \mathbf{s})\|, \tag{3.59}$$

where v is a scalar that can increase at each iteration to try to force the solution towards feasibility. If this merit function is smaller than it was on the previous iteration, then the step is taken.

Now the new control vector can be tested to see if the step tolerance, or optimality tolerance are met. Step tolerance is a relative bound comparing $\boldsymbol{\alpha}$ value to the previous one. Optimality tolerance measures how the objective function is varying in all feasible directions. If it is non-decreasing to within this tolerance, then the optimisation has been completed. These tolerances both show that a local minimum may have been found. If, however, the Hessian is not positive definite or the Jacobians are rank deficient, then the direct step algorithm is abandoned and instead a trust region step is tried.

Trust region step

The trust region step will always guarantee progress towards a stationary value, but is computationally too expensive to use exclusively. For each value α a neighbourhood N around it is considered. Within this region, a function $q(\alpha)$ is found which is a good approximation of $J(\alpha)$, but far less complicated. This then gives a sub-problem of finding a trial step s which minimises q. If $J(\alpha + s) < J(\alpha)$ then the control is updated to be $(\alpha + s)$. If this is not the case, then the trust region is shrunk until a suitable step is found.

N is found by using a preconditioned gradient process (Byrd et al., 1999). It is two dimensional in fmincon and spanned by a vector in the direction of the gradient (s_1) and either the direction of negative curvature (s_2 for $s_2^T Hns_2 < 0$) or an approximate Newton direction found by solving:

$$Hn.s_2 = -g. \tag{3.60}$$

The trust region sub-problem can be expressed as:

$$\min\{\frac{1}{2}s^T Hns + s^T \frac{\partial J}{\partial \alpha} : \|Ds\| \leq \Psi\}, \tag{3.61}$$

where D is a diagonal scaling matrix and Ψ is a positive scalar representing the trust region dimension. Algorithms to solve this sub-problem, rely on finding all of the eigenvalues of H_n and applying a Newton process to solve:

$$\frac{1}{\Psi} = \frac{1}{\|s\|}. \tag{3.62}$$

Again, once a step length is found, the merit function is used to see if this step is valid. If it is, then the new controls can be retrieved. If the tolerances mentioned previously are satisfied, then the iterations stop. If not a new iteration is attempted.

Iterations

Often the function is evaluated hundreds of times before a step is found that satisfies the merit function. Only once this happens can the new headings and airspeeds be tested using the tolerances. If tolerances are not met, then the next iteration is started, changing each α_i, by another small step. A limit is set to stop the search if a set number of function evaluations is exceeded.

3.3.3 Dynamic programming

Finally there is the method of dynamic programming, with the major advantage that the whole of the state space is searched for the optimal solution, meaning that a global optimum for the continuous problem is guaranteed. The method relies on the HJB equation which is based on the principle of optimality. This states that an optimal path from b to c, will form part of the optimal path from a to c, where b is between a and c. By solving the HJB equation we can use the value function and the HJB PDE to find an optimal feedback control, from whence the states and times can be retrieved. Dynamic programming requires spatial and time discretisation. Unfortunately this method suffers from the curse of dimensionality. This means that an increase in the number of states

present in a given problem, leads to an rapid growth in the number of calculations that must be completed. As well as guaranteeing a globally optimal solution this method does have the advantage that certain general cases can be solved by applying standard solutions.

In Chapter 7 dynamic programming is used to find minimum time cruise routes across the North Atlantic. The method is then adapted to find fuel minimal routes. In both cases mass is altered at each time step. To avoid the implications of the curse of dimensionality, the HJB equation is solved over a grid of state variables using a long time step. This information is then used to shrink the number of grid points to those directly around this initial route and the time step is reduced. The fuel burn results obtained in this way for the cruise phase of a flight can be compared with those of actual flights made between the 1st December, 2019 and the 29th February, 2020 as long as the actual flights are limited to a fixed cruise altitude.

Mathematical theory underpinning Dynamic Programming

Results are found for all times, rather than setting a fixed time period. In this way the minimum time is obtained for a cost functional based purely on time of flight. For fuel minimal routes the time of flight is retrieved from the number of time steps needed to reach the target in the final trajectory. The theory from this section can be found in the work of Bressan & Piccoli (2007), Falcone & Ferretti (2014), Bardi & Capuzzo-Dolcetta (1997) and Briani (2021). The notation used here follows Falcone & Ferretti (2014).
For the controlled system, the solutions of:

$$\begin{aligned} \dot{y} &= f(y(t), \alpha(t)) \quad t \in \mathbb{R}^+, \\ y(0) &= x, \end{aligned} \tag{3.63}$$

are the possible trajectories. The α values are the controls. Two key assumptions must be made here:

- The control assumption states that the set $A \subset \mathbb{R}^M$ must be closed and bounded. The set of admissible controls is defined as:

$$\mathcal{A} := \{\alpha : \mathbb{R}^+ \to A : \alpha(.) \text{ is measurable}\}, \tag{3.64}$$

- The dynamics assumption states that $f : \mathbb{R}^N \times A \to \mathbb{R}^N$ is continuous and must be both Lipschitz continuous:

$$\exists\ L_f > 0 : \quad |f(x,\alpha) - f(y,\alpha)| \leq L_f |x - y| \quad \forall x, y \in \mathbb{R}^N,$$

and bounded:

$$\exists\ M_f > 0: \quad |f(x,\alpha)| \leq M_f \quad \forall \alpha \in \mathcal{A}. \tag{3.65}$$

These assumptions show that a unique Lipschitz function solution of Equation (3.63) exists. This solution is a trajectory starting at state $x \in \mathbb{R}^N$ and controlled by $\alpha \in \mathcal{A}$:

$$y_x(.,\alpha) \in \mathrm{Lip}(\mathbb{R}^+;\mathbb{R}^N). \tag{3.66}$$

The cost functional to be minimised is:

$$J(x,\alpha) = \int_0^\infty l(y_x(t,\alpha),\alpha(t))e^{-\lambda t}\mathrm{dt}, \tag{3.67}$$

where l is the running cost and λ is a positive constant known as the discount factor. The running cost is assumed to be continuously bounded and continuous, so for a non-decreasing continuous function $\omega_l : \mathbb{R}^+ \to \mathbb{R}^+$ with $\omega_l(0) = 0$ and a constant M_l:

$$\begin{aligned} &|l(x,\alpha) - l(y,\alpha)| \leq \omega_l|x-y|, \\ &|l(x,\alpha)| \leq M_l \quad \forall\ x,y \in \mathbb{R}^N \quad \forall\ \alpha \in \mathcal{A}. \end{aligned} \tag{3.68}$$

To find the array of controls which leads to a minimum cost, the value function, $\nu : \mathbb{R}^N \to \mathbb{R}$, is introduced, with:

$$\nu(x) = \inf_{\alpha\in\mathcal{A}} J(x,\alpha). \tag{3.69}$$

If we accept assumptions Equations (3.64) and (3.65), Bellman's Dynamic Programming Principle (DPP) for infinite horizon problems is given in Proposition 8.2 of Falcone & Ferretti (2014) as:

$$\begin{aligned} &\forall\ x \in \mathbb{R}^d \text{ and } \tau > 0, \\ &\nu(x) = \inf_{\alpha\in\mathcal{A}} \left\{ \int_0^\tau g(y_x(s;\alpha),\alpha(s))e^{-\lambda s}\mathrm{ds} + e^{-\lambda\tau}\nu(y_x(\tau;\alpha)) \right\}. \end{aligned} \tag{3.70}$$

The idea here is that a minimum cost can be found by choosing an arbitrary control and letting the dynamical system evolve for a set period of time. The cost for the trajectory generated is added to the cost for continuing past this time to reach the target, using the best possible control. This sum is repeated for controls across the control set and then the minimum sum value is chosen.

The DPP is used to characterise the value function in terms of a non-linear partial differential equation, the HJB equation. As before we assume Equations (3.64), (3.65) and (3.68) all hold, then we can say that the value function, Equation (3.69), is a continuous viscosity solution of:

$$\lambda\nu(x) + H(x,D\nu) = 0 \quad x \in \mathbb{R}^N, \tag{3.71}$$

where $H : \mathbb{R}^N \times \mathbb{R}^N \to \mathbb{R}$ is the Hamiltonian defined as:

$$H(x,p) := \sup_{\alpha \in \mathcal{A}} \{-p.f(x,\alpha) - l(x,\alpha)\}. \tag{3.72}$$

It can be proved that $\nu(x)$ is the unique, viscosity solution of the HJB. Once the value function has been found, it can be used to synthesise the feedback control and from this the states belonging to the optimal trajectory can be obtained. The optimal control:

$$\alpha^* : [0,\infty) \to \mathcal{A}, \tag{3.73}$$

must be found which minimises the exponentially discounted cost functional:

$$J(y,\alpha) \doteq \int_0^{+\infty} e^{-\lambda t} h(x(t),\alpha(t))\mathrm{dt}, \tag{3.74}$$

where $\lambda > 0$ is a constant. The functions f and h are assumed to be bounded and Lipschitz continuous. It is also assumed that trajectory $t \to x(t;y,\alpha)$ is well defined at all $t \geq 0$. The value function is defined as:

$$\nu(y) \doteq \inf_{\alpha(.) \in \mathcal{A}} J(y,\alpha). \tag{3.75}$$

By the Dynamic Programming Principle (Theorem 8.8.2) for every $\tau > 0$ and $y \in \mathbb{R}^n$:

$$\nu(y) = \inf_{\alpha(.)} \left\{ \int_0^{\tau} e^{-\lambda t} h(x(t;y,\alpha),\alpha(t))\mathrm{dt} + e^{-\lambda \tau} \nu(x(\tau;y,\alpha)) \right\}. \tag{3.76}$$

Theorem 8.8.3 of Falcone & Ferretti (2014) states that ν is therefore the viscosity solution of the HJB equation:

$$-[-\lambda\nu + H(x,D\nu)] = 0 \quad x \in \mathbb{R}^n, \tag{3.77}$$

where the Hamiltonian is defined as:

$$H(x,p) \doteq \min_{\omega \in \mathcal{A}} \{f(x,\omega).p + h(x,\omega)\}. \tag{3.78}$$

The vector p contains the co-states as defined by PMP. This is analagous to Equation (3.72).

The minimum time problem

To minimise time to a set target, $C \subset \mathbb{R}^N$ we need to make an assumption that the target is closed with a compact boundary ∂C:

$$t_x(\alpha) := \begin{cases} +\infty & \text{if} \{t : y_x(t,\alpha) \in C\} = \emptyset, \\ \min\{t : y_x(t,\alpha) \in C\} & \text{otherwise.} \end{cases} \tag{3.79}$$

From here the minimal time function, $T : \mathbb{R}^N \to \mathbb{R}$ is used as the cost functional and can be expressed as:

$$T(x) := \inf_{\alpha \in \mathcal{A}} t_x(\alpha). \tag{3.80}$$

The set of all starting points from whence a trajectory can be found which reaches the target in finite time is called the reachable set, $\mathcal{R}$. This is defined as:

$$\mathcal{R} := \{x \in \mathbb{R}^N : \exists\ \alpha \in \mathcal{A} \text{ and } t \in \mathbb{R}^+ : y_x(t, \alpha) \in C\}. \tag{3.81}$$

Assuming Equations (3.64) and (3.65), Bellman's Dynamic Programming Principle for the minimum time problem says that for all $x \in \mathcal{R}$ and all $\tau \in (0, T(x))$:

$$T(x) = \inf_{\alpha \in \mathcal{A}} \{\tau + T(y_x(\tau; \alpha))\}. \tag{3.82}$$

The proof is given in Briani (2021).

The HJB equation for the minimum time problem states that given Equations (3.64), (3.65) and (3.68), the value function, Equation (3.80), is a viscosity solution of:

$$H(x, DT) = 1 \qquad \text{in } \mathcal{R} \setminus C, \tag{3.83}$$

where $H : \mathbb{R}^N \times \mathbb{R}^N \to \mathbb{R}$, the Hamiltonian, is given as:

$$H(x, p) := \sup_{\alpha \in \mathcal{A}} \{-p.f(x, \alpha)\}. \tag{3.84}$$

It is assumed that T is a continuous function in $\mathcal{R}$, bounded below and satisfying:

$$\begin{aligned} &T(x) = 0 && \text{on } \partial C, \\ &T(x) \to +\infty && \text{as } x \to x_0 \in \partial\mathcal{R}. \end{aligned}$$

To avoid having an infinite boundary condition, the Kružkow transform can now be applied. Let:

$$\nu(x) := \begin{cases} 1 & T(x) = +\infty, \\ 1 - e^{-T(x)} & T(x) < +\infty. \end{cases} \tag{3.85}$$

The value function, $\nu(x)$ is now that of an infinite horizon control problem with a running cost of 1 and a discount factor of 1, so that the cost functional is:

$$J_x(\alpha) = \int_0^{t_x(\alpha)} e^{-s} \mathrm{ds}, \tag{3.86}$$

giving:

$$\nu(x) = \inf_{\alpha \in \mathcal{A}} J_x(\alpha), \tag{3.87}$$

which is the unique solution of the HJB equation:

$$\nu(x) + \sup_{\alpha \in \mathcal{A}}\{-D\nu \cdot f(x,\alpha) - 1\} = 0 \quad x \in \mathbb{R}^N \setminus C, \tag{3.88}$$

$$\nu(x) = 0 \qquad \forall x \in C. \tag{3.89}$$

The optimal feedback control is found as:

$$\alpha^* = \arg\min_{\alpha \in \mathcal{A}}\{-D\nu \cdot f(x,\alpha) - 1\}. \tag{3.90}$$

Value iteration algorithm

Both the time minimal and infinite horizon problems can be solved numerically by using a Semi-Lagrangian scheme. This involves using a time discretisation which allows derivatives in the dynamical system to be approximated by an Euler scheme and the continuous cost functional by a discrete scheme, such as rectangular quadrature. So the Discrete Dynamic Programming Principle, characterises the time discrete value function as part of a semi-discrete approximation scheme. Next a spatial discretisation is applied, which replaces the time discretised value function with a polynomial interpolation. Thus a finite dimensional problem is obtained, which can be solved.

The method described here follows Falcone (1997). Consider a minimum time function:

$$T(x) = \inf_{\alpha \in \mathcal{A}}\{t > 0 : y_x(t,\alpha) \in C\}. \tag{3.91}$$

We will need to use both a Discrete Dynamic Programming Principle and a discrete version of the Kružkow transform to go from a partial differential equation to a functional equation. The first step is to apply a time discretisation with step size Δt. The derivative in the dynamical system is now approximated by an Euler scheme, so that $\dot{y} = f(y,\alpha)$ is replaced by its discrete approximation:

$$y_{k+1} = y_k + \Delta t f(y,\alpha). \tag{3.92}$$

The Euler scheme is used here, as where there are many evaluations of a formula, using a simple numerical method allows for considerable reductions in computation speed. Sensitivity analysis for time step length has been applied whenever the Euler scheme is used in this thesis, to ensure that an acceptable balance is maintained between stability and any errors caused by data truncation. The Kružkow Transform can be expressed as:

$$\hat{\nu} := \begin{cases} 1 & T(x) = +\infty, \\ 1 - e^{-\Delta t N(x)} & \text{else}, \end{cases} \tag{3.93}$$

where $N(x)$ is the number of timesteps to reach a time of $T(x)$. The time discretised value function can now be characterised as the solution of:

$$\hat{\nu}(x) = \min_{a \in \mathcal{A}} \{e^{-\Delta t} \nu^{\Delta t}(y + \Delta t f(y, a))\} + 1 - e^{-\Delta t} \qquad x \in \mathbb{R}^d \setminus C, \tag{3.94}$$

$$\hat{\nu}(x) = 0 \qquad x \in C. \tag{3.95}$$

It is worth noting that as $\Delta t \to 0$ so the solution of the discrete problem, $\hat{\nu}(x)$, tends to $\nu(x)$, the solution of the continuous problem.

Now a spatial discretisation is required to allow solution of the functional equation. A grid is set up to cover $\Omega \subset \mathbb{R}^d$ which contains all points we will consider for an admissible trajectory. For the moment we will assume that the grid has equal spacing of m in all dimensions. As $m \to 0$ so the solution of the time and space discretised problem $\hat{\nu}_m(x)$ tends to the solution of the time discretised problem, $\hat{\nu}(x)$. The value of $\hat{\nu}(x)$ is approximated by an array of values. For a two dimensional grid this would give a value of $\hat{\nu}_{i,j}$ at node (i, j). For simplicity, however, we number the nodes to appear in a vector $[\hat{\nu}]$. Now Equation (3.94) can be re-written to reflect the spatial discretisation:

$$[\hat{\nu}]_i = \min_{a \in \mathcal{A}} \{e^{-\Delta t} I[\hat{\nu}]_i(y_i + \Delta t f(y_i, a)\} + 1 - e^{-\Delta t}, \tag{3.96}$$

where the I represents a linear interpolation of the value function. As we obtain the value function at gridpoints, but $y_i + \Delta t f(y_i, a)$ may not get to a gridpoint, interpolation is necessary. This is now a fully discrete non-linear equation for $[\hat{\nu}]$, as long as we replace the continuous control function with sets of discrete control values. We can think of the equation as $\hat{\nu} = \omega(\hat{\nu})$ where ω is an operator. By Banach's fixed point theorem, as long as ω is a contraction and so satisfies:

$$||\omega(\hat{\nu}_1) - \omega(\hat{\nu}_2)|| \leq \delta ||\hat{\nu}_1 - \hat{\nu}_2|| \quad \text{where} \quad \delta < 1, \tag{3.97}$$

then a solution to the equation $\hat{\nu} = \omega(\hat{\nu})$ can be found by iterating $\hat{\nu}^{n+1} = \omega(\hat{\nu}^n)$ starting from any $\hat{\nu}^0$. Now this is a fully discretised scheme with a fixed point iteration. An algorithm for the method of solution is shown as Algorithm 1 and a flowchart of its application is given in Figure 3.7.

Once a value function and the corresponding controls have been found, the off-grid controls can be approximated by interpolation and then the trajectory traced using a numerical scheme, such as an Euler forward step method, to give states at each time step.

Algorithm 1 Algorithm to construct the value function

1: Estimate the starting value at each gridpoint. (Here we set the gridpoint values to 1, except for those in the target which we set to 0.)

2: Construct an interpolant $I[\hat{\nu}^m]$

3: Solve

$$[\hat{\nu}]_i^{m+1} = \min_{a \in \mathcal{A}} \{e^{-\Delta t} I[\hat{\nu}]_i^m (y_i + \Delta t f(y_i, a))\} + 1 - e^{-\Delta t} \tag{3.98}$$

for each combination of control variables and for all gridpoints.

4: If a particular set of controls takes the arrival point, $y_i + \Delta t f(y_i, a)$, outside of the specified grid, then set the value for this point to 1.

5: Find the minimum $[\hat{\nu}]_i^{m+1}$ values and the controls used to obtain them and store these.

6: Measure $||\hat{\nu}^{m+1} - \hat{\nu}^m||$. If it is below a pre-set tolerance, stop, if not go to next step.

7: Update so $\hat{\nu}^m = \hat{\nu}^{m+1}$

8: Reset target values to 0.

9: Go back to step 2.

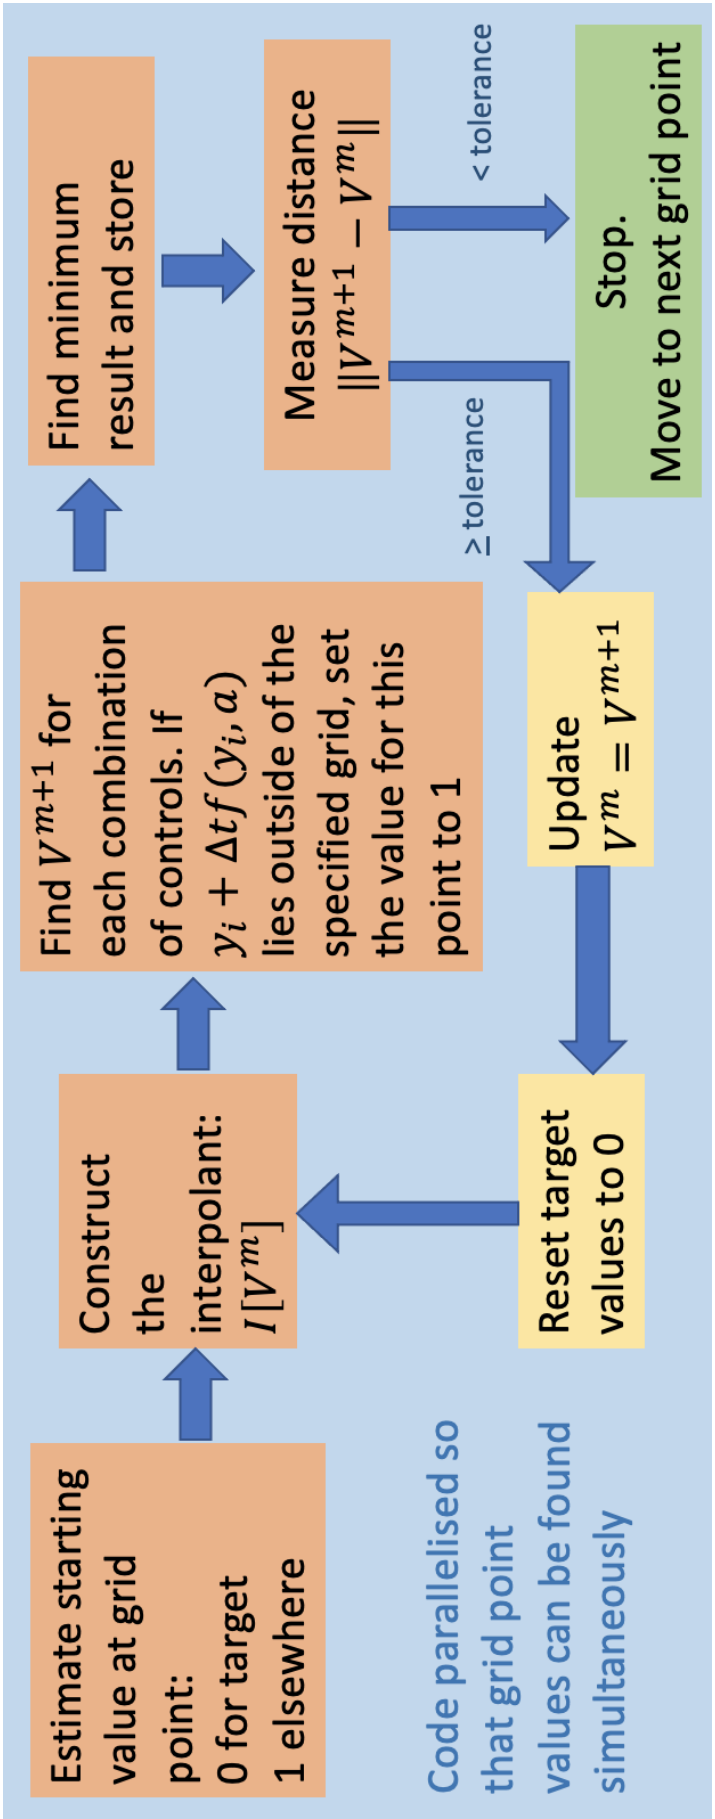

Figure 3.7: Flow chart to show how Algorithm 1 is used.

3.4 Fuel burn calculations

In the vast majority of research relating fuel burn to aircraft trajectories, EUROCONTROL's BADA method is used to model aircraft fuel flow. This is true of recent papers looking for climate optimised trajectories (Yamashita et al., 2020, 2021) as well as previous research into fuel optimal routing (García-Heras et al., 2014; Wickramasinghe et al., 2012; Soler et al., 2020). However, in this thesis a new method of calculating the fuel burn rate of commercial passenger aircraft is used, based on Poll & Schumann (2021a,b). Characteristic parameters for different models of aircraft are stated that allow the derivation of the fuel burn rate given current airspeed, temperature, mass and altitude. Certain values from the International Standard Atmosphere (ISA) and other standard constants are also used.

All fuel burn rate calculations made at any point in an aircraft's trajectory in the following chapters are from a Matlab code created from the formulae derived from these papers.

For the majority of the simulations, a Boeing 777-200ER aircraft is assumed to be used, as this is the model currently most frequently flown between LHR and JFK (Flightradar24, 2020). Where comparisons are made between actual and simulated flights in Chapter 7 the parameters for the aircraft used for that specific actual flight are used.

For the Boeing 777-200ER the aircraft specific parameters used are given in Table 3.1 (values given correct to appropriate accuracy as specified in Poll & Schumann (2021a,b)). A further table detailing aircraft specific parameters for all models of aircraft considered in this thesis is given in Appendix C. Some of the symbols and variable names used here are, for the purposes of describing the fuel burn rate calculation, different from those in the rest of the thesis.

Table 3.1: Table of aircraft specific constants for the Boeing 777-200ER as given in Poll & Schumann (2021b).

τ	0.122	coefficient in Equations (3.117), (3.118)
ψ_1	0.211	coefficient in Equation (3.117)
ψ_2	8.09	coefficient in Equation (3.118)
ψ_4	0.8106	approx. optimum Mach number with constant Reynolds number
ψ_5	1.266×10^8	coefficient in Equation (3.115)
ψ_6	0.631	coefficient in Equation (3.120)
$MTOM$	286900 kg	maximum take-off mass of aircraft
a	0.0269	skin friction parameter 1
b	0.14	skin friction parameter 2

The universal constants are given in Table 3.2.

Table 3.2: Table showing universal constants for the fuel burn function.

LCV	4.3×10^7 J kg^{-1}	lower calorific value of aircraft fuel
$(p_{TP})_{ISA}$	226.318 hPa	pressure at tropopause in hPa, ISA
T_0	288.15 K	sea level static temperature, given ISA
$(T_{TP})_{ISA}$	216.65 K	static temperature at the tropopause, ISA
$\mathfrak{R}$	287.05 J (kg K)$^{-1}$	air constant in ideal gas law
γ	1.4	ratio of constant pressure to volume specific heat for air

Other quantities found from the inputs and constants that are used subsequently in the derivation of the fuel burn rate are given in Table 3.3.

Table 3.3: Other quantities that must be calculated from inputs and constants.

a_∞	speed of sound at given altitude
$(a_{TP})_{ISA}$	speed of sound at the tropopause, given ISA
A	function dependent on ω derived in Equation (3.111) used in Equation (3.122)
B	function dependent on ω derived in Equation (3.112) used in Equation (3.122)
C_F^{ac}	mean skin friction coefficient of the aircraft
C_L	lift coefficient for current state of aircraft
$(C_L)_B$	lift coefficient for aircraft at best fuel efficiency for given Mach number
$(C_L)_B check$	C_{L_B} rounded to 1 dp for validity test
$(C_L)_{ratio}$	ratio of C_L to $(C_L)_B$
$\frac{\mathrm{d}m_f}{\mathrm{d}t}$	rate of fuel burn in kg s^{-1}
f_1	function dependent on ω derived in Equation (3.113) used in Equation (3.117)
f_2	function dependent on ω derived in Equation (3.114) used in Equation (3.118)
FL	Flight Level, given as one hundredth of altitude of aircraft in feet
h	altitude of aircraft in m
m	mass of aircraft at current time t in kg
m_{ratio}	fraction of maximum take-off mass remaining at current time t
M_∞	Mach number at current airspeed and Flight Level
p_∞	static pressure at current position in hPa
R_{ac}	Reynolds number for given inputs
T_∞	static temperature at given altitude in K
$(T_\infty)_{ISA}$	ISA static temperature at given altitude in K
V_∞	airspeed of aircraft in m s^{-1}
η_0	overall efficiency of propulsion system
$\eta_0 \frac{L}{D}$	maximum value of $\eta_0\times$ aircraft lift to drag ratio
$(\eta_0 \frac{L}{D})_B$	$\eta_0\times$ aircraft lift to drag ratio at best fuel efficiency for current Mach number
f_0	ratio between $\eta_0 \frac{L}{D}$ and $(\eta_0 \frac{L}{D})_B$
μ_∞	current value of dynamic viscosity in Pa s^{-1}
$(\mu_{TP})_{ISA}$	value of dynamic viscosity in Pa s^{-1}, given ISA
ϕ	atmospheric parameter
χ	ratio of pressure at tropopause to recorded pressure level
ω	ratio of M_∞ to ψ_4

3.4.1 Method to calculate fuel burn rate

The fuel burn rate is given in Poll & Schumann (2021a,b) as:

$$\frac{\mathrm{d}m_f}{\mathrm{d}t} = \frac{W \times V_\infty}{\eta_0 \frac{L}{D} \times LCV}, \tag{3.99}$$

where W is the current weight of the aircraft in N. Inputs to the formula are V_∞ (the airspeed), T_∞ (the temperature), p_∞ (the static pressure) and m (the mass). From these and the values in the previous section, the fuel burn rate, $\frac{\mathrm{d}m_f}{\mathrm{d}t}$, can be obtained following the working in Poll & Schumann (2021a,b). As these two papers provide full working and derivation for the fuel burn function, what follows here is a short summary.
First χ is calculated:

$$\chi = \frac{(p_{TP})_{ISA}}{p_\infty}. \tag{3.100}$$

The flight level can be found based on this value:

$$FL = \begin{cases} 1454.42(1 - 0.751865\chi^{-0.19026}) & \text{if } \chi < 1, \\ 360.8924 + 208.058 \ln \chi & \text{if } \chi \geq 1. \end{cases} \tag{3.101}$$

The altitude in metres, h, is obtained by multiplying the Flight Level (FL) by a conversion factor. This value is then used along with the sea level ISA static temperature (in degrees Kelvin), to find $(T_\infty)_{ISA}$:

$$h = 30.48FL, \tag{3.102}$$

$$(T_\infty)_{ISA} = \begin{cases} T_0 - 0.0065h & \text{if } h < 11000, \\ 216.65 & \text{if } h \geq 11000. \end{cases} \tag{3.103}$$

The speed of sound and dynamic viscosity at the tropopause, as well as each of these values at the current position, are needed to calculate ϕ:

$$(a_{TP})_{ISA} = \sqrt{\gamma \mathfrak{R} (T_{TP})_{ISA}}, \tag{3.104}$$

$$(\mu_{TP})_{ISA} = 0.000001458 \frac{((T_{TP})_{ISA})^{1.5}}{(T_{TP})_{ISA} + 110.4}, \tag{3.105}$$

$$a_\infty = \sqrt{\gamma \mathfrak{R} T_\infty}, \tag{3.106}$$

$$\mu_\infty = 0.000001458 \frac{(T_\infty)^{1.5}}{T_\infty + 110.4}, \tag{3.107}$$

$$\phi = \frac{a_\infty \mu_\infty}{(a_{TP})_{ISA} (\mu_{TP})_{ISA}}. \tag{3.108}$$

The speed of sound at the current position can now be used to find the corresponding Mach number:

$$M_\infty = \frac{V_\infty}{a_\infty}, \tag{3.109}$$

which in turn allows the value of ω to be calculated:

$$\omega = \frac{M_\infty}{\psi_4}. \tag{3.110}$$

The first test of input validity (as specified in Poll & Schumann (2021a,b)) can now be applied, $0.8 < \omega < 1.08$. If a value outside of this range is obtained an unrealistic fuel burn rate will result, so in the Matlab code the output is set to NaN (not a number). For ω values within the required range, values of A, B, f_1 and f_2 can now be calculated:

$$A = \begin{cases} -2.6 & \text{if } 0.8 < \omega < 0.975, \\ -(2.6 + 120(\omega - 0.975)^2 & \text{if } 0.975 \leq \omega < 1.08, \end{cases} \tag{3.111}$$

$$B = \begin{cases} -2.6 & \text{if } 0.8 < \omega < 0.975, \\ -(2.6 + 270(\omega - 0.975)^2 & \text{if } 0.975 \leq \omega < 1.08, \end{cases} \tag{3.112}$$

$$f_1 = \begin{cases} & 1 - 6(\omega - 1)^2 - 15(\omega - 1)^3 \\ \text{if} & 0.8 < \omega < 0.99, \\ & 1 - 5.8965(\omega - 1)^2 + 0.36024(\omega - 1)^3 - 31.684(\omega - 1)^4 - 53313(\omega - 1)^5 \\ \text{if} & 0.99 \leq \omega < 1.08, \end{cases} \tag{3.113}$$

$$f_2 = 1.05 - 14.8(\omega - 0.8)^3 + 116.75(\omega - 0.8)^4 - 370(\omega - 0.8)^5. \tag{3.114}$$

Calculation of the Reynolds number also requires the ω value and once found it can be used to obtain the mean skin-friction coefficient for the aircraft:

$$R_{ac} = \frac{\psi_5}{\phi} \times \frac{\omega}{\chi}, \tag{3.115}$$

$$C_F^{ac} = \frac{a}{R_{ac}^b}. \tag{3.116}$$

The best fuel efficiency for the Mach and Reynolds numbers calculated and the corresponding lift coefficient are now found:

$$(\eta_0 \frac{L}{D})_B = f_1 \psi_1 \left(\frac{1}{C_F^{ac}} \right)^{\left(\frac{1+\tau}{2}\right)}, \tag{3.117}$$

$$(C_L)_B = f_2 \psi_2 (C_F^{ac})^{\left(\frac{1-\tau}{2}\right)}. \tag{3.118}$$

The fraction of MTOM remaining at the current time is calculated next, in order to obtain the lift coefficient for the current state of the aircraft and the ratio between this and the $(C_L)_B$ value.

$$m_{ratio} = m/MTOM, \tag{3.119}$$

$$C_L = \frac{m_{ratio} \psi_6 \chi}{\omega^2}, \tag{3.120}$$

$$(C_L)_{ratio} = \frac{C_L}{(C_L)_B}. \tag{3.121}$$

This lift coefficient ratio provides the second input validity check. First the value is rounded to one decimal place and renamed $C_{Lratiocheck}$, before it is tested. Valid inputs will give $0.5 \leq (C_L)_{ratiocheck} < 1.3$, as specified in Poll & Schumann (2021a,b). Any answer obtained outside of this range will result in the fuel burn rate in the Matlab code being returned as NaN, to avoid erroneous values being used in later calculations.

From here the propulsion system efficiency ratio can be found using the unrounded value, C_{Lratio}, which leads to the current state value of propulsion system efficiency:

$$f_0 = 1 + \frac{A}{2}((C_L)_{ratio} - 1)^2 + \frac{B}{6}((C_L)_{ratio} - 1)^3, \tag{3.122}$$

$$\eta_0 \frac{L}{D} = (f_0 \times (\eta_0 \frac{L}{D})_B. \tag{3.123}$$

Finally, if all validity tests have been passed, the fuel burn rate in kg s^{-1} flown, with the aircraft and atmospheric conditions in their current state, is given by:

$$\frac{\mathrm{d}m_f}{\mathrm{d}t} = \frac{W \times V_\infty}{\eta_0 \frac{L}{D} \times LCV}, \tag{3.124}$$

where W is the current weight of the aircraft in N.

3.4.2 Finding the Take-off Mass

In order to track mass change across a trajectory, as is necessary in the third stage of the optimal control model, it is important to know the mass of the aircraft as it begins its cruise phase. This is shown in Poll & Schumann (2021a) to be 97.5% of the take-off mass. To find the take-off mass (TOM) requires the use of the following equation (also from (Poll & Schumann, 2021a)):

$$TOM \approx \frac{0.7 \times MZFM + 0.3 \times OEM}{\exp(-(0.014 + 1.015 \times \frac{g \times R_{t(air)}}{\eta_0 \frac{L}{D} \times LCV})) - 0.05}, \tag{3.125}$$

where $MZFM$ is the maximum permitted zero fuel mass, OEM is the operational empty mass, $R_{t(air)}$ is the air distance flown by the aircraft and all other variables and constants are as given in Tables 3.2 and 3.3. The values of $MZFM$ and OEM can be found in the aircraft characteristics documentation provided by both Boeing and Airbus (Airbus, 2021; Boeing, 2021).

For the fixed-time flights at a constant airspeed of 240 m s^{-1}, $R_{t(air)}$ can be taken as the scheduled time multiplied by the airspeed, thus for flights west $R_{t(air)} = 6\ 960$ km and for the shorter eastbound flights $R_{t(air)} = 5\ 280$ km.

To assess how sensitive to the $R_{t(air)}$ value the take-off mass is, the partial derivative of TOM with $R_{t(air)}$ was found:

$$\frac{\partial TOM}{\partial R_{t(air)}} = \frac{JK \exp -(0.014 + K \times R_{t(air)})}{(\exp{(-(0.014 + K \times R_{t(air)}))} - 0.05)^2}, \tag{3.126}$$

where:

$$J = 0.7 \times MZFM + 0.3 \times OEM, \tag{3.127}$$

$$K = \frac{1.015 \times g}{(\eta_o \frac{L}{D})_o \times LCV}. \tag{3.128}$$

To evaluate this the $(\eta_o \frac{L}{D})_o$ value is needed. It is found using:

$$(\eta_o \frac{L}{D})_o = \psi_1 \left(\frac{1}{C_F^{ac}} \right)^{\frac{1+\tau}{2}}, \tag{3.129}$$

where the ψ_1, τ and C_F^{ac} values are characteristic of a particular aircraft and C_F^{ac} can be calculated as shown in equation (3.116).

The rate of change of TOM with $R_{t(air)}$ was found for the range of air distances given across the winter period by the time optimal westbound model used in Chapter 4 and also its value for the wind-free air distance between the airports, the great circle distance. By comparing these answers for the different $R_{t(air)}$ values the sensitivity of the take-off mass to the air distance can be seen. The fact that the gradient of the rate of change is only $3{\times}10^{-10}$ kg m^{-1}, shows that the sensitivity is very small. The air distance range was then extended so that its first value is the assumed air distance for eastbound flights used in Chapter 6 and the final value is the assumed air distance for westbound flights used in Chapter 6. Both the variation of TOM with $R_{t(air)}$ and the rate of change of TOM with $R_{t(air)}$ are displayed in Figure 3.8.

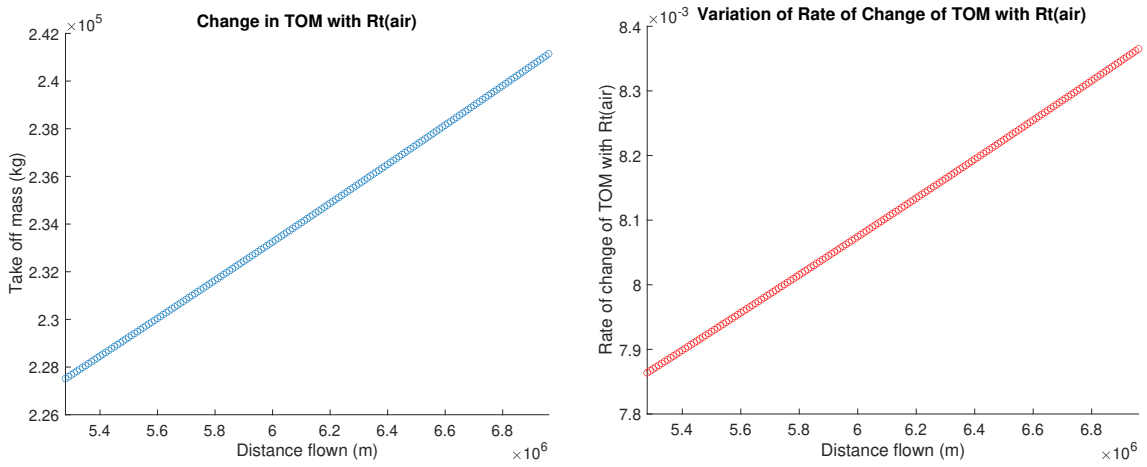

(a) Change in TOM

(b) Change in rate of change of TOM

Figure 3.8: Graphs showing how changing the air distance flown affects the take-off mass. These are linear relationships, as demonstrated by the data lying exactly along a straight line in each case. For every metre flown, the take-off mass is increased by just 8×10^{-3} kg, whilst the rate of change of take-off mass with air distance changes by 3×10^{-10} kg m^{-1}, demonstrating that the TOM is not greatly affected by the choice of $R_{t(air)}$.

3.5 Atmospheric state data

In Chapters 4 to 7 of this thesis zonal (eastward) winds, meridional (northward) winds and temperatures will be used to calculate both minimal time and minimal fuel burn routes. With the exception of the research in Chapter 5, all flights will be modelled as taking place between 1st December, 2019 and 29th February, 2020. This allows the full range of winter weather systems to be considered (Irvine et al., 2013), as the NAO has been shown to cause transatlantic routes to vary strongly (Woollings & Blackburn, 2012; Kim et al., 2016). In the future it is expected that cruise level winds in this region will continue to increase in velocity due to climate change (Williams, 2017; Storer et al., 2017; Simpson, 2016; Kim et al., 2020), so their inclusion in routing calculations seems set to become increasingly vital.

3.5.1 Re-analysis model

Using complex atmospheric state data in the simulations in this thesis is critical in ensuring their relevance to improving TBO.

Determination of fuel burn rate is dependent on temperature at any point in a trajectory. Numerical solution of the OCPs that arise from trying to find time and fuel minimal routes, necessitates obtaining speed and direction of the wind at any point across the North Atlantic. In the case where a varying altitude is included for estimates of actual flight fuel use, then these winds are required at a variety of different pressure levels.

In Chapters 4, 6 and 7 all atmospheric state data has been downloaded from the re-analysis data set provided by the National Center for Atmospheric Research (NCAR) (Kalnay et al., 1996). This is a data-driven model, relying on a large array of observations. Wind velocity and temperature values are given for a global grid of resolution 2.5° as daily averages. Linear interpolation is then applied to obtain atmospheric state data at specific points in a trajectory. This approach is justified as long haul flight routes have been shown to be largely insensitive to the resolution of weather data (Lunnon & Mirza, 2007) and the evolution of the jet stream at this altitude shows little variation across a 24 hour period (Mangini et al., 2018). Data from the 250 hPa pressure level is used where altitude is kept constant, as this is approximately equal to a standard cruise flight level of 34 000 feet. However, where altitude is varied during actual flights the wind and

temperature data is considered from the 150 hPa, 200 hPa, 300 hPa and 400 hPa pressure levels too. This is linearly interpolated to give winds and temperature information at any point in the three dimensional space used in the recorded flight data.

3.5.2 Atmospheric state at pressure altitude levels

Despite there being different weather patterns on different days, wind speeds on an individual day do not appear to be affected by altitude changes within the cruise phase altitude range. Temperature, which is also used in the fuel burn rate calculation shows more variation. Figures 3.9 and 3.10 show zonal and meridional winds and temperature at 150 hPa, 200 hPa, 250 hPa, 300 hPa and 400 hPa for the 1st December, 2019 and 8th February, 2020.

As current tracks are produced using the 250 hPa pressure altitude level, each set of weather data was tested by plotting a quantile-quantile plot comparing it with the 250 hPa level. A quantile-quantile plot assesses whether two sets of sample data come from the same distribution. The quantile values for the 250 hPa data set appear on the x-axis and the corresponding quantile values for the other pressure level appear on the y-axis. Figure 3.11 gives quantile-quantile plots where data points are shown as blue crosses and the line of best fit is the dashed red line. When quantile-quantile plots are linear, this means that data is likely to come from the same distribution. These graphs show that the wind distributions at 200 hPa and 300 hPa for both zonal and meridional wind are very similar to the distribution of wind at the 250 hPa level, with any differences being more noticeable at the extremes of the distribution. At 400 hPa the meridional wind distribution is also similar to that at 250 hPa. The temperature distributions at different altitudes vary more. Using a two sample Kolmogorov–Smirnov test at the 5% significance level, the temperature distributions at all other altitudes are shown to be significantly different from that at 250 hPa. Zonal wind distributions at 300 hPa and 250 hPa and at 200 hPa and 250 hPa are not significantly different, whilst those at 150 hPa and 250 hPa and at 400 hPa and 250 hPa are. Only the distribution of meridional winds at 150 hPa is significantly different from that at 250 hPa.

Now winds across all days of the winter period 2019-2020 are considered. The zonal and meridional wind speeds are interpolated to give winds every ten FLs. The difference between these values and the value at FL 340 (approximately corresponding to a pressure

level of 250 hPa) are calculated for all days from 1st December, 2019 to 29th February, 2020 for zonal and for meridional wind. The mean, standard deviation, median and quartiles of each data set are shown in Figures 3.12 and 3.13. As none of the wind fields followed a normal distribution when tested, the median and interquartile range are better measures of variability than the mean and standard deviation. Zonally the median difference between the winds is never greater than 1.8 m s^{-1}. Meridional winds show a similar pattern, with the median difference never being more than 1.2 m s^{-1}. This demonstrates that winds at different cruise range altitudes across the North Atlantic are very similar.

Testing winds on a gridpoint by gridpoint basis, using a Spearman's rank hypothesis test shows that there is no evidence to suggest that wind fields are significantly different at different altitudes within the cruise altitude range at the 5% significance level.
So we can conclude from this analysis that across the 2019 to 2020 winter season, the zonal and meridional winds are very similar at different pressure levels between 400 hPa and 150 hPa, but that temperature varies much more with static air pressure.

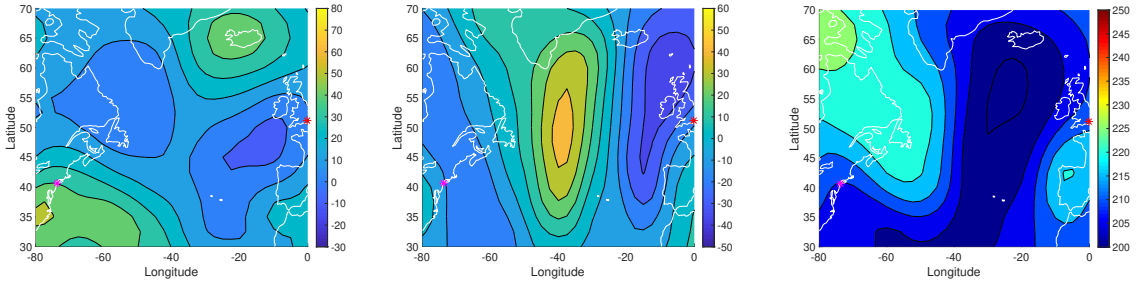

(a) 150hPa: Zonal wind m/s (b) Meridional wind m/s (c) Temperature deg K

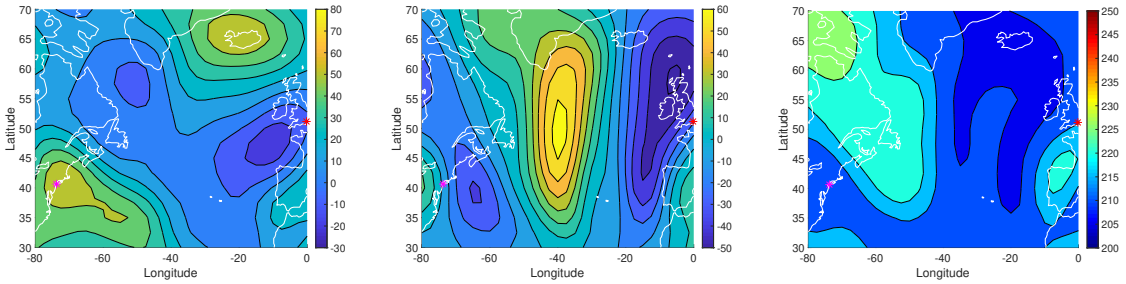

(d) 200 hPa: Zonal wind m/s (e) Meridional wind m/s (f) Temperature deg K

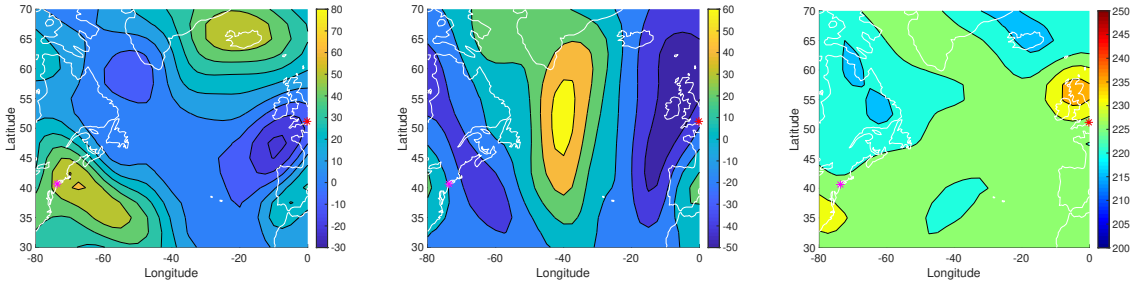

(g) 250 hPa: Zonal wind m/s (h) Meridional wind m/s (i) Temperature deg K

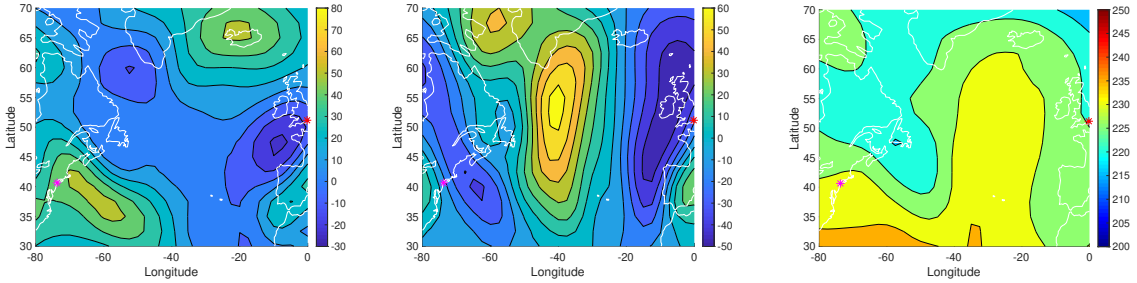

(j) 300 hPa: Zonal wind m/s (k) Meridional wind m/s (l) Temperature deg K

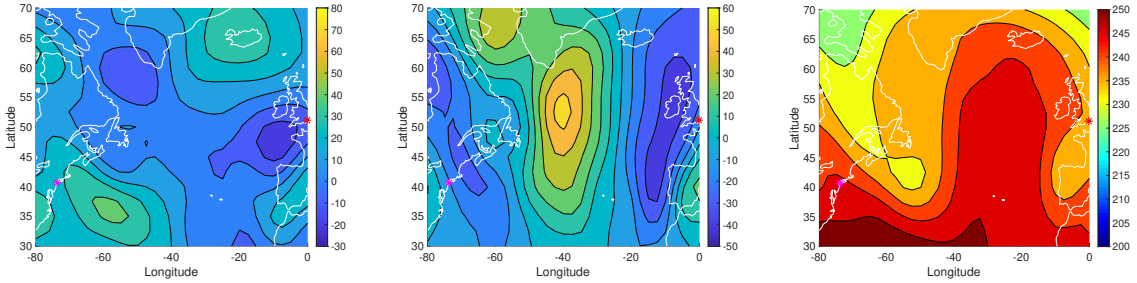

(m) 400hPa: Zonal wind m/s (n) Meridional wind m/s (o) Temperature deg K

Figure 3.9: Weather at different pressure altitudes on 1st December, 2019.

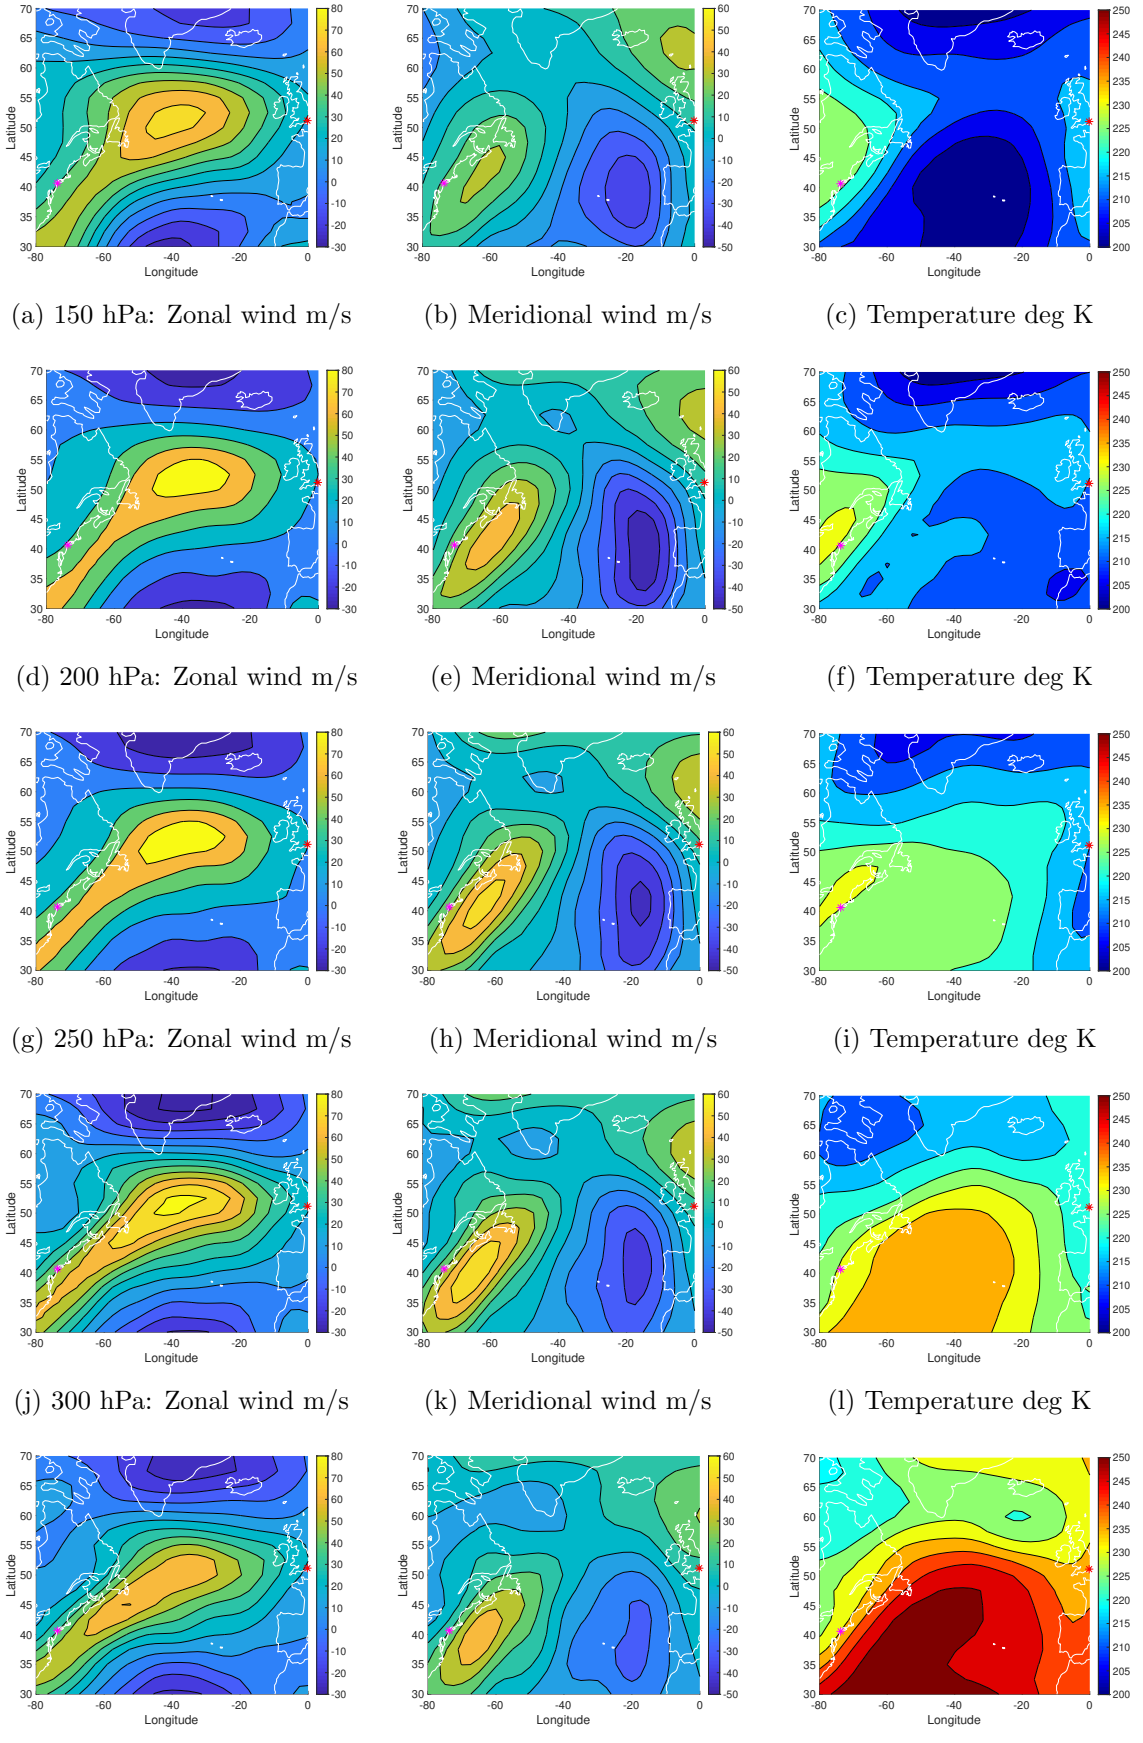

(a) 150 hPa: Zonal wind m/s (b) Meridional wind m/s (c) Temperature deg K

(d) 200 hPa: Zonal wind m/s (e) Meridional wind m/s (f) Temperature deg K

(g) 250 hPa: Zonal wind m/s (h) Meridional wind m/s (i) Temperature deg K

(j) 300 hPa: Zonal wind m/s (k) Meridional wind m/s (l) Temperature deg K

(m) 400 hPa: Zonal wind m/s (n) Meridional wind m/s (o) Temperature deg K

Figure 3.10: Atmospheric state at different pressure altitudes on 8th February, 2020.

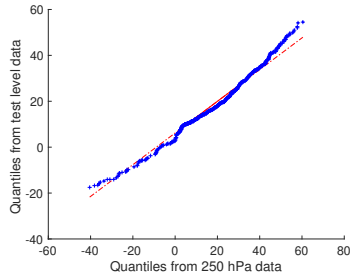

(a) Zonal wind quantiles 150 hPa v 250 hPa

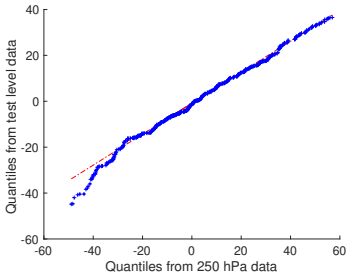

(b) Meridional wind quantiles 150 hPa v 250 hPa

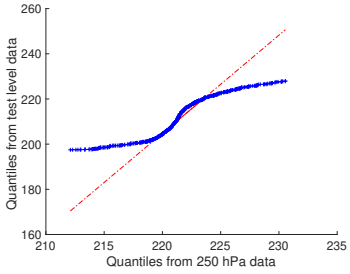

(c) Temperature quantiles 150 hPa v 250 hPa

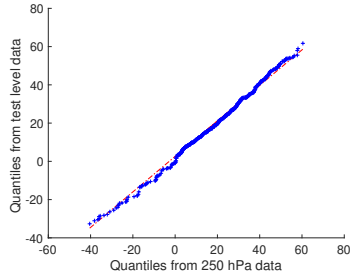

(d) Zonal wind quantiles 200 hPa v 250 hPa

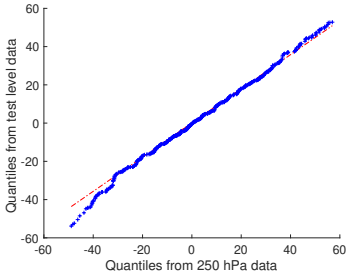

(e) Meridional wind quantiles 200 hPa v 250 hPa

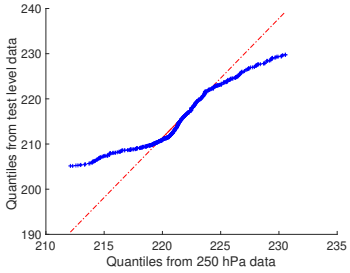

(f) Temperature quantiles 200 hPa v 250 hPa

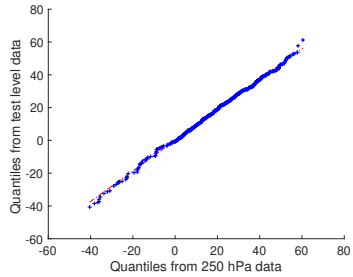

(g) Zonal wind quantiles 300 hPa v 250 hPa

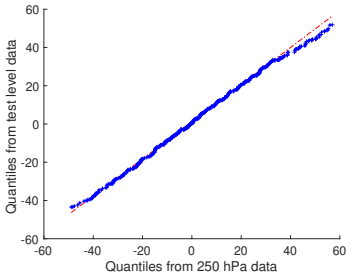

(h) Meridional wind quantiles 300 hPa v 250 hPa

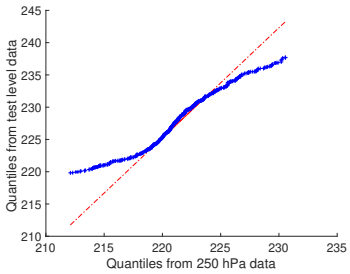

(i) Temperature quantiles 300 hPa v 250 hPa

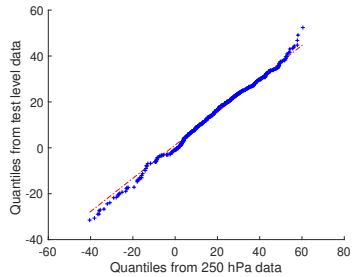

(j) Zonal wind quantiles 400 hPa v 250 hPa

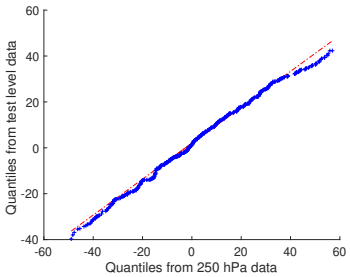

(k) Meridional wind quantiles 400 hPa v 250 hPa

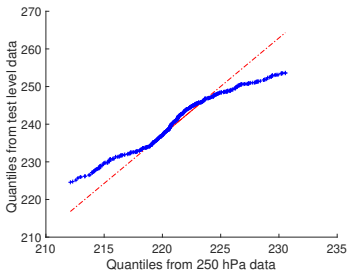

(l) Temperature quantiles 400 hPa v 250 hPa

Figure 3.11: Comparison of atmospheric state distributions compared to 250 hPa weather on 1st December, 2019.

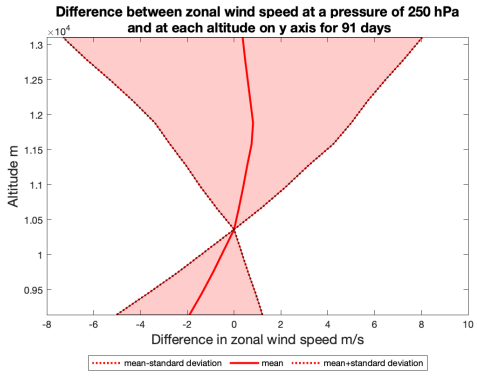

(a) Using mean and standard deviation

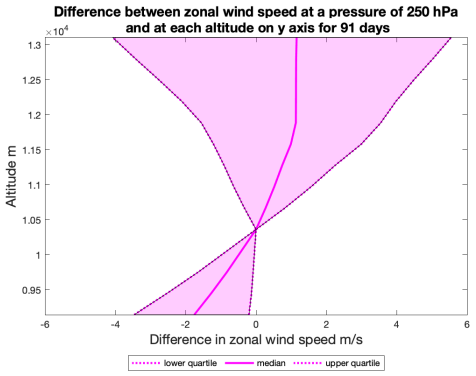

(b) Using median and quartiles

Figure 3.12: Differences between zonal winds at each point on a grid at different altitudes and the altitude corresponding to 250 hPa are found for a grid across the North Atlantic from -80 to 0 degrees longitude and from 30 to 70 degrees latitude. This is repeated for all days between 1st December, 2019 and 29th February, 2020. The mean value and one standard deviation each side of the mean for these figures are shown in (a), whilst the median and quartiles for the same data are shown in (b).

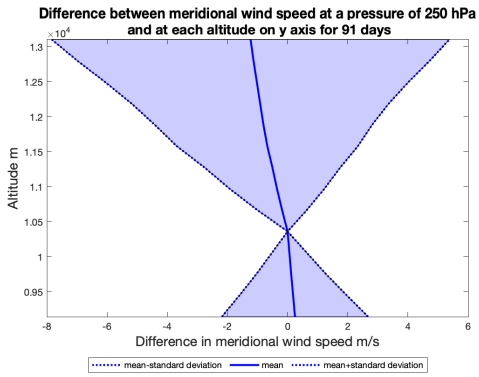

(a) Using mean and standard deviation

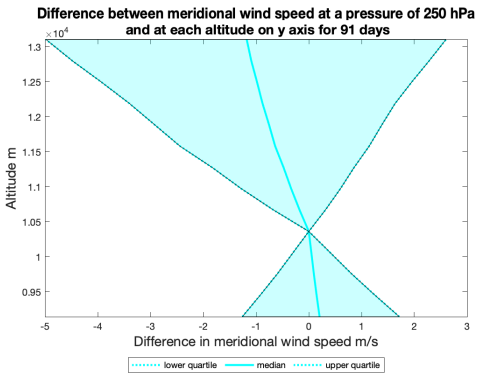

(b) Using median and quartiles

Figure 3.13: The same process is repeated with meridional winds, generating the corresponding charts across the same time period and spatial grid.

3.6 Track data

Aircraft crossing the North Atlantic have been constrained to fly along routes that form part of an OTS since the 1960s, with tracks first being published in 1965, although they were used by ATM as early as 1961.

These tracks were adopted due to this being both the busiest oceanic airspace in the world and also suffering from a lack of radar coverage across the central section of the region. To ensure safety, aircraft are kept separated, both horizontally and vertically on set tracks dictated ANSPs. NATS in the United Kingdom provide the westbound paths and NAV CANADA those for eastbound flights.

Now that the Aireon low Earth orbit satellites can be used to give 100% global coverage it can be argued that such measures are no longer necessary and in Chapter 4 the possible time and thus fuel and emissions savings that could be made by flying time minimal routes, rather than along the OTS, are discussed.

Each day a set of tracks is published in each direction, giving waypoints that must be flown through on a route. Westbound traffic on the tracks must pass a longitude of 30°W between 1130 and 1900 UTC (Co-ordinated Universal Time), whilst aircraft flying eastbound must pass the same point between 0100 and 0800 UTC. These timings have been chosen to accommodate the noise level rules at the airports and the needs of air travellers.

Tracks take into account the wishes of the airlines who send PRMs to the ANSPs a few hours before a flight. The track waypoints used in this thesis have been downloaded from the website: https://blackswan.ch/northatlantictracks.

Westbound tracks use labels from A to K, where A is always the northernmost route and eastbound tracks are labelled from N to Z with Z being the southernmost route. There are different numbers of tracks each day, depending on the wind field and the number of aircraft flying, so although westbound tracks are always labelled from A whilst eastbound are labelled in reverse from Z, tracks with the same label on different days are not necessarily similarly efficient.

The ATM tracks are used at a variety of different altitudes, but they are optimised at the 250 hPa iso-bar, since, on average, this is where the jet stream is strongest (Mangini et al., 2018). An example of daily tracks is given in Chapter 4.

Since the publication of Wells et al. (2021), on which Chapter 4 is based, NATS have removed the OTS on 21 days, starting on the 9th March, 2021, allowing airlines to decide on their preferred trajectory (Ahlgren, 2022). This has been possible due to reduced demand during the Covid pandemic, as well as the fact that the Aireon satellites update aircraft positions every 7 to 8 seconds, reducing the necessary safety margin between aircraft from 40 to 14 nautical miles. At and below FL330, all tracks have been removed permanently from 1st March, 2022.

3.7 Chapter summary

In this chapter we have defined how optimal control can be used in finding optimal flight paths across the North Atlantic. The mathematical theory underpinning the three main ways of solving OCPs has been summarised, so that future use of an indirect method based on PMP in Chapters 4 and 5, a direct method incorporating Matlab's fmincon function used in Chapter 6 and the dynamic programming method used in Chapter 7 should be familiar to the reader.

The fuel burn function used to calculate the rate of fuel burn in Chapters 6 and 7 has been introduced, together with the reanalysis data model used in Chapters 4, 6 and 7 and the track data used in Chapter 4.

In Chapter 1 we asked some questions about re-routing flights, we have now justified our chosen methods by way of an extensive literature review given in Chapter 2, showing that indirect methods for solving an OCP work best with time minimisation at a fixed mass and airspeed, direct methods work best for fuel minimisation when time of flight is fixed and dynamic programming is of most use when minimising fuel burn, but with a free-time for trajectories and multiple control variables. In Chapter 3 each of those methods has been introduced in general terms, so that in the next four chapters we can start to tackle the questions from Chapter 1 in depth.

Chapter 4

Reducing transatlantic flight emissions by fuel-optimised routing

4.1 Introduction

By routing aircraft based on daily wind fields, can we reduce the amount of air distance that they cover during a flight between LHR and JFK compared to the OTS supplied by NATS and NAV CANADA? This is the key research question answered in this chapter. As pressure to reduce global greenhouse gas emissions continues to increase (IPCC, 2019), aviation must respond at least as ambitiously as the other transport sectors. Currently aviation is responsible for approximately 2.4% of all anthropogenic sources of CO_2 (Graver et al., 2019; Grewe et al., 2019; Lee et al., 2020), but this figure is growing (Ryley et al., 2020; Grewe et al., 2017; Graver et al., 2019). The ICAO has already established a policy of improving the fuel efficiency of international flights by 2% annually (ICAO, 2016), through improvements to aircraft technology, sustainable fuels and ATM and operations (ICAO, 2019). However, greater savings are needed. Additionally, 192 nations agreed to CORSIA in 2016, pledging to use offset schemes to maintain net emissions at the 2020 level (Timperley, 2019). CORSIA only provides short term alleviation, as there are difficulties in ensuring that genuine net emissions reduction takes place.

If the global economy fails to decarbonise sufficiently rapidly, there may be significant consequences for aviation from the ensuing climate change. These consequences include

increased turbulence as the jet stream becomes more sheared (Williams & Joshi, 2013; Williams, 2017; Storer et al., 2017; Lee et al., 2019; Kim et al., 2015), modified flight routes and journey times as the prevailing high-altitude winds shift and strengthen (Karnauskas et al., 2015; Irvine et al., 2016; Williams, 2016; Kim et al., 2020), and take-off weight restrictions as warmer air reduces lift and thrust on the runway (Coffel & Horton, 2015; Gratton et al., 2020). Therefore, aviation is not only a contributor to climate change, but may also suffer from its adverse effects increasingly in future.

Various alternative solutions to offsetting have been suggested, from greater use of synthetic bio-fuels to replacing the entire air transport fleet of approximately 31 000 aircraft (CAPA, 2018) with updated models (Monbiot, 2007). However, those technologies that have the potential to produce significant reductions in fuel use are high-risk, high-cost and have implementation timescales measured in decades (Jensen et al., 2015). By contrast, improvements to current operational procedures, such as routing flights more efficiently, have the potential to provide immediate, low-cost, low-risk and significant reductions (Németh et al., 2018). Fuel saving through more efficient operations would be a benefit to both the airlines through reduced fuel expenditure and to the environment through reduced emissions.

Historically flight routes across the North Atlantic have been constrained by the large volume of air traffic and the absence of radar coverage in mid-ocean (Dhief, 2018). However, a new network of low Earth orbit satellites being tested currently will improve situational awareness dramatically (Aireon, 2020). With aircraft able to transmit and receive accurate information continuously, it is now possible to consider the implementation of fuel-optimised routes (NATS, 2019). These new routes would take greater advantage of the prevailing eastward winds when flying east and reduce the negative impact of these same air currents when flying west. Air distance, the distance flown by an aircraft relative to the surrounding air, will be used in this study as a measure of the efficiency of a flight path. Ground distance, in contrast, is not related to fuel burn in a simple way, because of the conveyor effect of the winds. As fuel burn is directly proportional to air distance and as emissions, including carbon dioxide, are directly proportional to fuel burn (Henderson et al., 2012; Green, 2009), any saving made in air distance is a valuable step towards meeting the ICAO target.

Currently flight tracks in the North Atlantic's OTS are created on a daily basis by ANSPs, NATS in the United Kingdom for the westbound paths and NAV CANADA for those going east. These are based primarily on the need to separate aircraft safely, whilst taking some account of the winds. Airlines request their preferred tracks by submitting PRMs in the hours before a flight and the ANSPs create a daily track system that reflects the airlines' wishes as closely as possible. An example of these tracks can be seen in Figure 4.1. By calculating the air distances for these ATM tracks and comparing them with the optimised minimum values, the potential savings can be assessed. These savings are characteristic of the route flown and not of the aircraft being used.

Recent research has focused on limiting energy output, rather than time (Franco & Rivas, 2011; Burrows, 1983; Pierson & Ong, 1989; Murrieta Mendoza et al., 2020). Other strands of route optimisation have considered turbulence avoidance (Jardin & Bryson, 2012b,a; Kim et al., 2015) and balancing the reduction of climate effects with time of flight (Grewe et al., 2019, 2017). This research, however, is the first to identify fuel and emissions savings for transatlantic traffic by calculating the excess air distance flown along the OTS relative to the minimum air distance route. Thus the focus of this thesis is CO_2 reduction. Fuel optimised routes are not necessarily climate optimised (Grobler et al., 2019), as additional effects such as contrail formation, documented in other sources (Teoh et al., 2020; Poll & Schumann, 2021b), are not taken into account.

This research is set out in five sections. In Section 4.2, the different datasets are described. Section 4.3 explains the analysis method, shows how time optimisation impacts fuel use, how time optimised routes are found and how to compare these with the current tracks. Results showing potential air distance savings are set out in Section 4.4. Finally, the results are summarised and discussed in Section 4.5.

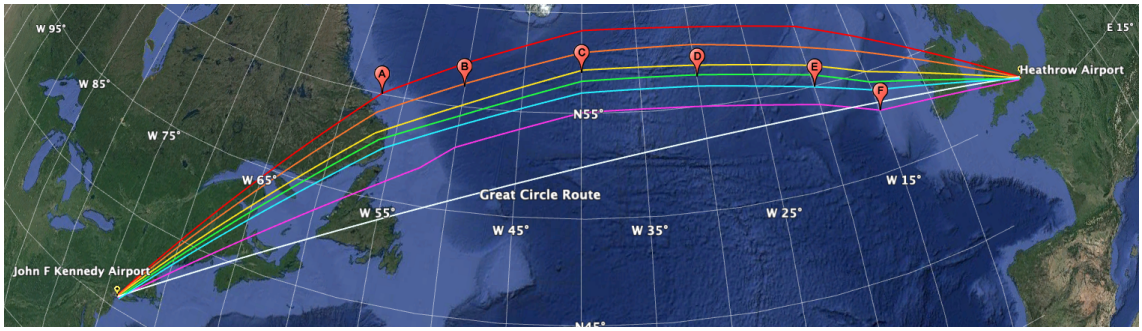

(a) Westbound ATM tracks

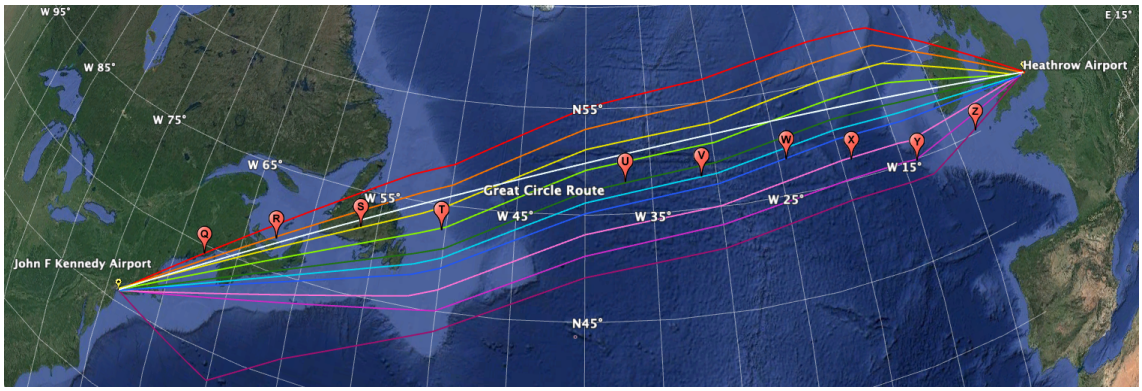

(b) Eastbound ATM tracks

Figure 4.1: Figure 4.1a shows all westbound tracks between LHR and JFK on the 3rd December, 2019. The Great Circle path (GC Path), the shortest distance along the ground between the airports, is shown in white. The six tracks are labelled from A to F and lie predominantly North of the GC path to avoid the prevailing jet stream air currents. They were valid for all flights reaching 30W from 11:00 to 19:00(UTC). Figure 4.1b shows all eastbound tracks between JFK and LHR on the 3rd December, 2019. The GC path is again shown in white. The ten tracks are labelled from Q to Z and lie both sides of the GC path. They were valid for all flights reaching 30W from 01:00 to 08:00(UTC). Map data: Google, Data SIO, NOAA, U.S. Navy, NGA, GEBCO, Image IBCAO, Image Landsat/Copernicus.

4.2 Data sources

Flights between LHR (51.5° N, 0.5° W) and JFK (40.6° N, 73.8° W) are modelled both eastbound and westbound. Although the trajectory prediction methods used by airlines are commercially sensitive (Cheung, 2018), the resulting tracks are in the public domain. Past flight tracks were downloaded from the https://blackswan.ch/northatlantictracks website, which provides archived eastbound and westbound way points for all paths flown over the last year. Westbound tracks use labels from A to K, where A is always the northernmost route and eastbound tracks are labelled from N to Z with Z being the southernmost route. There are different numbers of tracks each day, depending on the wind field and the number of aircraft flying, so although westbound tracks are always labelled from A whilst eastbound are labelled in reverse from Z, tracks with the same label on different days are not necessarily similarly efficient. The ATM tracks are used at a variety of different altitudes, but they are optimised at the 250 hPa iso-bar, since, on average, this is where the jet stream is strongest (Mangini et al., 2018). This pressure level corresponds approximately to FL340, which is an altitude of 34 000 feet in the ISA. Mangini et al. (2018) found that altering altitude between 200 and 300 hPa when flying the ATM tracks, made less than a 1% difference to total route time. Therefore, all simulations in this study have been run at the single pressure altitude of 250 hPa.

Flights from 1st December, 2019 to 29th February, 2020, have been considered as the wind fields in the winter months tend to be at their strongest and most variable. This can be seen when seasonal weather patterns are split into categories; in winter five are used, whereas in summer only three are necessary (Irvine et al., 2013). This is (at least partly) because transatlantic flight routes in winter vary strongly in response to the NAO (Kim et al., 2016). The average wind speed at typical cruise altitudes in the North Atlantic flight corridor in winter is expected to increase in future, because of climate change, meaning that winds will play an increasingly important role in flight routing (Williams, 2016; Simpson, 2016; Kim et al., 2015, 2020).

Past wind field data are obtained from the re-analysis data set provided by NCAR (Kalnay et al., 1996). Average daily horizontal wind velocities are given every 2.5 degrees of both latitude and longitude, at a range of pressure levels. Given that there is little variation in the wind field at this altitude on timescales of one day (Mangini et al., 2018), a daily

average is sufficiently accurate for these routes. Linear interpolation is used to obtain meridional (northward) and zonal (eastward) winds in between the grid points. Following Lunnon and Mirza (Lunnon & Mirza, 2007) the calculated optimal routes are largely insensitive to the wind grid resolution for journeys of this length, with use of high resolution data making at most a few seconds difference on a transatlantic flight. Air distances for flights made at the 200 and 300 hPa pressure levels were also considered, but these were shown to be very similar to values at the 250 hPa level, so only this level is considered in this chapter.

Since the flights are long haul, the time and distance covered during climb and descent are small compared to the cruise phase, with flights between LHR and JFK spending about 92% of the ground distance in cruise (Flightradar24, 2020). Therefore, we neglect the climb and descent phases in this analysis.

4.3 Analysis methods

Here methods used to obtain optimised fuel-efficient flight paths are described. The air distance of each route is considered. This is the distance travelled relative to the wind field and so can be calculated by multiplying the time in the air by the airspeed. By contrast, ground distance is the length of the route that the aircraft is observed to have flown from the ground. Thus in a zero wind field, the air distance and the ground distance will be identical and the Great Circle path — the path giving the shortest distance along the ground between two points on a sphere — will be the minimum time route. However, if a wind field is added, then maintaining the Great Circle path (GC path) as the ground track involves a change to the air distance flown, resulting in a flight time that is no longer the minimum.

In this section, the relationship between fuel use and air distance is first established, before the time optimisation method used is outlined.

4.3.1 Aircraft Fuel Use

For a given aircraft, the mass of fuel burned per unit time, $\frac{\mathrm{d}m_f}{\mathrm{d}t}$, depends on the weight of the aircraft, the true airspeed and the altitude (Poll, 2018; Poll & Schumann, 2021a,b). Hence, if the weight, airspeed and altitude are held constant, as we assume here, the

quantity of fuel burned per unit distance travelled through the air, $\frac{\mathrm{d}m_f}{\mathrm{d}s}$, is also constant. In this case the total fuel required for a journey is the product of $\frac{\mathrm{d}m_f}{\mathrm{d}s}$ and the total air distance flown, where total air distance is the product of the constant airspeed and the total flight time. Therefore, the total fuel requirement for a flight at a given airspeed is minimised when the total flight time is minimised. Finding the minimum flight time at constant airspeed through a specified wind field is a classical problem in aeronautics, first addressed by Zermelo (Zermelo, 1930; Levi-Civita, 1931).

Furthermore, for an aircraft of a given weight at a specified altitude, there is a particular airspeed at which $\frac{\mathrm{d}m_f}{\mathrm{d}s}$ is also a minimum (Poll, 2018; Poll & Schumann, 2021b). For a modern large, long-range, turbofan-powered aircraft, this optimum airspeed is in the region of 240 m s^{-1} (Mach$\approx$ 0.82). Hence, for a given aircraft travelling between a given airport pair, the absolute minimum fuel requirement is found by calculating the airspeed that minimises the product of $\frac{\mathrm{d}m_f}{\mathrm{d}s}$ for the aircraft and the minimum air distance through the wind field. Therefore, the complete problem depends on both the aircraft and the wind field. In this research the mass of the aircraft has been kept constant at 80% of the take-off mass, as this gives an approximate average for the aircraft mass as fuel is burned during the flight.

4.3.2 Time Optimisation

Time minimisation (and therefore air distance minimisation) is achieved by solving a time OCP, where the dynamics are given by the Zermelo equations (see Appendix A) mapped conformally onto a sphere (see Appendix B) (Zermelo, 1930; Arrow, 1949; Bryson & Ho, 1975):

$$\frac{\mathrm{d}\lambda}{\mathrm{d}t} = \frac{u + V\cos\theta}{R\cos\phi}, \tag{4.1}$$

$$\frac{\mathrm{d}\phi}{\mathrm{d}t} = \frac{v + V\sin\theta}{R}, \tag{4.2}$$

$$\frac{\mathrm{d}\theta}{\mathrm{d}t} = -\frac{Wind2D}{R\cos\phi}, \tag{4.3}$$

where the $Wind2D$ term is:

$$\begin{aligned} Wind2D = &-\sin\theta\cos\theta\frac{\partial u}{\partial\lambda} + u\cos^2\theta\sin\phi + \cos^2\theta\cos\phi\frac{\partial u}{\partial\phi} \\ &+ v\sin\theta\cos\theta\sin\phi + \cos\theta\sin\theta\cos\phi\frac{\partial v}{\partial\phi} + V\cos\theta\sin\phi - \sin^2\theta\frac{\partial v}{\partial\lambda}. \end{aligned} \tag{4.4}$$

Here t is time from departure, $\lambda(t)$ and $\phi(t)$ are the aircraft's longitude and latitude (radians), $u(\lambda,\phi)$ and $v(\lambda,\phi)$ are zonal and meridional wind speeds (m s^{-1}), R is the radius

of the Earth, taken as 6 371 km, V is the constant airspeed (m s^{-1}) and θ is the aircraft's heading angle, the direction in which the nose of the aircraft is pointing, here measured anticlockwise from due east. This is illustrated in Figure 3.1. Initial conditions $\lambda(0)$ and $\phi(0)$ are given by the departure airport's longitude and latitude respectively. The initial value for $\theta(0)$ is unknown and to be determined. Zonal and meridional components of both the wind velocity and the velocity of the aircraft relative to the air are combined to give Equations (4.1) and (4.2), from which Equation (4.3) is derived, using first-order optimality conditions (Bryson & Ho, 1975; Ng et al., 2011, 2014).

This process of deriving Equation (4.3) from Equations (4.1) and (4.2) using PMP is described here. For more information on PMP please refer to Section 3.3.1. Initially we also include aircraft mass, M and the variable fuel burn rate, g, as given by the fuel burn function in Section 3.4. The first two lines of the dynamical system follow Zermelo's equations as expressed in Zermelo (1930) very closely, but this planar system has been scaled conformally to fit the spherical model:

$$\frac{\mathrm{d}\lambda}{\mathrm{d}t} = \frac{u + V\cos\theta}{R\cos\phi}, \tag{4.5}$$

$$\frac{\mathrm{d}\phi}{\mathrm{d}t} = \frac{v + V\sin\theta}{R}, \tag{4.6}$$

$$\frac{\mathrm{d}M}{\mathrm{d}t} = -g(\lambda(t), \phi(t)). \tag{4.7}$$

It is assumed that thrust is equal to drag. The airspeed is kept constant, with the aircraft starting from the departure airport at $t = 0$ and arriving within a certain radius of the destination airport when $t = t_f$. The control variable is the heading angle.

The cost functional aims to minimise time, giving a Lagrangian of 1, so the Hamiltonian of the system becomes:

$$H = p_\lambda \left(\frac{V\cos\theta + u}{R\cos\phi} \right) + p_\phi \left(\frac{V\sin\theta + v}{R} \right) + 1 + p_M(-g), \tag{4.8}$$

where p_λ, p_ϕ and p_M are the adjoint variables. As mass is kept constant at 80% of take-off mass during this chapter, to allow fuel burn to be proportional to air distance, we now remove Equation (4.7), as the rate of change of mass will be assumed to be zero. The last term of the Hamiltonian can therefore be ignored giving:

$$H = p_\lambda \left(\frac{V\cos\theta + u}{R\cos\phi} \right) + p_\phi \left(\frac{V\sin\theta + v}{R} \right) + 1. \tag{4.9}$$

From here we look at the partial derivatives of H with respect to each variable:

$$-\dot{p}_\lambda = \frac{\partial H}{\partial \lambda} = \frac{p_\lambda}{R\cos\phi}\frac{\partial u}{\partial \lambda} + \frac{p_\phi}{R}\frac{\partial v}{\partial \lambda}, \tag{4.10}$$

$$-\dot{p}_\phi = \frac{\partial H}{\partial \phi} = \frac{p_\phi}{R}\frac{\partial v}{\partial \phi} + \frac{p_\lambda}{R\cos\phi}\frac{\partial u}{\partial \phi} + \frac{p_\lambda}{R\cos\phi}\tan\phi(u + V\cos\theta) \tag{4.11}$$

$$\frac{\partial H}{\partial \theta} = \frac{-p_\lambda V\sin\theta}{R\cos\phi} + \frac{p_\phi V\cos\theta}{R}. \tag{4.12}$$

For an extremum to exist this last expression must be zero, giving:

$$\tan\theta = \frac{p_\phi\cos\phi}{p_\lambda}. \tag{4.13}$$

For optimality, H must also be zero. We use this value and rearrange Equation (4.13), to make p_ϕ the subject and then substitute this into the equation for H:

$$0 = \frac{p_\lambda(V\cos\theta + u)}{R\cos\phi} + \frac{p_\lambda\tan\theta(V\sin\theta + v)}{R\cos\phi} + 1. \tag{4.14}$$

This can now be rearranged to make p_λ the subject of the formula:

$$p_\lambda = \frac{-R\cos\theta\cos\phi}{V + u\cos\theta + v\sin\theta}. \tag{4.15}$$

Following the same method, but finding p_ϕ in terms of V, u, v and θ this time gives us:

$$0 = \frac{p_\phi\cos\phi(V\cos\theta + u)}{R\cos\phi\tan\theta} + \frac{p_\phi(V\sin\theta + v)}{R} + 1,$$

$$p_\phi = \frac{-R\sin\theta}{V + u\cos\theta + v\sin\theta}. \tag{4.16}$$

These adjoint expressions can now be substituted back into Equations (4.10) and (4.11):

$$\dot{p}_\lambda = -\frac{\partial H}{\partial \lambda} = \left(\frac{\cos\theta}{V + u\cos\theta + v\sin\theta}\right)\frac{\partial u}{\partial \lambda} + \left(\frac{\sin\theta}{V + u\cos\theta + v\sin\theta}\right)\frac{\partial v}{\partial \lambda}, \tag{4.17}$$

$$\dot{p}_\phi = -\frac{\partial H}{\partial \phi} = \left(\frac{\sin\theta}{V + u\cos\theta + v\sin\theta}\right)\frac{\partial v}{\partial \phi} + \left(\frac{\cos\theta}{V + u\cos\theta + v\sin\theta}\right)\frac{\partial u}{\partial \phi}\ldots \tag{4.18}$$

$$+ \left(\frac{\cos\theta}{V + u\cos\theta + v\sin\theta}\right)\tan\phi(V\cos\theta + u).$$

Our aim is to find the rate of change of θ so that we can plot a path, minimising time. So using implicit differentiation on Equation (4.13) gives:

$$\frac{\mathrm{d}\theta}{\mathrm{d}t}\sec^2\theta = \frac{\mathrm{d}\phi}{\mathrm{d}t}\left(\frac{-p_\phi\sin\phi}{p_\lambda}\right) + \frac{\dot{p}_\phi}{p_\lambda}\cos\phi - \frac{\dot{p}_\lambda p_\phi}{p_\lambda^2}\cos\phi. \tag{4.19}$$

Substituting Equations (4.6) and (4.15)-(4.18) into Equation (4.19) we obtain:

$$\frac{\mathrm{d}\theta}{\mathrm{d}t} = \cos^2\theta\left[\frac{\mathrm{d}\phi}{\mathrm{d}t}\left(\frac{-p_\phi\sin\phi}{p_\lambda}\right)\right] + \cos^2\theta\left[\frac{\dot{p}_\phi}{p_\lambda}\cos\phi\right] - \cos^2\theta\left[\frac{\dot{p}_\phi p_\phi}{p_\lambda^2}\cos\phi\right]. \tag{4.20}$$

To clarify the working, the coloured sections from the equation above, will be simplified in turn. The first section can be simplified as:

$$\begin{aligned}\cos^2\theta\left[\frac{\mathrm{d}\phi}{\mathrm{d}t}\left(\frac{-p_\phi\sin\phi}{p_\lambda}\right)\right] &= \cos^2\theta\frac{\mathrm{d}\phi}{\mathrm{d}t}\left(\frac{\tan\theta}{\cos\phi}\right)(-\sin\phi)\\ &= -\frac{\mathrm{d}\phi}{\mathrm{d}t}\tan\phi\sin\theta\cos\theta\\ &= -\left(\frac{v+V\sin\theta}{R}\right)\tan\phi\sin\theta\cos\theta\\ &= -\frac{v\sin\phi\sin\theta\cos\theta}{R\cos\phi}-\frac{V\sin^2\theta\cos\theta\sin\phi}{R\cos\phi}.\end{aligned} \tag{4.21}$$

The second section becomes:

$$\begin{aligned}\cos^2\theta\left[\frac{\dot{p}_\phi}{p_\lambda}\cos\phi\right] &= \left[\left(\frac{\sin\theta}{V+u\cos\theta+v\sin\theta}\right)\frac{\partial v}{\partial\phi}+\left(\frac{\cos\theta}{V+u\cos\theta+v\sin\theta}\right)\frac{\partial u}{\partial\phi}+\ldots\right.\\ &\quad\left.\left(\frac{\cos\theta}{V+u\cos\theta+v\sin\theta}\right)\tan\phi\,(V\cos\theta+u)\cos\phi\left(\frac{V+u\cos\theta+v\sin\theta}{-R\cos\theta\cos\phi}\right)\right]\cos^2\theta\\ &= -\left(\frac{\sin\theta\cos\theta}{R}\right)\frac{\mathrm{d}v}{\mathrm{d}\phi}-\left(\frac{\cos^2\theta}{R}\right)\frac{\mathrm{d}u}{\mathrm{d}\phi}-\frac{\tan\phi\,(V\cos\theta+u)}{R}\cos^2\theta\\ &= -\left(\frac{\sin\theta\cos\theta}{R}\right)\frac{\mathrm{d}v}{\mathrm{d}\phi}-\left(\frac{\cos^2\theta}{R}\right)\frac{\mathrm{d}u}{\mathrm{d}\phi}-\frac{V\cos^3\theta\sin\phi}{R\cos\phi}-\frac{u\cos^2\theta\sin\phi}{R\cos\phi}.\end{aligned}$$

Finally the third part gives:

$$\begin{aligned}&-\cos^2\theta\left[\frac{\dot{p}_\lambda p_\phi}{p_\lambda^2}\cos\phi\right]\\ &= -\cos^2\theta\left(\frac{\cos\theta\frac{\partial u}{\partial\lambda}+\sin\theta\frac{\partial v}{\partial\lambda}}{V+u\cos\theta+v\sin\theta}\right)\frac{\tan\theta}{\cos\phi}\cos\phi\left(\frac{V+u\cos\theta+v\sin\theta}{-R\cos\theta\cos\phi}\right)\\ &= \left(\frac{\sin\theta\cos\theta}{R\cos\phi}\right)\frac{\partial u}{\partial\lambda}+\left(\frac{\sin^2\theta}{R\cos\phi}\right)\frac{\partial v}{\partial\lambda}.\end{aligned} \tag{4.22}$$

Putting these three sections together gives:

$$\begin{aligned}\frac{\mathrm{d}\theta}{\mathrm{d}t} &= -\frac{v\sin\phi\sin\theta\cos\theta}{R\cos\phi}-\frac{V\sin^2\theta\cos\theta\sin\phi}{R\cos\phi}\\ &\quad-\left(\frac{\sin\theta\cos\theta}{R}\right)\frac{\mathrm{d}v}{\mathrm{d}\phi}-\left(\frac{\cos^2\theta}{R}\right)\frac{\mathrm{d}u}{\mathrm{d}\phi}-\frac{V\cos^3\theta\sin\phi}{R\cos\phi}-\frac{u\cos^2\theta\sin\phi}{R\cos\phi}\\ &\quad+\left(\frac{\sin\theta\cos\theta}{R\cos\theta}\right)\frac{\partial u}{\partial\lambda}+\left(\frac{\sin^2\theta}{R\cos\phi}\right)\frac{\partial v}{\partial\lambda}.\end{aligned} \tag{4.23}$$

Rearranging by using $\cos^2\theta+\sin^2\theta=1$ and $\sin^2\theta=1-\cos^2\theta$ we obtain the expression for $Wind2D$, as used in the simulation:

$$\begin{aligned}\frac{\mathrm{d}\theta}{\mathrm{d}t} = -\frac{1}{R\cos\phi}\Bigg[&-\sin\theta\cos\theta\frac{\partial u}{\partial\lambda}+u\cos^2\theta\sin\phi+\cos^2\theta\cos\phi\frac{\partial u}{\partial\phi}-\frac{\partial v}{\partial\lambda}+\cos^2\theta\frac{\partial v}{\partial\lambda}\ldots\\ &+v\sin\theta\cos\theta\sin\phi+\sin\theta\cos\theta\cos\phi\frac{\partial v}{\partial\phi}+V\cos\theta\sin\phi\Bigg].\end{aligned} \tag{4.24}$$

Along with Equations (4.5) and (4.6), the expression for $\frac{d\theta}{dt}$ forms the dynamical system used in the simulation.

4.3.3 Numerical Algorithm

We present a numerical method to estimate the unknown initial heading angle $\theta(0)$. Each $\theta(0)$ value will have an associated trajectory, which will be time optimal for this particular set of initial conditions ($\lambda(0), \phi(0)$ and $\theta(0)$), but which will intersect the arrival airport only for certain values of $\theta(0)$. The method we use seeks these values of $\theta(0)$ by locating trajectories which satisfy the optimality system shown in Equations (4.1) to (4.3) and also pass within 200 m of the destination airport. These are obtained using extreme initial heading angles of 50 degrees either side of the GC path joining the departure and destination airports. The GC path initial heading from JFK to LHR is 38.7° and from LHR to JFK is 162.1°. Subsequent headings along the route are found by advancing the Equations (4.1) -(4.3) using the Euler forward step method (Williams, 2016). At each aircraft position, the wind field is given by linear interpolation of the re-analysis data. The flight path is calculated by advancing in time in 1 s intervals. The integration duration is taken as 1.8 times the time taken to cover the GC path at the chosen airspeed in still air, i.e. when airspeed is equal to ground speed. This duration ensures that, even in the strongest headwinds, the calculated paths will always pass through the destination's meridian.

From the paths for the two extreme initial heading angles, the distance to the destination airport from each point on each route (d_H) is calculated using the Haversine formula (Veness, 2019):

$$a_H = \sin^2(\Delta lat/2) + \cos(lat1)\cos(lat2)\sin^2(\Delta lon/2), \tag{4.25}$$

$$c_H = 2\text{atan2}(\sqrt{a_H}, \sqrt{1-a_H}), \tag{4.26}$$

$$d_H = Rc_H, \tag{4.27}$$

where $lat1$ is the latitude of the first point, $lat2$ is the latitude of the second point, Δlat is the difference in latitudes between the two points, Δlon is the difference in their longitudes and R is the radius of the Earth. The atan2 function is the four-quadrant inverse tangent of the two real values in the function bracket. (The first of these values dictates the y position and the second the x position of a point on Cartesian axes. The returned value is the angle, between $-\pi$ and π, swept out from the positive x axis to a line joining the origin to the given position.)

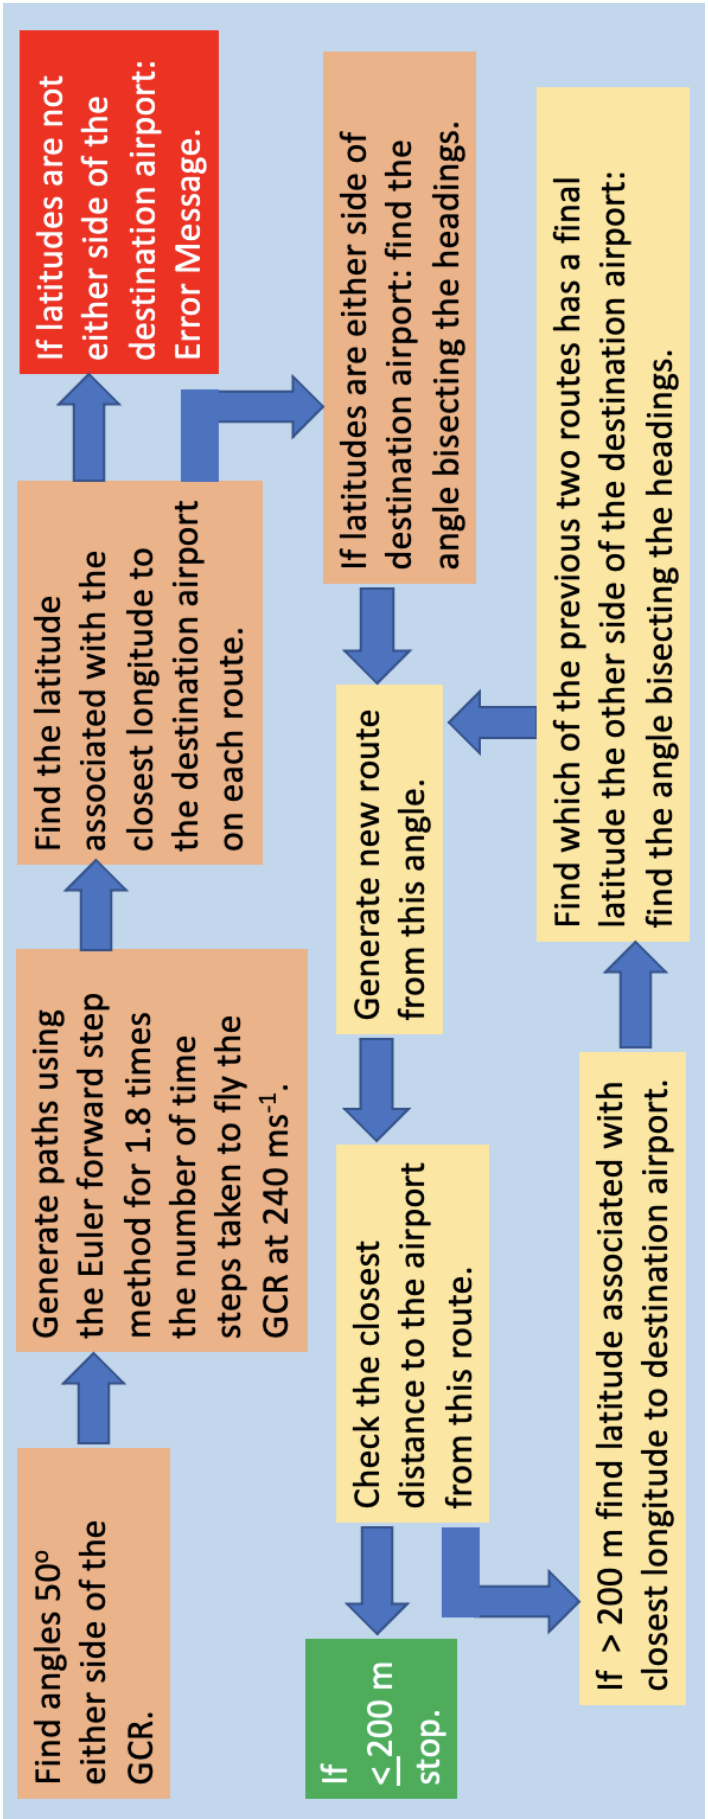

Figure 4.2: Flow chart to show how the bisection method is used to find the intial heading angle of the time minimal route.

If no point within 200 m of the destination airport is found along the two paths, a bisection method is applied to the initial heading angles and the algorithm above is repeated until this criterion is achieved. The method is summarised in Figure 4.2 and a set of diagrams showing intermediate results and the final trajectory is shown in Figure 4.3. Minimum time trajectories for airspeeds between 200 m s^{-1} and 270 m s^{-1} can be found in this way for all eastbound flights and most westbound flights.

For a small number of westbound flights, the bisection method finds multiple optimal time paths which pass within 200 m of the destination airport, due to routes from more than one initial heading angle fulfilling the success criterion. In these cases the time optimal route with the shortest duration is chosen. For example, paths for 1st December, 2019 and 8th February, 2020 are plotted in Figure 4.4 over a quiver plot showing the wind velocity at each grid point. It can be seen in Figure 4.4a that on 1st December, 2019, the bisection method is successful in finding a single time optimal path. The daily wind field leads to routes with similar initial heading angles changing heading gradually and almost identically, so that they do not cross. In contrast, Figure 4.4b shows that on 8th February, 2020, where the region of stronger adverse wind in the mid-North Atlantic causes flights to change heading more rapidly, three trajectories will reach the destination airport. Only one of these will, however, be a global (within the initial 100°-wide search range) minimum time route for the journey, as the flight times differ. For the three time optimal routes shown in Figure 4.4b these flight times are 8 hours 53 minutes for an initial heading angle of 167.5 degrees, 8 hours 23 minutes for an initial heading angle of 196.3 degrees and 7 hours 25 minutes for an initial heading angle of 121.7 degrees.

Thus using the Euler forward method applied to the Zermelo dynamical system, mapped conformally onto a sphere, the minimal time paths can be plotted between airport pairs across a range of air speeds for all 91 days in winter 2019-2020. These paths are called OFW routes. From these trajectories the air distance can be found and compared with the air distance of an aircraft with a ground track following the GC path and with the air distances associated with each of the daily ATM tracks in the OTS. Due to the daily wind field data, the air distance for the GC path will not be equal to the GC ground track distance.

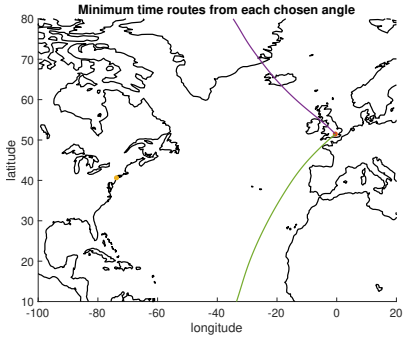

(a) Initial angles

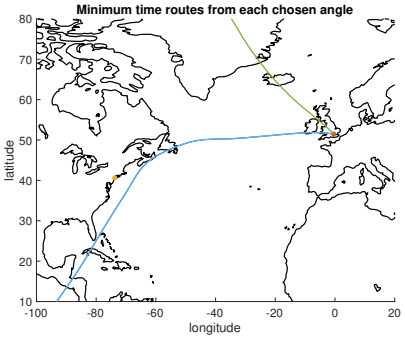

(b) 1st bisection

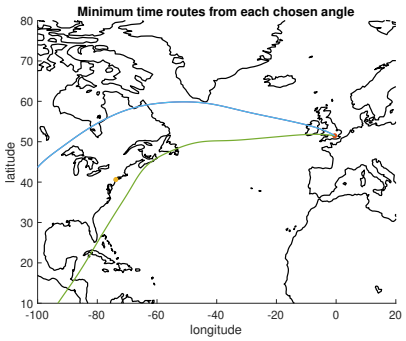

(c) 2nd bisection

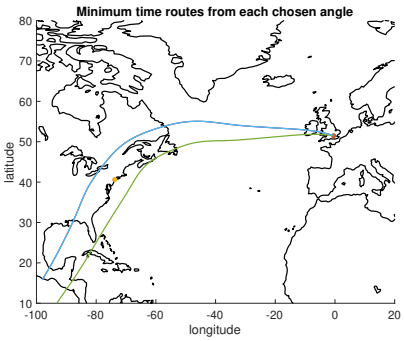

(d) 3rd bisection

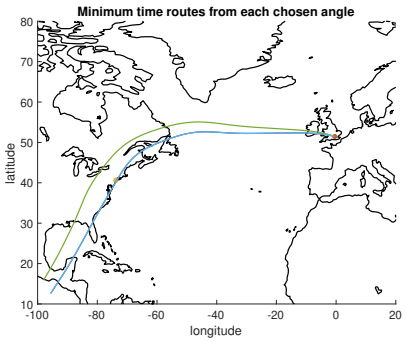

(e) 4th bisection

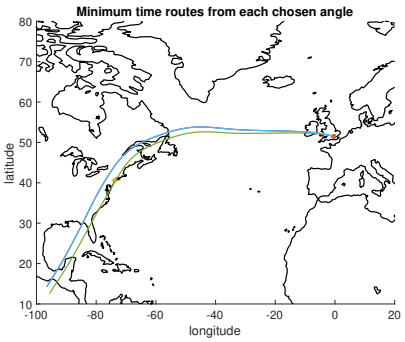

(f) 5th bisection

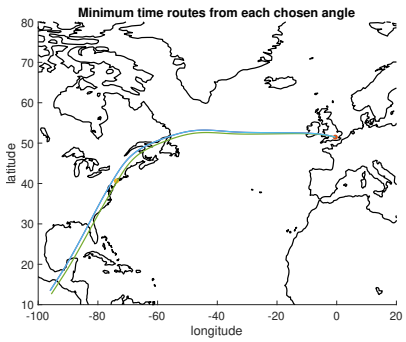

(g) 6th bisection

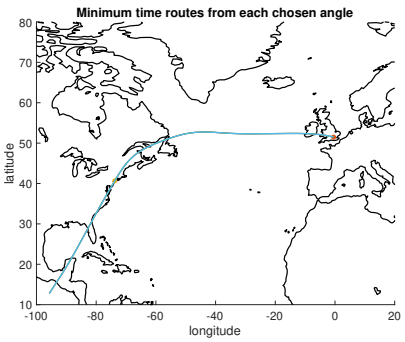

(h) 13th bisection

Figure 4.3: Initial routes and stages 1-6 and 13 of the bisection method for 1st December, 2019, westbound. New routes at each stage are in blue.

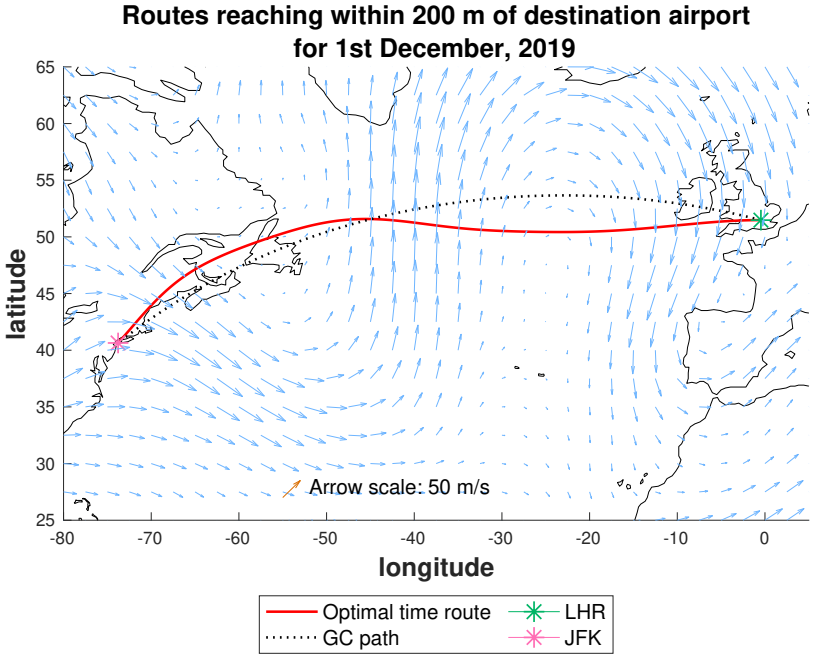

(a) Westbound Quiver Plot 1st December, 2019

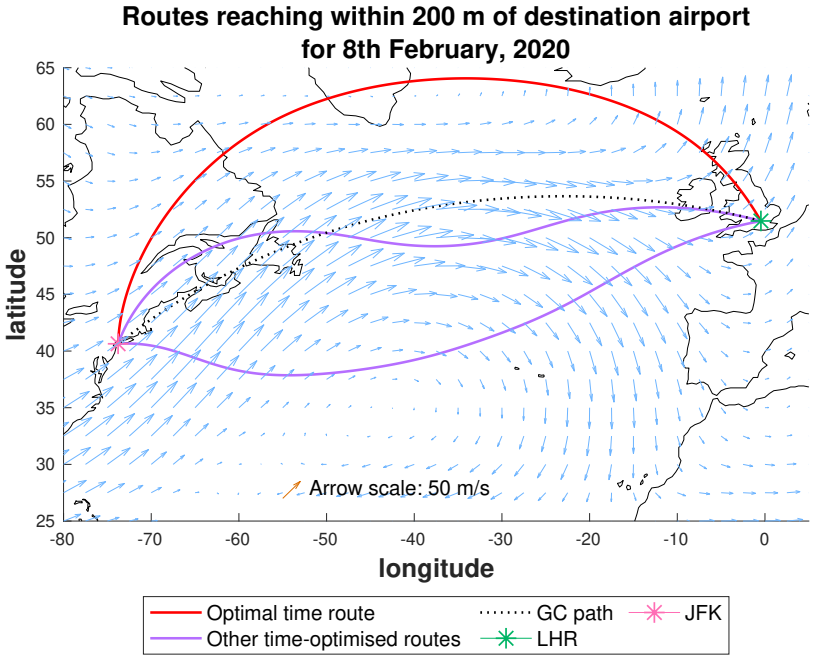

(b) Westbound Quiver Plot 8th February, 2020

Figure 4.4: Quiver plots showing values for wind velocity in m s^{-1} on 1st December, 2019 and 8th February, 2020 across the North Atlantic. The westbound time optimal flight paths for each day's wind field, from LHR to JFK (indicated by green and pink asterisks respectively) are shown in red. The two routes on 8th February, 2020 which also reach the airport, but have longer flight times are shown in purple. The GC path is marked as a dotted black line on each plot.

4.3.4 ATM tracks and the GC path

The minimised flight times are used to compute air distances, as a function of airspeed, for all flights between JFK and LHR. Mean time for westbound OFW routes, plus or minus standard deviation, was 415 ± 18.4 minutes, whilst eastbound it was 323.5 ± 13.7 minutes. The OFW route air distances are compared with those covered by aircraft following as ground tracks the GC path and the ATM tracks (produced by NAV CANADA and NATS), through the same wind fields. The corresponding air distances for these journeys are calculated as follows.

To validate comparisons between the GC path and the OFW routes, the GC path is split into 25 000 way points. For each pair of consecutive way points, the rhumb line bearing, a single heading that will take an aircraft from the first waypoint to the second around a sphere, is found and the ground speed for each interval calculated by solving the vector equation:

$$\mathbf{u}_{wind} + \mathbf{u}_{air} = \mathbf{u}_{ground}, \tag{4.28}$$

where $\mathbf{u}_{wind}$ is a vector comprised of zonal and meridional wind components, $\mathbf{u}_{air}$ is a vector giving the aircraft's velocity relative to the air and $\mathbf{u}_{ground}$ is a vector giving the aircraft's velocity relative to the ground (see Figure 4.5).
The sum of the rhumb line distances gives the total ground distance. The rhumb line distances for each interval are divided by the corresponding ground speeds to give the interval times, which are summed to give total flight time. Multiplying the flight time by the airspeed gives the air distance travelled in each case.

Air and ground distances for the ATM tracks are calculated in a similar way. The journey between each of the ATM track way points is divided into smaller intervals reflecting the ratio between the length of the journey between these way points and the whole journey distance. Again a total of 25 000 steps is used. Once these intervals are generated, the methods used for the GC path are applied.

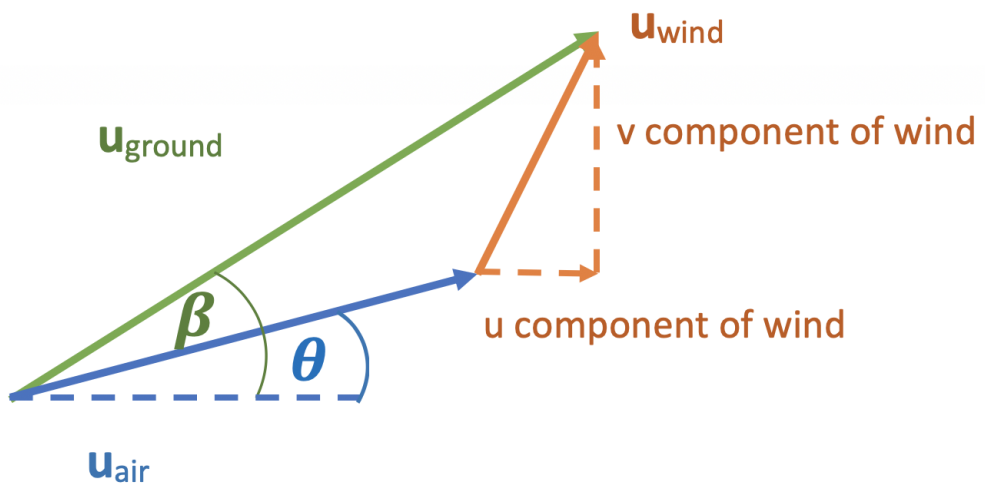

Figure 4.5: Diagram showing vector triangle of aircraft velocity relative to the ground, $\mathbf{u}_{ground}$, aircraft velocity relative to the air, $\mathbf{u}_{air}$, and wind velocity, $\mathbf{u}_{wind}$. Here θ is the aircraft's heading angle and β is the rhumb line or ground-track angle, which is the single heading an aircraft in wind-free conditions would take to reach the next waypoint.

4.4 Results

4.4.1 Air distance analysis

Air distances for the ATM tracks, GC path and OFW route are now compared. An example of the ground tracks for these paths is given in Figure 4.6. This shows that on the 1st December, 2019, the OFW routes going east are very similar to the most efficient ATM track, but that flying west there is more of a discrepancy. All routes except for the westbound ATM track cross the GC path. The savings in air distance made, can be seen in Figure 4.7, which gives air distances of each route for a range of airspeeds.

As expected the OFW route has the shortest air distance for all airspeeds both westbound and eastbound. Air distance reduces as airspeed increases for the westbound tracks, whilst increasing for eastbound flights, since in a strong wind field, airspeed has less effect on air distance when flying with the prevailing wind than against it. As air distance is airspeed multiplied by time, if time is greatly reduced by flying faster, the increase in airspeed does not mean necessarily that air distance is increased. This can be seen in westbound results. Conversely, when aircraft fly east, extra airspeed does not change the time the

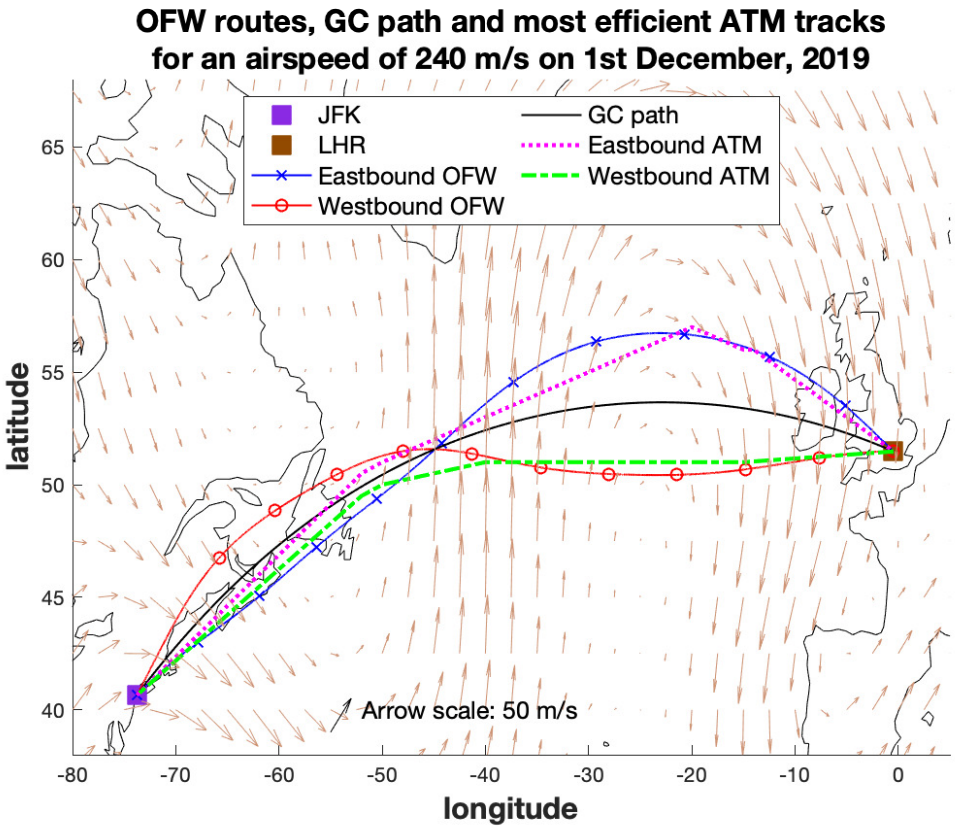

Figure 4.6: Plots to show the ground tracks of OFW routes and the most efficient ATM tracks flying both east and west between LHR and JFK. The GC path is also shown.

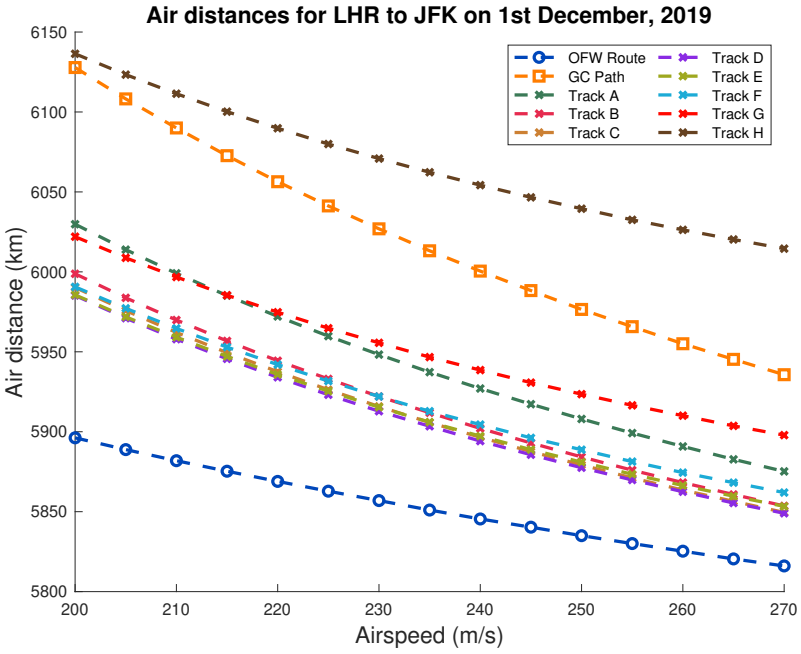

(a) Air distances 1st December, 2019 westbound

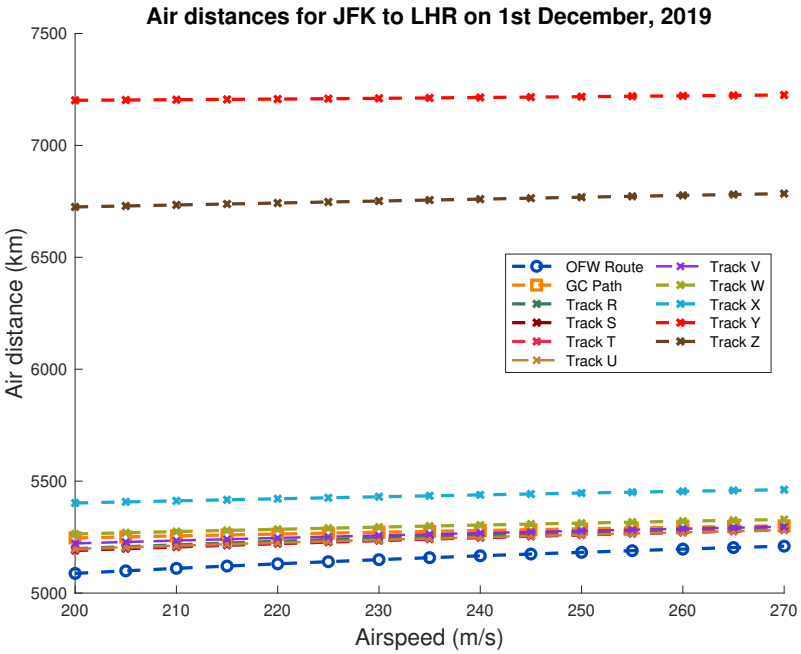

(b) Air distances 1st December, 2019 eastbound

Figure 4.7: The variation of air distance with airspeed for OFW routes, GC paths and ATM Tracks on 1st December, 2019.

route takes as dramatically, meaning that the airspeed factor in the air distance product becomes dominant and the air distance increases with airspeed. Both the GC path curves in Figure 4.7 approach the great circle distance between the two airports (5.6×10^6 m) as the airspeed increases, which is the expected behaviour, because the wind speed becomes less important and so the air distance tends to the ground distance.

The excess air distances that are incurred by flying the ATM tracks and GC path instead of the OFW route, averaged across the whole winter period, are shown in Figure 4.8. Here both westbound and eastbound air distance savings decrease with increasing airspeed, with the exception of savings made when the OFW route is compared with the least efficient eastbound ATM track. In this case the savings increase slightly as airspeed increases, demonstrating that once a track is very far from the advantageous eastbound winds, flying at a higher airspeed will only burn even more fuel and produce even more emissions, without reducing air distance. However, in all other cases, there is less of a saving in air distance as airspeed increases, meaning that increased airspeeds are needed for less advantageous winds.

Taking the results for an airspeed of 240 m s^{-1} and averaging savings in air distance between the most efficient ATM track and the OFW route across all 91 days of winter 2019-2020 for flights from JFK to LHR, gives an air distance saving of 37 km, but the saving for the least efficient ATM track is over 931 km. The average saving for all ATM tracks is 232 km. In the opposite direction, flying from LHR to JFK, 54 km of air distance are saved by using the optimised route compared with the most efficient ATM track and air distance is reduced by 502 km compared with the least efficient ATM track. The average across all tracks gives a saving of 173 km. Westbound ATM tracks are normally closer on average to the most efficient track than those going east, explaining the larger eastbound average savings in comparison to each day's least efficient track.

Averaged over all 91 days of data, the difference in air distance between the OFW route and each ATM track is statistically significant at the 95% level using a one-tailed t-test (Student, 1908) with unequal variances, for each airspeed. In all cases the adoption of the OFW routes significantly reduces air distance. Percentage improvement in air distance, for flights from JFK to LHR flown at 240 m s^{-1} show a saving of 0.7% when the OFW route is compared to the most efficient ATM tracks, but 16.2% when compared to the

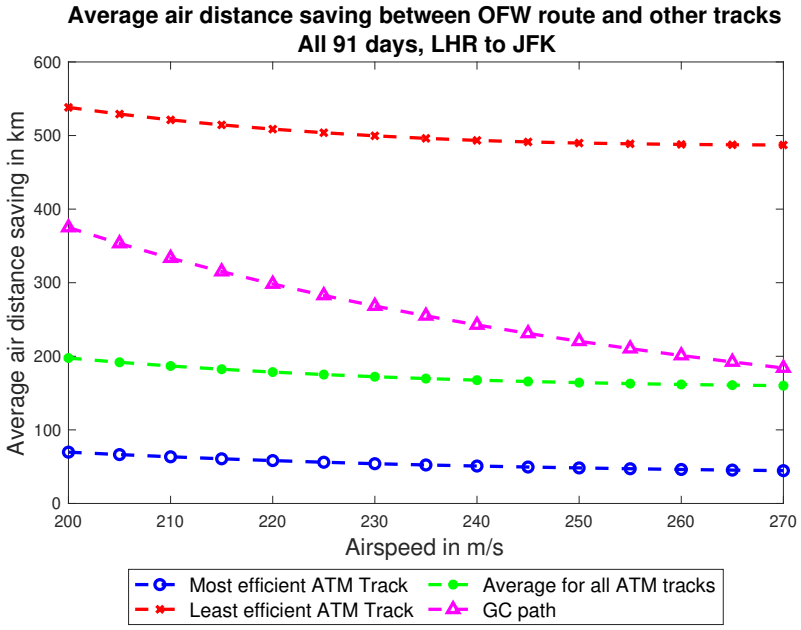

(a) Average air distance westbound

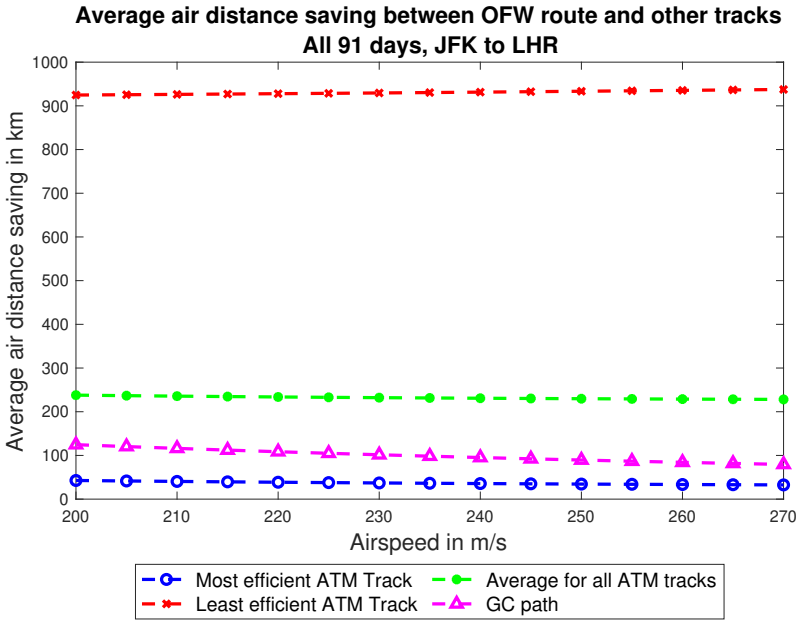

(b) Average air distance eastbound

Figure 4.8: Plots of air distance saved when the OFW route is compared to the most efficient ATM track, the least efficient ATM track and the average ATM track for each day at each airspeed. Savings compared with the GC path through the same wind field are also shown.

least efficient ATM tracks. The average across all tracks gives a saving of 4.7%. For the same journey in reverse the range of savings is from 0.8% to 7.3%, with an average of 2.9%. Table 4.1 shows all percentage savings across both routes, for a range of airspeeds. Given the large number of flights using these routes every day, adopting the OFW routes would save a significant amount of air distance and thus a significant amount of fuel and greenhouse gas emissions.

4.4.2 Observed usage of ATM tracks

As shown in the previous section, air distance savings depend to a large extent on the ATM track flown, so track usage statistics are crucial when estimating potential savings. We have obtained such statistics from NATS, giving the number of times each track was flown in each direction each day. A limitation is that all mid-North Atlantic flights are included, not just those between JFK and LHR. Also data for 14th and 15th December, 2019 are not available, so these days have been omitted from the calculations. Nevertheless, using a weighted average, in which tracks are used in the same ratio as in the NATS daily track usage figures, to look at differences in air distance between OFW routes and ATM tracks, provides some insight into potential percentage savings. Figure 4.9 shows that if airlines use all provided tracks equally their flights are less fuel efficient than the track usage figures from NATS imply. Therefore airlines already have a good idea of where the most fuel and time efficient routes are each day. The percentage saving obtained, using this weighted average for flights at a constant airspeed of 240 m s^{-1}, is 2.5% for eastbound flights and 1.7% for those flying west. Thus with more flexibility allowed in the track system, fuel savings could be made, enabling an important reduction in emissions.

Table 4.1: Table showing percentage savings in air distance for routes in both east and west directions at a range of airspeeds when the OFW route is compared to least (L) and most (M) efficient ATM tracks over the 91 day period between 1st December, 2019 and 29th February, 2020.

Airport		**Airspeed in m/s**															
Dept	**Dest**	**200**		**210**		**220**		**230**		**240**		**250**		**260**		**270**	
		M	L	M	L	M	L	M	L	M	L	M	L	M	L	M	L
LHR	JFK	1.08	7.78	0.99	7.59	0.92	7.46	0.85	7.37	0.81	7.32	0.77	7.29	0.74	7.30	0.72	7.31
JFK	LHR	0.90	16.45	0.85	16.37	0.80	16.29	0.77	16.23	0.73	16.18	0.70	16.13	0.68	16.10	0.66	16.06

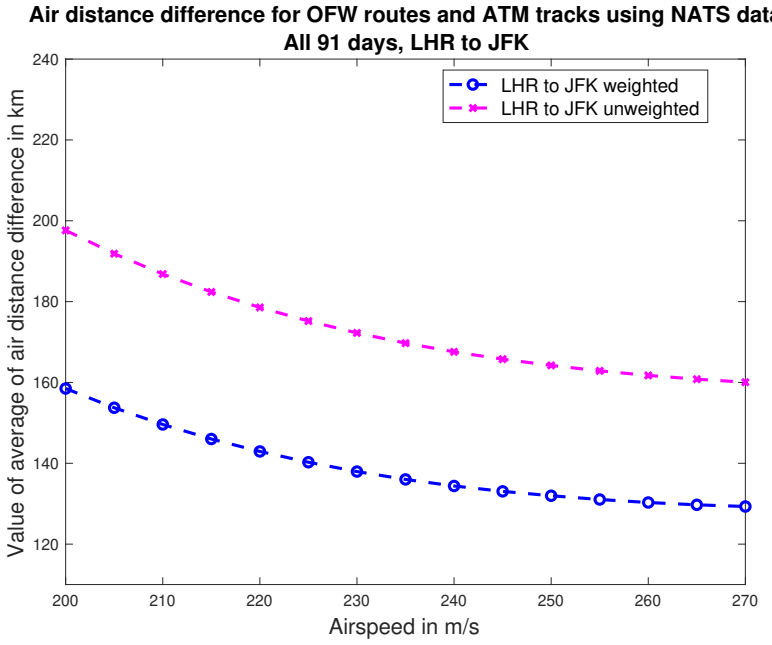

(a) Average air distance difference westbound

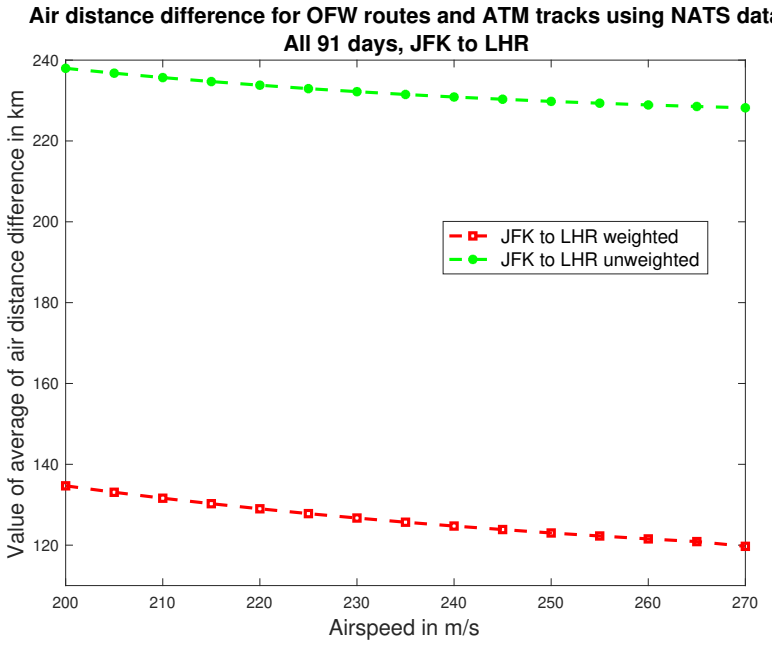

(b) Average air distance difference eastbound

Figure 4.9: Plots to show the savings to be made on average at each airspeed for aircraft flying both east and west along the route between LHR and JFK. Averages are calculated as if all tracks are used equally for the unweighted results and then as if tracks are used in the same ratio as in the NATS daily track usage figures for all mid-North Atlantic crossings, for the weighted results.

4.4.3 Limitations

The limitations of the current approach lie in the simplification of the underlying model. Aiming for the minimum time for a route does not necessarily fit the reality of a timetable of scheduled flights. However, other assumptions such as maintaining the same altitudes and airspeeds for the cruise phase could be adjusted in future work to allow flight times to be fixed. This would allow aircraft to benefit from different wind fields at different altitudes. In Figure 4.10a average actual flight times from LHR to JFK are compared with those found by time optimisation across each daily wind field, for all flights made from 1st December, 2019 to 29th February, 2020. The actual flights have varying airspeeds, with cruising airspeeds of between 230 m s^{-1} and 250 m s^{-1}. As the actual flights are not entirely in cruise phase and are not optimised for time, most actual flight times are longer than the OFW route times, as is shown by the data points lying above the red line. However, on a handful of days the opposite is true, as flights are not constrained to a single airspeed or altitude. The green and pink lines on the plot show the shortest and longest scheduled times for a crossing from LHR to JFK. As data is widely spread from these boundaries, this reveals that current scheduling includes much taxiing and contingency time, timetables having been written long before airlines know the nature of the daily wind fields. However, during this period, eleven westbound flights did exceed their scheduled time, presumably due to exceptionally strong winds.

In Figure 4.10b the same patterns are shown for the journey from JFK to LHR. In this case, all flights have shorter OFW route times. All eastbound flights during the 2019-2020 winter period landed well within their scheduled times.

Both OFW route times and average actual times are found to be normally distributed in a single-sample Kolmogorov–Smirnov test, at a 95% level of significance, so a product moment correlation coefficient can be calculated. This shows that there is a strong positive correlation between datasets, as shown by the turquoise line of best fit in Figure 4.10. The shallow gradient of this line, which is less than one for both eastbound and westbound flights, shows that as flights take longer, the difference between the actual flight times and the optimised flight times becomes smaller, as taxi, take-off and landing phases take a set time and so for longer flights they become a smaller percentage of the whole trip. The line of best fit for the westbound journeys is steeper than that for the eastbound journeys, as the OFW routes save more air distance in this direction.

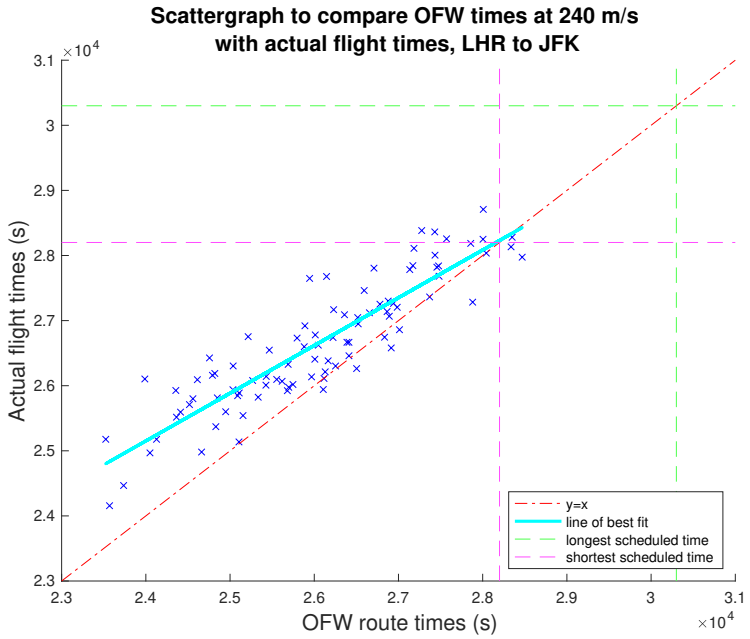

(a) Westbound flights

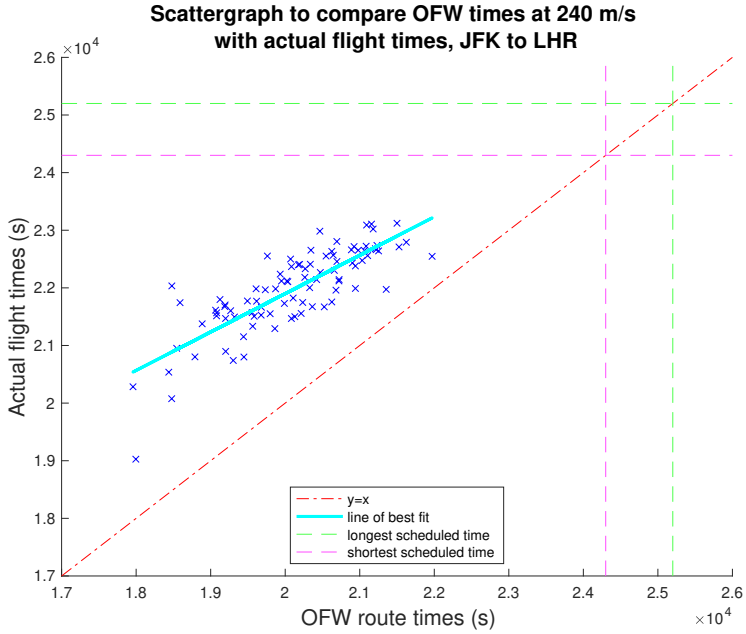

(b) Eastbound flights

Figure 4.10: Figure 4.8a shows how the optimal westbound times calculated for each day's wind field varied with the average actual flight time recorded each day from JFK to LHR. Figure 4.8b shows the corresponding information for the eastbound route from LHR to JFK.

4.5 Chapter summary

Using wind-optimised tracks would reduce North Atlantic air distances significantly, even in comparison to the most efficient current ATM tracks. This research has been undertaken at the constant static pressure level of 250 hPa, as this is the level at which the tracks are optimised (Mangini et al., 2018). At present, the ATM tracks are heavily constrained by safety considerations driven by the poor situational awareness available over the mid-Atlantic. This results in a track system that is sub-optimal for fuel use and thus airlines are obliged to produce excess emissions. Airlines also currently choose routes which minimise the CI of a flight (the ratio of the time cost to the fuel cost), not the fuel consumption or emissions. With more reliable and high-resolution situational awareness available, greater track flexibility should be possible enabling worthwhile fuel savings to be made by the airlines and reducing emissions. This work demonstrates a way to compare the current North Atlantic ATM Tracks with routes generated by a time optimisation method based on solutions of the spherical version of the Zermelo equations. Air distance is used as a measure of the efficiency of ATM tracks. Savings of between 0.7 and 16.4% in air distance can be made by adopting time-optimised routes through each daily wind field, with level of savings dependent on flight direction and chosen ATM track.

To estimate the potential CO_2 emissions savings over a whole winter period, consider the 3 833 701 seats provided between New York and London in 2019 (OAG, 2020). According to the ICAO carbon emissions calculator (ICAO, 2020), an economy class return flight between London and New York generates 670 kg of CO_2 per passenger. Taking an airspeed of 240 m s^{-1}, an average saving of 1.7% can be assumed for the 479 333 passengers flying west over the winter period and an average saving of 2.5% can be assumed for the 479 093 passengers flying east. These figures are derived from the weighted averages for the air distance savings as discussed in Section 4.4.2 and from assuming that one quarter of the annual flight figures provided by OAG (OAG, 2020) pertain to winter flights. This gives a potential saving of over 6.7 million kg of CO_2 emissions across the winter period of each year alone.

Here we investigated optimal flight trajectories given a deterministic wind field from re-analysis data. In Chapter 5 we will consider atmospheric data from climate models allowing us to compare historic and future time minimal flight durations.

Chapter 5

Future changes to minimum flight times for transatlantic trajectories

5.1 Introduction

Just as the carbon dioxide emitted by aircraft is contributing to climate change, so the warming climate is likely to affect aviation. These changes could take many forms, such as increased turbulence due to amplification of the shearing of the jet stream (Williams & Joshi, 2013; Williams, 2017; Storer et al., 2017; Lee et al., 2019; Kim et al., 2015) or take-off weight limitations based on the reduced lift afforded by warmer air (Coffel & Horton, 2015; Gratton et al., 2020). In this chapter the effect on time minimal transatlantic flight duration, between LHR and JFK caused by the predicted future changes to both the position and strength of the jet stream will be investigated. We use average daily wind fields, so consideration of turbulence and convection is not possible. For both of these effects atmospheric data would need very high spatial and temporal resolution.

Previously increases in round trip flight time across the Pacific have been the subject of research by Karnauskas et al. (2015). By comparing time optimal trajectories between Hawaii and continental North America in 2006, to those predicted in 2100, the research showed that the predicted strengthening of winds at 300 hPa could increase round-trip flight times by just under a minute, leading to the potential production of an extra 4.6 million kg of CO_2 per year. This work was based on 34 climate models from CMIP5 (PCMDI, 2018). The data used is from the RCP8.5 scenario, which assumes a rise in average global temperature of 4.3°C by 2100, compared to pre-industrial temperatures.

Optimal transatlantic flight cruise times under climate change were scrutinised by both Williams (2016) and Irvine et al. (2016). In Williams (2016), the average round trip time increase for routes between JFK and LHR was found to be 1 minute 18 s. This amounts to 70 million kg of extra CO_2 emissions annually. Here flight times were considered for 20 years of pre-industrial CO_2 levels and compared with those from 20 years of doubled CO_2. All simulated flights were run at an airspeed of 250 $\mathrm{m\,s^{-1}}$ and an altitude of 200 hPa. Wind data was from the single model GFDL CM2 from CMIP3 (PCMDI, 2007).

Results from Irvine et al. (2016) showed that, on average, eastbound routes in the period from 2073 to 2099 were likely to be just under a minute shorter than those flown between 1979 and 2005, but that westbound routes would be about a minute longer. Here a constant airspeed of 250 $\mathrm{m\,s^{-1}}$ was again assumed, but the cruise altitude was 250 hPa. Simulations were based on five different climate models from CMIP5, including GFDL ESM2G and MPI-ESM-MR (see Table 5.1).

To quantify the more immediate effects of climate change on commercial flights it was decided to compare minimum flight times for a historic period, 1986-2005 and a future period, 2036-2055. An ensemble of eight different CMIP5 models was used under the RCP8.5 scenario. Both the summer season (1st June to 31st August) and the winter season (1st December to 28th February) were considered, giving 20 summer and 19 winter seasons in each time period. We also compare average wind speeds at different latitudes in these two time periods to explain our flight duration results.

5.2 Choice of climate models

Eight different climate models (shown in Table 5.1) are considered in finding minimal time trajectories. The flight times for aircraft crossing the North Atlantic between LHR and JFK recorded for each model are averaged to give final times for analysis. A Met Office climate specialist chose the ensemble members based on previous research and full details of all of the models can be found in (Lowe et al., 2018). As each model works with different zonal (x) and meridional (y) grid points, the data was re-meshed (using linear interpolation) onto a grid with 35 grid points between the longitudes of the airports. The grid point latitudes were calculated to give a square mesh. Numbers of grid points in each direction in the original data are given in Table 5.1.

Table 5.1: Ensemble of climate models from CMIP5

Model	Grid points		Description
	x	y	
ACCESS1-3	192	144	General circulation climate model ARC Centre of Excellence for Climate System Science Australia
CanESM2	128	64	Fourth generation coupled global climate model Canadian Centre for Climate Modelling and Analysis Canada
CMCC-CM	480	240	Coupled atmosphere–ocean circulation model Centro Euro-Mediterraneo sui Cambianeati Climatica Italy
CNRM-CM5	256	128	5 existing Earth system models coupled by OASIS Centre National de Recherches Météorologique France
GFDL-ESM2G	144	90	Coupled atmosphere–ocean model with biochemistry NOAA Geophysical Fluid Dynamics Laboratory US
IPSL-CM5A-MR	144	143	Atmosphere–land–ocean–sea-ice full Earth system Institute Pierre Simon Laplace Climate Modelling Center France
MPI-ESM-MR	192	96	Couples atmosphere, ocean and land surface Max-Planck Institut für Meteorologie Germany
MRI-CGCM3	320	160	Atmosphere–land, aerosol, ocean-ice, cloud microphysics Meteorological Research Institute Japan

5.3 Optimal control formulation

As in Chapter 4, a dynamical system detailing the zonal and meridional progress of an aircraft at constant altitude through a wind field is used. This is based on the Zermelo equations mapped conformally onto a spherical shell (Arrow, 1949; Bryson & Ho, 1975):

$$\frac{\mathrm{d}\lambda}{\mathrm{d}t} = \frac{u + V\cos\theta}{R\cos\phi}, \tag{5.1}$$

$$\frac{\mathrm{d}\phi}{\mathrm{d}t} = \frac{v + V\sin\theta}{R}, \tag{5.2}$$

$$\frac{\mathrm{d}\theta}{\mathrm{d}t} = -\frac{Wind2D}{R\cos\phi}, \tag{5.3}$$

where the $Wind2D$ term is:

$$\begin{aligned} Wind2D = &-\sin\theta\cos\theta\frac{\partial u}{\partial\lambda} + u\cos^2\theta\sin\phi + \cos^2\theta\cos\phi\frac{\partial u}{\partial\phi} \\ &+ v\sin\theta\cos\theta\sin\phi + \cos\theta\sin\theta\cos\phi\frac{\partial v}{\partial\phi} + V\cos\theta\sin\phi - \sin^2\theta\frac{\partial v}{\partial\lambda}. \end{aligned} \tag{5.4}$$

Here t is time from departure in s, $\lambda(t)$ and $\phi(t)$ are the aircraft's longitude and latitude in radians, $u(\lambda,\phi)$ and $v(\lambda,\phi)$ are zonal and meridional wind speeds in $\mathrm{m\,s^{-1}}$, R is the radius of the Earth, taken as 6 371 km, V is the airspeed in $\mathrm{m\,s^{-1}}$ and θ is the aircraft's heading angle in radians. Heading angle is taken as the direction in which the nose of the aircraft is pointing, measured anticlockwise from due east. A detailed derivation of this system can be found in Chapter 4.

As before, the dynamical system is solved numerically using an Euler forward stepping scheme, but this time Mach number (ratio of airspeed to speed of sound) rather than airspeed is kept constant, resulting in an extra parameter update at each timestep:

$$V = 0.82 * 340.3 * \sqrt{T/288.15}, \tag{5.5}$$

where T is temperature in degrees Kelvin at the start of a time step (TSCM, 2020). The constant Mach number used is 0.82, with 340.3 $\mathrm{m\,s^{-1}}$ giving the speed of sound at 288.15 K. This equation is drawn from the fact that airspeed is equal to the product of Mach number and speed of sound at the point under consideration. As there is very little variation in temperature at 250 hPa across the North Atlantic, this has very little effect on the airspeed, which remains very nearly constant. This modification was included to align with Met Office practices, as the material in this chapter was researched during a placement with the Aviation Group of the Met Office.

5.4 Obtaining optimal flight times and their associated trajectories

By integrating Equations (5.1) to (5.3) forward in time, trajectories can be traced from a departure point by choosing an initial heading angle. Paths obtained in this way will be time minimal for the given heading angle, but there is no guarantee that a path will reach the destination airport. A shooting method is used to improve estimates for the initial heading angle in order to reach the final airport. Initially trajectories are found using angles 60° either side of the first heading angle of the Great Circle Route between the airports. Using a bisection method the heading angle is improved until the route reaches the destination airport. In over 95% of cases this is an adequate approach. However, where trajectories are complicated by unusual wind patterns, multiple local minima mean that the bisection method is no longer valid. In these instances a new approach is applied. Trajectories are sent from angles in the original range at 1° increments. A hexagon is then created joining the closest point to the airport on a trajectory with the points 100 s before and after it and the corresponding three points on the next trajectory. An example showing a complicated wind pattern in which the destination airport lies in multiple hexagons is given in Figure 5.1. This is based on climate data from the MPI-ESM-MR model, for the 12th December, 1988. A sketch illustrating the new approach is shown in Figure 5.2.

Finding all of the initial angle pairs with trajectories forming part of a hexagon containing the airport, allows us to apply the method repeatedly using increasingly small angle intervals and increments, until paths are found with endpoints within 200 m of the destination airport. The path with the shortest associated flight time is adopted as the time minimal path. To demonstrate the method more clearly a less complicated set of trajectories has been chosen, as shown in Figure 5.3a. This example is based on weather data from the CNRM-CM5 model, for the 9th January, 1987.

By plotting a path due east from the destination airport (as shown in Figure 5.3b) and counting how many sides of each hexagon it intercepts, it is possible to locate any hexagons containing the destination. (In the case of a point lying within a polygon, a horizontal line projected from the interior of the polygon to the right will always intersect an odd number of sides.) The initial heading angles generating paths that produce each of these hexagons can now be examined further. The interval between each pair of such angles is

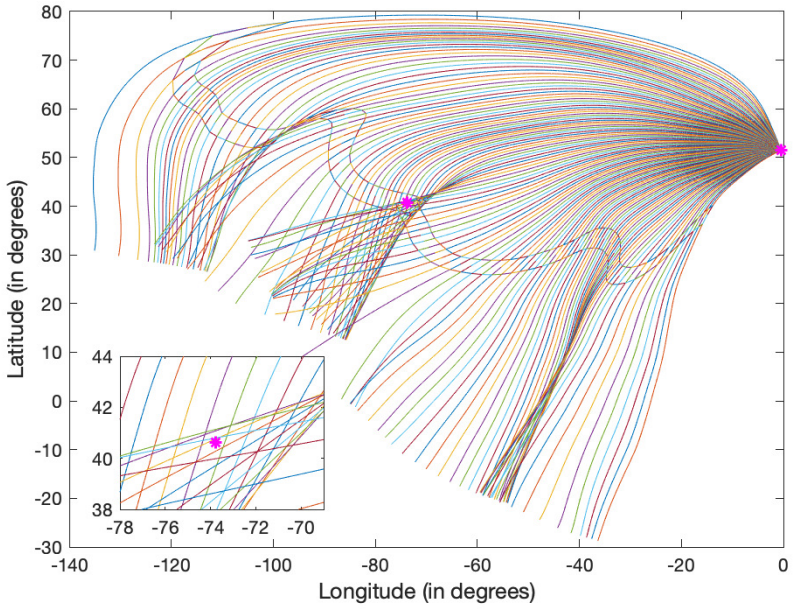

Figure 5.1: Plot showing optimal trajectories from a range of initial heading angles from -60 to 60 degrees either side of the GCR path. The departure and destination airports are shown by stars.

divided into smaller increments and the new trajectories and hexagons plotted (see Figures 5.3c and 5.3d). This can be repeated until a trajectory comes within 200 m of the destination airport (see Figures 5.3e and 5.3f).

This new method is more computationally expensive than the bisection method, so is used only as necessary. Working to four decimal places of a degree ensures that a path close enough to the destination is found where this exists, which normally requires three levels of angle increments being used. Although Pontryagin's Maximum Principle gives only necessary and not sufficient conditions for such a path to exist, in this particular research, time minimal paths were found for all climate models, on all days, in both directions.

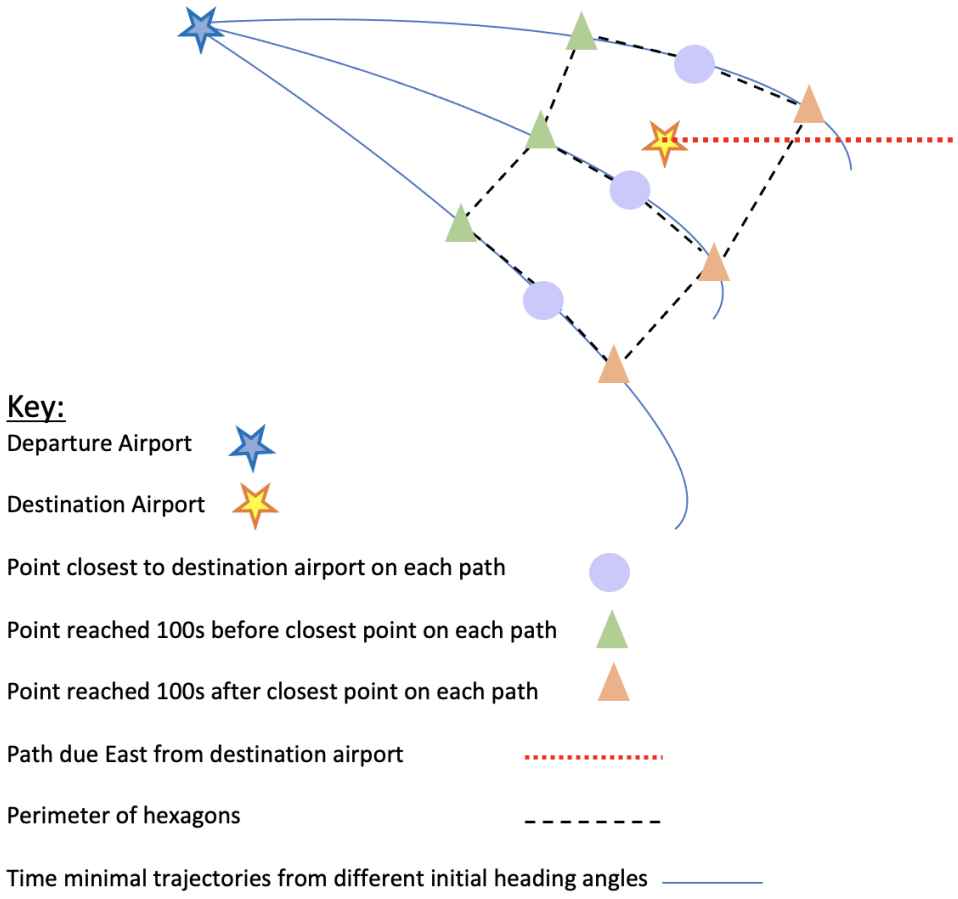

Figure 5.2: Sketch showing how the pair of trajectories between which the destination airport lies can be found by creating hexagons on neighbouring paths and seeing how many sides of each hexagon a line due east from the destination airport crosses.

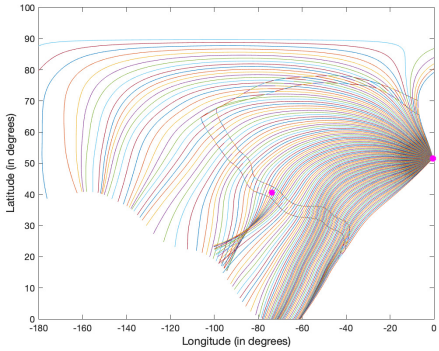

(a) Step 1

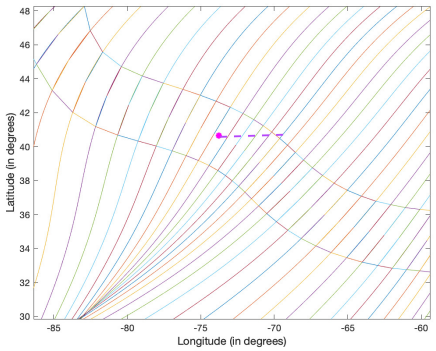

(b) Step 1 zoomed in

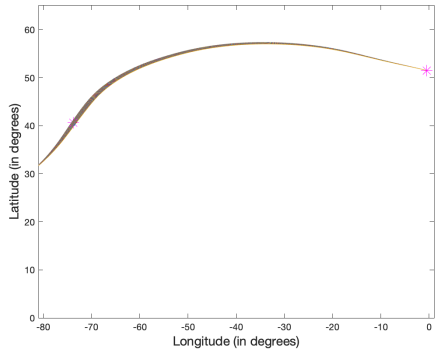

(c) Step 2

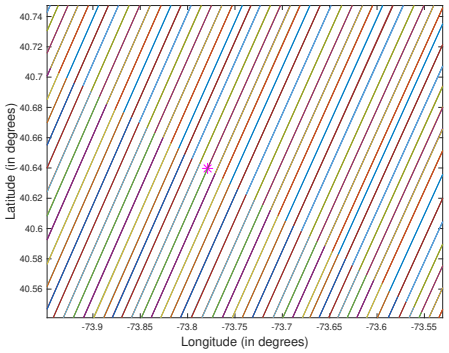

(d) Step 2 zoomed in

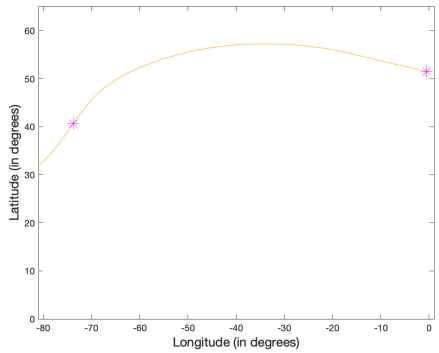

(e) Step 3

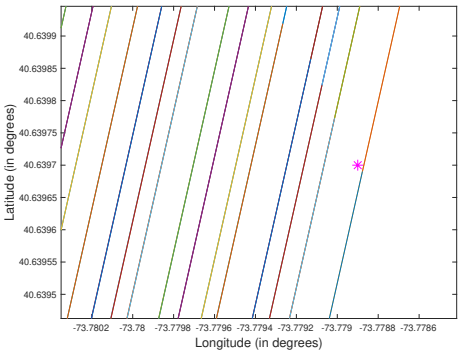

(f) Step 3 zoomed in

Figure 5.3: Plots of trajectories for a number of different initial heading angles. At each step the choice of angles becomes more focused and the interval between pairs of angles smaller, so longitude scale changes between Figures 5.3b, d and f. Changes in line colour show where each side of the hexagons have been drawn over the flight paths.

5.5 Results

Minimum times for each day, of each season of each year, for each model were collected, together with co-ordinates for each minimum time trajectory. Times were averaged across the ensemble of models to give one minimum time for each day of weather data. These results are shown as relative frequency histograms in Figures 5.4 and 5.5.

It can be seen that the future minimum times seem slightly shorter than those from the historic period in all cases, although this is more marked with eastbound journeys. A Student's t-test at a 5% significance level was run for the difference between mean times for corresponding historic and future minimum times. This showed that there is a significant decrease in time in all cases, but that this decrease is no more than one or two minutes, depending on season and direction of flight. Confidence intervals for the differences are shown in Table 5.2. The shorter flight times lead to an average reduction in round trip duration of 3 minutes in the summer and 2.6 minutes in the winter in the future, which is inconsistent with both Williams (2016) and Irvine et al. (2016), but agrees with findings in a new report from the Met Office, in which a different method is used to find minimum time paths, based on Djikstra's algorithm (Eurocontrol, 2021b; Cheung, 2018).

Table 5.2: Confidence intervals (at 5% two-tailed significance level) for difference in mean minimum flight times between historic and future time periods.

Season	Direction	Confidence Interval (s)
Summer	West	(38.3, 75.1)
	East	(104, 140)
Winter	West	(45.6, 98.0)
	East	(64.5, 106)

The trajectories, shown in Figures 5.6, 5.7, 5.8 and 5.9 give all points on all routes from each model. These are not averaged in any way, but instead include 27 360 flights in winter and 29 440 flights in summer. Plotting all of these routes allows the area of the North Atlantic flown over in each direction in each season to be seen. Flying West, aircraft do cover a wide range of latitudes, but most of the routes lie further North than the Great Circle Route (GCR) (shown in green on each plot). Eastbound, aircraft trajectories have less meridional range and cover either side of the GCR symmetrically in summer, with more of a tendency to fly south of the GCR in winter.

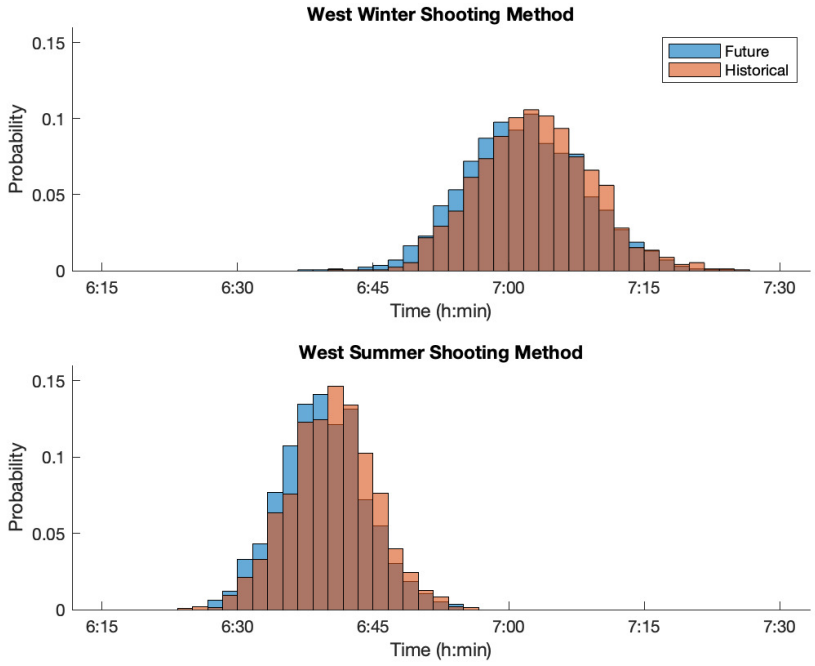

Figure 5.4: Relative frequency histograms showing minimum flight duration between JFK and LHR. Y axis values display the probability of the given range of flight times occurring.

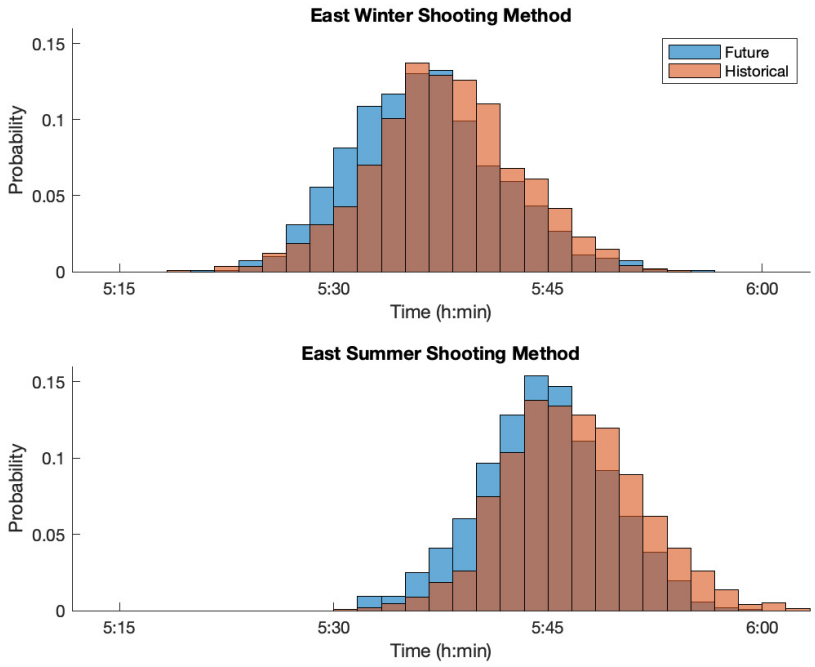

Figure 5.5: Relative frequency histograms for the minimum times for trajectories between LHR and JFK.

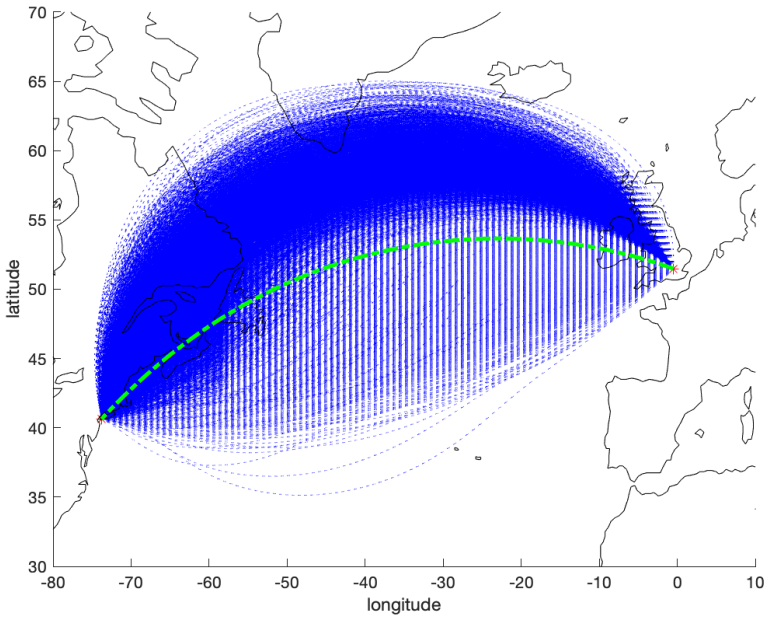

(a) Winter westbound, historic

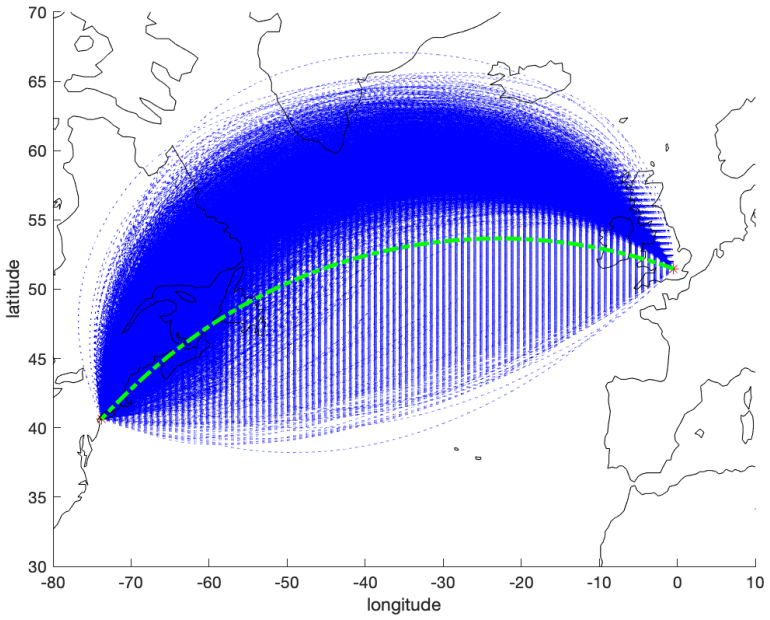

(b) Winter westbound, future

Figure 5.6: Plots showing trajectories of all westbound winter flights for wind fields provided by all climate models over the historic and future periods. The GCR between LHR and JFK is shown in green.

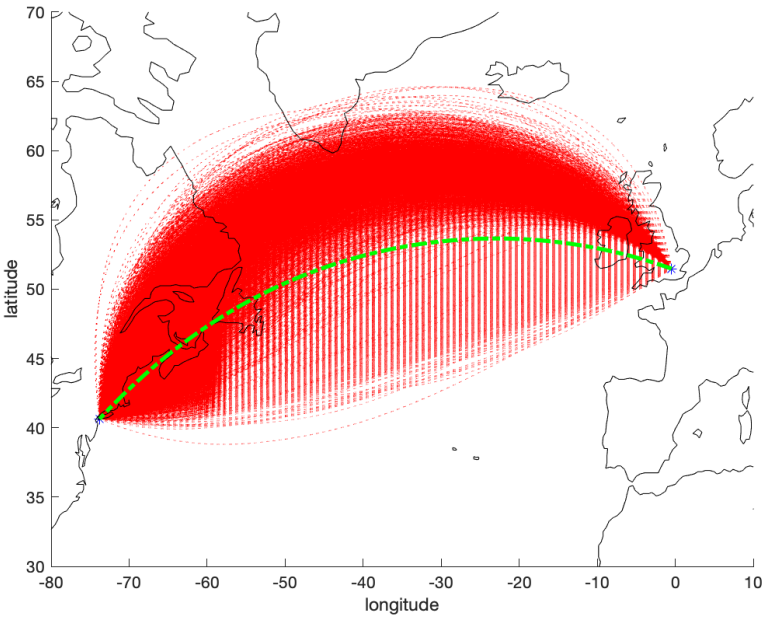

(a) Summer westbound, historic

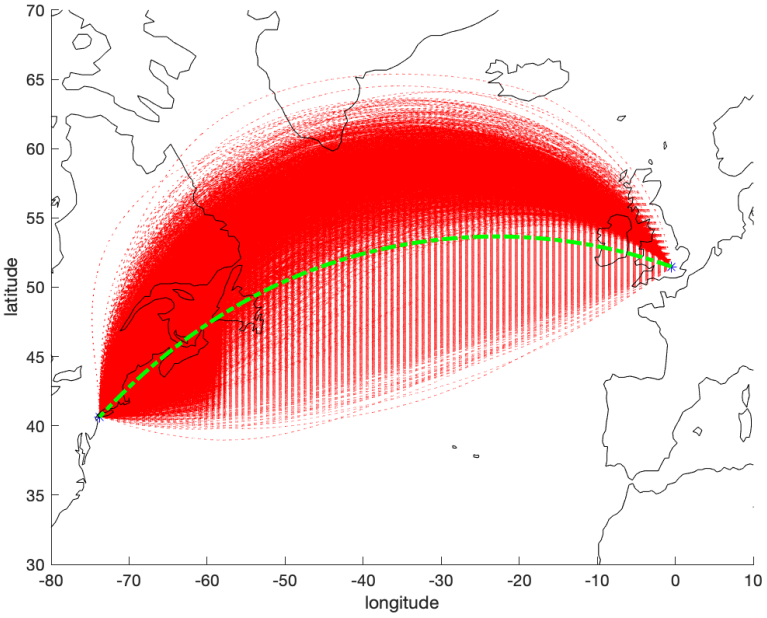

(b) Summer westbound, future

Figure 5.7: Plots showing trajectories of all westbound summer flights for wind fields provided by all climate models over the historic and future periods. The GCR between LHR and JFK is shown in green.

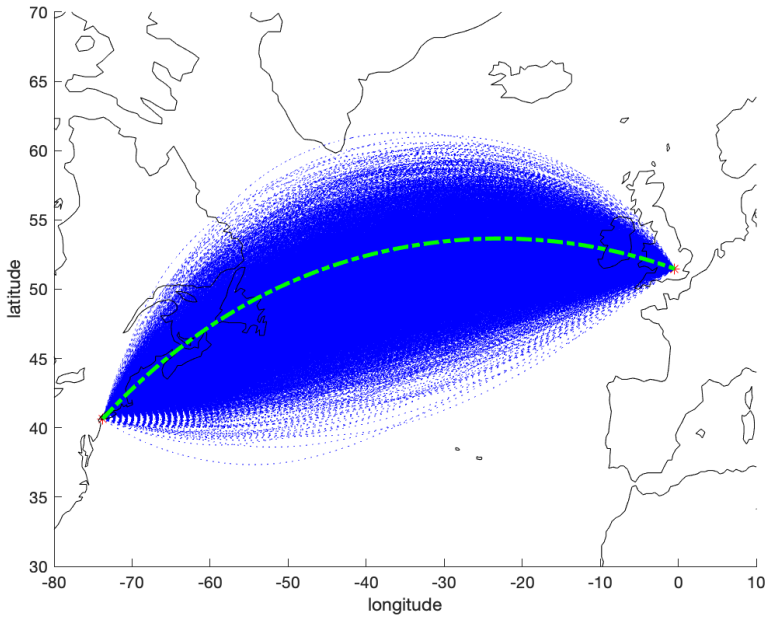

(a) Winter eastbound, historic

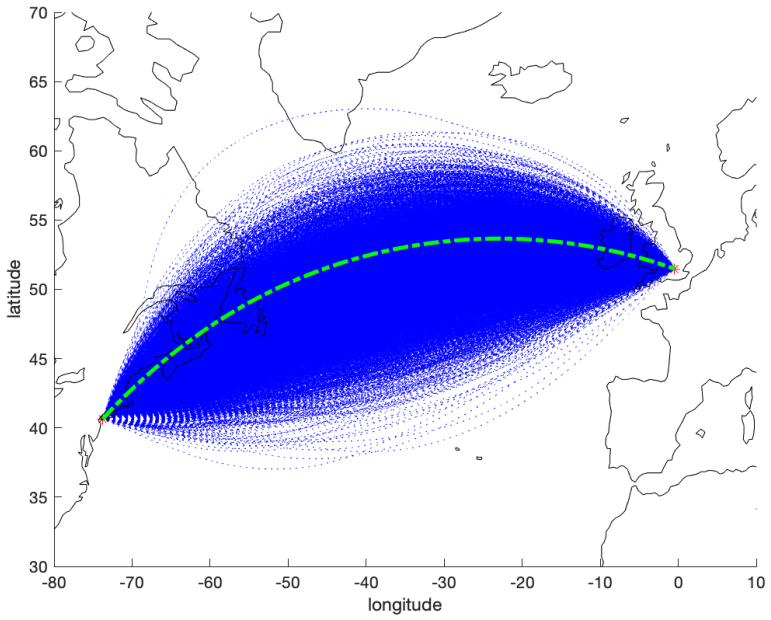

(b) Winter eastbound, future

Figure 5.8: Plots showing trajectories of all eastbound winter flights for wind fields provided by all climate models over the historic and future periods. The GCR between JFK and LHR is shown in green.

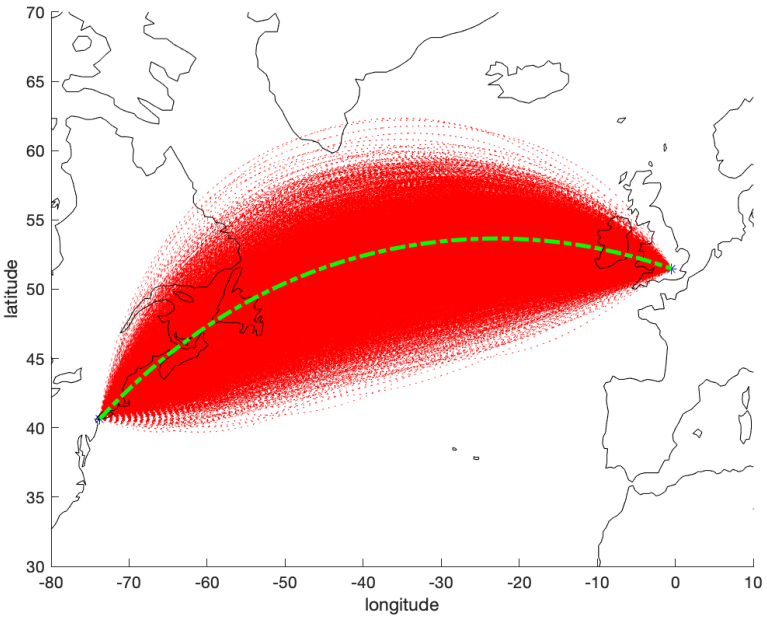

(a) Summer eastbound, historic

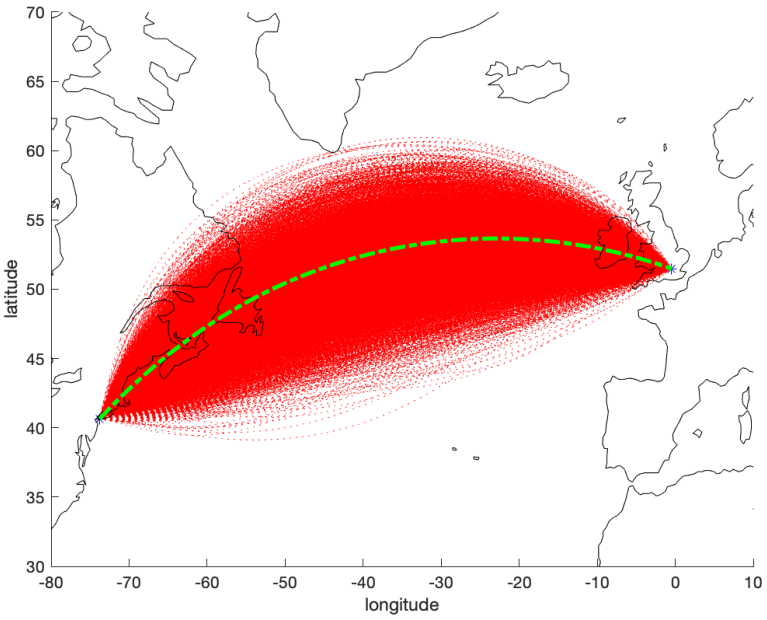

(b) Summer eastbound, future

Figure 5.9: Plots showing trajectories of all eastbound summer flights for wind fields provided by all climate models over the historic and future periods. The GCR between JFK and LHR is shown in green.

5.6 Discussion of results

Why should transatlantic flight durations become shorter in the future in both directions? This seems counter intuitive as faster winds blowing east would seem to imply faster eastbound flights, but slower journeys westbound. However, it can be explained by the patterns shown in the jet stream position and strength from the ensemble of climate models. The average zonal wind speed at each latitude was found for each season and for each time period.

Figure 5.10 shows that, in the winter, the models predict a faster jet stream core slightly further South than before, but a weaker jet stream at higher latitudes. As westbound flights avoid the core already, the weaker winds toward the poles allow for quicker journeys detouring to the North. Eastbound flights will benefit from the increased speed of the core, as the prevalent current is in the correct direction and the core itself is in the region of the shortest ground distance path. In the summer the jet stream core is expected to be of similar strength at the centre, but extend meridionally, both pole-wards and equator-wards, as shown in Figure 5.10. This results in westbound flights staying more central and not deviating so far North on longer ground tracks. Flying east the jet core is very similar, but the added meridional spread, means that tailwinds can be accessed across a greater portion of the selected route.

To illustrate these points further, graphs were plotted showing the latitude of each trajectory at the point where the longitude was midway between the departure and destination airports.

Figure 5.11a shows that, as predicted, in the winter more westbound flights detour North in the future period to avoid the stronger jet stream core. Flying east (Figure 5.11b) more flights take advantage of the stronger core across the central latitudes. Results for the summer trajectories are shown in Figures 5.11c and 5.11d. Fewer flights detour so far North when flying West, as the jet stream core is too spread to make this an advantage. Eastbound, aircraft are detouring further North and South, which corresponds to the wider jet stream. So the trajectories seem to reinforce the ideas given to explain how minimum time flights can be of shorter duration both flying east and west in the future.

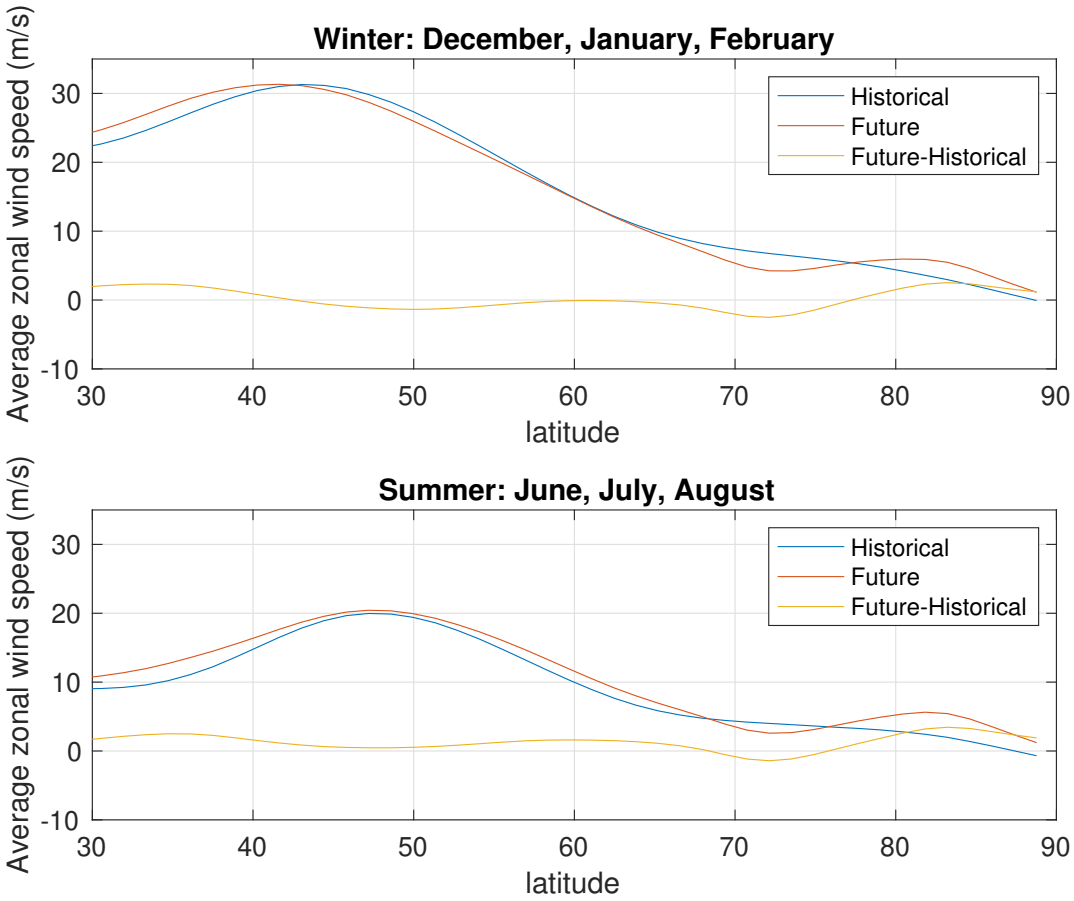

Figure 5.10: Plots showing average zonal winds at each latitude both for the historic and future climate model data. The winter season is illustrated above and the summer season below.The difference between historical winds and projections of future winds is also shown.

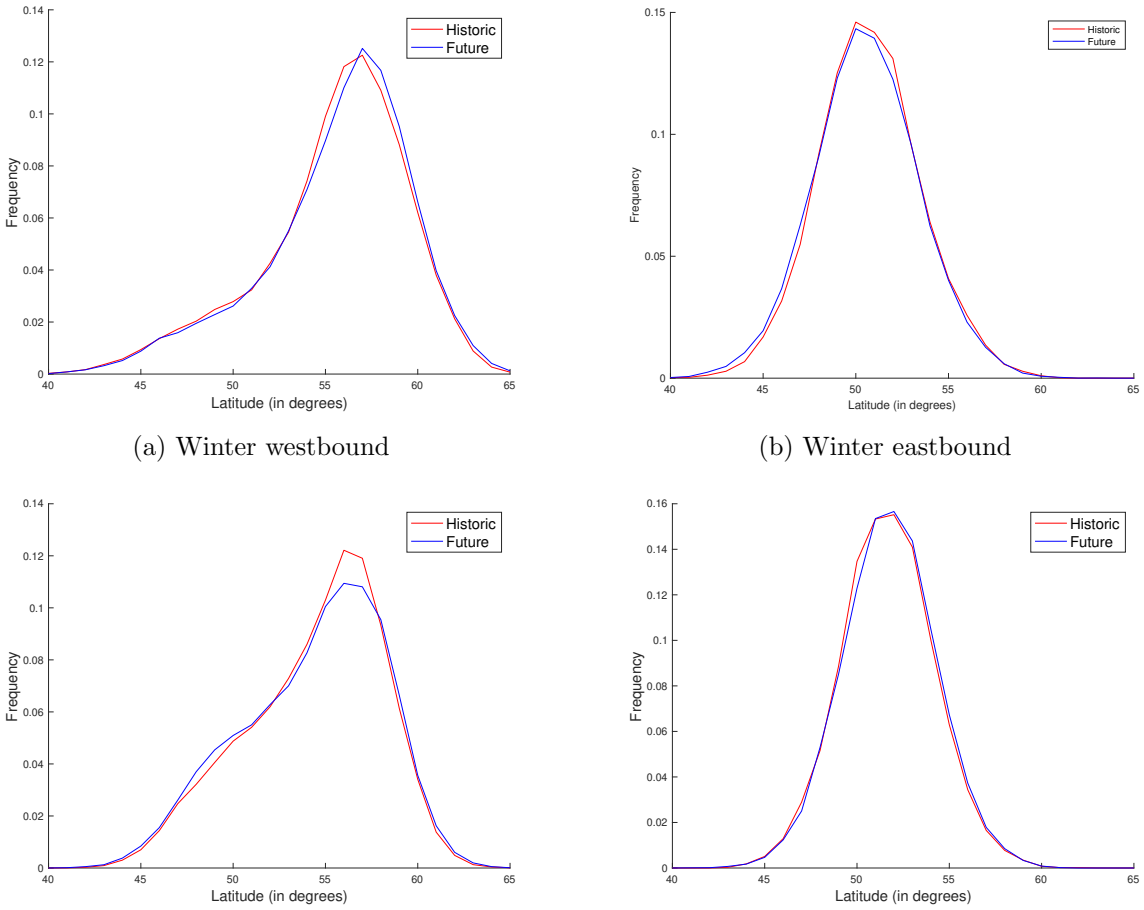

(a) Winter westbound

(b) Winter eastbound

(c) Summer westbound

(d) Summer eastbound

Figure 5.11: Plots showing the average latitude of aircraft when they reach a longitude halfway between the departure and destination airports.

5.7 Why the difference between these results and previous research?

There are two key reasons why results from this research differ from previous studies.

Firstly, in the current research, two relatively close time periods are being discussed. There is only a three decade gap between the historic and future time periods, whilst in Williams (2016) the weather data is pre-industrial to double pre-industrial, which would reflect a significantly longer interval between the historic and future periods. In Irvine et al. (2016) there was a seventy year interval between the two time periods considered. It is unlikely that changes to winds as the climate warms will be linear. Thus comparative results between historic and future periods can be expected to show different patterns.

In addition to the time period changes, there is also the reliance on different climate models. By using an ensemble, differences between models are to some extent averaged out, so results from an ensemble and results based on a single model are not really comparable. In Figure 5.12 the percentage flight time change from the historic to the future time period is plotted, for each of the climate models used for this research.

All models are shown to produce a faster average flight time eastbound in the summer, whereas flying west in both seasons and east in the winter some models show an increase in average journey duration. In Irvine et al. (2016) the GFDL ESM2G climate model formed part of the ensemble of five CMIP5 models used. Figure 5.12b shows an increase in average flight time for westbound flights in the winter for this climate model in this research too. If this pattern were repeated in the other four ensemble members used, it could explain why flights travelling west appear to be slower in the future.

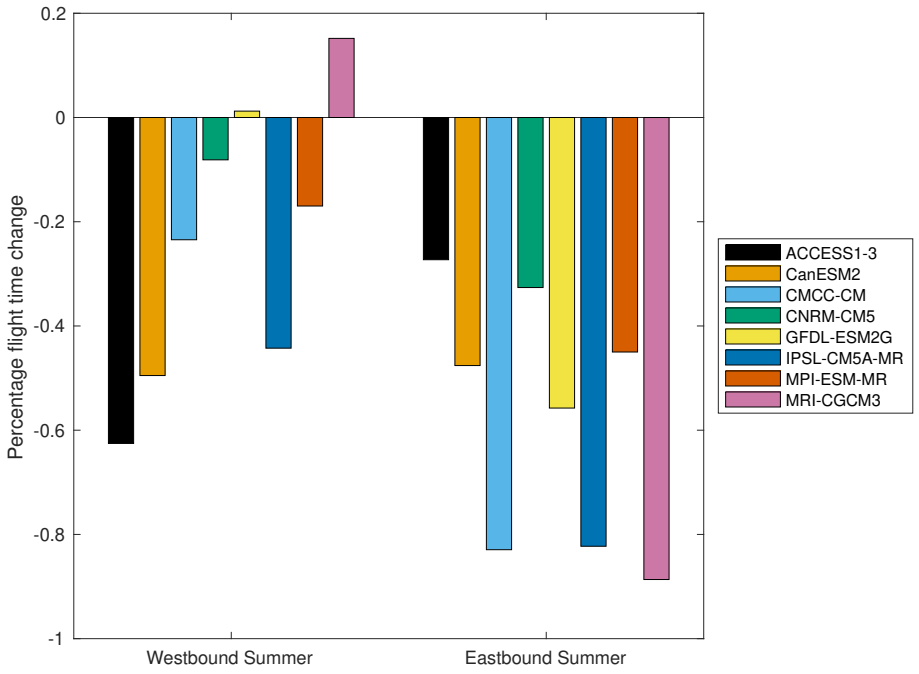

(a) Summer

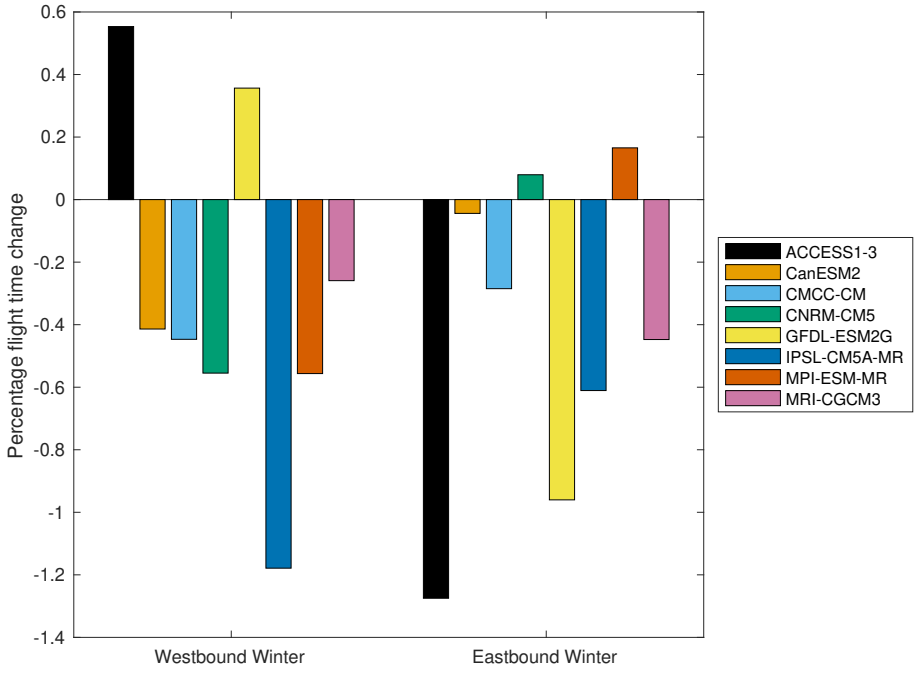

(b) Winter

Figure 5.12: Barcharts to show percentage change to flight time going from Historic to Future periods for each ensemble model.

5.8 Chapter summary

An ensemble of eight climate models is used to compare minimum flight times for the summer and winter seasons of a historic time period (1986-2005), with a future time

period (2036-2055). Results show that there is very little difference to flight durations, where routes have been optimised for time, across the North Atlantic in the next three decades. Any change to the future climate has been shown to reduce flight duration slightly in both directions. However, research from Williams (2016) and Irvine et al. (2016) shows that beyond this period, as the climate continues to warm, westbound flights could slow significantly.

The climate model data used in this research, looks at average zonal and meridional winds across a 24 hour period, so it does not capture local turbulence or convection, which are localised effects on a shorter timescale (Storer et al., 2017). As the climate warms, it is expected that both of these atmospheric phenomena will become more widespread, so although our analysis predicts faster transatlantic flights based on daily wind fields, there may be other effects that will need to be avoided in the future, adding to route length and thus journey times (Storer et al., 2017; Williams & Joshi, 2013; Williams, 2017; Lee et al., 2020; Kim et al., 2015).

Chapter 6

The role of airspeed variability in fixed-time, fuel-optimal aircraft trajectory planning

6.1 Introduction

In Chapters 4 and 5 minimal time routing was applied, but in this chapter it is explained that such a method is not necessarily suitable in a commercial setting.
With airlines, ATM and the ICAO keen to reduce the environmental impact of commercial flight (ICAO, 2020; Molloy, 2020; AirlinesUK, 2019), it has been shown that TBO provide the key to improving overall efficiency (Wickramasinghe et al., 2012; García-Heras et al., 2014). This approach focuses on designing individual flight routes accounting for weather conditions in order to optimise efficiency with regard to a number of different factors. For long haul flights, such as those between LHR and JFK, for which the cruise phase makes up around 92% of the ground distance of the flight (as shown in Wells et al. (2021)), TBO promises the greatest rewards (Girardet et al., 2014).

Time optimal routing has been shown to reduce fuel burn when altitude and airspeed are fixed, compared to the tracks provided across the North Atlantic by NATS and NAV CANADA (Wells et al., 2021). However, this option is not necessarily the most practical for airlines and airports. For and financial reasons airlines also need to adhere to their published timetables. Flights arriving early, create additional costs, through extra fuel burned in holding patterns, missed connecting flights, blocked gates and additional

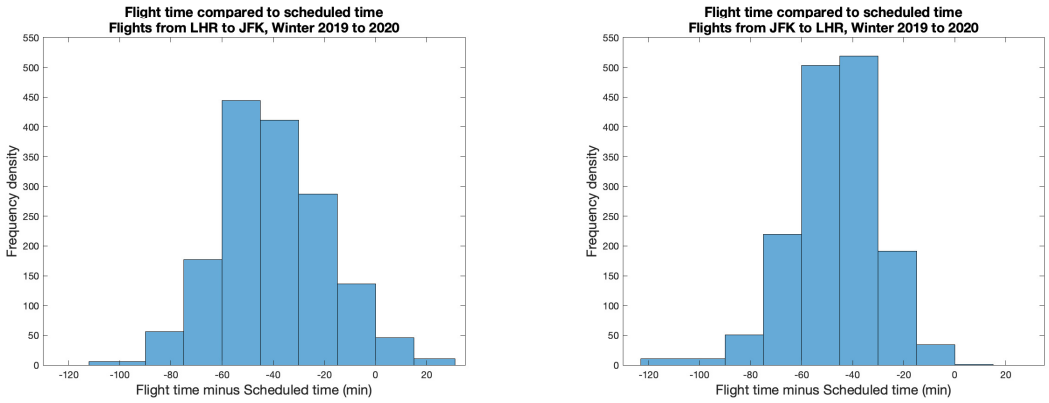

(a) Difference in times: West

(b) Difference in times East

Figure 6.1: Histogram showing differences between scheduled time and flight time for the 3 000 flights with available flight data, between LHR and JFK made from 1st December, 2019 to 29th February, 2020.

crew time. Customer dissatisfaction is also a key issue for airlines when delays occur. If flight trajectories are planned to ensure a fixed flight time, whilst minimising fuel burn, then the additional costs to both the airlines and the environment associated with early arrival can be saved.

Figure 6.1 shows the difference between scheduled times and flight times for the flights for which full data was available, flown between LHR and JFK from 1st December, 2019 to 29th February, 2020. Flight time is measured as the time from the aircraft wheels leaving the tarmac at the start of the journey, to them touching down on the destination airport runway. This is not the same as the scheduled time for a flight, as this is composed of the gate to gate time, thus including time between the gate to the runway at both ends of the journey and some contingency time.

Flights that arrive more than 15 minutes before or after they are scheduled to be at the arrival gate are considered early or late. Trying to judge this from flight times is difficult, as scheduled time does include what is known as “padding” to help to prevent delays (OAG, 2022), but this can result in passengers spending time in the aircraft once landed, waiting for a gate to be assigned. Taking into account taxiing times at both ends of a flight, a fair assumption is that any aircraft with a flight time 60 minutes or more shorter than its scheduled time is probably inconveniently early.

In the same way, any aircraft with a flight time less than 10 minutes shorter than its scheduled gate to gate time is probably unacceptably late, as not enough time has been left for movement between the runway and the gate. Using these rules, 15.2% of westbound flights and 18.5% of eastbound flights shown in Figure 6.1 arrived early and 7.5% of westbound flights and 1.0% of eastbound flights would be classed as late.

A very early arrival, at a busy time for an airport, can result in the aircraft being stacked for as much as 30 to 40 minutes. This circling of the airport vicinity at low altitude is particularly expensive in terms of fuel use and thus emissions produced, as the aircraft is being forced to fly at sub-optimal airspeeds and altitudes (Yamashita et al., 2015). In 2001, more than half of air traffic delays were caused by ATM (Lane, 2001). Since then a number of new initiatives have led to some improvements, despite runway capacity still being one of the biggest issues in ensuring flights arrive as scheduled (Soomer & Franx, 2008). Tactically enhanced arrivals mode at LHR has even seen some aircraft landing on the departure runway to reduce stack holding delays (Molloy, 2020).

If flight trajectories are planned to ensure a fixed flight time, whilst minimising fuel burn, then the costs to both the airlines and the environment associated with early or late arrival can be saved. For late arrivals it is clear that missed transfers, departing flight delays due to crew and aircraft not being available, ground crew rescheduling and crew overtime payments are all costly (Soomer & Franx, 2008). It is perhaps more surprising to learn that early arrivals can be just as costly in terms of missed aircraft rotations which could have generated extra revenue (OAG, 2022).

In order to use a fixed time across the whole of the winter 2019 to 2020 period, it was necessary to choose a journey time that would allow all flights to reach their destinations regardless of the winds encountered en route. The results described in Chapter 4 show that time optimal flight times westbound between LHR and JFK from 1st December, 2019 to 29th February, 2020 varied in duration between 6 hours 32 minutes (23 524 s) and 7 hours 54 minutes (28 469 s), so a fixed time of 8 hours 3 minutes (29000 s) is chosen based on the maximum duration recorded for time minimal flights, to ensure flights can be completed in any of the daily wind fields encountered. The time optimal model and the current fuel optimal model consider the entire route as cruise phase, so these times are slightly shorter than actual flight times. Scheduled times for actual flights on this route

over the same winter period are between 7 hours 50 minutes (28 200 s) and 8 hours 25 minutes (30 300 s), with an average of 8 hours 7 minutes (29 205 s) (Flightradar24, 2020).

Eastbound flight times are shorter due to the generally more favourable wind field. Time optimal routes from JFK to LHR vary between 4 hour 59 minutes (17 958 s) and 6 hours 6 minutes (21 971 s). Flights are scheduled for between 6 hour 45 minutes (24 300 s) and 7 hours (25 200 s), with an average time of 6 hours 54 minutes (24 819 s). The fixed time for eastbound flights is set to 6 hours 7 minutes (22 000 s) based on the time minimal maximum time duration, to allow enough time for flights across each day's wind field, when the entire flight is considered to be in the cruise phase. This is shorter than scheduled times, as we are considering cruise conditions across the whole trajectory.

Optimal control theory (Kirk, 1970; Macki & Strauss, 1982; Bressan & Piccoli, 2007) is applied here to formulate a problem leading to fixed-time, fuel minimal trajectories subject to arrival constraints, an aerodynamic fuel burn model and a data-driven wind field. The fuel usage of trajectories generated by two finite horizon optimal control formulations are compared. In the first, a single control variable is given as a set of position-dependent aircraft headings. The second formulation varies both the headings and the airspeed of the aircraft. Fuel consumption is modelled with a new physics-driven fuel burn function, which is aircraft type specific (Poll & Schumann, 2021a,b). Optimal trajectories are found numerically using a reduced gradient approach.

Long range cruise trajectories minimising fuel for a fixed mass aircraft have previously been researched by others using control variables of thrust and flight path angle (Schultz & Zagalsky, 1972; Speyer, 1973) or lift and thrust (Schultz, 1974; Menon, 1989), but these have not accounted for the wind field. In these papers the fuel flow is minimised for a set distance of flight. Wind has been factored into fixed range cruise calculations which compare fixed and free thrust as airspeed varies (Erzberger & Lee, 1980), but here a direct operating cost is minimised, with this cost dependent on both time of flight and fuel use. Similar comprehensive research into trajectory optimisation for hybrid UAVs, incorporating weather data and optimising for energy has been completed recently, but these routes are time variable (Dobrokhodov et al., 2020).

Fixed-time trajectories minimising direct operating cost have also been designed, but

unlike the current research, the optimal control model is set up with a free final time (Sorensen & Waters, 1981; Burrows, 1983). The trajectory best fitting the imposed arrival time constraint is then selected. Other methods have been used to solve a minimum fuel cruise at constant altitude with fixed arrival time (García-Heras et al., 2014; Franco & Rivas, 2011), but these either did not consider the effect of wind at all or simplified the wind fields. Optimal routes through realistic wind fields are often very different from their wind-free equivalents (Girardet et al., 2014).

The minimum fuel optimisation problem for a fixed-time journey has, therefore, been examined previously. The novelty in the approach shown here lies in:

- Applying a recently developed, novel, analytic aerodynamic fuel-burn model (Poll & Schumann, 2021a,b) that is quicker and easier to use in complex computation schemes than the standard EUROCONTROL BADA (García-Heras et al., 2014; Wickramasinghe et al., 2012; Soler et al., 2020; Yamashita et al., 2021, 2020). This new method is open source and fully transparent.
- A numerical assessment across 91 days, covering the entire winter period from 1st December, 2019 to 29th February, 2020, using real weather data.

This research is set out in five sections. In Section 6.2, the system dynamics, a description of the optimal control formulation and an outline of the approximate synthesis of optimal trajectories are discussed. Section 6.3 contains model data specific to the the data-driven wind field and the fuel burn function. In Section 6.4 fuel savings by incorporating both heading angle and airspeed control, rather than just heading angle, are quantified and the effect of differing wind fields on results is analysed. Finally, the results are summarised and discussed in Section 6.5.

6.2 Mathematical modelling for trajectory planning

In this section the dynamical system governing the trajectory of an aircraft is presented. An optimal control formulation is stated to minimise fuel burn along a fixed time horizon with a fixed departure point. The destination target comprises all points within 1 km of the point $(\lambda_{dest}, \phi_{dest})$. Two different optimal control formulations are presented, representing fixed and variable airspeed models.

We then redefine this continuous system as a set of discrete paths, making a piecewise continuous route, in order to render the solution more practical in an sense. The numerical method applied to provide an approximate solution to this discrete problem is then discussed. The nomenclature for variables used throughout the chapter is defined in Table 6.1.

6.2.1 System dynamics for the trajectory of an aircraft

A simplified model for the trajectory of an aircraft is considered. The model is a variation of Zermelo's navigation problem transformed conformally onto a sphere, but here the objective is to minimise fuel burn rather than flight time (Zermelo, 1930). The aircraft travels the surface of a spherical shell with the same radius as the Earth. Although the aircraft will actually travel at a fixed altitude h above this shell, as $h \ll R$ this can be approximated by R in line with previous research (Arrow, 1949; Jardin & Bryson, 2012b; Ng et al., 2014; Kim et al., 2016; Williams, 2016). The aircraft trajectory, is characterised by longitude, λ, latitude, ϕ, and mass, M, at any time t. The rate of change of these states depends both on airspeed and wind speed. The aircraft is controlled through its heading angle, θ, and its airspeed, V. By varying these quantities, the path of the aircraft will be altered, thus resulting in different regions of winds being encountered. The airspeed V will determine which winds it is possible to access within the given fixed time of flight. This initial continuous problem assumes that:

1. The aircraft's heading angle and airspeed are varied continuously.

2. Airspeed will remain within bounds throughout the flight.

The equations of motion for the controlled aircraft are given by:

$$\dot{\lambda} = \frac{1}{R\cos\phi(t)}(V(t)\cos\theta(t) + u(\lambda(t),\phi(t))), \tag{6.1}$$

$$\dot{\phi} = \frac{1}{R}(V(t)\sin\theta(t) + v(\lambda(t),\phi(t))), \tag{6.2}$$

$$\dot{M} = -g(\lambda(t),\phi(t),M(t),V(t)). \tag{6.3}$$

Equation (6.1) is derived from the zonal components of both airspeed and wind speed, mapped conformally onto a sphere. Equation (6.2) is the sum of the meridional components of airspeed and wind speed, again multiplied by the correct transformation factor to allow for all (λ,ϕ) co-ordinate pairs to lie on a spherical shell. Equation (6.3) is the rate

Table 6.1: Nomenclature

ϕ	latitude in radians
ϕ_{dept}	latitude in radians of departure airport
ϕ_{dest}	latitude in radians of destination airport
λ	longitude in radians
λ_{dept}	longitude in radians of departure airport
λ_{dest}	longitude in radians of destination airport
M	mass of aircraft in kg
M_{dept}	mass of aircraft in kg at start of trajectory
M_{ref}	scaling constant based on nominal aircraft take-off mass 235 113 kg
θ	heading angle in radians
u	zonal (eastward) wind in $\mathrm{m\,s^{-1}}$
v	meridional (northward) wind in $\mathrm{m\,s^{-1}}$
R	radius of Earth ≈6 371 000 m
h	altitude of aircraft above Earth in m ($h \ll R$)
V	airspeed in $\mathrm{m\,s^{-1}}$
V_{ref}	scaling constant based on nominal cruise airspeed 240 $\mathrm{m\,s^{-1}}$
t	time in s
t_f	final time in s
g	fuel burn in $\mathrm{kg\,s^{-1}}$
d_H	distance around the sphere between $(\lambda_{dest}, \phi_{dest})$ and $(\lambda(t_f), \phi(t_f))$ in m
δ_{dest}	radius of target around destination airport in m

of decrease of fuel over time, according to fuel burn rate function g which is discussed in detail in Section 6.3.2. However, in order to avoid scaling issues, as the state variables are of very different orders, it is necessary to non-dimensionalise this dynamical system. Of the state variables, only M, the mass of the aircraft has a dimension, as angle measures are considered dimensionless. In the same way airspeed, V, time, t, fuel burn rate, g, zonal wind speed, u and meridional wind speed, v must be similarly scaled. The following definitions will be used:

$$M^* = \frac{M}{M_{ref}}, \qquad V^* = \frac{V}{V_{ref}}, \qquad t^* = \frac{V_{ref}t}{R},$$
$$g^* = \frac{Rg}{M_{ref}V_{ref}}, \qquad u^* = \frac{u}{V_{ref}}, \qquad v^* = \frac{v}{V_{ref}},$$

leading to the non-dimensionalised dynamical system of:

$$\frac{\partial\lambda}{\partial t^*} = \frac{1}{\cos\phi(t^*)}(V^*(t^*)\cos\theta(t^*) + u^*(\lambda(t^*),\phi(t^*))), \tag{6.4}$$
$$\frac{\partial\phi}{\partial t^*} = V^*(t^*)\sin\theta(t^*) + v^*(\lambda(t^*),\phi(t^*)), \tag{6.5}$$
$$\frac{\partial M^*}{\partial t^*} = -g^*(\lambda(t^*),\phi(t^*),M^*(t^*),V^*(t^*)). \tag{6.6}$$

This describes a continuous problem, but ly the trajectory is viewed as a piecewise continuous route. It comprises rhumb-line tracks, the paths with constant heading angle around the globe between two given points, between waypoints corresponding to the start and end of equal time intervals. We now discuss this discrete optimal control formulation and its approximate synthesis.

6.2.2 Optimal control formulation and approximate synthesis

The objective of the dynamic optimisation procedure is to compute the airspeeds and headings that will ensure fuel burn across the trajectory is minimised. At first, the control variables are not constrained. However, it is assumed that airspeed, where this is allowed to vary, will lie within a specified range.

In formulating the OCP a constant altitude flight is considered, starting from the initial position, $(\lambda_{dept}, \phi_{dept})$ at a time $t = 0$ to a final target that encompasses all points within a certain distance, δ_{dest}, of $(\lambda_{dest}, \phi_{dest})$ at a final time of $t = t_f$, with a fixed time of flight.

An aircraft arriving within $\delta_{dest} = 1000$ m of the destination airport is assumed to have

completed its journey successfully. This terminal constraint is based on the Haversine formula (Veness, 2019):

$$a_H = \sin^2(|\phi_{dest} - \phi(t_f)|/2) + \cos(\phi_{dest})\cos(\phi(t_f))\sin^2(|\lambda_{dest} - \lambda(t_f)|/2),$$
$$d_H = 2R\,\text{atan2}(\sqrt{a_H}, \sqrt{1-a_H}), \tag{6.7}$$

where atan2 is the four-quadrant inverse tangent of the two real values in the function bracket, d_H is the spherical distance to the destination airport in metres and R is the radius of the Earth.

The mass, M(t), is not constrained, but is assumed to lie between the start of cruise mass (taken to be 97.5% of the take-off mass and denoted here as SOCM) and the OEM. At $t = 0$, $M(0) = M_{dept}$ which is the SOCM. For a Boeing 777-236ER aircraft, this is calculated as 221 826 kg flying east and 235 112 kg flying West, using the method developed in Poll & Schumann (2021b).

The formulation involving two control variables is referred to subsequently as OCP2. The first control is the heading angle in radians, measured anti-clockwise from due East, in line with the original derivation of the dynamical system in Zermelo (1930). The true airspeed, $V(t)$, is the second control and is assumed to stay within the boundaries of 199 to 252 $\mathrm{m\,s^{-1}}$. A second formulation, referred to as OCP1, is a simplified version of OCP2, involving just a single control variable, the heading angle $\theta(t)$. In both cases, as flights progress, at each time step mass will be reduced by the amount of fuel burned during the preceding time interval.

In the continuous problem with two control variables, outlined in Section 6.2.1, the optimal heading angle $\theta(t)$ and airspeed $V(t)$ are found by minimising the cost functional:

$$J^* = \min_{\theta(\cdot),V(\cdot)} \int_0^{t_f} g(\lambda(t), \phi(t), M(t), V(t))\,dt\,, \tag{6.8}$$

where $g(\lambda, \phi, M, V)$ is the physics-driven fuel burn function discussed further in Section 6.3.2, and subject to Equations (6.1)-(6.3).

In the case of OCP1, where a single variable only is applied, the system dynamics remain as described in Equations (6.1)-(6.3), but taking a constant airspeed, i.e. $V(t) = V$.

The cost functional now becomes:

$$J^* = \min_{\theta(\cdot)} \int_0^{t_f} g(\lambda(t), \phi(t), M(t), V)\, dt. \tag{6.9}$$

However, for applicability, we do not wish to solve the continuous problem itself, but the discretisation of the continuous problem.

The construction of a non-linear optimisation approach continues with the discretisation in time of the system dynamics. A uniform temporal grid with N time steps of size Δt is set. The spatial states at the start and end of each time interval are considered as waypoints in the flight and the temperature and wind conditions at each of these $N+1$ waypoints are used in updating the dynamical system as the trajectory progresses. The third state, the mass of the aircraft at the start of each time interval is used in evaluating the fuel burn across each time interval. This is done using a Runge-Kutta 4th order method with fixed time step for Equations (6.1)-(6.3), providing a set of discrete state variables $\{\lambda(i\Delta t^*), \phi(i\Delta t), M(i\Delta t^*)\}_{i=0}^{N}$, which we denote by $(\lambda(i), \phi(i), M(i))$. In Chapters 4, 5 and 7 the Euler Scheme is used for numerical approximations, but here the Runge-Kutta 4th order method is applied as in this case it does not unduly slow the computational time. This leads to the following non-linear optimisation problem for the non-dimensionalised system:

$$J^*_{\Delta t^*} = \min_{\lambda(\cdot),\phi(\cdot),M^*(\cdot)\in\mathbb{R}^N} \Delta t^* \sum_{i=1}^{N} g^*\big(\lambda(i), \phi(i), M^*(i), V^*(i)\big)\,, \tag{6.10}$$

subject to:

$$\lambda(i+1) = \lambda(i) + \Delta t^* \Phi_\lambda(\lambda(i), \phi(i), M^*(i), \theta(i), V^*(i), \Delta t^*)\,, \tag{6.11}$$

$$\phi(i+1) = \phi(i) + \Delta t^* \Phi_\phi(\lambda(i), \phi(i), M^*(i), \theta(i), V^*(i), \Delta t^*)\,, \tag{6.12}$$

$$M(i+1) = M(i) + \Delta t^* \Phi_M(\lambda(i), \phi(i), M^*(i), \theta(i), V^*(i), \Delta t^*)\,, \quad i = 0 \ldots N\,, \tag{6.13}$$

with the initial conditions:

$$(\lambda(0), \phi(0), M^*(0)) = (\lambda_{dept}, \phi_{dept}, \frac{M_{dept}}{M_{ref}})\,, \tag{6.14}$$

and the constraints:

$$d \leq \delta_{dest}\,, \tag{6.15}$$

$$M^*(i) \in \left[\frac{\text{OEM}}{M_{ref}}, \frac{\text{SOCM}}{M_{ref}}\right]\,, \tag{6.16}$$

$$V^*(i) \in \left[\frac{V_{min}}{V_{ref}}, \frac{V_{max}}{V_{ref}}\right]\,. \tag{6.17}$$

The function Φ denotes an explicit RK4 time integration expressed as a single-stage update. Controls are considered as piecewise constant, as across each 100s time step controls and atmospheric conditions are taken to be constant. This is a justifiable assumption, as a 100 s time step is small in comparison to the whole journey of about six hours flying east and eight hours flying West. Equations (6.16) and (6.17) of this program include constraints on the aircraft mass and airspeed.

There are different numerical optimisation solvers which can deal with the solution of the non-linear programming problem Equations (6.10)-(6.17). We chose not to use a specialised commercial solver here, but rather one widely used throughout the scientific community to allow for easier replication of results, with the data, formulae and processes all in the public domain. Matlab's fmincon solver was used within the multistart solver from Matlab's Global Optimisation Toolbox (Waltz et al., 2006; Byrd et al., 1999, 2000; Ugray et al., 2007). The multistart solver generates a specified number of starting points randomly within the bounds of the given problem structure. It then filters to feasible points before running these through the fmincon function, as a local solver. Once all points have been run, the local solver output with the lowest cost is found. Fmincon uses an interior point algorithm to alter heading angles and airspeeds at each iteration to ensure that the cost functional is reduced. This continues until either the step tolerance or the optimality tolerance fall below given bounds. Step tolerance is a relative bound which compares the size of a control value to the size of the previous one. Optimality tolerance is a measure of how the objective function is varying in all feasible directions. Its value is taken as the infinity norm of the sum of the partial derivatives of J^* with respect to the control variables at each timestep $i = 1, 2, ...N$:

$$\sum_{i=1}^{i=N} \frac{\partial J^*}{\partial \psi(i)} + \sum_{i=1}^{i=N} \frac{\partial J^*}{\partial V^*(i)}. \tag{6.18}$$

However, gradients of J^* are computed using finite differences, avoiding an explicit numerical calculation of gradients for the wind field and the fuel burn model. This is a standard procedure, as encountering objective functions where the evaluation of derivatives is not practical is not unusual in real-world systems (Colson & Toint, 2001; Naresh Kumar et al., 2018). As fmincon finds local minima of a problem, there is no guarantee that results will be globally optimal, but by using fmincon within the multistart solver the method becomes more robust.

6.3 Weather data and fuel burn modelling

In order to solve the discrete problem described in Equations (6.10)-(6.17), data regarding zonal (eastward) and meridional (northward) winds and temperatures from daily averaged re-analysis data is needed, in addition to a model for the aerodynamic fuel burn function. These are detailed fully in this section.

6.3.1 Data-driven wind and temperature fields

The use of realistic atmospheric data in these simulations is critical in ensuring their relevance to improving TBO. Plots of two sample wind fields are shown in Figure 6.2.

Determination of fuel burn rate is dependent on temperature at any point along a trajectory, the mass of the aircraft and also on the airspeed of the aircraft. The temperature is required in the calculation of both the Mach number and the dynamic viscosity of the air, both of which values are used to find the current fuel efficiency of the aircraft at any point on a trajectory. Numerical solution of Equations (6.1) -(6.2), necessitates obtaining speeds for the zonal and meridional winds at any point across the North Atlantic. Thus both records of wind speed and temperature are needed in order to find admissible trajectories. Contour plots of the two temperature fields corresponding to the wind fields in Figure 6.2 are shown in Figure 6.3.

In this research all weather data has been downloaded from the re-analysis data set provided by the NCAR (Kalnay et al., 1996). This comprises an atmospheric model and a large array of observations combined via data assimilation to produce a weather hindcast. Wind velocity and temperature values are given for a global grid of resolution 2.5° as daily averages. Linear interpolation is then applied to obtain wind components and temperature at specific points in a trajectory. This approach is justified as long haul flight routes have been shown to be largely insensitive to the resolution of weather data (Lunnon & Mirza, 2007) and the evolution of the jet stream at this altitude shows little variation across a 24-hour period (Mangini et al., 2018). The use of daily wind data is in line with methods used in previous transatlantic trajectory research (Wells et al., 2021; Kim et al., 2020; Williams, 2016; Mangini et al., 2018).

Wind fields can be considered to be smooth, as there are no flow discontinuities in the at-

mosphere. Molecular viscosity prevents discontinuities from occurring by smoothing them out over the Kolmogorov scale (which is typically a few millimeters). There are no flow fluctuations smaller than this scale.

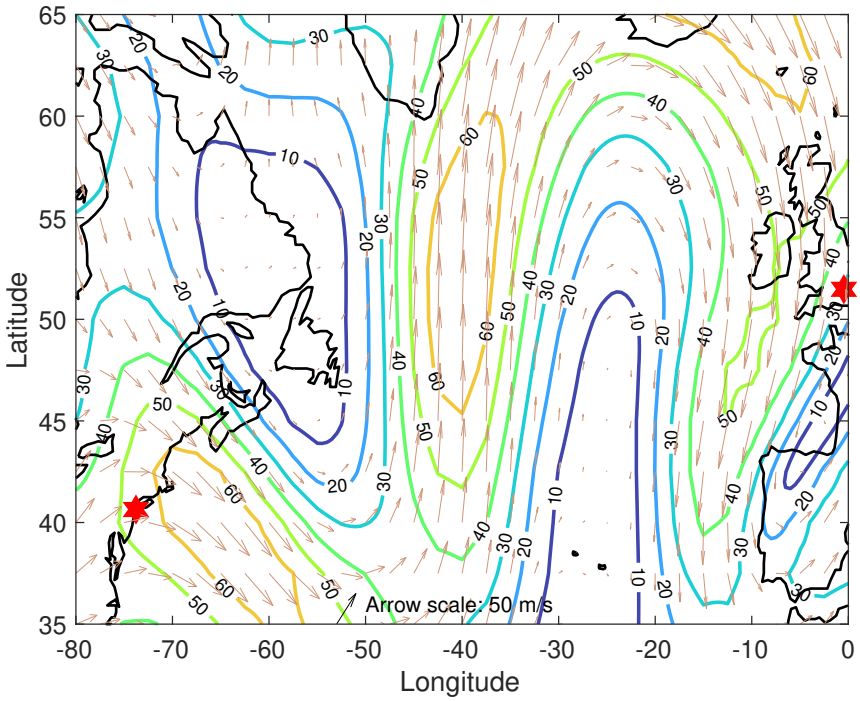

(a) Plot of Norm of Winds 1st December, 2019.

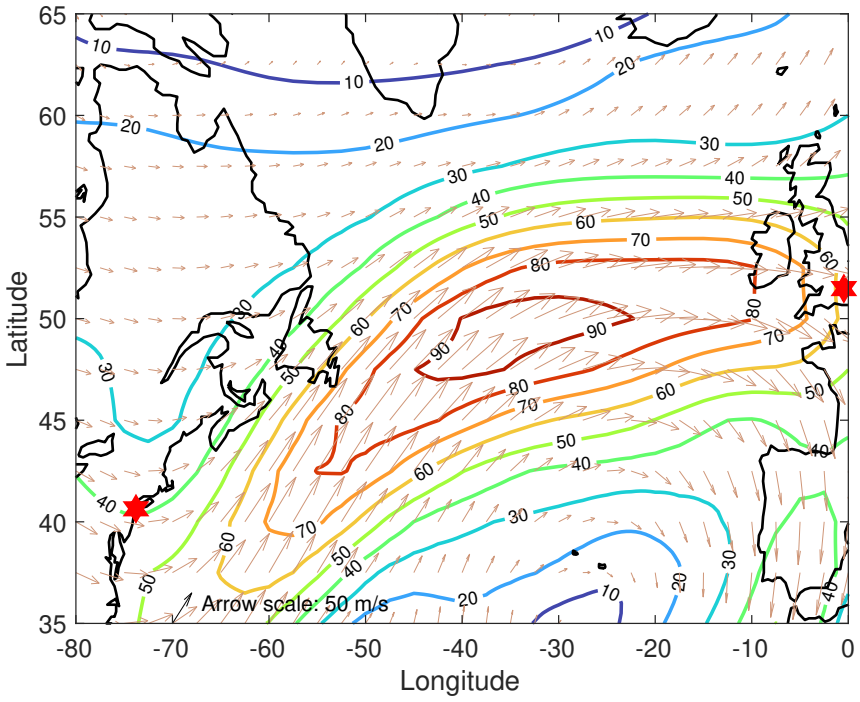

(b) Plot of Norm of Winds 9th February, 2020.

Figure 6.2: Contour plots showing strength of winds in $\mathrm{m\,s^{-1}}$ across the North Atlantic on 1st December, 2019 and 9th February, 2020. Behind these contour plots is a quiver plot of the same winds to show the direction of flow. These two wind fields are representative of the winter period. The positions of LHR and JFK are represented by red stars.

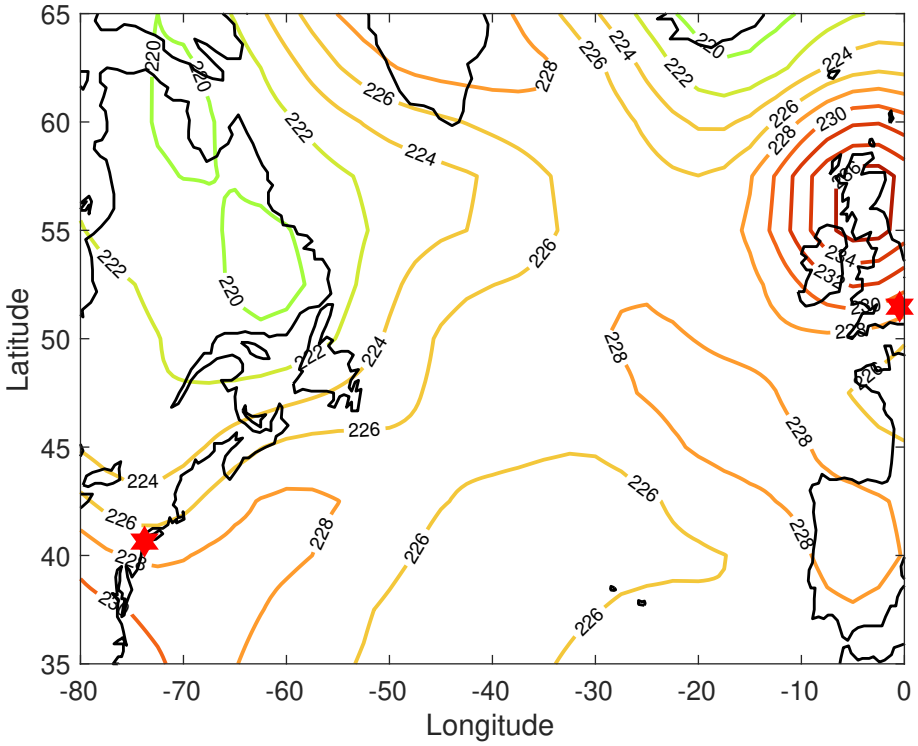

(a) Contour Plot of Air Temperature at 250 hPa, 1st December, 2019.

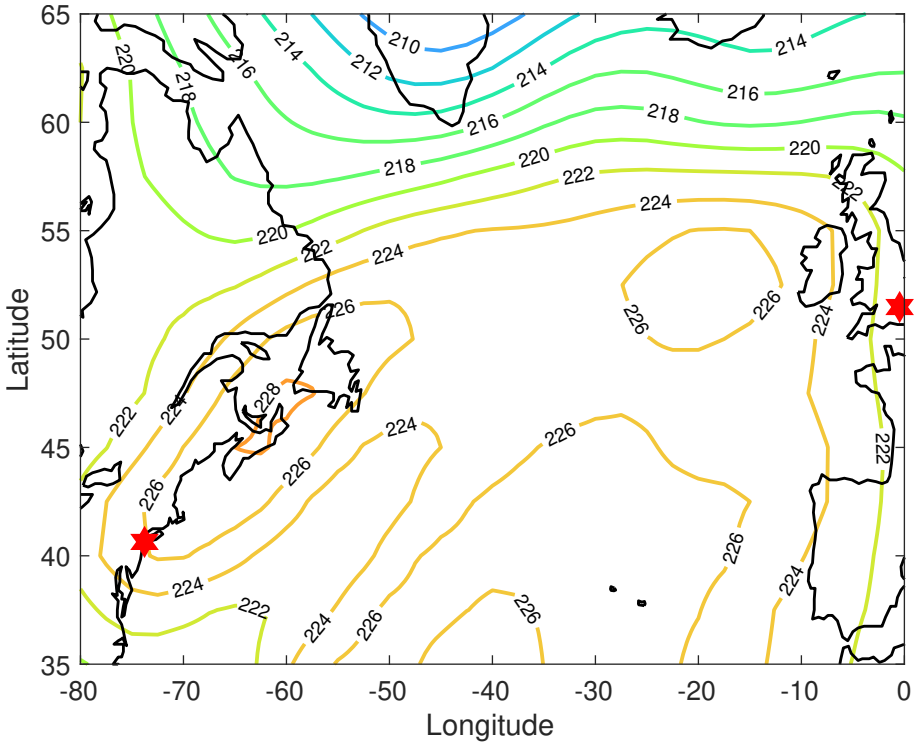

(b) Contour Plot of Air Temperature at 250 hPa, 9th February, 2020.

Figure 6.3: Contour plots showing temperature variation in degrees K across the North Atlantic on 1st December, 2019 and 9th February, 2020. These two temperature fields correspond to the wind fields shown in Figure 1. The positions of LHR and JFK are represented by red stars.

6.3.2 Fuel burn model

In the vast majority of research relating fuel burn to aircraft trajectories, EUROCONTROL's BADA Version 4 method is used to model aircraft fuel flow (Eurocontrol, 2021a). This is true for recent papers looking for climate optimised trajectories (Matthes et al., 2020; Yamashita et al., 2020, 2021) as well as previous research into fuel-optimal routing (García-Heras et al., 2014; Wickramasinghe et al., 2012; Soler et al., 2020). However, we have chosen to use a new analytic method for estimating the fuel burn rate of commercial passenger aircraft due to Poll & Schumann (2021a,b). This method is more ideally suited to our research, being quicker and easier to use than BADA, which is proprietary data, as part of a complex computation scheme. As this method is open source, there is no need for a licence to access it and there are no restrictions on its use. The derivation of the method has been set out in refereed journals, ensuring that its validity has received appropriate endorsement.

The fuel burn rate (in $\mathrm{kg\,s^{-1}}$) can be expressed as:

$$g = \frac{VW}{\left(\eta_0 \frac{L}{D}\right) LCV}, \tag{6.19}$$

where W is the weight of the aircraft in N, V denotes the airspeed in $\mathrm{m\,s^{-1}}$, LCV is the lower calorific value of aircraft fuel (43 $\mathrm{MJ\,kg^{-1}}$ for kerosene) and $\eta_0 \frac{L}{D}$ is the maximum value of overall efficiency of the propulsion system, multiplied by the lift-to-drag ratio. Full details of this function can be found in Chapter 3.

Obtaining the fuel burn rate for each step of a trajectory depends on aircraft specific parameters and ISA parameters. Here it is assumed that a Boeing 777-236ER aircraft is used, as this is the model currently flown most frequently in transatlantic routes between LHR and JFK (Flightradar24, 2020). In addition, the fuel burn is also dependent, at each point along a trajectory, on the altitude and airspeed of the aircraft, and on the environment temperature $T = T(\lambda, \phi)$. The term $\eta_0(L/D) = \eta_0(L/D)(T)$ has a non-linear physical dependence on the temperature field, which has been modelled in Poll & Schumann (2021a,b) as shown in Section 3.4.1.

The aircraft are assumed to fly along the 250 hPa isobar, which corresponds to an altitude of approximately 34 000 feet. This is close to the average cruise altitude for flights across the North Atlantic and the flight level at which the OTS is currently calculated

(Mangini et al., 2018). Aircraft on this route rarely change altitude and so this is an acceptable simplification.

The fuel burn rate can be seen in Figure 6.4 for a range of airspeeds and temperatures. The airspeeds are determined by the model of aircraft flown, in this case the Boeing 777-236ER, and the temperatures are representative of those recorded across the North Atlantic during the winter period from 2019 to 2020.

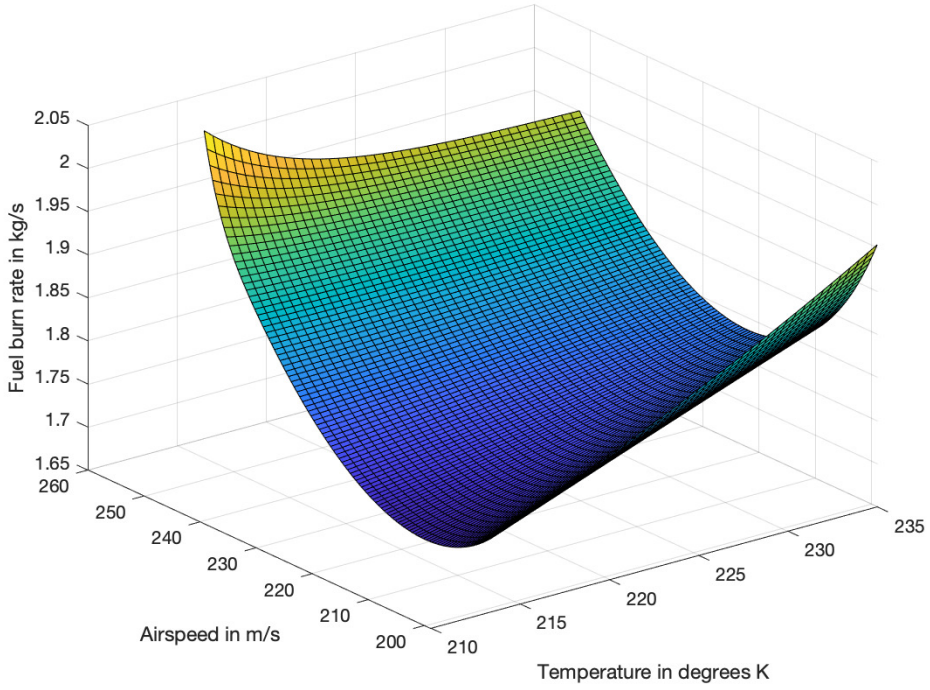

Figure 6.4: Surface plot showing fuel burn rate variation with airspeed and temperature. The range of airspeeds shown are those practical for a Boeing 777-236ER flying at a cruise altitude of 34 000 ft. The temperatures cover a realistic range for the winter season over the North Atlantic.

Where both heading angle and airspeed are controlled, fuel burn rate will vary with airspeed, mass and temperature, whereas in the second case, where airspeed remains constant, fuel burn rate will vary only due to the mass of the aircraft and the temperature encountered along the route, which can be seen to be a small effect.

6.4 Results

In this section, the particular parameters for the numerical models used are specified and then results, given the relevant atmospheric data, are presented and discussed.

6.4.1 Parameters of the system

Trajectories are modelled between LHR (51.5° N, 0.5° W) and JFK (40.6° N, 73.8° W) both eastbound and westbound. This particular route has been chosen as it is not only one of the busiest, but goes through the slowly evolving background wind field provided by the jet stream, the prevailing eastbound nature of which causes the challenges of flying in each direction to be quite different.

All flights modelled occur between 1st December, 2019 and 29th February, 2020. This allows the full range of winter weather systems to be considered (Irvine et al., 2013), as the NAO has been shown to cause transatlantic routes to vary strongly (Woollings & Blackburn, 2012; Kim et al., 2016).
Further model parameters include time step length, number of time steps in each direction, search algorithm tolerances, chosen airspeeds for the initialization of the OCP2 formulation, the fixed airspeed for OCP1 and the initial estimate for the heading angle at each time step for both OCP1 and 2.

The time step length was chosen following a sensitivity analysis and a time step of $\Delta t =$ 100 s allows a stable and consistent application of the numerical method, whilst not unduly increasing truncation errors. This also, in practice, allows time for adjustments to heading and airspeed to be made. Having a practical knowledge of the situation under consideration is vital in choosing a reasonable time step (Rumpfkeil & Zingg, 2010).

Travelling from LHR to JFK against a headwind, the fixed final time is set to 29 000 s, giving $N = 290$, whilst in the opposite direction this is reduced to just 22 000 s, with $N = 220$. These times have been chosen following research into time optimal routes across the same winter period (Wells et al., 2021). They allow flights on all days enough time to reach their destination, whilst also lying between the longest and shortest scheduled flight times given by the airlines for this route between 1st December, 2019 and 29th February, 2020.

For the interior point algorithm the step tolerance was set to 1×10^{-6}, and the optimality tolerance to 1×10^{-3}. These values were shown to allow efficient convergence across all wind fields considered.

Given the time restriction on trajectories, the fixed airspeed yielding the minimum fuel use will depend on the daily wind field. For this reason, fuel use for trajectories flown at all airspeeds in increments of $1\,\mathrm{m\,s^{-1}}$, within the constraints, was calculated and the trajectory results for the airspeed associated with the lowest fuel use for each day adopted. In some cases, this was the lowest airspeed allowing the trajectory to reach the destination target, whereas on days where winds were more favourable, the airspeed chosen depended on the most efficient airspeed for the model of aircraft used.

The initial airspeed across each time interval in the numerical model for OCP2, $V(t)$, was chosen to be the airspeed obtained from OCP1, but by using the multistart solver, a range of other initial airspeeds was also applied. In the original formulation of OCP2, the airspeed is assumed to be $V(t) \in \mathbb{R}$. However, there are obvious boundaries to an aircraft's airspeed in the practical setting. In order to apply the fuel function based on the work of Poll & Schumann (2021b) across the range of temperatures recorded across the North Atlantic at cruise altitudes, it was necessary to have $V(t) \in [199, 252]$. This control constraint was applied in judging if a feasible solution had been found. In the case of westbound flights, optimised airspeeds varied between 199 and $242\,\mathrm{m\,s^{-1}}$ and for eastbound flights this range was from 199 to $240\,\mathrm{m\,s^{-1}}$, so in all cases a vector of feasible airspeeds was retrieved from the numerical optimisation.

In solving both OCP1 and OCP2 an initial vector of heading angles is required to produce the first trajectory. The GCR is now considered for the journey between the airports. The GCR between the departure and destination airports is divided into N equal length intervals for flights and the rhumbline angle that would take an aircraft from the start to the end of each interval in a no-wind scenario is calculated. (Rhumbline angles provide a single heading on which to travel between two points on the surface of a sphere and are calculated here using the Matlab mapping toolbox.) These angles form the estimate for θ.

6.4.2 Daily results from each numerical model

Results for the numerical solution shown in Section 6.2 are obtained as a vector of optimal headings and a value for the fuel used along the trajectory for each of the 91 days from 1st December, 2019 to 29th February, 2020. The headings are used to generate the states. The area spanned by these states and eight example routes can be seen for all days in Figure 6.5 and Figure 6.6. The shaded area shows the extent of the most extreme trajectory positions across the time period; not all points within this area will have formed part of a trajectory. Westbound flight routes obtained using just the heading angle as a control cover an area that stretches further South than those where time of arrival can be guaranteed by changes to both airspeed and heading angle. The extent of flight routes with varying airspeed also covers an area further North at the start of the flight, benefiting from the variable airspeed to allow a deviation into lower headwind regions. Eastbound flights with variable airspeed can be seen to avoid diversions to both North and South as they approach LHR. However, the area spanned by the trajectories is very similar whether just heading angle or heading angle and airspeed vary.

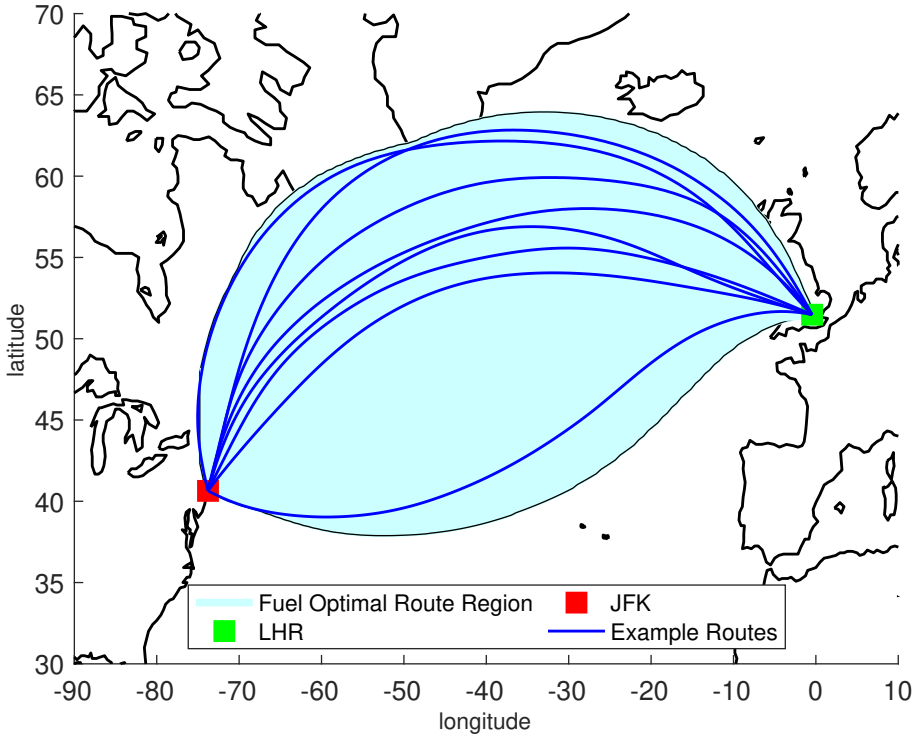

(a) All days West

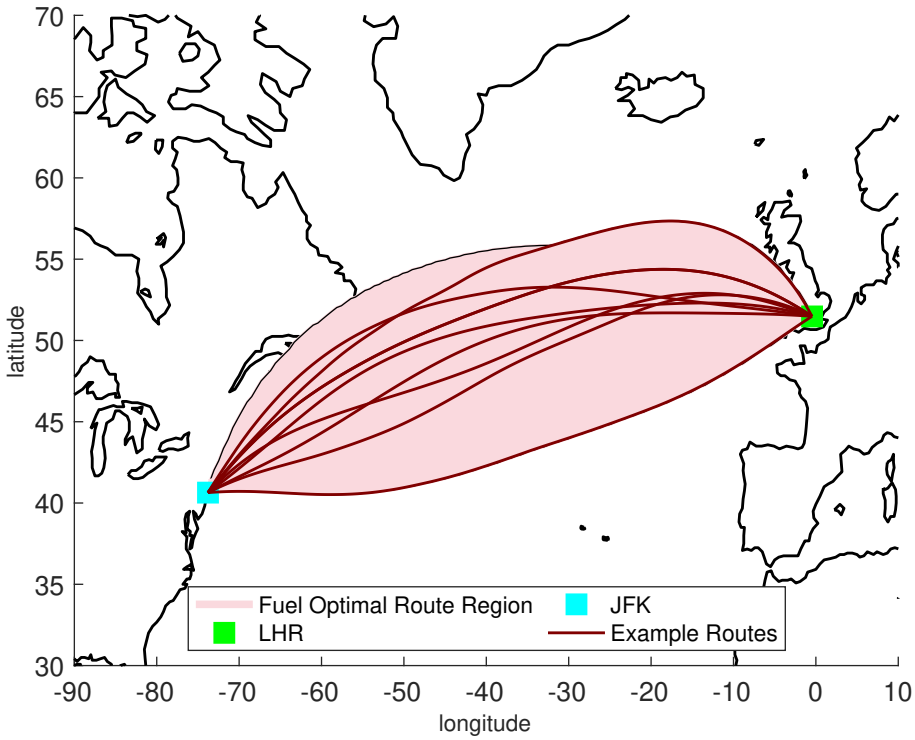

(b) All days East

Figure 6.5: Plots showing region spanned by all fuel optimal trajectories across the North Atlantic for 1st December, 2019 to 29th February, 2020 found from the solution of OCP1, with 8 example routes in each direction shown. The shaded area shows the extent of trajectories, but not all points within the area are part of a route.

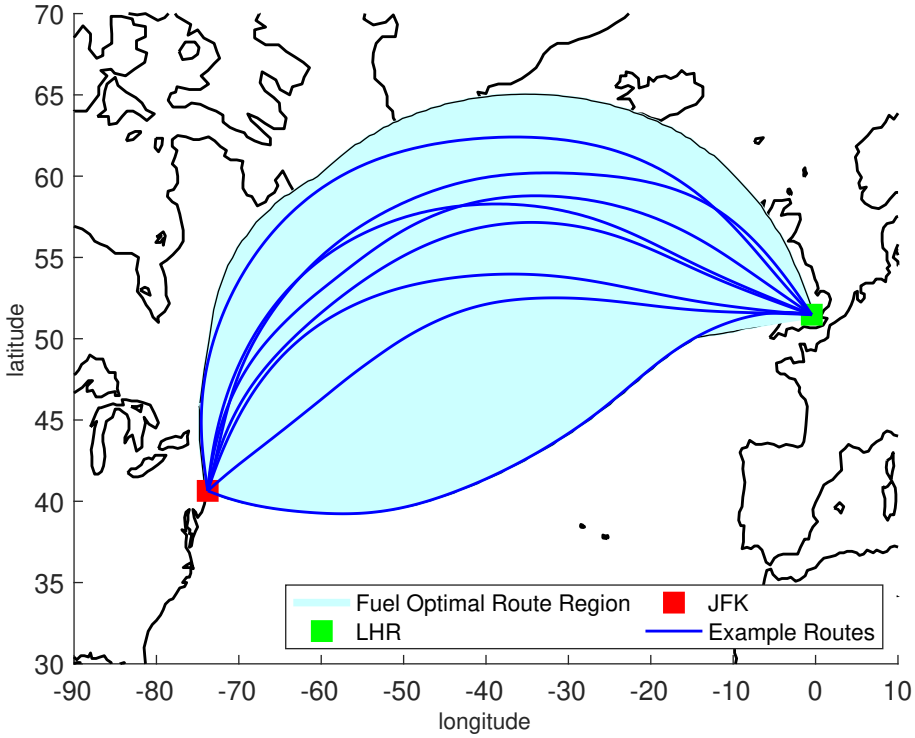

(a) All days West

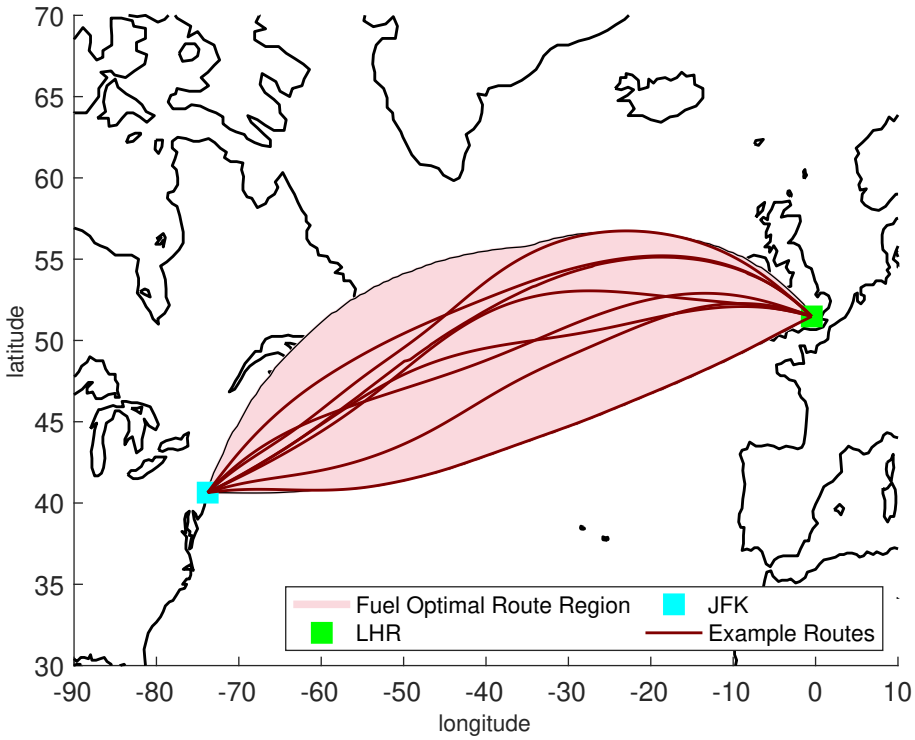

(b) All days East

Figure 6.6: Plots showing region spanned by all fuel optimal trajectories across the North Atlantic for 1st December, 2019 to 29th February, 2020 found from the solution of OCP2, with 8 example routes in each direction shown. The shaded area shows the extent of trajectories, but not all points within the area are part of a route.

Percentile plots, shown in Figures 6.7 and 6.8, display the percentage of flights passing to the South of each particular area of the North Atlantic. Outlying waypoints, those that have longitudes passed through by fewer than three of the flights, are shown as yellow crosses. All but one or two flights stayed within a longitude range of $1 - 73$ °W.
These graphs are created by taking longitudes between the maximum and minimum longitudes of all flight trajectories generated in 1° increments and then interpolating to find the corresponding latitudes for each daily trajectory. Percentiles for the latitudes at each longitude are found and these are plotted. This demonstrates the distribution of latitude positions across all routes and allows for easier comparison between the trajectories resulting from OCP1 and 2.

Flying west (Figure 6.7) it can be seen that the range of latitudes at each longitude is wider for OCP1, but that the interquartile ranges are very similar. More of the OCP2 flights tend to fly further North across the middle of the Atlantic as they can vary their airspeed to counter headwinds and then slow down in more favourable regions of the wind field in order to adhere to the fixed journey time and save on fuel. The 1% percentile for the OCP1 flights demonstrates that at all latitudes certain journeys must go very far south to avoid headwinds to balance the need for both a prompt arrival at the target and a low fuel burn rate.

Figure 6.8 shows that the vast majority of flights going east follow a path very close to the GCR. Using two controls allows a small percentage of flights to go further North or South than the flights constrained by a fixed airspeed. The interquartile range of latitudes is also wider for flights with variable airspeed.

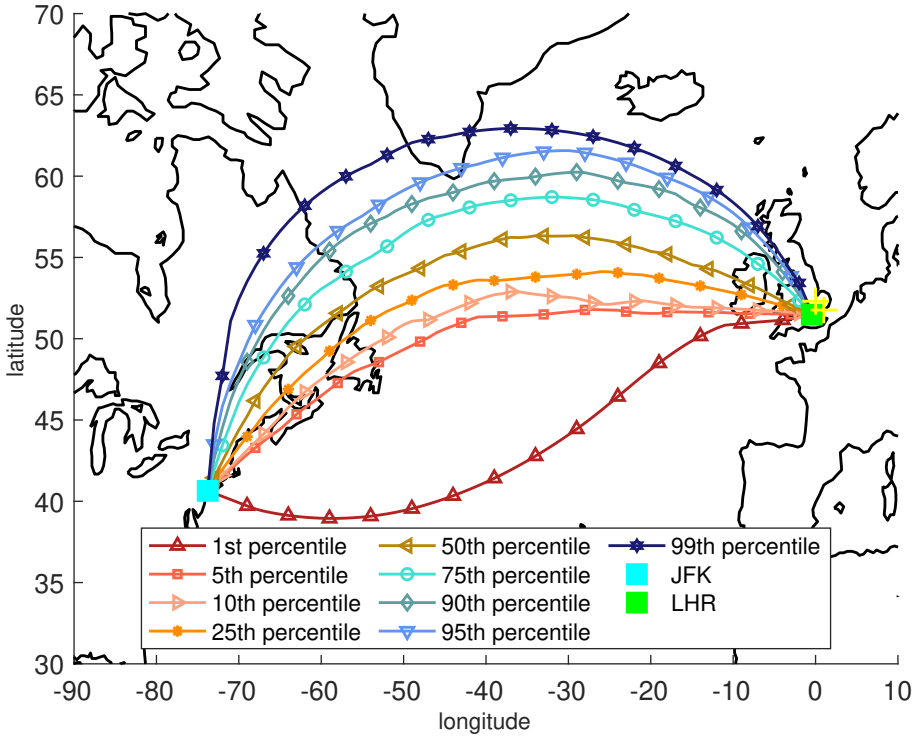

(a) All days West OCP1

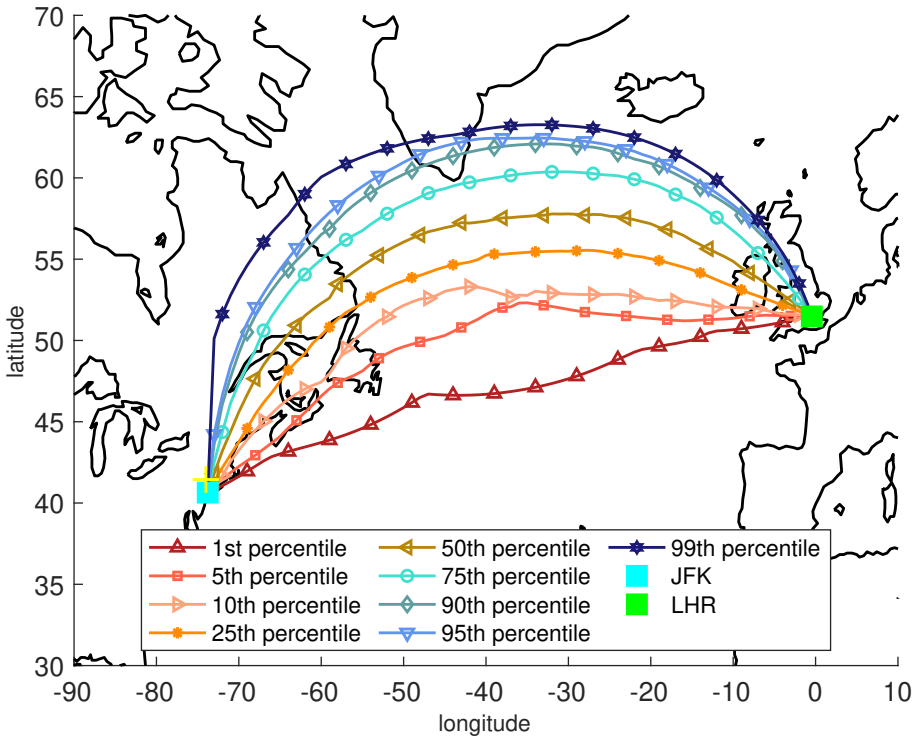

(b) All days West OCP2

Figure 6.7: Plots showing percentiles of latitude positions for a set of integer longitude positions ranging between the furthest easterly and westerly points in the westbound trajectories obtained from OCP1 and 2. Waypoints at longitudes not accessed by all flights are shown in yellow. Percentiles show proportion of flight routes passing south of each waypoint.

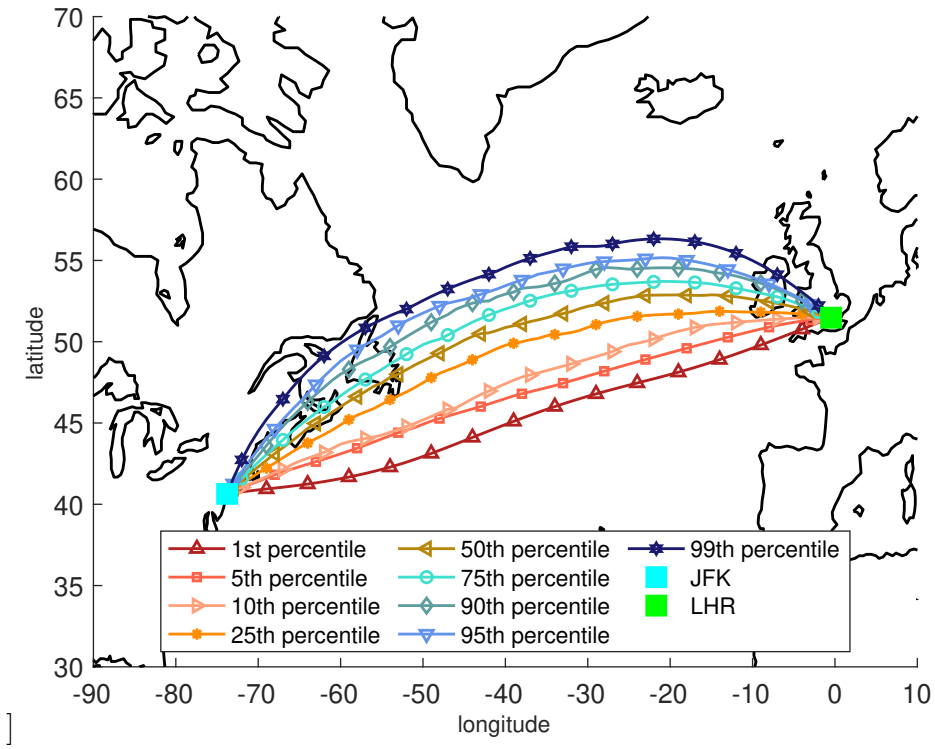

]

(a) All days East OCP1

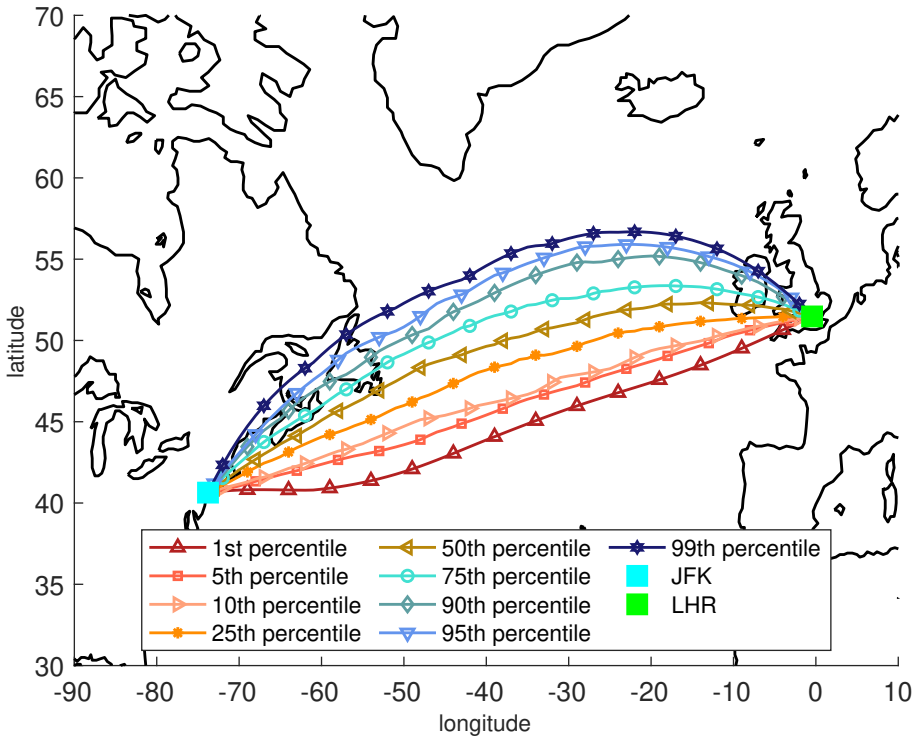

(b) All days East OCP2

Figure 6.8: Plots showing percentiles of latitude positions for a set of integer longitude positions ranging between the furthest easterly and westerly points in the westbound trajectories obtained from OCP1 and 2. All longitudes were accessed by all flights. Percentiles show proportion of flight routes passing south of each waypoint.

6.4.3 Comparison of fuel use between OCP1 and OCP2

If airspeed is fixed at a level that allows all daily flights to arrive at the target in the appropriate fixed time, then clearly more fuel will be used than by flights with a variable airspeed. In this case the highest airspeed used by any of the eastbound flights in the OCP1 formulation was 240 $\mathrm{m\,s^{-1}}$ and for the westbound flights it was 241 $\mathrm{m\,s^{-1}}$. Interestingly the most efficient fixed airspeed for the fixed time crossing between the two airports in a zero wind field, with a constant ISA temperature of 221°, would be 252 $\mathrm{m\,s^{-1}}$ flying east in the 22 000 s time window and 200 $\mathrm{m\,s^{-1}}$ flying west in a fixed time of 29 000 s. Given that one of the key defining points of this research is to create trajectories with a fixed time across an entire winter season, these results demonstrate the importance of adapting to the daily wind field within this system. If, instead of varying the OCP1 airspeed on a daily basis, the maximum OCP1 airspeeds were used for all days in the 2019 -2020 winter season, by comparison the variable airspeed model used in OCP2 would save an average of 8% of the fuel. However, this can be viewed as unnecessarily high, as the fixed airspeed in OCP1 can be lowered to navigate the wind field specific to the day of flight. By analysing data from the solution of the approximate numerical methods for both OCP1 and OCP2, we can provide an estimate of the improvement in fuel efficiency when comparing the most efficient fixed airspeed flights each day to the corresponding most efficient variable airspeed flights.

Comparing results from the solution of OCP1 and OCP2, there are relative differences of 0.5% across the winter period possible. On particular days, these savings are just under 4%. Figure 6.9 shows the percentage fuel saving made each day by using both airspeed and heading angle as control variables in each direction rather than just heading angle. The savings vary on a daily basis, so it is clear that the wind field is instrumental in dictating the importance of varying the airspeed for a fixed-time flight.

When viewed as a box and whisker plot as in Figure 6.10, the distribution of savings can be compared across the eastbound and westbound routes. Whilst median values are similar, the range of eastbound savings is smaller (excluding outliers), as is expected from a route that is most often benefiting from tailwinds, rather than fighting headwinds.

Figure 6.11 shows the distribution of actual fuel savings made each day by using both airspeed and heading angle. Here it can be seen that in absolute terms the median fuel

saving of eastbound flights is lower than that of westbound flights. Flying west just under 23 tonnes of fuel could be saved over the 91 days, whilst flying east fuel savings amount to over 15 tonnes. This gives a total fuel saving across the winter period for nineteen flights in each direction per day of 723 tonnes. These savings are a clear indication that when planning fixed-time trajectories, varying the airspeed and the heading angle will give more efficient results in terms of fuel usage and thus emissions reduction, than controlling heading angle alone.

A comparison of routes generated from the time optimal model used in Wells et al. (2021), the fixed-time, fuel optimal model with both a single control variable and with two control variables developed in this chapter and actual flight routes flown on the 12th December, 2019 is shown in Figure 6.12. The 12th December, 2019 was chosen as a typical day on which a large number of flights were scheduled. From these plots it appears that most airlines prefer to use a path close to the time optimal route and the fuel optimal route, based on two controls. The fuel optimal route which depends purely on altering the heading angles can be seen to be very different from these. In this simplified model, if strong tailwinds are encountered the aircraft must deviate from a more direct path to avoid an early arrival at the destination airport, as the airspeed is fixed throughout the trajectory. These trajectories are representative of those obtained across the winter season considered. Airlines currently try to minimise their operating cost, which is made up of both fuel and time factors. However, trying to adhere to the timetable is also important. On this day 14 of the 19 westbound flights were more than 15 minutes early, with 10 of these being more than 30 minutes early and 3 being more than 45 minutes early. Eastbound all 18 flights were more than 45 minutes early, with 13 being over an hour early and 1 arriving an hour and a half early. Although airlines prefer to be early than late, such early arrival times can lead to added costs. Often aircraft are slowed and stacked as they approach the airport or extra waiting time is spent on the tarmac before accessing a gate, which is unpopular with customers and can lead to compensation claims (John, 2020). From an operations point of view it can lead to blocked gates and ground crew, baggage handlers and fuel bowsers being in the wrong place.

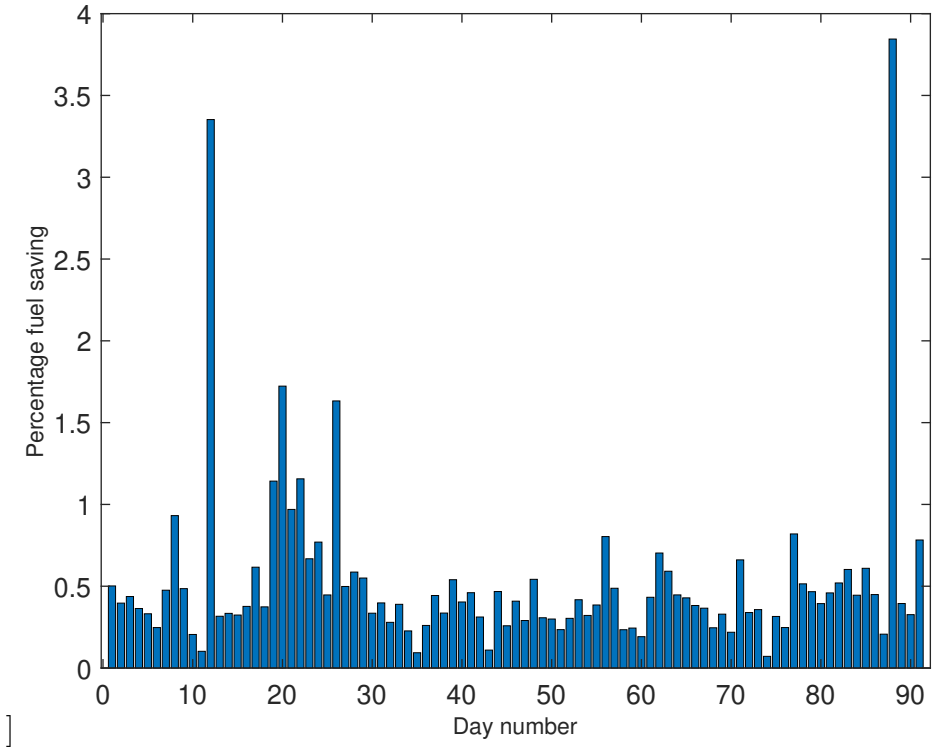

]

(a) Savings made on each day West

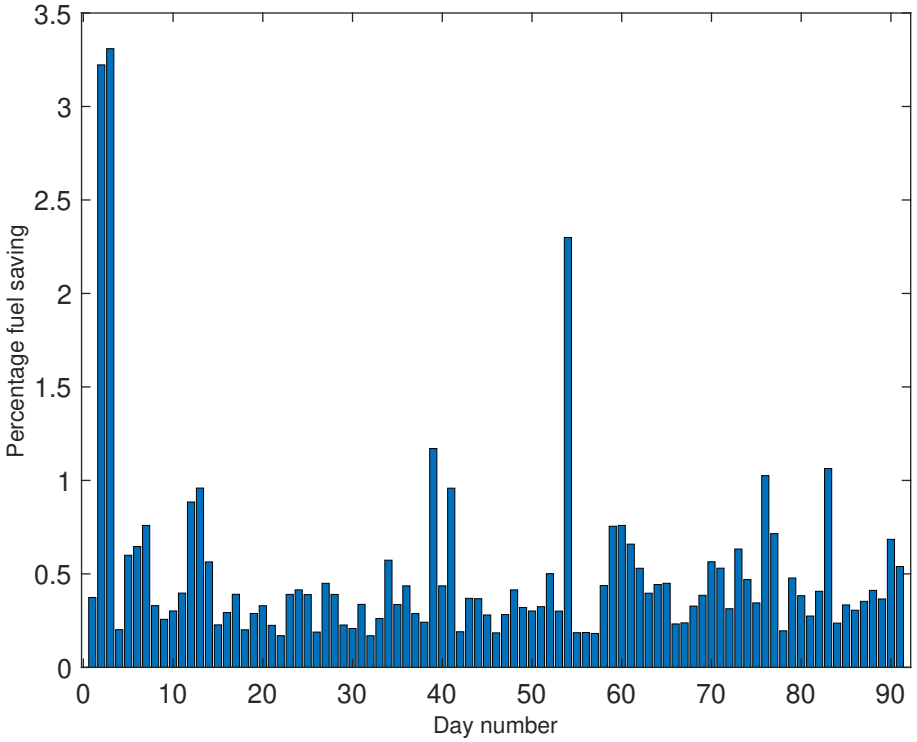

(b) Savings made on each day East

Figure 6.9: Bar charts to show percentage of fuel from using the single control variable of heading angle (corresponding to OCP1) saved by using two control variables, both heading angle and airspeed (corresponding to OCP2).

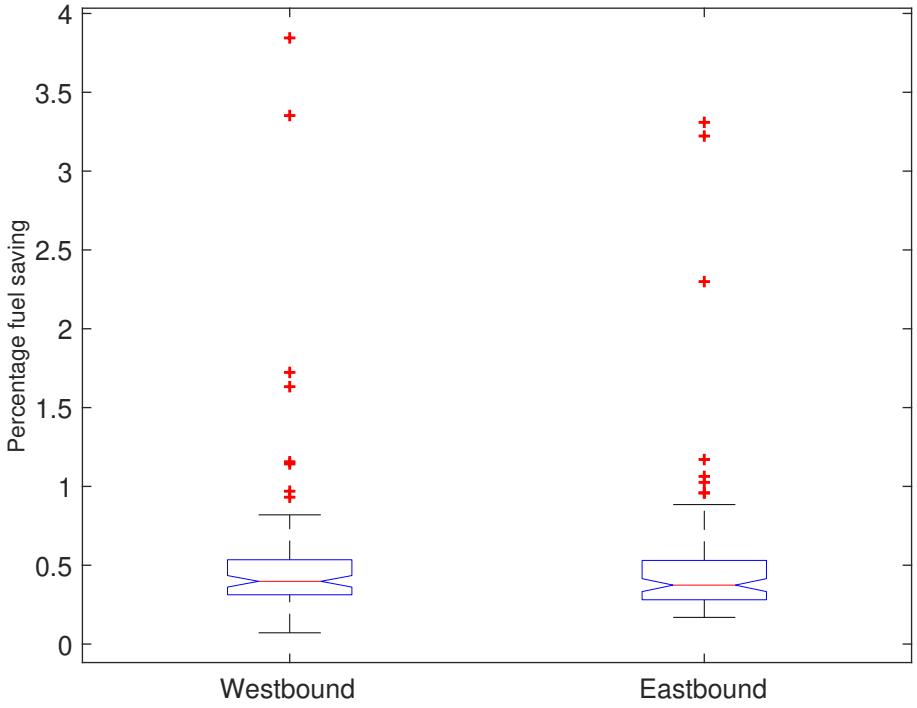

Figure 6.10: Box and whisker plot to compare distribution of percentage savings made made by using OCP2 as the problem formulation rather than OCP1, for all 91 days of the winter of 2019 to 2020, for both eastbound and westbound flights.

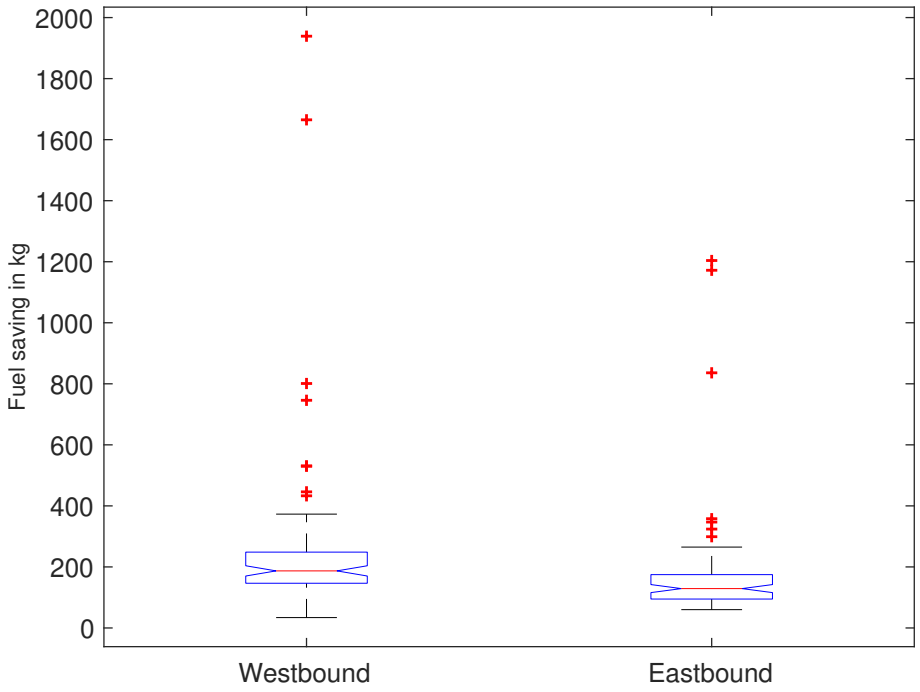

Figure 6.11: Box and whisker plot to compare distribution of absolute savings made by using OCP2 as the problem formulation rather than OCP1, for all 91 days of the winter of 2019 to 2020, for both eastbound and westbound flights.

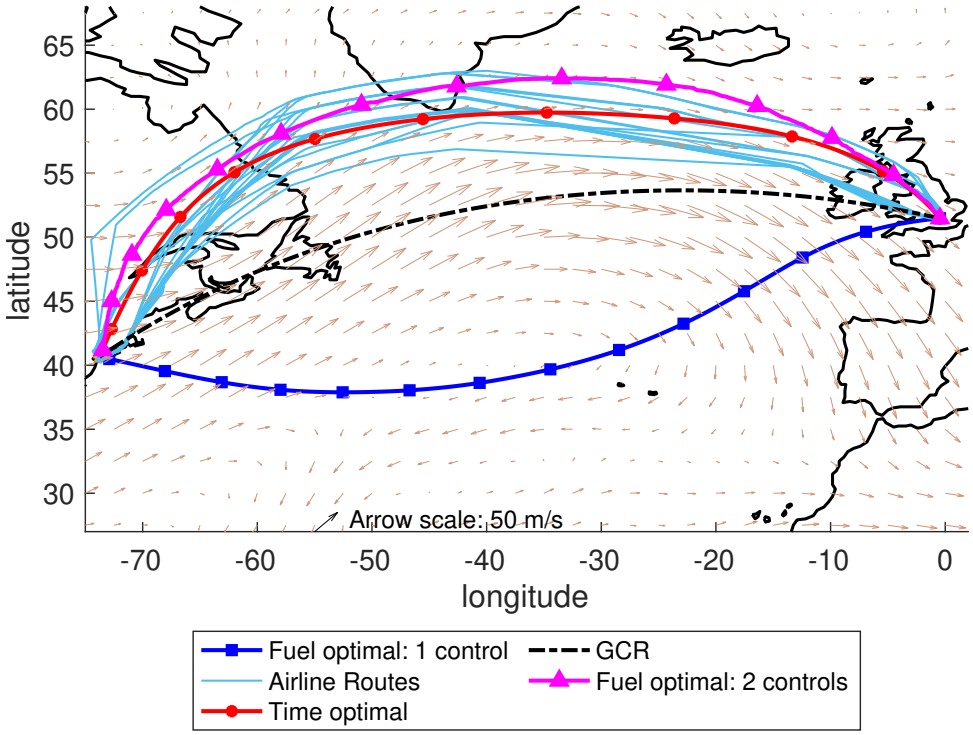

(a) Westbound routes on 12th December, 2019

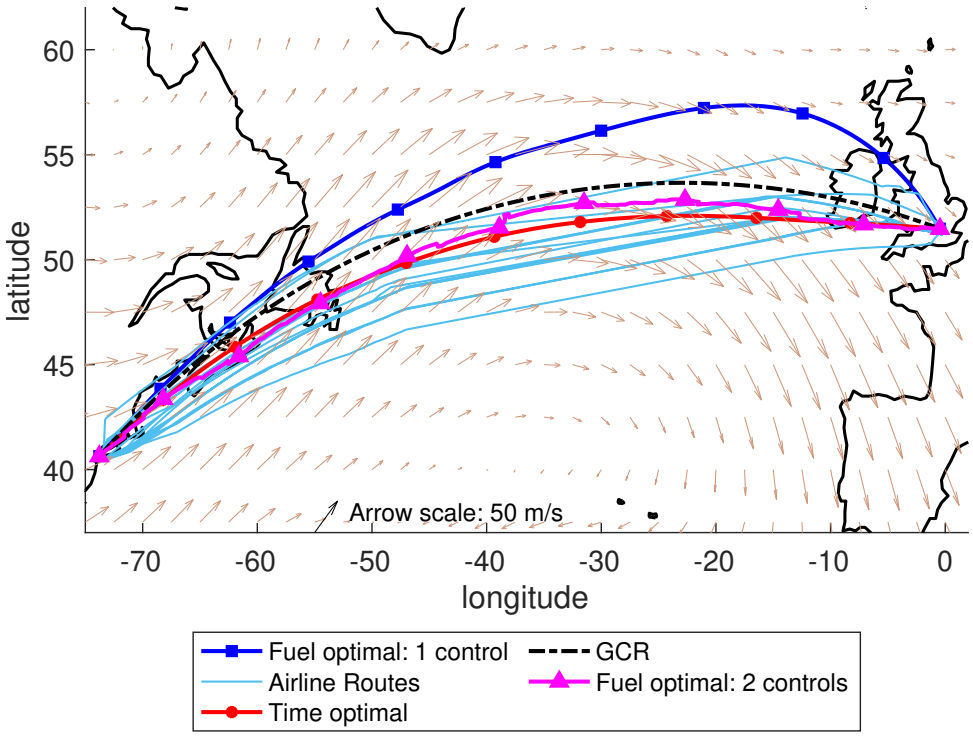

(b) Eastbound routes on 12th December, 2019

Figure 6.12: Route maps showing constant airspeed time optimal trajectory, constant airspeed fuel optimal trajectory (OCP1), varying airspeed fuel optimal trajectory (OCP2), the great circle route (GCR) and actual flights for the 12th December, 2019.

The time optimal routes are simulated as flown at $250\,\mathrm{m\,s^{-1}}$. The fuel optimal routes with fixed airspeed are simulated as flown at $238\,\mathrm{m\,s^{-1}}$ going west and $204\,\mathrm{m\,s^{-1}}$ going East, which are the most fuel efficient airspeeds for a fixed-time flight from the OCP1 formulation on that day. All simulated routes are considered at a fixed altitude of approximately 34 000 feet, with a variable mass. The fuel optimisation is calculated for a Boeing 777-236(ER). The actual flights are not fixed time, mass, altitude or airspeed and use a variety of aircraft. However, they are restricted to the OTS across the North Atlantic, so the flight path taken does not always reflect the airline's chosen route. At a particular flight time, aircraft are given waypoints across the North Atlantic, which must define their routes, in order to maintain a safe separation between all aircraft in the vicinity, given the limited situational awareness before the advent of the new satellite communications. These waypoints can lead to aircraft following less time and fuel efficient routes.

6.4.4 Effect of wind field on optimised airspeed

As the airspeed control allows quite significant fuel savings on some days compared to optimisation by controlling heading angle alone, but far smaller savings on other days, a link between daily wind conditions and airspeed variation is sought. Whilst it is true that mass and airspeed are also connected, with airspeed reducing with mass across a single trajectory in the absence of winds, this pattern is accounted for in the averaging of daily airspeed in the current analysis.

The winds along the GCR are used as a measure of likely headwinds and tailwinds experienced during a flight. These are calculated by splitting the GCR between LHR and JFK into 290 intervals for a westbound flight and 220 intervals for an eastbound flight. The rhumbline angle, β, needed to fly directly between each pair of waypoints is calculated. The wind field for each day is interpolated to give zonal and meridional winds, u and v, at each waypoint. The sum of the components of these lying along the rhumbline angle for each interval of the GCR gives the tailwind, τ at each waypoint:

$$\tau = \begin{bmatrix} uv \end{bmatrix} \begin{bmatrix} \cos\beta \\ \sin\beta \end{bmatrix}. \tag{6.20}$$

By plotting airspeed for the OCP1 simulations and average airspeed for the OCP2 simulations each day against average daily tailwind, the effect of the wind field on the airspeed can be seen.

For westbound flights, shown in Figures 6.13a and 6.13c, as the headwind strengthens, the aircraft must fly faster to make up time, so there is a negative correlation at the 5% significance level, with a product moment correlation coefficient of -0.8. For eastbound flights, shown in Figures 6.13b and 6.13d, there is a similar negative correlation at the 5% significance level, but with product moment correlation coefficients of -0.9. This shows that the aircraft fly more slowly as they have a set time to reach their destination and the tailwind can help to minimise fuel burn, by avoiding the need for higher, less fuel efficient airspeeds.

The variation in airspeeds across a trajectory is also likely to be directly affected by the wind field. In Figure 6.14a the positive correlation, with product moment correlation coefficient of 0.5, between average tailwind along the GCR and the standard deviation of optimal airspeeds from OCP2 used each day, shows that as headwinds become stronger flying West, the variety of wind speeds used to make a fuel optimal crossing decreases. This shows that higher airspeeds are being used as necessary to counter the headwinds to allow prompt arrival, but on days when a lower airspeed is adequate for some parts of the journey these are used to save fuel. For eastbound flights the strength of the wind around the GCR only accounts for about 30% of the variation in optimal airspeeds and so the link between range of airspeeds and GCR average wind speed is less strong.

The airspeed changing along a trajectory in the context of a fixed-time flight, has been shown to be a result of the winds encountered at each time step of the journey. To illustrate this effect further, the variation in airspeeds used along each daily fuel optimal trajectory is considered. Flying east on the 28th January, 2020 this variation is at its smallest and on the 25th January, 2020 it is at its most. On the 8th December, 2019 the westbound route has the largest variation in airspeeds, with the smallest variation on the 12th February, 2020.

In Figures 6.15a and 6.15b the eastbound trajectories are shown, colour coded for airspeed, against a quiver plot of the wind field. On the 28th January, 2020, the route largely follows the wind direction from the middle of the North Atlantic, so airspeed can be reduced as the flight moves closer to LHR. However, at the start of the trajectory a higher airspeed is needed to fly perpendicular to the wind. This can be reduced as the

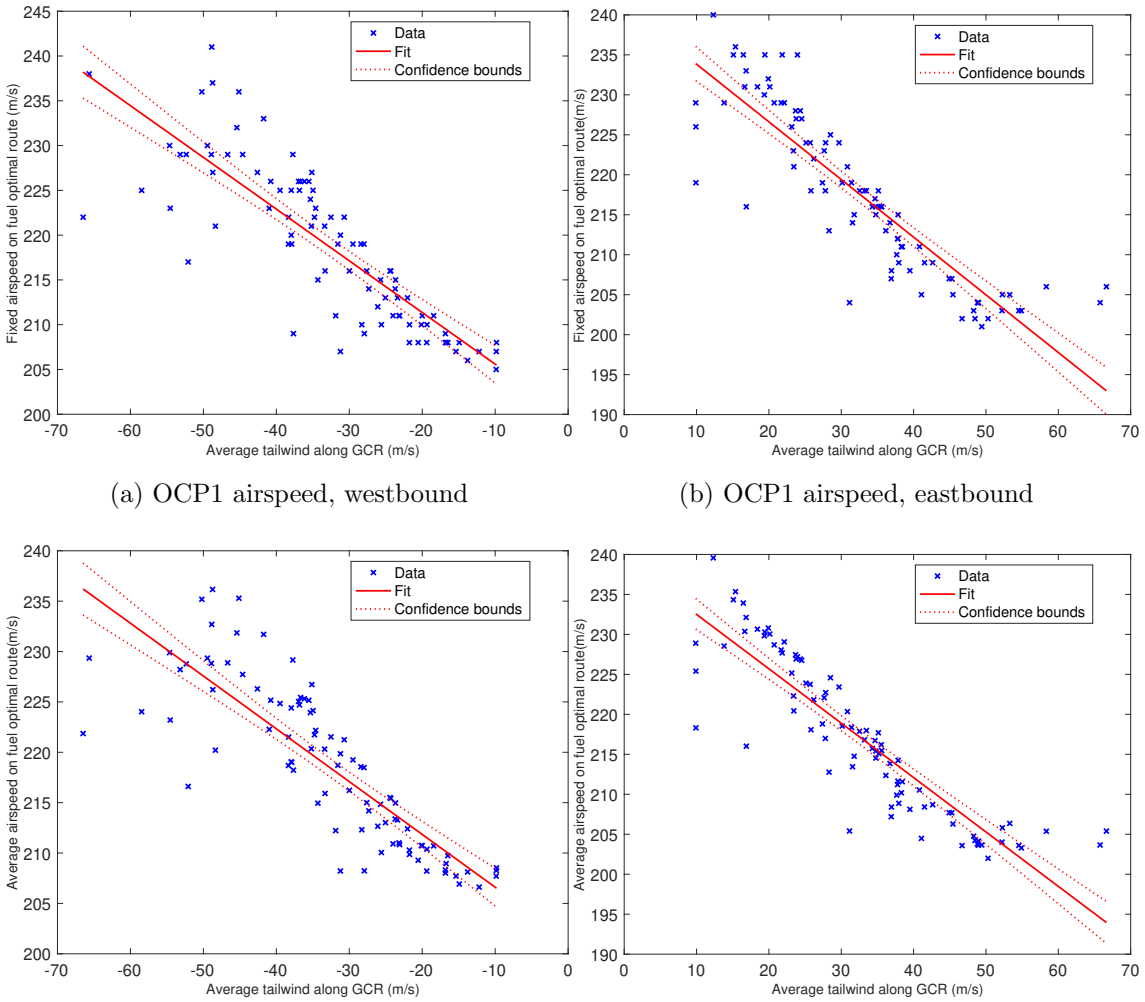

(a) OCP1 airspeed, westbound

(b) OCP1 airspeed, eastbound

(c) OCP2 average airspeed, westbound

(d) OCP2 average airspeed, eastbound

Figure 6.13: Scatter graphs to show variation of OCP1 trajectory airspeed and average OCP2 trajectory airspeed each day with average tailwind along the GCR. Least squares regression line is shown with a 5% confidence interval. These show that the airspeeds used depend largely on the wind field along the GCR.

route and the winds begin to come more into line between longitudes of 60°W and 50°W. On the 25th January, 2020 a smaller range of airspeeds is used as the flight is almost always flying in a direction that is not parallel to the wind vectors shown. This means that to arrive promptly, the aircraft must use higher airspeeds throughout. The only noticeable slowing occurs between the longitudes of 25°W and 10°W when the aircraft uses the vertical component of the wind to increase its latitude ready for the approach into LHR. This reinforces the previous finding, that stronger winds in the correct direction do mean that it is most efficient to use a larger range of airspeeds.

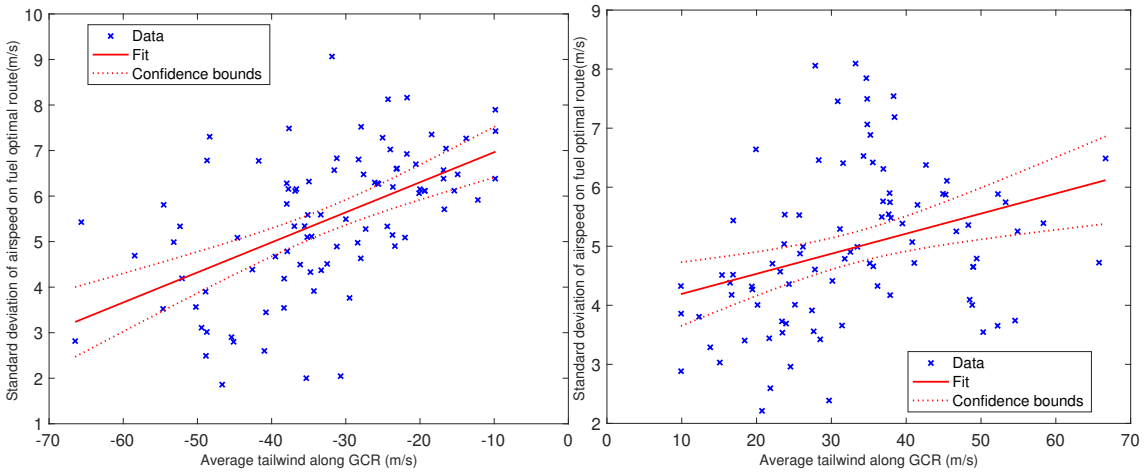

(a) Winds along GCR, westbound

(b) Winds along GCR, eastbound

Figure 6.14: Scatter graphs to show variation of standard deviation of OCP2 trajectory airspeed each day with average tailwind along the GCR with least squares regression line and 5% confidence interval. These show that there is a weak link between variation in airspeed and the daily windfield along the GCR.

We now compare the westbound trajectories in Figures 6.15c and 6.15d, on the 8th December, 2019 and the 12th February, 2020. During the first of these flights, a large range of airspeeds is used as the aircraft is able to fly further North and thus reach a patch of very weak winds for the majority of the flight. Having made good time across this central section, the aircraft can then slow down as it reaches the US coast. On the 12th February, 2020 the headwinds are stronger both as the aircraft leaves LHR and as it approaches JFK. This means that even across the centre of the North Atlantic a high airspeed must be maintained. From these examples the pattern from Figure 6.14 is evident, with higher headwinds around the GCR leading to a smaller range of airspeeds being used to ensure fuel optimality.

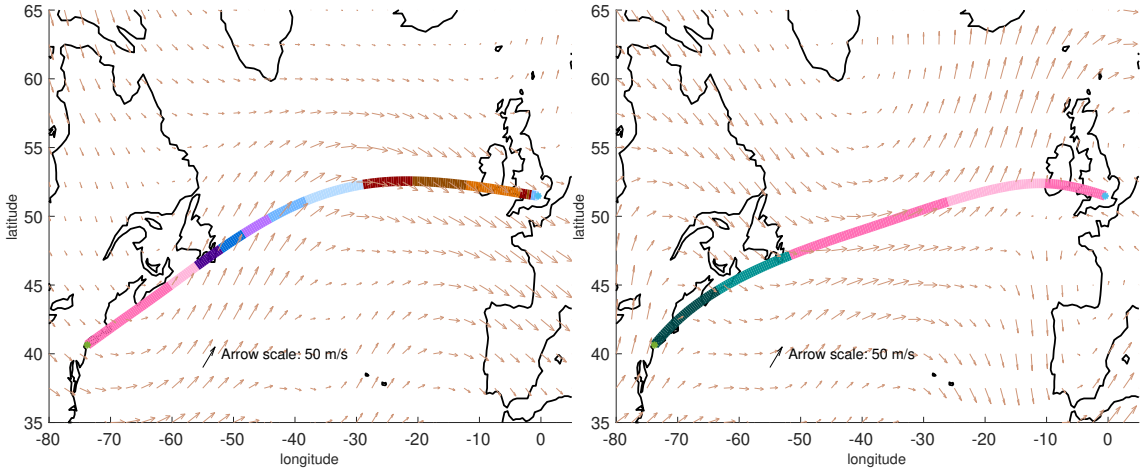

(a) 28th January, 2020 eastbound

(b) 25th January, 2020 eastbound

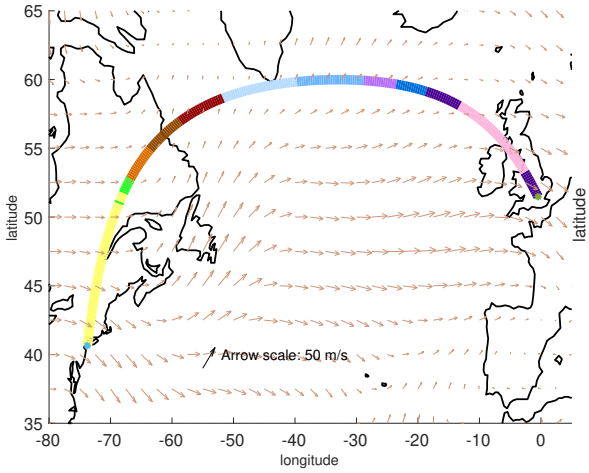

(c) 8th December, 2019 westbound

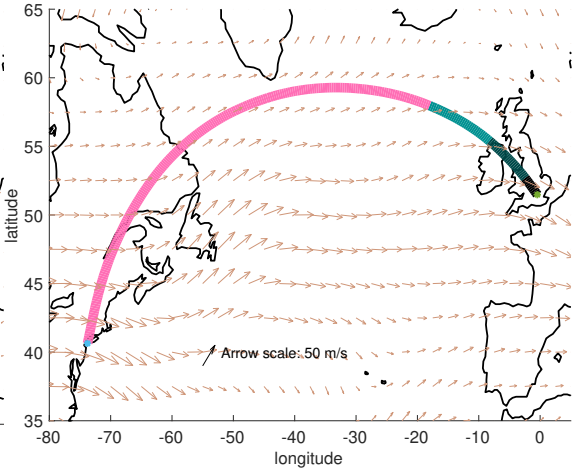

(d) 12th February, 2020 westbound

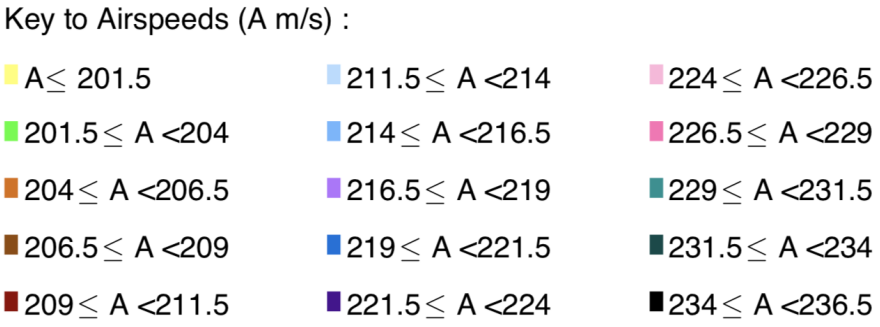

(e) OCP2 average airspeed, westbound

Figure 6.15: Fuel optimal trajectory maps showing airspeed change, in $\mathrm{m\,s^{-1}}$.

6.5 Conclusions

In this chapter a method is created to optimise fuel burn across a deterministic wind field for a fixed time of flight, using a novel fuel burn function. The time of flight is discretised into time steps. The altitude of the aircraft is assumed to be constant. This method involves controlling just the heading angle of the aircraft, or both the heading angle and the airspeed.

Flights are between LHR and JFK, with fixed times chosen to be slightly longer than the longest time optimal flight in each direction between these airports. The whole trajectory is assumed to be completed in cruise phase, as this makes up the vast majority of all transatlantic flights.

The fuel burn rate is calculated using a new physics based method (Poll & Schumann, 2021a,b) and the flights are assumed to move through the deterministic wind fields supplied by the NCAR re-analysis data. Wind and temperature data at each time step is found by linearly interpolating the grid of weather data, which is at a 2.5° resolution.

Results show that by including true airspeed as a second control up to 4% less fuel is used than when the flight is flown at constant airspeed. Over the course of all 91 days of the winter period, if nineteen flights were made each day (as was the case in 2019 to 2020), then just under 723 tonnes of fuel could be saved.

Links between airspeeds obtained from the fuel optimisation using both heading and airspeed as controls and the daily wind conditions are established. Daily wind conditions are used to generate an average tailwind speed along the GCR. It is seen that as this tailwind increases, the average airspeed used falls. Although the range of airspeeds used along a trajectory is not strongly linked to the average tailwind around the GCR, there is a weak positive correlation which is further demonstrated using airspeed patterns for four example days.

Airspeed use along a flight path is plotted for the days with the highest and lowest standard deviations of airspeed flying both east and west. These flights are considered against the backdrop of the day's wind field to show how the airspeed is adapted to make the best

use of available winds.

As the time of flight is fixed for an entire season, on some days it will be far longer than for the corresponding time optimal trajectory. This means that despite using more fuel efficient airspeeds, fixing the time of flight will consume more fuel over all, than using a time optimal route. This does assume, however, that time optimal flights are able to land and passengers disembark immediately they arrive at their destination, which in practice is not always the case. As tarmac delays are treated identically for late and early arrivals at an airport, arriving too early can also be costly for airlines.

Future research should incorporate more diverse routes and aircraft models, as well as turbulence avoidance, given the projected increase in more severe turbulence with climate change (Williams & Joshi, 2013; Williams, 2017; Storer et al., 2017; Lee et al., 2019). The use of dynamic programming to ensure sufficient conditions for an optimal route is also of interest as the current method uses a numerical global search function and so cannot guarantee optimality.

6.6 Chapter summary

In this chapter the need for fuel minimal, fixed-time flights has been established and the choice of fixed flight times in each direction justified. A direct method has then been applied to solve two different OCPs, one involving taking heading and airspeed as controls and one using just heading. By using two control variables up to 4% of fuel can be saved, but this depends on the wind field encountered each day. The link between fuel use and wind field is established as the average optimal airspeed each day is shown to have a strong negative correlation with the average wind speed along the GCR on that day. There is a weak positive correlation between standard deviation of airspeeds and wind speeds, demonstrating that there is more variation of airspeeds when wind speeds are higher.

In Chapter 7 dynamical programming will be used to solve the fuel minimal free-time problem, as fixing the time for a whole winter period means using extra fuel, just for predictability. Landing times can be varied on a daily basis, without this causing fuel penalties, with a more flexible approach to landing schedules at airports.

Chapter 7

Dynamic programming

7.1 Introduction

Instead of a direct or indirect approach to the OCP, which both have the issue of not guaranteeing a global minimum, here dynamic programming is used. This involves finding a value function to describe the optimality of each point we may encounter on a trajectory. To do this we apply Bellman's Dynamic Programming Principle, as shown for the infinite horizon problem in Equation (3.70) and this leads to the expression given in Equation (3.76). By solving this HJB equation using the method shown in Section 3.3.3, given as Algorithm 1, across the whole state space, we create a map of the value function. From this grid of values an optimal route is retrieved by identifying the controls associated with the smallest values at the start of each time interval and thus the corresponding states.

In the first instance a minimal time problem is considered, but in the second part of this work, the fuel burn is minimised, with time unconstrained. As change in mass does have an effect on the rate at which fuel is burned by an aircraft, mass is included as the third state variable in the fuel minimal formulation.

In Figures 7.1a and 7.1c the fuel burn rate for the nine different aircraft that flew between LHR and JFK during the winter of 2019 to 2020 is shown for a set of masses representative of those corresponding to a transatlantic flight. The airspeed and flight level are fixed at 240 m s^{-1} and FL340 respectively. The ISA temperature corresponding to this flight level is also used in the fuel burn calculation which is based on Poll & Schumann (2021a,b). Aircraft model specific parameters drawn from Poll & Schumann (2021b) to be included

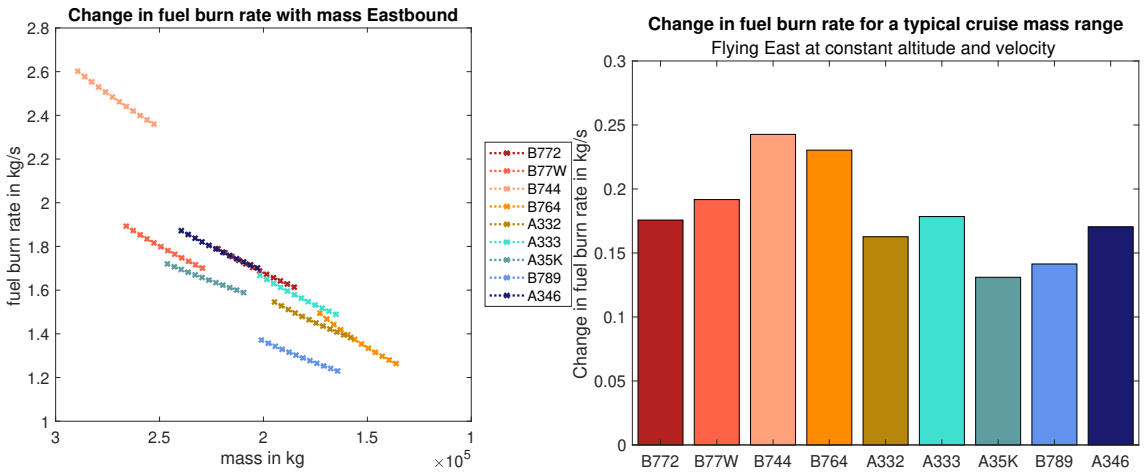

(a) Fuel burn rate

(b) Change in fuel burn rate

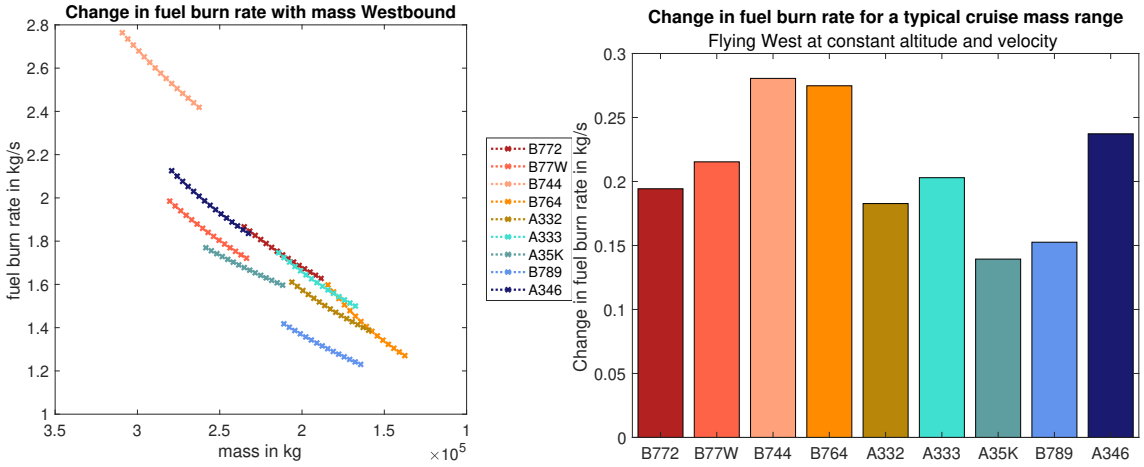

(c) Fuel burn rate

(d) Change in fuel burn rate

Figure 7.1: Plots showing values for fuel burn rate at a variety of different masses for the nine aircraft types flown between LHR and JFK during winter 2019-2020. Airspeed, temperature and altitude are fixed. Differences between eastbound and westbound results reflect the different take-off masses of aircraft. Differences depend on aircraft model parameters as given in Appendix C.

in the fuel burn model are given in Appendix C. In Figures 7.1b and 7.1d the change in fuel burn rate is shown, illustrating the need to include mass changes in the model where fuel burn is considered. To allow a fair comparison of results from time and fuel minimal versions of the problem, the fuel burn across each interval for the time minimal approach will also be included in the working and the mass change incorporated into the model after optimisation.

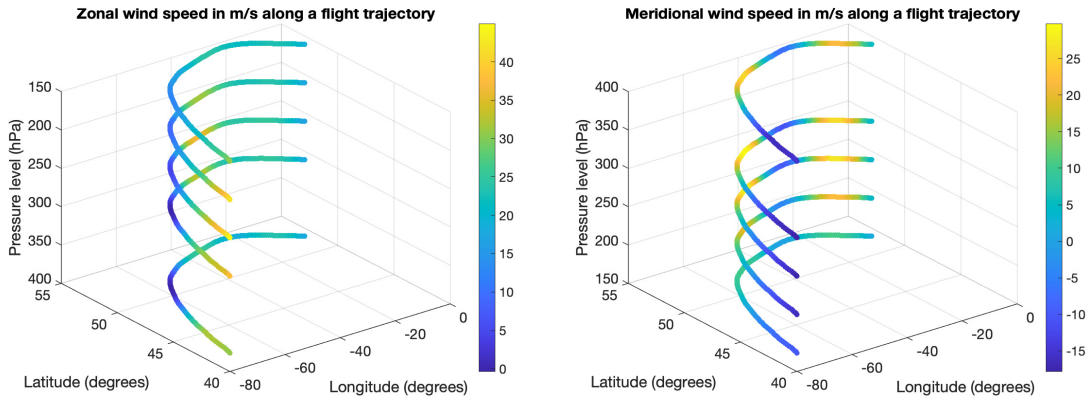

(a) Zonal wind speeds

(b) Meridional wind speeds

Figure 7.2: Scatter plots showing values for wind speed in m/s at a number of different pressure levels on 2nd January, 2020, along the horizontal trajectory of flight VS137.

As was shown in Chapter 3, the zonal and meridional winds across the North Atlantic are very similar at all cruise altitudes, so step climbs during flight will not give any substantial benefit from the point of view of atmospheric conditions, except in avoiding unexpected turbulence. This is illustrated in Figure 7.2, in which both the zonal and meridional wind speeds are shown for each time interval along the recorded horizontal flight path of a Virgin Airways flight, VS137, on the 2nd January, 2020. Data is from the Flightradar24 website, as discussed in Section 7.4. Wind speeds are shown at each of the pressure levels given in the re-analysis model data, as described in Section 3.5. As cruise altitudes are likely to be between 190 and 300 hPa, it can be seen that both zonal and meridional winds are very similar along the route at altitudes in this range. Changes in altitude to reduce drag as fuel is burned are not included in the initial analysis, but are considered in Section 7.7.1.

7.2 Method to retrieve optimal trajectories

In this section parameters and specifications for the grid and target are detailed. Discretisation of variables is chosen after sensitivity analysis involving small changes to one parameter in isolation and then a comparison of results. In addition the dynamical systems for both the time and fuel minimal formulations are non-dimensionalised and the relevant cost functionals specified.

7.2.1 Time minimal formulation

The system is modelled with two state variables: longitude, λ and latitude, ϕ. The state variables are given boundary conditions of -80 to 10° for longitude (in radians: $\lambda \in (\frac{-4\pi}{9}, \frac{\pi}{18})$) and 30 to 70° for latitude (in radians: $\phi \in (\frac{\pi}{6}, \frac{7\pi}{18})$), in order to include all feasible positions of a transatlantic flight. A horizontal grid is set up across this airspace with a resolution of 2.5° or $\frac{\pi}{72}{}^c$. This discretisation was chosen as atmospheric data is obtained at this level of accuracy and making a finer grid would require the same interpolation as is already inherent in reading the optimal route from the value function map.

Although mass was recorded at each time step based on the fuel burned in moving from one waypoint to the next, it did not form part of the optimisation itself, as fuel use is only calculated once a time minimal path has been found.

The aircraft begins its flight at the departure airport and continues until it is within 225 km of the destination airport. This target has been chosen to reflect the end of the cruise phase. Once the route has been found, the first few time intervals are removed from the route to ensure that only the part of the aircraft flight between 225 km from each airport is included. This allows for a fair comparison with actual cruise phase data.

In order for an aircraft to finish cruise within 225 km of the destination airport, we can use a query circle of all points distance d from a target to generate box constraints. Use of the conformal mapping from a planar to a spherical system is shown in Appendix B. The angular radius of the query circle is first found by dividing the radius of the query circle by the radius of the Earth:

$$r = \frac{d}{R} = \frac{225}{6371}. \tag{7.1}$$

As one radian of latitude corresponds to the same distance anywhere on the globe, the latitude bounds are found by just using a difference of r either side of the target airport:

$$latmin = \phi_{dest} - r, \tag{7.2}$$

$$latmax = \phi_{dest} + r. \tag{7.3}$$

To find the corresponding longitude boundaries, the fact that the length associated with a radian of longitude changes with latitude must be taken into consideration. So here Δlon

is calculated as:

$$\Delta lon = \text{asin}\left(\frac{\sin r}{\cos \phi_{dest}}\right). \tag{7.4}$$

This gives boundaries for the final longitude as:

$$lonmin = \lambda_{dest} - \Delta lon, \tag{7.5}$$

$$lonmax = \lambda_{dest} + \Delta lon. \tag{7.6}$$

So only points inside this bounding rectangle could possibly be in the destination target. Now only these need to be checked via the Haversine function to see if they are within the target. Here the bounding rectangle constraints can be given as:

$$\lambda_{dest} - \text{asin}\left(\frac{\sin 0.0353}{\cos \phi_{dest}}\right) \leq \lambda(t_f) \leq \lambda_{dest} + \text{asin}\left(\frac{\sin 0.0353}{\cos \phi_{dest}}\right), \tag{7.7}$$

$$\phi_{dest} - 0.0353 \leq \phi(t_f) \leq \phi_{dest} + 0.0353. \tag{7.8}$$

The value of 225 km is obtained by considering the 3114 flights between LHR and JFK from 1st December, 2019 and 29th February, 2020 for which route data is available. The distance between the end of the cruise phase for each flight and the destination airport is calculated. The 90th percentile for this distance for westbound flights is 224.4 km and for eastbound flights is 224.6 km. By choosing a radius of 225 km, this covers 90% of actual flights being considered.

The control variables are the heading angle, θ which varies from 0 to 360° in 2° intervals (0-2π^c with $\frac{\pi}{90}{}^c$ resolution) and the airspeed, V which is restricted to the acceptable flying range of 200 to 250 m s^{-1}, in steps of 2 m s^{-1}. As expected in a time minimal optimisation, the optimal airspeed obtained is always the maximum value of this range. Altitude is restricted to a single value corresponding to a pressure of 250 hPa, which approximates to FL340. The OCP is formulated as for Sections 3.2.2 and 3.2.3 and then the value iteration algorithm shown in Section 3.3.3, Algorithm 1, is applied.

By non-dimensionalising the dynamical system, any issues with the relative size of the variables is overcome. Longitude, latitude and heading angle are dimensionless, but airspeed, wind speed and time all have dimensions. The following definitions are used:

$$V^* = \frac{V}{V_{ref}}, \qquad u^* = \frac{u}{V_{ref}},$$

$$v^* = \frac{v}{V_{ref}}, \qquad t^* = \frac{V_{ref} t}{R},$$

where V_{ref} is the nominal cruise speed of $240\,\mathrm{m\,s^{-1}}$ and R is the approximate radius of the Earth, 6 371 000 m. Using these definitions the new dimensionless set of state equations is:

$$\frac{\mathrm{d}\lambda}{\mathrm{d}t^*} = \frac{1}{\cos\phi}(V^*\cos\theta + u^*), \tag{7.9}$$

$$\frac{\mathrm{d}\phi}{\mathrm{d}t^*} = V^*\sin\theta + v^*. \tag{7.10}$$

In the time minimal problem the continuous summation of the time for each section of the trajectory gives the running cost:

$$J(\lambda, \phi, \theta, V^*) = \int_{t_0^*}^{t_f^*} 1\mathrm{d}t^* = t_f^*, \tag{7.11}$$

which is minimised over the control variables, giving:

$$\min_{\theta, V^*} \int_{t_0^*}^{t_f^*} 1\mathrm{d}t^* = t_f^*, \tag{7.12}$$

where t_f^* is the time at which the aircraft first reaches the target circle of radius 225 km around the destination airport.

7.2.2 Fuel minimal formulation

The system when minimising fuel is modelled with three state variables, longitude, λ, latitude, ϕ and mass, M, with mass non-dimensionalised to give:

$$M^* = \frac{M}{M_{ref}}, \tag{7.13}$$

where M_{ref} is taken to be the mass of the aircraft at the start of the cruise phase, which is 97.5% of the take-off mass (TOM). The boundary values for the aircraft mass are taken as $M_{min} = OEM$ and $M_{max} = M_{ref}$. The first of these values is available from the Aircraft Characteristics documentation provided by both Boeing and Airbus (Airbus, 2021; Boeing, 2021) and a method to find the TOM is given in Section 3.4.2. The fuel burn rate, g, is also non-dimensionalised as:

$$g^* = \frac{Rg}{M_{ref}V_{ref}}. \tag{7.14}$$

The state mesh is split into cuboids that are:

$$\delta\lambda \times \delta\phi \times \delta M^* = \frac{\pi}{72} \times \frac{\pi}{72} \times \frac{3333\ kg}{M_{ref}}. \tag{7.15}$$

The mass discretisation is chosen, after a sensitivity study, to give a fine enough third dimension to the value function for meaningful results, without unduly slowing the computation time. The departure point, altitude and target are the same as for the time

minimal formulation. The controls also remain as heading angle and airspeed, although in this formulation airspeed is found to vary considerably across a trajectory. Due to the extra state, the dynamical system becomes:

$$\frac{\mathrm{d}\lambda}{\mathrm{d}t^*} = \frac{1}{\cos\theta}(V^*\cos\theta + u^*), \tag{7.16}$$

$$\frac{\mathrm{d}\phi}{\mathrm{d}t^*} = V^*\sin\theta + v^*, \tag{7.17}$$

$$\frac{\mathrm{d}M^*}{\mathrm{d}t^*} = -g^*. \tag{7.18}$$

The aim of the optimisation is to minimise a cost based on fuel burned during the journey:

$$\min J(\lambda, \phi, M^*, V^*, \theta) = \int_{t_0^*}^{t_f^*} g^*(\lambda, \phi, M^*, V^*)dt^*, \tag{7.19}$$

such that the aircraft arrives within a target set as a circle of radius 225 km around the destination airport at time t_f.

7.2.3 Adapting the value function for the fuel minimal case

To solve the fuel minimal, but free-time example, requires a different value function from the time minimal case, as the running cost is no longer 1, but the fuel burn function, g. The Hamiltonian is adapted from the infinite horizon problem in Section 3.3.3, but includes an extra term, involving the terminal cost ω:

$$H(x, \omega, p) := \sup_{\alpha\in\mathcal{A}}\{-p.f(x, \alpha) - g(x, \alpha) + (g(x, \alpha) - 1)\omega\}. \tag{7.20}$$

Here x represents the state variables, α the control variables, f the dynamical system and p the adjoints, as in Section 3.3.3. (This expression for the Hamiltonian is drawn from Cristiani & Martinon (2009).) Taking the terminal cost at each step to be the value function of the previous iteration, gives us a value function of:

$$\nu(x) - \inf_{\alpha\in\mathcal{A}}\{g(x, \alpha) + f(x, \alpha)\cdot\Delta\nu(x) - (g(x, \alpha) - 1)\nu(x)\} = 0 \quad x \in \mathbb{R}^n \setminus C, \tag{7.21}$$

$$\nu(x) = 0 \qquad\qquad x \in C, \tag{7.22}$$

where C is the target set.

Including the initial $\nu(x)$ in the infimum we get:

$$-\inf_{\alpha\in\mathcal{A}}\{g(x, \alpha) + f(x, \alpha)\cdot\Delta\nu(x) - \nu(x)g(x, \alpha) + \nu(x) - \nu(x)\} = 0, \tag{7.23}$$

which is the same as:

$$\sup_{\alpha\in\mathcal{A}}\{-f(x, \alpha)\cdot\Delta\nu(x) + g(x, \alpha)(\nu(x) - 1)\} = 0. \tag{7.24}$$

This can be approximated to first order by, :

$$\sup_{\alpha\in\mathcal{A}}\left\{\frac{\hat{\nu}(x)-\hat{\nu}(x+\Delta tf)}{\Delta t}+g(x,\alpha)(\hat{\nu}(x)-1)\right\}=0, \tag{7.25}$$

where Δt represents a small change in time and Δx is approximated to Δtf. This in turn leads to:

$$\sup_{\alpha\in\mathcal{A}}\{\hat{\nu}(x)-\hat{\nu}(x+\Delta tf)+\Delta tg(x,\alpha)(\hat{\nu}(x)-1)\}=0, \tag{7.26}$$

which means that:

$$\hat{\nu}(x)=\min_{\alpha\in\mathcal{A}}\{\hat{\nu}(x+\Delta tf)+\Delta tg(x,\alpha)(1-\hat{\nu}(x))\}, \tag{7.27}$$

is the value function for which the optimal feedback control can be found as:

$$\alpha^*=\arg\min_{\alpha\in\mathcal{A}}\{\hat{\nu}(x+\Delta tf)+\Delta tg(x,\alpha)(1-\hat{\nu}(x))\}. \tag{7.28}$$

Using a Semi-Langrangian Scheme which incorporates both a temporal and spatial discretisation, as in Section 3.3.3, the value policy update is found to be:

$$[\hat{\nu}]_i^{m+1}=\min_{a\in\mathcal{A}}\{I[\hat{\nu}]_i^m(x_i+\Delta tf(x_i,a))+\Delta tg(x_i,\alpha)(1-\hat{\nu}_i^m(x_i))\}, \tag{7.29}$$

where the subscript i refers to a grid and the superscript m to a fixed point iteration used to approximate the solution of Equations (7.21) and (7.22).

7.3 Aircraft specific fuel burn

Ten different aircraft models were used in making the transatlantic crossing between LHR and JFK between 1st January, 2019 and 29th February, 2020. These are shown in Table 7.1, together with the airline that flew them and their model code. It should be noted that Aircraft 6 and 9 have the same parameters and thus their model code and fuel burn formula are identical. All model specific parameters are shown in Appendix C.

Using the formula described in Section 3.4 and the aircraft specific parameters for each of these models of aeroplane, as given in Poll & Schumann (2021a,b), the amount of fuel used across each time interval of the cruise phase of the minimum time routes is found. These values are summed to give total cruise fuel, approximating the continuous summation from the cost functional shown in Equation (7.12) by a discrete summation.

In the case of the fuel minimal flights, the formula is used as part of the optimisation process, with each set of possible control variables being applied at each state grid point via the value function to build the value map.

Table 7.1: Aircraft models flying between LHR and JFK in the winter from 1st December, 2019 to 29th February, 2020, their aircraft codes and the airlines that flew them.

Aircraft	Model	Model code	Airlines
1	Boeing 777-236(ER)	B772	American Airlines, British Airways
2	Boeing 777-323(ER)	B77W	American Airlines, British Airways
3	Boeing 747-436	B744	British Airways
4	Boeing 767-432(ER)	B764	Delta Air Lines
5	Airbus 330-223	A332	Delta Air Lines
6	Airbus 330-323	A333	Delta Air Lines, Virgin Atlantic
7	Airbus 350-1041	A35K	Virgin Atlantic
8	Boeing 787-9 Dreamliner	B789	Virgin Atlantic
9	Airbus 330-343	A333	Virgin Atlantic
10	Airbus 340-642	A346	Virgin Atlantic

7.4 Flight data

Recorded sets of timed aircraft positions for all flights between LHR and JFK from 1st December, 2019 to 29th February, 2020 are used to estimate fuel burned by actual flights. The data needed is recorded on transponders in the aircraft themselves and gathered from the Flightradar24 ADS-B (Automatic Dependent Surveillance Broadcast) network of over 20 000 connected receivers (Flightradar24, 2020). Transponders on planes transmit signals to the receivers based either terrestrially or, where this is not possible, to satellites equipped with ADS-B receivers. Transponders are positioned in most commercial aircraft making the transatlantic crossing. The information sent to the Flightradar24 receivers includes: (i) position as degrees of longitude and latitude, (ii) direction as a bearing, ground speed in knots and (iii) altitude in feet. However, we recalculated ground speeds as those recorded are not always consistent with given positions and times.

Data gives a step by step record of all flights, as detailed in Table 7.2. If data is missing for all or part of a flight, then this is removed from the dataset.

Table 7.2: Flight numbers for aircraft flying between LHR and JFK in the winter from 1st December, 2019 to 29th February, 2020. Odd numbers are westbound flights, for all but Delta Airlines, whose eastbound flights have odd numbers. Eastbound flights have the corresponding even (or in the case of Delta Airlines, odd) numbers.

American Airlines	British Airways	Delta Air Lines	Virgin Atlantic
AA101, AA100	BA113, BA112	DL1, DL2	VS3, VS4
AA105, AA104	BA115, BA114	DL3, DL4	VS9, VS10
AA107, AA106	BA117, BA116		VS25, VS26
AA141, AA142	BA173, BA172		VS45, VS46
	BA175, BA174		VS137, VS138
	BA177, BA176		
	BA179, BA178		
	BA183, BA182		

To find airspeeds for each stage of each flight, as these provide the basis for fuel burn calculations, the ground speed and the wind field for each day are used. By taking the ICAO definition of the ISA and interpolating between the altitude levels with a given pressure conversion in the hypothetical model, the pressure at each point reached is estimated. This is paired with the co-ordinates of the point. Using linear interpolation of a three dimensional array, the zonal and meridional wind components and the temperature at every point in an aircraft's cruise phase are approximated. Once the wind speeds are known, the wind in the direction of the plane is calculated. These can be subtracted from the recorded ground speed, to give the airspeed required in the fuel burn calculation. The flights are modelled to have constant airspeed between recorded positions, unless the time step is longer than two minutes.

Across each trajectory, there is usually a large time interval for which no extra information is recorded due to signalling issues. In view of this any time intervals longer than two minutes, are split into smaller steps, with equal duration of approximately two minutes. The aircraft positions are then modelled as following a shortest ground distance route between the co-ordinates recorded for the start and end of the original interval. Al-

titude is assumed to stay the same throughout these new time steps, but the atmospheric conditions are adjusted for the new positions to best model the actual aircraft airspeed and thus fuel use.

In some cases, the recorded positions when coupled with the re-analysis data do not give feasible values for airspeed, as the maximum Mach number for the aircraft is exceeded or the minimum Mach number subceeded. In such cases, the airspeeds are replaced by the closest feasible values and the time intervals adjusted accordingly. These occurrences are most likely to be due to inaccuracies in the data recorded. Although there may have been slight discrepancies between the re-analysis data atmospheric conditions and those actually encountered, these would be minimal given the large amount of observational data incorporated in the re-analysis model, with some of that data being recorded directly by aircraft (Kalnay et al., 1996).

In order to minimise fuel use, aircraft will change altitude to travel at the most efficient flight level for their current mass. This means that when data from the simulations is compared to estimated actual fuel use there are some non-optimised flights which appear to use less fuel than the simulated fuel minimal flights. In view of this, the actual fuel use was calculated both for the altitudes recorded and for horizontal paths to allow for a fairer comparison. Ideally the simulation would include a fourth state of altitude, but this was found to cause issues with data overloading, due to what Richard Bellman described as "the curse of dimensionality" (Bellman, 1957). This is when adding extra state variables causes the complexity of a problem to increase rapidly due to the increasing number of possible combinations of states and controls.

In Figure 7.3 it can be seen that the fuel burn rate is not affected equally by altitude changes for all models of aircraft, but that it does have an affect. As the total number of seconds for each flight is so large (approximately 20 000), these differences are magnified. here the airspeed has been fixed to 240 $\mathrm{m\,s^{-1}}$.

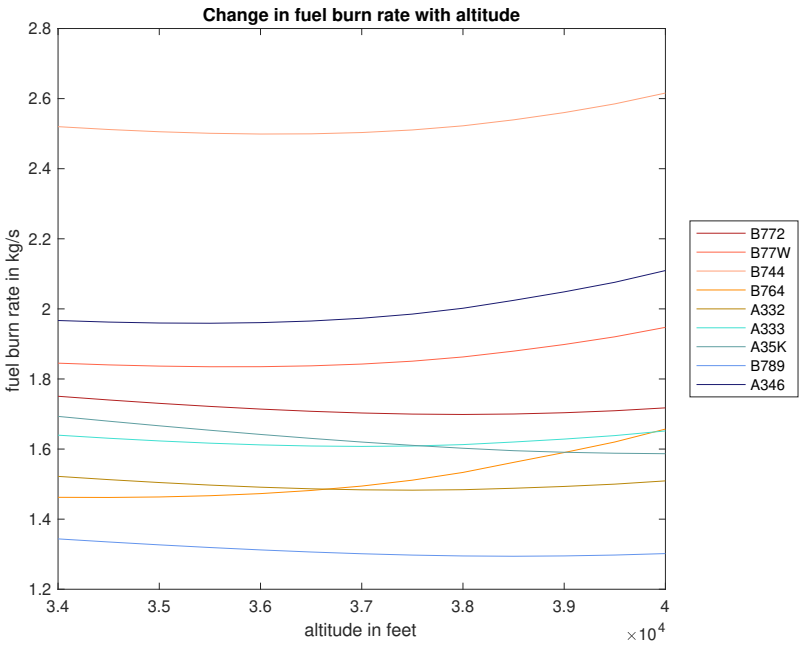

Figure 7.3: Plot showing how fuel burn rate varies with altitude for each aircraft model when half of the flight fuel has been used, at a point midway across the North Atlantic. Temperatures at each altitude are taken from a typical day, the 26th January, 2020.

7.5 Practical computational considerations

A time step of 125 s was chosen to allow for the practical adjustment of heading angle and airspeed. Flight data recorded for transatlantic flights on Flightradar24 has changing time steps which average to this value (Flightradar24, 2020).

Whilst the time minimal code can be run across the whole state space with a time step of 125 s in a timely manner, including the extra state of mass for the fuel minimal code leads to lengthy optimisation times of over an hour. In view of this the fuel minimal code is first run with a time step of 500 s. Once a route is retrieved from the value map given, the state space is reduced to only include grid boxes including the route and the grid boxes immediately adjacent to these. In order to work in this way, any grid points outside of the new range are given a value of 1 at each iteration. The next application of the code has a time step of 250 s. This process is repeated once more, to arrive at the time step of 125 s. A diagram illustrating this method of reducing the state space is given for a two dimensional example in Figure 7.4.

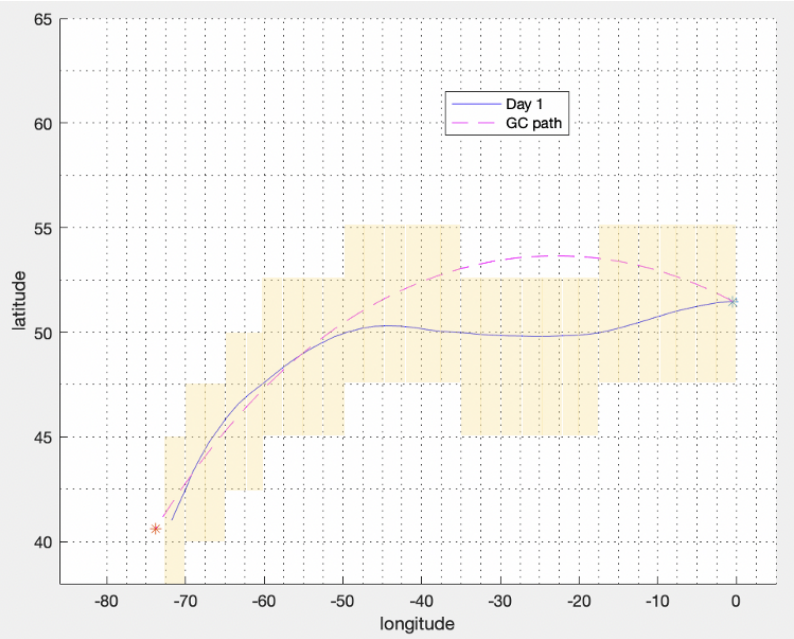

Figure 7.4: 2D example showing how the state space is reduced after the first stage of the simulation on 1st December, 2019 flying West. In this case the diagram shows just a horizontal state space, although in the actual simulation mass is used as a third dimension. The nodes in the shaded region are taken as the new state space after the initial route has been found. The Great Circle (GC) path is the shortest distance path around a sphere between the airports.

The accuracy of this method was confirmed first by comparing time minimal routes to final fuel minimal routes, as these had not been calculated with a larger initial time step or on a reduced state space. Next some examples of fuel minimal calculations were run for an extended computation time using just the 125 s time step across the whole state space. Results for this method were identical to those for the computationally much faster version. Code run with a 10 s timestep across the whole region and with a finer spatial discretisation confirmed that results were accurate, but took up to seven days to run. By reducing the state space considered and using a more practical timestep, computational time to create the initial value map was reduced to five minutes.

7.6 Results

Here the data from the time minimal and fuel minimal simulations is compared to actual flight data that has been limited to a pressure level of 250 hPa, corresponding to an

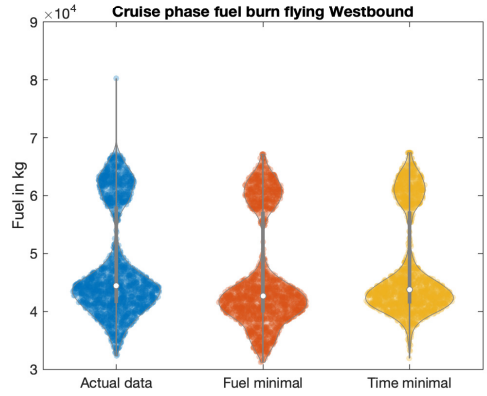

(a) Fuel use West

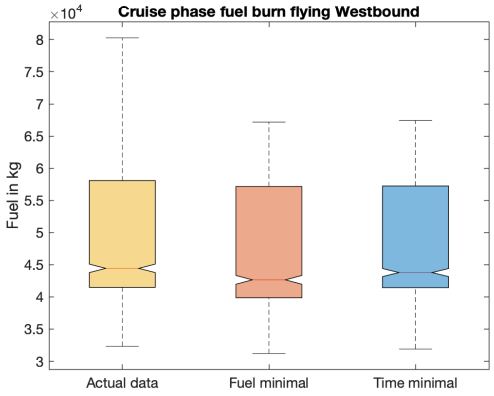

(b) Fuel use West

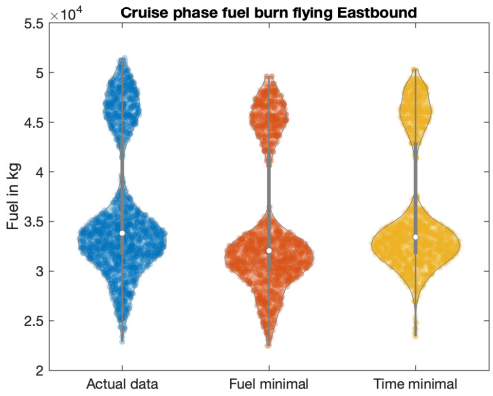

(c) Fuel use East

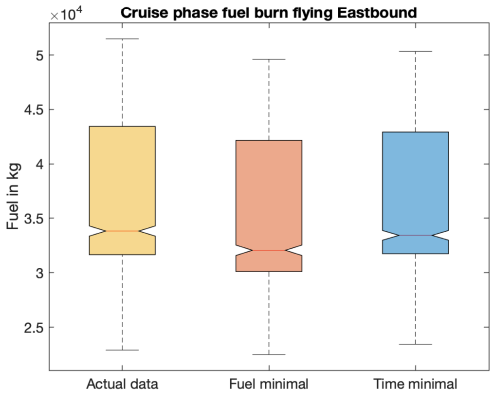

(d) Fuel use East

Figure 7.5: Violin and box plots showing values for fuel burned during the cruise phase of flights between LHR and JFK during winter 2019-2020. Fuel burn is considered for actual flights, simulated fuel minimal trajectories and simulated time minimal trajectories.

altitude of approximately 34 000 ft.

7.6.1 Comparing fuel and time data: all flights

In total, data was available for 1567 flights travelling west and 1547 flights travelling East, between the 1st December, 2019 and the 29th February, 2020. Time and fuel minimal routes were simulated using the same atmospheric and aircraft data as was applied to finding the fuel burn for these actual flights.

In Figure 7.5 both violin plots and notched box plots are shown for the fuel use for

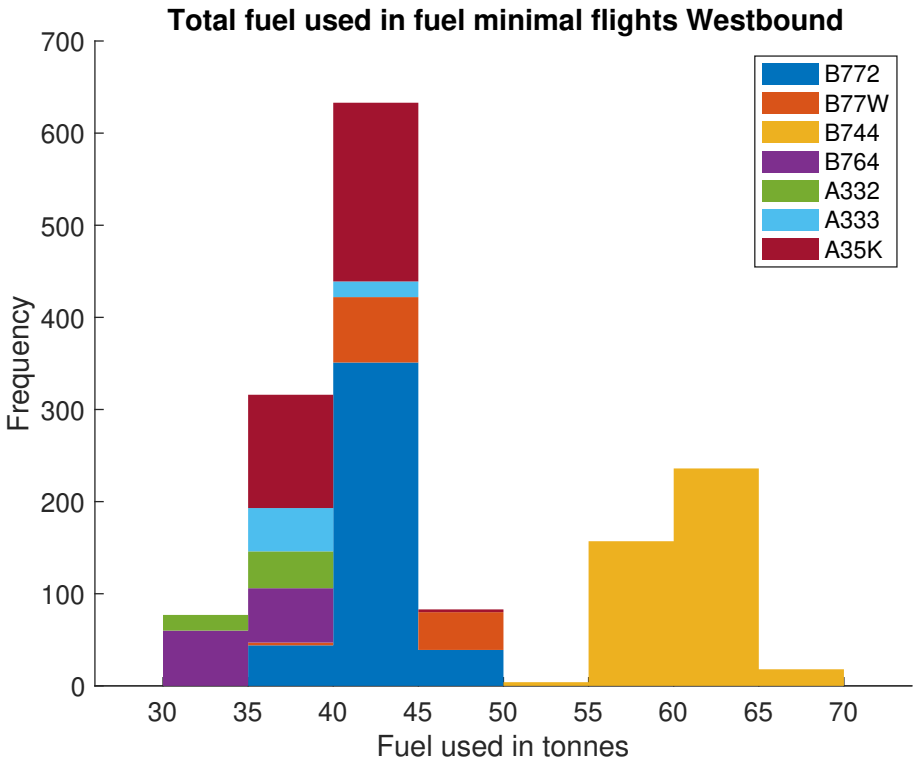

Figure 7.6: Plot showing how much fuel is used for the fuel minimal trajectories for the 7 most used models of aircraft for routes between LHR and JFK from 1st December, 2019 to 29th February, 2020.

the routes in each direction. It is clear from the violin plots that the data is not normally distributed, but follows a bi-modal distribution, indicating that there are two distinct groups within the data. The first of these has a global maximum and shows fuel use for the majority of routes flown. The second mode is lower, appearing to show that for a smaller group of flights more fuel is needed. Although at first glance this might be thought to be due to the atmospheric conditions encountered on different days, in actual fact it is entirely due to the different models of aircraft flown. In Figure 7.6, the fuel use for each fuel minimal trajectory is plotted on a stacked bar chart, clearly showing that the model of aircraft is causing the second mode. In this chart only the 7 models of aircraft flying the most routes are included, for clarity. Although the Boeing 747-436 is responsible for far greater fuel use than the other aircraft, it does also carry more passengers, with capacity for 524 in a two class layout, compared with 335 for the Airbus 350-1041. In Section 7.6.2 the possible savings for different models of aircraft will be discussed fully.

All sets of routes have a similar distribution of fuel use, but from the notched box-plots it can be seen that the median of the fuel minimal routes is significantly lower than

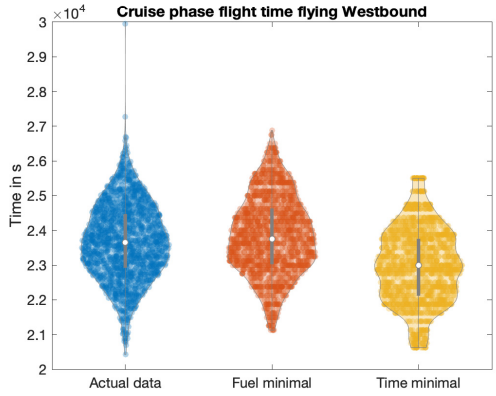

(a) Time of cruise West

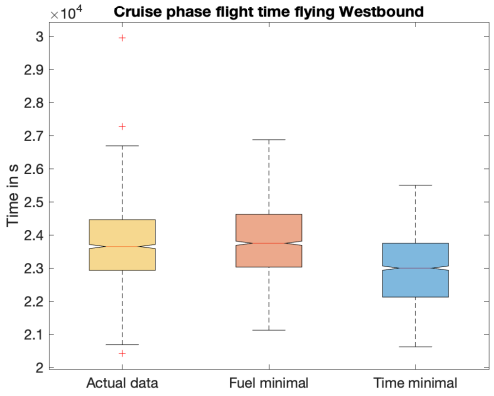

(b) Time of cruise West

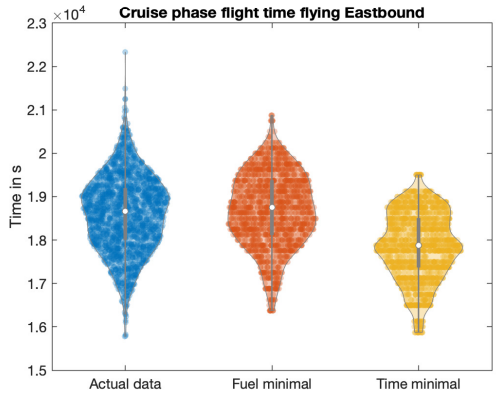

(c) Time of cruise East

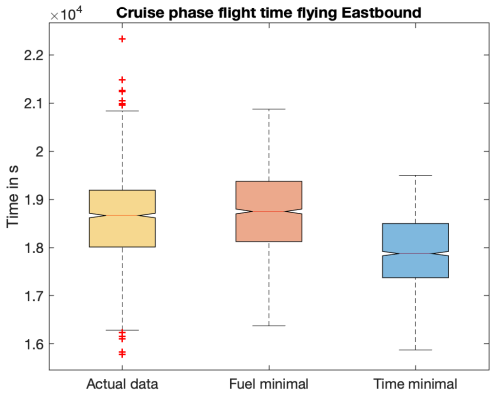

(d) Time of cruise East

Figure 7.7: Violin and box plots showing duration of the cruise phase of flights between LHR and JFK during winter 2019-2020. Time is considered for actual flights, simulated fuel minimal trajectories and simulated time minimal trajectories.

that for the time minimal routes, as might be expected and also for routes based on the actual flight data. The difference between the fuel use of the fuel minimal routes and actual flight routes is tested using a one tailed t-test at a 95% level of significance (Student, 1908). The pairwise difference between these data sets has a mean of more than 0, showing that the fuel minimal routes are significantly more efficient in terms of fuel burn. These patterns are the same for flights in both directions with more pronounced differences between eastbound data sets, but one westbound flight does use far more fuel than the rest, as can be seen in Figure 7.5a.

In Figure 7.7 the same charts are plotted for the time of flight for each type of trajec-

tory. It is unsurprising that the time minimal routes had a significantly shorter duration than the fuel minimal routes and the actual flight routes. However, the difference in the distribution of these times is surprising. The violin plots show that the actual cruise flight times and the fuel minimal cruise flight times are distributed in a far more similar way than the time minimal values. Flying west the cruise flight durations of the time minimal routes have a negative skew and both the other two route types show a slight positive skew. Heading east this pattern is reversed.

A one-tailed t-test is again applied and showed that the time minimal routes are significantly faster than the actual routes and the fuel minimal routes (Student, 1908). It might have been expected that the fuel minimal route would take significantly longer than the actual routes in order to burn fuel more efficiently, but in fact this is not the case and there is no evidence at the 95% level of significance that the fuel minimal trajectories are in fact slower than the actual routes flown.

Overall, by flying a fuel minimal route, it was found that the airlines flying between LHR and JFK could have saved approximately 2.5 million kg of fuel flying east and 2.8 million kg flying West, as is illustrated in Figure 7.8. Using the Environmental and Energy Study Institute's estimate that every 1 kg of aviation fuel burned produces 3.16 kg of CO_2, this would amount to a reduction of almost 16.6 million kg of CO_2 emissions across just one winter period (Overton, 2022).

One of the reasons why actual cruise phase flights may be using more fuel than fuel minimal flights, is down to the airspeeds that the aircraft are flown at. If these are too high, the fuel burn rate will rise as the efficiency of the propulsion system decreases, but if they are too low, the cruise phase will take longer and as fuel burn rate is measured in kg s^{-1}, more fuel will be burned over all. To check whether airspeed does have an important effect on fuel efficiency, Figure 7.9 shows box plots of average airspeed per flight for the fuel minimal and actual routes. The time minimal trajectories were all flown at 250 m s^{-1} (as the time optimisation chose the highest control value available for airspeed at all time steps), so a single line on the chart represents all of the data for these routes. It is obvious from Figure 7.9 that the actual flight routes are flown faster than is optimal. A one-tailed t-test at a 95% significance level confirmed that along the actual routes mean airspeeds are significantly higher than those for the fuel optimal trajectories (Student, 1908). This

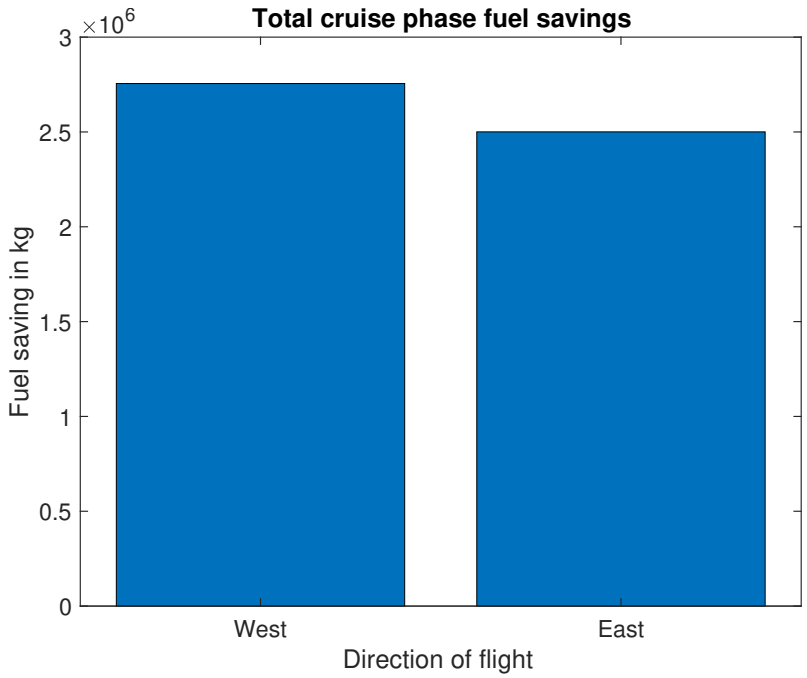

Figure 7.8: Plot showing how much fuel could be saved if all aircraft flew fuel minimal routes between LHR and JFK from 1st December, 2019 to 29th February, 2020.

is true despite some outlying actual route airspeed values being far lower than expected. If the overall average airspeed per route is significantly lower in the fuel minimal simulations, but these trajectories are not of significantly longer duration, then clearly the fuel decrease is also due in part to a better route having been chosen via the heading angle control.

7.6.2 Comparing data: by route, airline and aircraft

The patterns shown in the previous results section do not make a distinction between subsets of data, that could be relevant to our understanding of how improvements in routing could best be implemented.

In Figure 7.10 we look first at percentage differences in fuel and time use between actual flights and fuel minimal flights separated by flight number. There are 19 different flight codes for the journey between JFK and LHR in each direction over the winter of 2019 to 2020. However, on any given day a particular flight may or may not happen or the flight data may not have been recorded fully. The box plots in Figure 7.10c show the distribution of complete sets of flight data taken across the whole winter period.

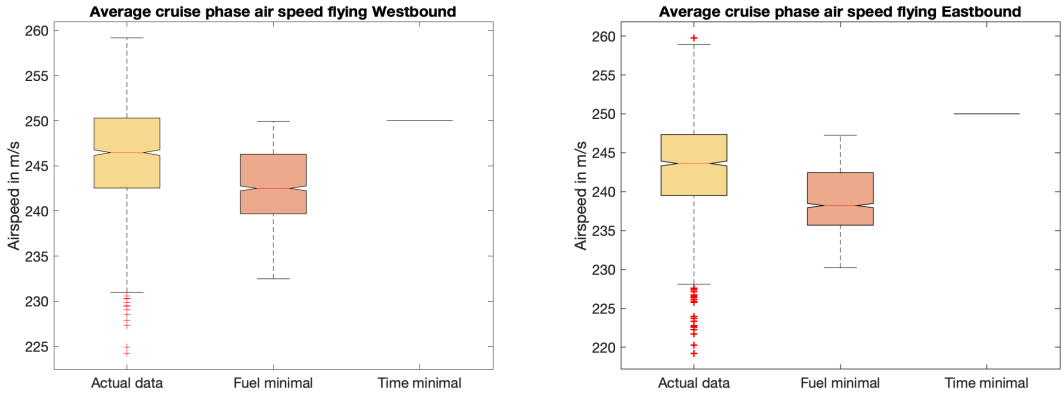

(a) Average airspeeds West

(b) Average airspeeds East

Figure 7.9: Box plots showing average airspeed for each aircraft during the cruise phase of flights between LHR and JFK during winter 2019-2020. Average airspeed is considered for actual flights and simulated fuel minimal trajectories. Time minimal trajectories have a constant airspeed of 250 m s^{-1}, as this highest airspeed level is chosen as the optimal feedback control in the optimisation.

The charts show that in order to make a fuel saving, it is often necessary to incur a time penalty on a particular route. This is most marked in the case of the Virgin Atlantic (VS) flight routes. Most of the other routes show a very small time penalty for a much larger fuel reduction, with the largest percentage savings in fuel to be made flying East. This reflects the fact that the OTS prevents all flight routes from accessing the most favourable winds.

Having seen that the patterns in fuel and time savings between different flight routes seemed to be largely based on the particular airline flying these routes, it is also useful to group results just by airline. Now it is more obvious if certain transatlantic service providers are already flying more efficient routes than others. In Figures 7.11a and 7.11c the percentage of time and fuel saved by flying the fuel minimal route rather than the actual recorded flight route is shown. To help to gauge the relative importance of these effects a pie chart showing the proportion of flights flown by each airline is also included in Figures 7.11b and 7.11d. From these charts it is clear that the majority of the flights are flown by British Airways, who would be able to reduce fuel by between 2 and 3% with

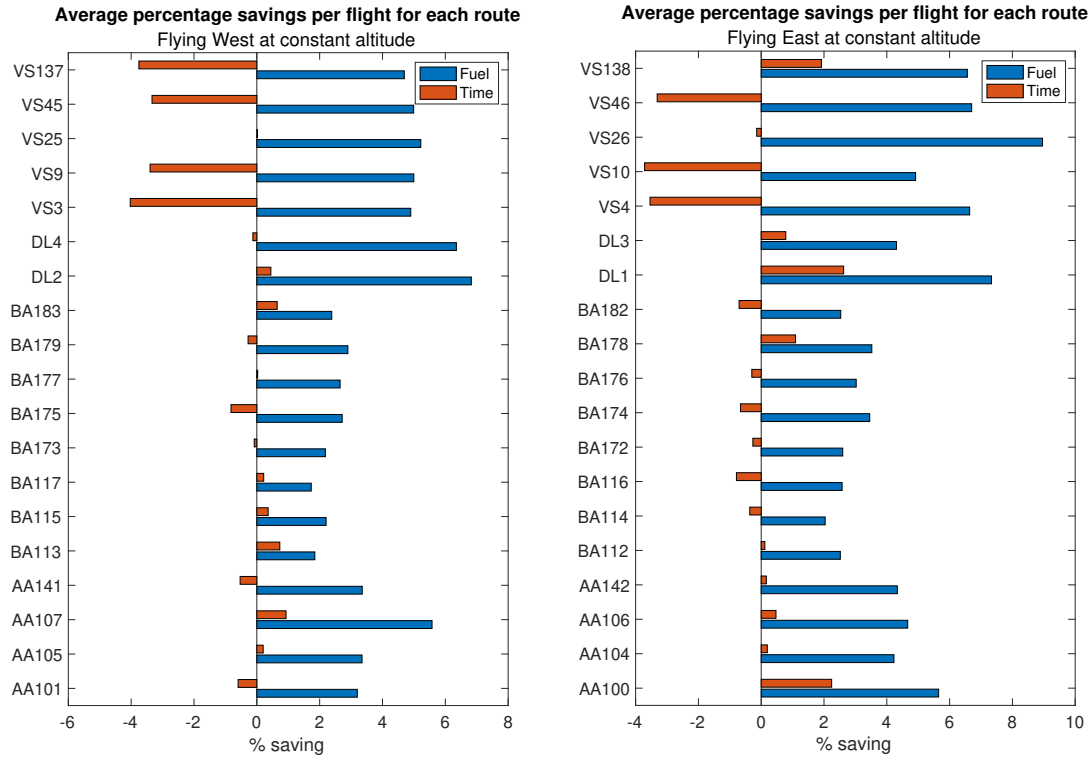

(a) Percentage saving West

(b) Percentage saving East

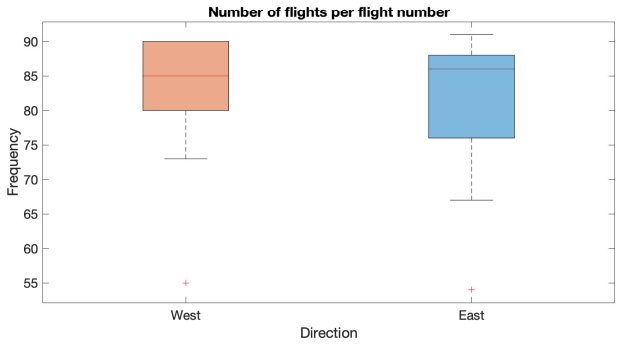

(c) Distribution of flights

Figure 7.10: Bar charts showing percentage savings in time and fuel made by using fuel minimal trajectories rather than actual routes during the cruise phase of flights between LHR and JFK during winter 2019-2020. The box plot shows the distribution of the number of flights per route code.

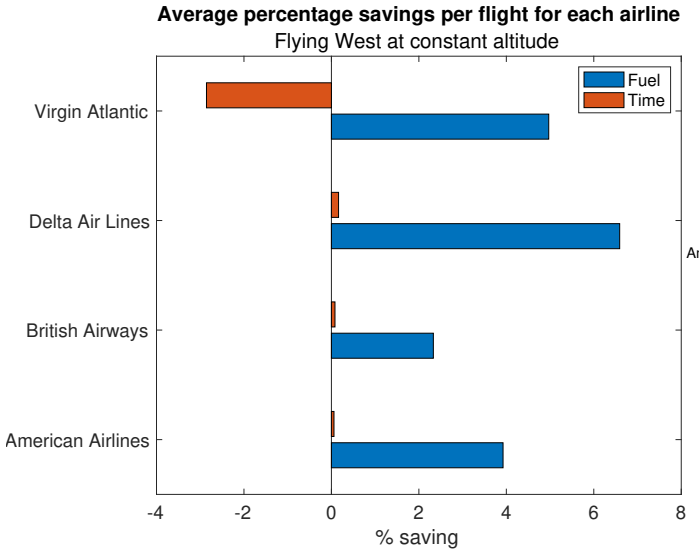

(a) Percentage savings West

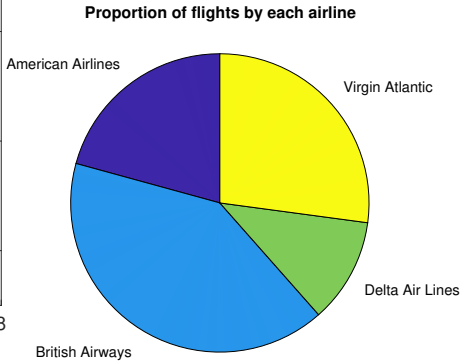

(b) Proportion per airline West

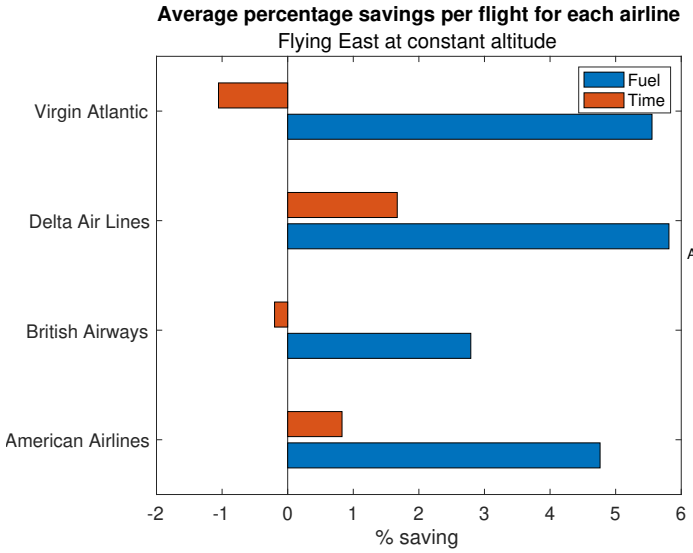

(c) Percentage savings East

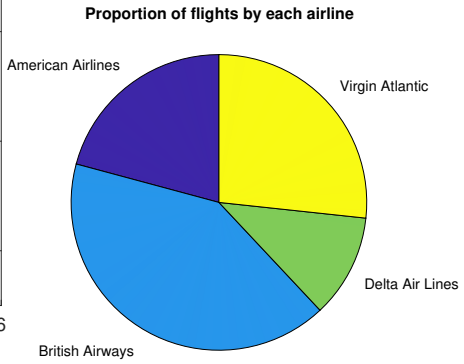

(d) Proportion per airline East

Figure 7.11: Bar charts showing percentage savings in time and fuel made by using fuel minimal trajectories rather than actual routes during the cruise phase of flights between LHR and JFK during winter 2019-2020. The pie charts show the proportion of the flights made by each airline.

only a very minimal time penalty on eastbound flights. This may mean that both their routes and airspeeds would become more fuel efficient as part of the simulations. In contrast, Virgin Atlantic could save far more fuel, but would also have longer duration flights, showing that the fuel savings will mainly come from reducing airspeed. Delta Airlines and American Airlines have a smaller share of the routes between LHR and JFK, but could save both time and fuel by moving away from the OTS, so may have been flying on the less efficient tracks in the past.

Finally, as aircraft parameters do have an effect on fuel burn rate, data from each type of aircraft detailing average fuel saved per cruise flight is plotted in Figure 7.12. It can be seen from Figures 7.12a and 7.12c that the largest savings are made by the changes to the routes and airspeed of the Airbus 340-642, 330-323 and 330-343 when flying eastbound, with the wind and by changes for the Boeing 767-432 and Boeing 787-9 when flying westbound against the wind. The Boeing 747-436 has the smallest average absolute savings in both directions, but as the flights for this model were all British Airways routes, it could be argued that it was generally flown on the more efficient tracks.

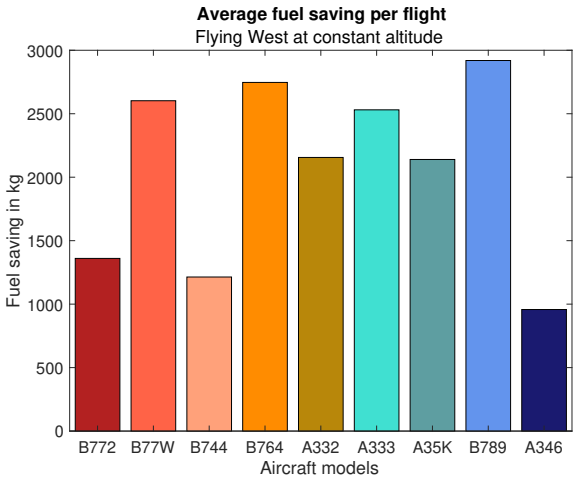

(a) Absolute average fuel savings West

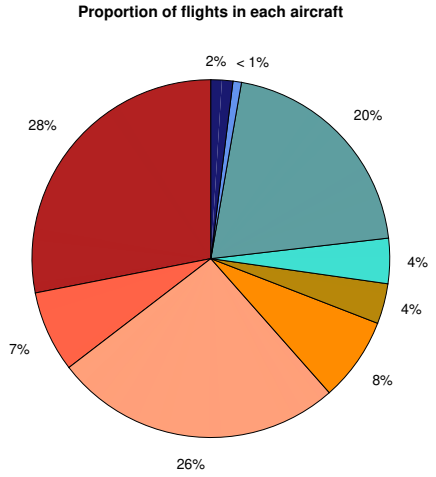

(b) Proportion per aircraft West

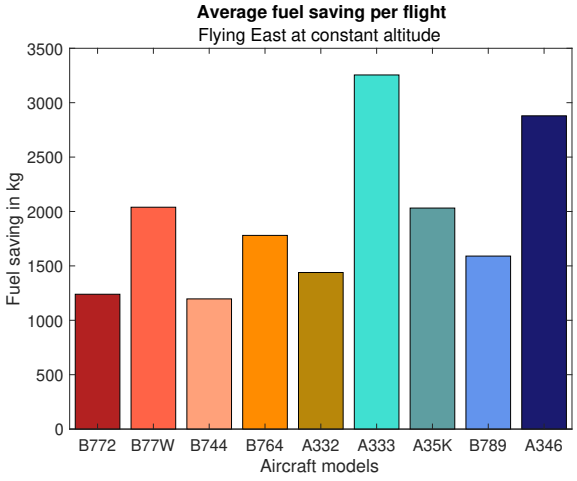

(c) Absolute average fuel savings East

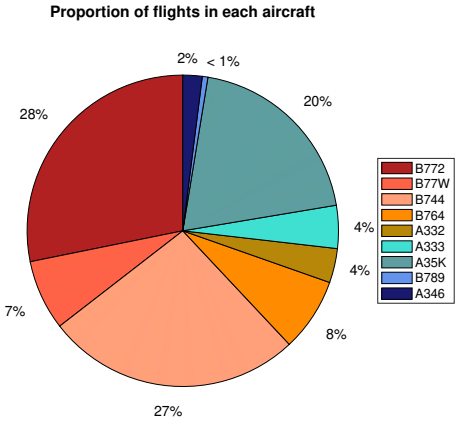

(d) Proportion per aircraft East

Figure 7.12: Bar charts showing absolute savings in fuel made by using fuel minimal trajectories rather than actual routes during the cruise phase of flights between LHR and JFK from 1st December, 2019 to 29th February, 2020. The pie charts show the proportion of the flights made by each type of aircraft.

7.7 Discussion

The results from comparing fuel minimal trajectories with actual routes flown look very favourable, but there are limitations to these calculations which need further discussion.

7.7.1 Altitude changes

It could be argued that comparing simulation results with routes that have never been flown at the given altitude may lead to unrepresentative findings. There are two ways to tackle this idea.

The first is to compare fuel use for the fixed altitude simulations with the actual variable altitude flights. This immediately gives an advantage to the flight routes that are allowed to change altitude to ensure more efficient fuel burn. However, even given this extra bonus overall absolute fuel savings of 2.4 million kg are still possible and a t-test at the 95% level of significance still shows a significant reduction in fuel when results are compared pairwise between fuel estimates from actual data and fuel minimal simulation results (Student, 1908). Figure 7.13 shows the fuel usage from these two methods across all days. Although the fuel minimal data is visibly weighted to the left of the charts, this is less marked for the westbound flights.

By allowing the actual flights to change altitude and keeping the simulations at fixed altitude reduces the eastbound savings by 47% and the westbound savings by 58%. Although across all airlines, savings are still made by flying the simulated routes at fixed altitude, the Airbus 350-1041 actually uses 1% more fuel over the 317 flights it makes going West. This aircraft has a take-off mass which lies just above the median for all models, but the design of the propulsion system makes fuel efficiency reduce with altitude more than for the other aircraft, as shown in Figure 7.3.

The second method is to allow the simulated routes to fly at the most efficient altitudes and model the wind fields as identical to that at 250 hPa at all cruise altitudes. This is not completely accurate, but as was explained in Section 7.4, including the altitude as part of the optimisation was prohibited by computational restrictions. Figure 7.2 shows that our assumption of the same wind field at each altitude is not entirely spurious. Temperatures, however, are different, as shown in Section 3.5.2, so the temperature used in the fuel burn

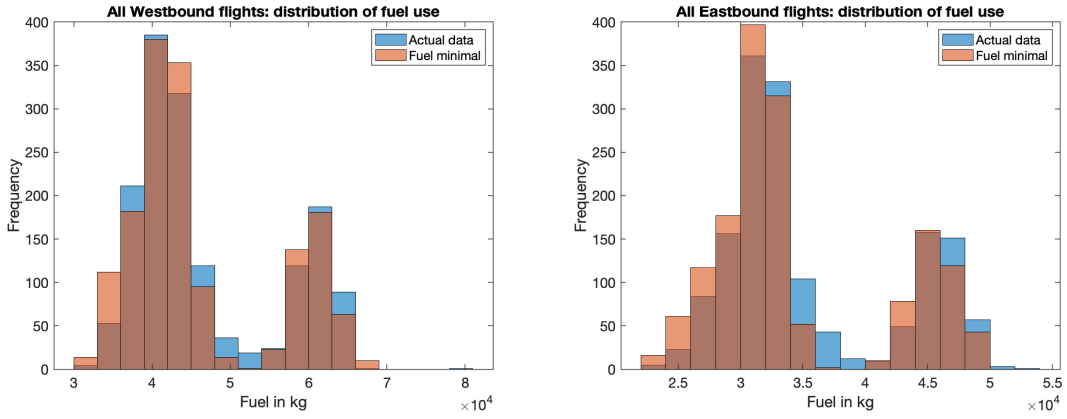

(a) Fuel use West

(b) Fuel use East

Figure 7.13: Histograms showing the distribution of fuel used between LHR and JFK for all days between 1st December, 2019 and 29th February, 2020. Results for the estimated fuel used by actual flights changing altitude and the fuel used in simulated fuel minimal fixed altitude flights is displayed.

calculations will be taken from the appropriate altitude. By using the horizontal paths found in Section 7.6 and adjusting the fuel burn calculation to reflect altitude changes we are also ignoring the fact that the mass state formed part of the original optimisation. The state at each point is an input to the creation of the value map which in turn allows the optimal feedback control to be obtained. However, this gives at least an approximation of the results that could be obtained with more computing power from an optimisation with four states and three controls.

The altitude was limited to FL 300 at the first and last steps of the trajectory, up to FL 330 at the second and penultimate steps and up to FL 350 at the third and third to last steps. Between these states, altitudes were allowed to take any value from FL 300 to FL 400, with the altitude giving the lowest fuel burn being chosen. In this way cruise was restricted so that the aircraft would be in the correct positions for the end of the ascent phase and the start of the descent phase. The actual flights did not always follow these restrictions, but on average it allows for a fair comparison.

Results for the routes between LHR and JFK show similar, but slightly higher total season savings than those obtained from limiting the actual flights to FL340. 2.8 million kg

of fuel savings are possible flying westbound and 3.0 million kg eastbound. This would reduce CO_2 emissions by 18 million kg across the 2019-2020 winter season. The results for separate aircraft and airlines follow those shown in Section 7.6.2.

7.7.2 Days where results show an increase in fuel use

On a very small number of days, estimates for the actual fuel use for certain routes is lower than the amount required in the fuel minimal model. These "negative fuel savings" are included in all of the summary statistics calculated and the previous overall results. However, this issue is worthy of further review. In this section we consider the original data in which all flight altitudes are fixed to approximately FL340.

Flying east there are 9 flights where this is observed out of the total of 1547 flights. These flights take place on 7 different days and involve 5 different models of aircraft. All increases in fuel use are very small, the smallest being 0.0001% and the largest still under 0.5% of the estimated actual total fuel used. Looking at the routes chosen by the aircraft, these are very similar to optimised trajectories on these days and the difference is explained by the distance between the ends of cruise and the airports. In each case the first and last positions are those just within 225 km of the airports. However, as the actual data has varying time steps, this means that towards the beginning and end of the cruise the time steps are sometimes less than half of the simulation time step of 125 s. This has the effect that although inside the target circles at either end of a trajectory, actual flights still have further to go in the ascent and descent phases than the fuel minimal trajectories, which will change the overall fuel use. We can conclude from this that flying eastbound the simulated routes are always as fuel efficient, if not more fuel efficient, than those currently flown. Indeed in over 99% of cases the simulated routes lead to obvious fuel savings.

Flying westbound there are more instances where the new routes do not lead to fuel savings. This is unsurprising, given that there will be more of a choice of routes when avoiding strong headwinds, than when trying to make use of favourable tailwinds. The 24 flights (out of 1567) where more fuel is required by the simulations than the estimates for actual flights, are flown by 2 different models of aircraft, the Boeing 777-236 and the Boeing 747-436, on 12 different days. Of these, 16 instances show small fuel differences (less than 0.5%) that can be attributed to the distance from the airports as explained above. This does, however, leave 8 flights (0.51% of the westbound flights). In the case of

these routes the guaranteed minimum fuel use offered by applying dynamic programming is limited by the initial choices of time step and control variable resolution. On certain days, atmospheric conditions may have meant that a less efficient route has been plotted, due to these restrictions. This can be seen clearly in Figure 7.14 which shows plots of routes flown on the 19th February, 2020. Here the simulation takes a route South of the

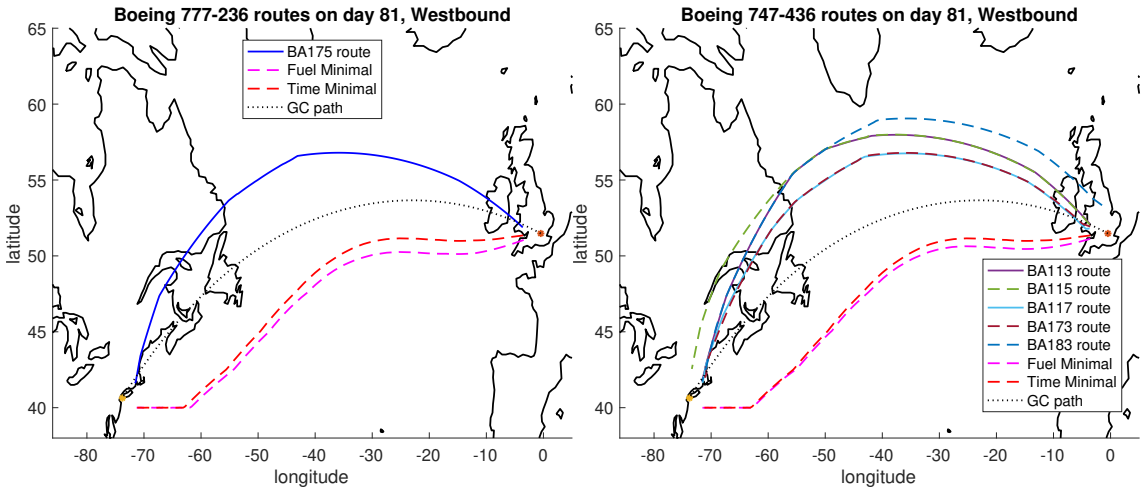

(a) Boeing 777-236 (b) Boeing 747-436

Figure 7.14: Plots showing routes taken by actual flights and the time and fuel minimal routes simulated for the 19th February, 2020.

GC path, whilst the British Airways routes go North. Only flights BA113 and BA115 use significantly less fuel (here taken as 0.5%) than the fuel minimal route, with BA175 using over a tonne of fuel more. So clearly the wind field on the 19th February, 2020 allows for similar fuel burn from two very different paths, an effect noted in Chapter 4. The optimised route flying South of the GC path is more fuel efficient than 88% of the chosen flight paths taken by all of the aircraft making the journey from LHR to JFK through this wind field.

Looking at all cases where more fuel is used by the simulated routes there does not seem to be a particular wind pattern that would warn route planners that this could occur. As can be seen in Figure 7.15 the winds all follow different patterns. On these days the majority of routes which are taken by the actual aircraft are not as fuel efficient as the fuel minimal routes, but one or two use less fuel.

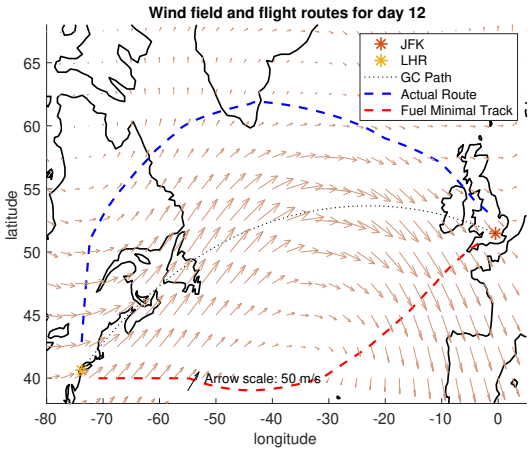

(a) 12th December, 2019

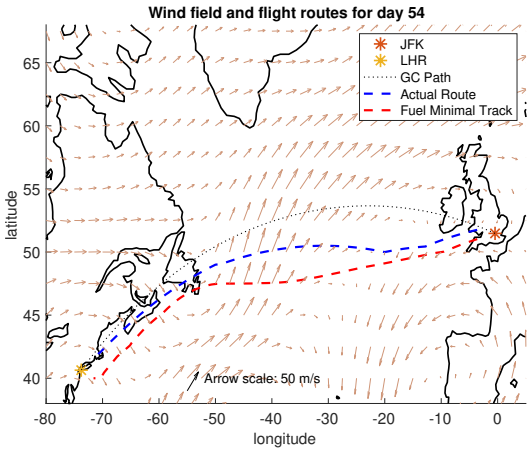

(b) 23rd January, 2020

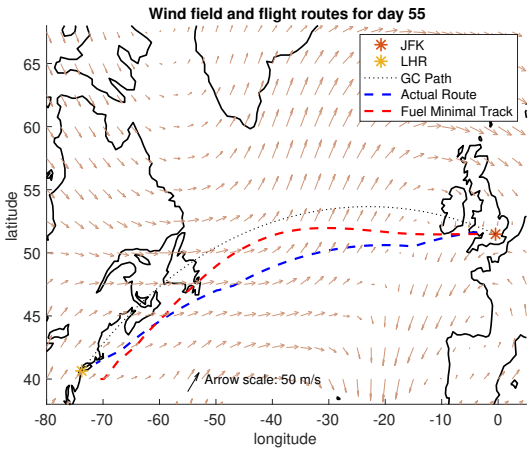

(c) 24th January, 2020

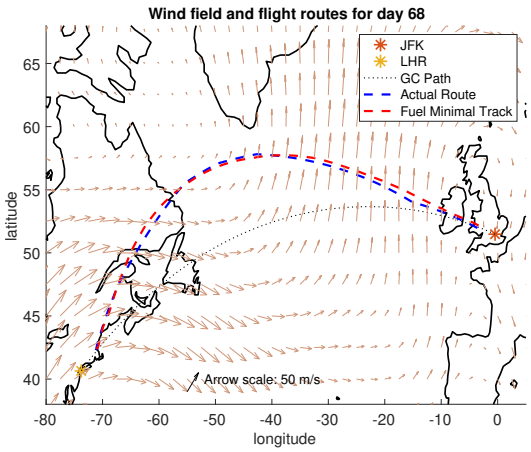

(d) 6th February, 2020

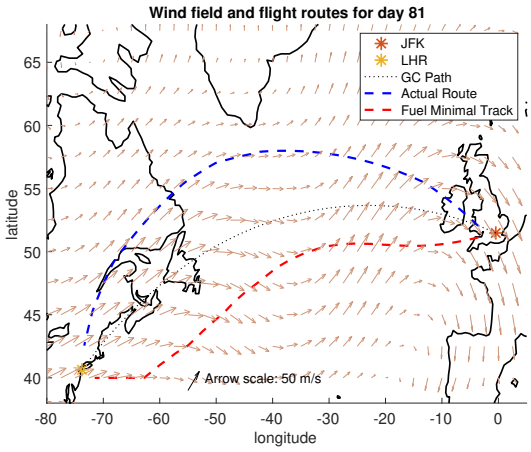

(e) Day 19th February, 2020

Figure 7.15: Plots showing fuel minimal simulated routes and actual routes that did not make a fuel saving. Arrows show the wind field for flights.

7.8 Chapter summary

Dynamic programming is applied to create both time and fuel minimal routes. This involves using a grid of horizontal positions in the case of the time minimal formulation, whilst a third dimension of mass is included for the fuel minimal formulation. The dynamical system for each formulation requires information on atmospheric conditions at each step and these are taken from a reanalysis model. Cost functionals depending on the total journey duration or the total fuel burn can be minimised by varying the heading angle and the airspeed.

A numerical solution to the relevant HJB equations is found by discretising both spatially and temporally. This solution provides a value map across the state space. From here both the optimal feedback control and the optimised trajectory can be retrieved. Only the cruise phase, here defined as from within 225 km of LHR and JFK is considered. In the fuel minimal case the fuel use can be obtained directly from the change in mass state. In the time minimal case the fuel used during each time step is calculated after optimisation.

Fuel use for optimised trajectories is compared with estimated fuel use of actual cruise phase routes, based on recorded data. In calculating this, the same atmospheric data and fuel burn rate formula are applied.

Results show that the fuel use across the cruise phase of flights between LHR and JFK, from 1st December, 2019 to 29th February, 2020 can be reduced. Looking at actual flights as restricted to a single altitude of approximately 34 000 feet, the average fuel reduction flying east is 4.6% and flying west is 3.8%. Overall this amounts to a reduction of 16.6 million kg of CO_2 emissions for this single winter season.

Time minimal flights are shown to have a significantly shorter duration than actual flights and fuel minimal flights, but use more fuel. Comparing fuel minimal trajectories to actual flights flown, the fuel use is significantly reduced, without a significant increase in flight duration. Average airspeeds are significantly lower for fuel optimised trajectories than for actual flights, which lowers the fuel burn rate. However, by choosing optimal heading angles, this does not have a significant impact on the fuel minimal trajectory durations.

Whilst this research shows the value of taking action to plan TBO around a fuel burn equation that takes atmospheric and model specific data into account, it is limited by the resolution of the control variables, as is shown by the days where an improved route is not found.

Increased computational power to simulate flights with an extra state of altitude and a finer resolution of controls would result in fuel savings on every day for all models of aircraft. The compromise of optimising altitude for a set horizontal route, although not accounting for changes to winds at different altitudes, at least shows that it is important to include changes to altitude in this model. This is due to the propulsion efficiency of the aircraft themselves, rather than atmospheric conditions.

Comparing just the cruise phases of routes optimised from the departure airport may also be limiting the optimality of the trajectories, so future work should include the ascent, cruise and descent phases.

Chapter 8

Conclusion

8.1 Answering research questions

In Section 1.1 a number of questions were asked, all of which have been answered to some extent during the course of the last seven chapters. Here we summarise these findings.

8.1.1 Time minimal flights versus the OTS

In the introduction we asked the question:

- How much difference would horizontal time optimal trajectory planning make to fuel use and thus carbon dioxide emissions, in fixed airspeed transatlantic flights, compared with the OTS?

By simplifying the problem to look just at flights between LHR and JFK, at a variety of fixed airspeeds and at a pressure level of 250 hPa (approximately FL340), we are able to show that significant reductions in fuel use are possible, by moving from the OTS to TBO. Using the atmospheric conditions given by re-analysis data the air distance for flights going along the tracks was compared with the air distance given by time minimal trajectories. These were found by applying PMP to the dynamical system describing the aircraft movement zonally and meridionally.

Compared to the worst tracks a maximum reduction of 16.4% of the air distance and thus of the CO_2 emissions is possible. However, this is an extreme result as most aircraft will travel on the more efficient tracks. To gain an estimate of how much CO_2 could actually be prevented by re-routing aircraft, data from the air service navigation provider,

NATS, showing how many aircraft used each track each day was considered. This gave the more realistic result that for a fixed airspeed of 240 m s^{-1} emissions could be reduced by 2.5% eastbound and 1.7% westbound across an entire winter period, amounting to a total of 6.7 million kg of emissions when all flights between LHR and JFK are taken into account.

These results were communicated as part of a paper and directly to NATS through an invited presentation (Wells et al., 2021). Since this research was published NATS have referred to it in justifying their decision to try disbanding the tracks for certain days in 2021 (Young, 2021). Following this the tracks were removed up to FL330 from March 2022, but more changes are likely in the future.

8.1.2 The effect of climate change on trajectory duration

We next asked:

- To what extent is climate change likely to affect minimum times of flight within the next 35 years?

By comparing the results of time minimal trajectories simulated across the North Atlantic between LHR and JFK for two different time periods in both the summer and the winter, we were able to show that there would be no large changes to flight duration in either direction as a result of climate change. Using an ensemble of eight CMIP 5 climate models for a historic time period (1986-2005) and a future time period (2036-2055), it was shown that the round trip flight duration would on average be reduced by between one and two minutes, depending on the season. This could be explained by the expected changes to the core position and intensity of the jet stream predicted by the climate models. This result was contrary to previously published research, based on earlier climate models and a longer interval between time periods, that had suggested a net increase in round trip duration.

So we can conclude that for the chosen ensemble of models and selection of time periods there is very little change to flight duration, but given different models and time periods the result could change. The data used to answer this question came from part of a joint project with Jacob Cheung of the Met Office Aviation Group, in which time minimal results obtained by solving a dynamical system numerically were compared with those found by applying the Met Offices' A Star discrete mathematics method. The findings of

the whole project have been accepted for publication in the Meteorological Applications journal.

8.1.3 Fuel minimal fixed-time flights

- Can fixed-time flights be planned for a whole winter season to ensure that fuel is minimised?

Here the cost functional of the OCP was changed to minimise the fuel burned across a trajectory, using the fuel burn model and aircraft parameters given in Poll & Schumann (2021a,b). This change necessitated a new way to solve the problem. A direct method was adopted, in which the system was first discretised and then optimised. The Matlab fmincon function was encased in a global search mechanism, to hunt for a trajectory giving a global minimum for fuel use in each case.

By solving an OCP based on minimising fuel use for a fixed-time flight it is possible to plan fuel minimal flights for a whole season. However, in restricting these flights to a schedule that allows all atmospheric conditions across the whole of the winter to be taken into account, more fuel is being burned in flight than would be on days when a shorter duration trajectory is possible. If however, there was no contingency for landing at variable times and shorter duration flights would remain in low altitude stacks burning more fuel, then this approach would save fuel.

8.1.4 The benefit of controlling airspeed in addition to heading angle

- In minimising fuel burn of transatlantic flights, can extra benefits result from controlling the airspeed in addition to the heading angle?

In Chapter 6 two different formulations of the OCP were considered, the first where both heading angle and airspeed were controlled and the second where heading angle was the only control variable. In this second formulation airspeed was fixed across each trajectory. Even allowing the fixed airspeed to take the most efficient value for the particular winds encountered each day, using variable airspeeds yielded up to an extra 4% in fuel savings. This showed that in answering the final research question including airspeed as a control variable was necessary.

It was also found that the airspeed chosen for each time step of the simulated fuel minimal route was linked to the wind field. This was clear as there was a significant negative correlation between the average airspeed of each trajectory and the average tailwind around the Great Circle path between LHR and JFK each day.

8.1.5 Comparison of fuel minimal, time minimal and actual flight trajectories

- Would minimising fuel and having a free time in a fixed altitude flight produce trajectories that are more fuel efficient than those currently flown and to what extent would the flight duration be changed?

In solving the time minimal and fuel minimal routing problem for free-time routes at a fixed altitude, dynamic programming was used. This method involves solving the HJB equation to create a map of the value function. In moving to a solution, it is necessary to discretise both temporally and spatially, so we obtain a solution to the discrete problem, via a Semi-Lagrangian method, that as the time step length tends to zero, approximates to the continuous solution. For the time minimal optimisation there were the states of longitude and latitude to consider, but for the fuel minimal optimisation a third state of mass was required. In each case the control variables were the heading angle and airspeed.

Results from both of these optimisations were compared with estimates of fuel use from actual flight routes, calculated using the same atmospheric data and fuel burn function as applied in obtaining the dynamic programming solutions. The actual flight data was given as a timed horizontal route and treated as if flown at the same altitude as the simulations.

The fuel minimal simulated trajectories were shown to use significantly less fuel than the actual routes flown and led to a reduction in carbon dioxide emissions of 16.6 million kg across the winter of 2019 to 2020 for flights between LHR and JFK. The average airspeed of simulated routes was significantly lower than that for the actual routes, but the flight duration was not significantly longer.

This means that it is possible to produce more fuel efficient trajectories than are currently being flown, that do not result in significant time penalties. It would however,

necessitate a more fluid approach to timing landing slots, as each trajectory takes a different amount of time. As these could be planned 1-2 days in advance, due to the slow evolution of the wind field across the North Atlantic, this would give airports time to fine tune a landing schedule.

8.1.6 Evolution of methodology

In answering these questions a range of techniques has been used and at each stage choices for parameters, discretisation of variables and numerical methods have been made. In all cases sensitivity analysis was employed to test for stability, computational time, data storage issues and errors due to truncation of multiple time steps. A TPBVP was also derived based on moving exactly between the two airports using a simplified analytic windfield and an approximate fuel burn model. This confirmed the magnitude of the results, but as co-states had to be found by trial and improvement, incurred its own numerical errors. In each case the methods employed were checked in a wind-free scenario to ensure that the shortest distance route, around the GCR, was selected.

In Chapters 4 and 5 a time minimal route was required, which meant that the best method was the indirect approach, using PMP. This is true as in this case the adjoint variables cancel out from the dynamical system, leaving only an initial control variable to be found (see Section 4.3.2). The flights were assumed to be entirely in cruise phase as the OTS are designed to be most efficient at one particular altitude (Mangini et al., 2018). In this part of the research the route was simulated to arrive within 200 m of the destination airport, as it was being compared with routes along the OTS that were modelled as going directly between LHR and JFK. In Chapter 4 the time step was set to 1 s in finding both the time minimal and OTS route air distances, as this research was a proof of the concept that the OTS tracks were not as efficient as flying time minimal routes rather than a system designed to be used operationally. By setting a very small time step piecewise continuous routes tend to continuous routes. The airspeed, mass and altitude were fixed to allow the air distance to be a suitable comparative metric, as it allowed fuel burn to be directly proportional to air distance. A similar system was used in Chapter 5, but the time intervals were now dictated by the constraints of the A star method as this project was used to compare the two approaches. Computational time (on a Macbook Pro, with 2.3 GHz Dual-Core Intel Core i5 and 8GB memory) for generating the OFW routes was approximately 45 minutes for a single route, varying with the climatic

conditions encountered. In Chapter 4 routes were recorded for 91 days, at 15 different airspeeds in 2 different directions, making a total of over 2700 OFW routes. Although only one Mach number was used for simulations in Chapter 5 over twenty times more simulations were completed as flight paths were found for both winter and summer over 20 years, using 8 different sets of atmospheric projections for both historic and future time periods. This meant that code was run a huge number of times. The largest proportion of computational time was spent on generating routes for the iterations of the bisection method and so using the Euler forward step numerical method kept run-time manageable.

In Chapter 6 the focus was on trying to fix the time for flights between the two airports for a whole winter season and then test if controlling just the heading angle, or controlling both the heading angle and the airspeed was most effective, in terms of minimising fuel burn. Given that this was a straight comparison between two simulations, the choice of parameters was based on practical flight considerations as wells as a compromise between accuracy and computational time and memory constraints. The original function using fmincon was able to include a Runge-Kutta 4th order method with fixed time step of 100 s. As initial control vectors were estimated and then improved in the direct optimisation method, the code ran in 8 minutes. However, when the global function was used in order to provide a more rigorous framework by way of including extra sets of initial controls, the computational time went up to three hours. The direct method was most useful here as the time was fixed and we were minimising fuel burn via a function that could not be differentiated or effectively approximated. Routes were modelled to arrive within 1 km of the destination airport, as runways themselves are often up to 2 km long and the wider target improved the convergence rate of the code.

In Chapter 7 dynamic programming was used, as it was important to find globally optimal routes that could be compared with actual flights. The change in mass could be factored into the mesh, but this did make running times far longer than the dynamic programming model which just optimised flight duration and so did not need the third mesh dimension. The mesh discretisation of 3333 kg was chosen after sensitivity analysis showed that the choice of discretisation for mass made very little difference to the final fuel use. It also allowed the mesh arrays to remain within the preset size constraints of Matlab. Finding value maps took approximately 1 hour using the 3 state variables and 2 control variables, with a 125 s time step, but once a value map was available, retrieving the associated

controls and states took only 2 minutes. Again it was most time efficient to use an Euler method within the algorithm as this numerical calculation had to be applied to every combination of state and control variables. The time step of 125 s was shown to allow a stable and consistent application of the numerical method, whilst not unduly increasing truncation errors. It is also of a practical length for pilots and agreed with the average step length for available flight data. As the main goal of this chapter was to compare with actual flights, the target was set at 225 km from the destination airport to reflect the end of the cruise phase as shown in the actual flight data. Data up to 225 km away from the destination airport was also removed from results to make the start of the cruise phase comparable to recorded data.

Therefore, in each case, it was important to use the most suitable method for solving the optimisation problem and to adjust parameters to reflect the context of the research.

8.2 Future directions

The questions answered by this research show potential for change in transatlantic aircraft routing. However, there are a few limitations of this research that could be addressed in future studies.

In using deterministic atmospheric data, an ability to predict winds and temperatures accurately is assumed. In future work a more probabilistic approach to atmospheric data, by using ensembles to capture forecast uncertainty, would not only enhance the level of route detail possible, but also enable different routes to be planned based on different possible weather scenarios.

In the previous chapters only the main wind effects have been taken into account, with no allowance for localised phenomena such as convective weather patterns, super saturated icy regions and turbulence. As weather forecasts improve, including such information will be very useful. With the advent of machine learning enhanced forecasts and the design of new super computing hardware, which by working to lower numerical precision is able to perform faster calculations, future forecasts will have greater spatial and temporal resolution, both of which are necessary for tactical and strategic aviation route planning

(Dueben, 2019; METOffice, 2020). According to WMO (2021) the future aim for forecasters is to provide updates on the minute scale. This will be possible using high performance computing to include increased observations and incorporate more processes into Earth system models. A full characterisation of the uncertainty of forecast details will also be provided. With this new data, accurate avoidance of turbulence and convection can be factored into route optimisation.

Currently only the route between JFK and LHR has been considered, but with airlines looking to create new longer haul routes, such as Qantas' Project Sunrise flights between Australia and destinations such as New York, Rio de Janeiro, Frankfurt, Capetown, and London it will be essential to fly fuel minimal routes across a variety of different wind fields (Curran, 2021). During this research a side project was undertaken which involved creating time minimal routes across the Pacific Ocean to inform research into the effects ash clouds have on aviation (Capponi et al., 2022). This shows that the theory can be easily adapted to different oceanic regions. In plotting trajectories over land other issues, however, such as airspace borders and orography could prove challenging.

One of the main limitations to the routes generated during this research, is that in each case just one flight is considered at a time. Although the increased situational awareness across the North Atlantic will make scheduling by departure time possible for the nineteen flights a day between JFK and LHR, there are obviously many other airport pairs that would need to be considered in creating a comprehensive route plan for the North Atlantic. If the idea of flocks of aircraft is used, then having one fuel minimal route to be followed by the lead aircraft would be essential, but this is still an idea requiring further research.

Another area that could prove very exciting is the use of artificial intelligence, specifically deep learning, in route planning. Such methods have already been explored by Zhang et al. (2018) in finding the positions of fighter planes and Jiang et al. (2021) in predicting trajectories of aircraft currently in the air, but as yet no attempt has been made to work from atmospheric data and previously planned fuel or time minimal trajectories to create a neural network capable of route planning for future winds and temperatures. A system of this type could be trained offline to ensure almost instant results in the air if conditions changed.

The questions answered in this thesis, aim to show that TBO are a possibility in the immediate future and have the potential to reduce aviation CO_2 emissions substantially. The future adoption of either time minimal or fuel minimal routing will depend on the adaptability of airports and ATM as well as on the number of flights to be scheduled.

Bibliography

Ahlgren, L. (2022). Why are North Atlantic Tracks below FL330 being scrapped? https://simpleflying.com/low-north-atlantic-tracks-scrapped/. Accessed: 2022-02-18.

Ahmed, K., Bousson, K., & Coelho, M. (2021). A modified dynamic programming approach for 4d minimum fuel and emissions trajectory optimization. *Aerospace*, **8**(5).

Airbus (2020). Airbus reveals new zero emission concept aircraft. https://www.airbus.com/en/newsroom/press-releases/2020-09-airbus-reveals-new-zero-emission-concept-aircraft. Accessed: 2021-09-21.

Airbus (2021). Commercial Aircraft. https://www.airbus.com/aircraft.html. Accessed: 2021-03-12.

Airbus (2022). Wake-retrieval to boost environmental performance. https://www.airbus.com/en/innovation/disruptive-concepts/biomimicry/fellofly. Accessed: 2022-03-30.

Aireon (2020). Operations overview. https://aireon.com/resources/overview-materials/operations-overview/. Accessed: 2020-05-23.

Airlines (2019). The History of Flight. http://www.century-of-flight.freeola.com/new_site/commercial/Transatlantic_Services.htm. Accessed: 2020-07-29.

AirlinesUK (2019). Responding to the carbon challenge. https://airlinesuk.org/wp-content/uploads/2017/01/Airlines-UK-Responding-to-the-Carbon-Challenge.pdf. Accessed: 2020-02-12.

Archer, D., Eby, M., Brovkin, V., Ridgwell, A., Cao, L., Mikolajewicz, U., Caldeira, K., Matsumoto, K., Munhoven, G., Montenegro, A., & Tokos, K. (2009). Atmospheric lifetime of fossil fuel carbon dioxide. *Annual Review of Earth and Planetary Sciences*, **37**(1), 117–134.

Arrow, K. (1949). On the use of winds in flight planning. *Journal of Meteorology*, **6**(2), 150–159.

Athans, M. & Falb, P. (1966). *Optimal Control:An Introduction to the Theory and its Applications*. McGraw-Hill Book Company, New York.

Avila, D., Sherry, L., & Thompson, T. (2019). Reducing global warming by airline contrail avoidance: A case study of annual benefits for the contiguous united states. *Transportation Research Interdisciplinary Perspectives*, **2**, 100033.

Bardi, M. & Capuzzo-Dolcetta, I. (1997). *Optimal Control and Viscosity Solutions of Hamilton–Jacobi–Bellman Equations*. Birkhauser, Boston.

Barrett, S., Britter, R., & Waitz, I. (2010). Global mortality attributable to aircraft cruise emissions. *Environmental Science & Technology*, **44**(19), 7736–7742.

Barron, E. & Jensen, R. (1986). The Pontryagin Maximum Principle from Dynamic Programming and Viscosity Solutions to First Order Partial Differential Equations. *Transactions of the American Mathematical Society*, **298**(2), 635–641.

Bellman, R. (1957). *Dynamic Programming*. Princeton University Press.

Bertsekas, D. (1999). *Nonlinear Programming*. Athena Scientific.

Boeing (2021). Boeing Commercial Airplanes. `https://www.boeing.com/commercial/`. Accessed: 2021-03-12.

Boucher, O., Randall, D., Artaxo, P., Bretherton, C., Feingold, G., Forster, P., Kerminen, V.-M., Kondo, Y., Liao, H., Lohmann, U., Rasch, P., Satheesh, S. K., Sherwood, S., Stevens, B., & Zhang, X. Y. (2013). *Clouds and aerosols*, pages 571–657. Cambridge University Press, Cambridge, UK.

Bressan, A. & Piccoli, B. (2007). *Introduction to the mathematical theory of control*. American Institute of Mathematical Sciences, Springfield.

Briani, A. (2021). An introduction to optimal control theory and hamilton-jacobi equations. http://www.lmpt.univ-tours.fr/~briani/AppuntiCorsoBriani.pdf. Accessed: 2021-10-23.

Bryson, A. & Ho, Y.-C. (1975). *Applied Optimal Control*. Routledge.

Buis, A. (2019). The atmosphere: Getting a handle on carbon dioxide. https://climate.nasa.gov/news/2915/the-atmosphere-getting-a-handle-on-carbon-dioxide/#:~:text=The_concentration_of_carbon_dioxide,it_was_near_20370_ppm. Accessed: 2021-09-23.

Burrows, J. (1983). Fuel-optimal aircraft trajectories with fixed arrival times. *Journal of Guidance, Control, and Dynamics*, **6**(1), 14–19.

Byrd, R., Hribar, M., & Nocedal, J. (1999). An Interior Point Algorithm for Large-Scale Nonlinear Programming. *SIAM Journal on Optimization*, **9**(4), 877–900.

Byrd, R., Gilbert, J., & Nocedal, J. (2000). A trust region method based on interior point techniques for nonlinear programming. *Mathematical Programming*, **89**(1), 149–185.

CAPA (2018). Record global aircraft deliveries in 2017: Boeing ahead of Airbus again, but behind on order backlog. https://centreforaviation.com/analysis/reports/record-global-aircraft-deliveries-in-2017-393914. Accessed: 2019-11-12.

Capponi, A., Harvey, N., Dacre, H., Beven, K., Saint, C., Wells, C., & James, M. (2022). Refining an ensemble of volcanic ash forecasts using satellite retrievals: Raikoke 2019. *Atmospheric Chemistry and Physics*, **22**(9), 6115–6134.

Cartis, C., Massart, E., & Otemissov, A. (2021). Global optimization using random embeddings. *arXiv*. Accessed: 2022-03-10.

Chakravarty, A. (1985). Four-dimensional fuel-optimal guidance in the presence of winds. *Journal of Guidance, Control, and Dynamics*, **8**(1), 16–22.

Chen, Y. (2020). COVID-19 Pandemic Imperils Weather Forecast. *Geophysical Research Letters*, **47**(15).

Cheung, J. (2018). Flight planning: node-based trajectory prediction and turbulence avoidance. *Meteorological Applications*, **25**(1), 78–85.

CitytoSea (2022). Ethical travel. https://www.citytosea.org.uk/campaign/plastic-free-travel/ethical-travel/. Accessed: 2022-03-30.

Coffel, E. & Horton, R. (2015). Climate change and the impact of extreme temperatures on aviation. *Weather, Climate, and Society*, **7**(1), 94–102.

Colson, B. & Toint, P. (2001). Exploiting Band Structure in Unconstrained Optimization Without Derivatives. *Optimization and Engineering*, **2**(4), 399–412.

Cristiani, E. & Martinon, P. (2009). Initialization of the Shooting Method via the Hamilton-Jacobi-Bellman Approach. https://arxiv.org/abs/0910.0521. Accessed: 2021-09-30.

Cummins, N. (2019). The evolution of transatlantic flight. https://simpleflying.com/transatlantic-flight-evolution/. Accessed: 2020-07-29.

Curran, A. (2021). Qantas eyes greenlighting project sunrise flights within 12 months. https://simpleflying.com/qantas-project-sunrise-capa/. Accessed: 2021-09-23.

Dalmau, R., Prats, X., Ramonjoan, A., & Soley, S. (2020). Estimating fuel consumption from radar tracks. https://upcommons.upc.edu/bitstream/handle/2117/183452/Pre2020_Dalmau_CANJ_Fuel-estimation_pre-print.pdf. Accessed: 2020-09-03.

Dancila, B., Beulze, B., & Botez, R. (2019). Flight phase and altitude-dependent geometrical vertical flight plan optimization minimizing the total number of vertical plan segments. *Proceedings of the Institution of Mechanical Engineers, Part G: Journal of Aerospace Engineering*, **233**(13), 4825–4838.

Dhief, I. (2018). *Optimization of aircraft trajectories over the North Atlantic Airspace*. Ph.D. thesis, Université Paul Sabatier.

Dobrokhodov, V., Walton, C., Kaminer, I., & Jones, K. (2020). Energy-Optimal Trajectory Planning of Hybrid Ultra-Long Endurance UAV in Time-Varying Energy Fields. https://calhoun.nps.edu/handle/10945/64615. Accessed: 2022-04-17.

Doctor, D., Budd, T., Williams, P., Prescott, M., & Iqbal, R. (2022). Modelling the effect of electric aircraft on airport operations and infrastructure. *Technological Forecasting and Social Change*, **177**, 121553.

Dueben, P. (2019). Will machine learning replace conventional weather prediction models. `https://www.ecmwf.int/en/about/media-centre/science-blog/2019/will-machine-learning-replace-conventional-weather`. Accessed: 2020-06-23.

Dunstone, N., Smith, D., Scaife, A., Hermanson, L., Eade, R., Robinson, N., Andrews, M., & Knight, J. (2016). Skilful predictions of the winter north atlantic oscillation one year ahead. *Nature Geoscience*, **9**.

Economon, T., Copeland, S., Alonso, J., Zeinali, M., & Rutherford, D. (2011). Design and optimization of future aircraft for assessing the fuel burn trends of commercial aviation.

English, S., Shapland, M., & Armitage, J. (2020). Business leaders call for air bridge between london and new york. Evening Standard, 21/9/20.

Erzberger, H. & Lee, H. (1980). Constrained optimum trajectories with specified range. *Journal of Guidance and Control*, **3**(1), 78–85.

ETS, E. (2021). Reducing emissions from aviation. `https://ec.europa.eu/clima/eu-action/transport-emissions/reducing-emissions-aviation\_en`. Accessed: 2021-07-26.

Eurocontrol (2021a). Base of aircraft data. `https://www.eurocontrol.int/model/bada`. Accessed: 2021-06-13.

Eurocontrol (2021b). Climate Change Risks for European Aviation 2021 Annex 5 Impact of changes in wind patterns on flight operations. `https://www.eurocontrol.int/sites/default/files/2021-09/eurocontrol-study-climate-change-risk-european-aviation-annex-5.pdf`. Accessed: 2021-09-28.

Fahroo, F. & Ross, I. (2002). Direct trajectory optimization by a chebyshev pseudospectral method. *Journal of Guidance, Control, and Dynamics*, **25**(1), 160–166.

Falcone, M. (1997). *Optimal Control and Viscosity Solutions of Hamilton–Jacobi–Bellman Equations*, chapter Appendix A: Numerical Solution of Dynamic Programming Equations, pages 471–504. Birkhauser, Boston.

Falcone, M. & Ferretti, R. (2014). *Semi-Lagrangian Approximation Schemes for Linear and Hamilton–Jacobi Equations*. SIAM, Philadelphia.

Fan, Y., Yang, L., Li, Q., Nong, C., Zheng, Z., & Xue, F. (2020). Cost index-based cruise flight trajectory optimization. In *2020 IEEE 4th Information Technology, Networking, Electronic and Automation Control Conference (ITNEC)*, volume 1, pages 103–110.

Flightradar24, A. (2020). Flight data. `https://www.flightradar24.com/data/flights`. Accessed: 2022-08-12.

Franco, A. & Rivas, D. (2011). Minimum-cost cruise at constant altitude of commercial aircraft including wind effects. *Journal of Guidance, Control, and Dynamics*, **34**(4), 1253–1260.

Franco, A. & Rivas, D. (2013). Analysis of optimal aircraft cruise with fixed arrival time including wind effects. *Aerospace Science and Technology*, **32**.

Freedman, A. (2020). The coronavirus pandemic and loss of aircraft data are taking a toll on weather forecasting. `https://www.washingtonpost.com/weather/2020/05/12/weather-forecasting-coronavirus-flights/`. Accessed: 2020-05-20.

García-Heras, J., Soler, M., & Sáez, F. (2014). A Comparison of Optimal Control Methods for Minimum Fuel Cruise at Constant Altitude and Course with Fixed Arrival Time. *Procedia Engineering*, **80**, 231–244. 3rd International Symposium on Aircraft Airworthiness (ISAA 2013).

Gardi, A., Sabatini, R., & Ramasamy, S. (2016). Multi-objective optimisation of aircraft flight trajectories in the ATM and avionics context. *Progress in Aerospace Sciences*, **83**(0376-0421), 1–36.

Girardet, B., Lapasset, L., Delahaye, D., & Rabut, C. (2014). Wind-optimal path planning: Application to aircraft trajectories. In *2014 13th International Conference on Control Automation Robotics Vision (ICARCV)*.

Godbold, J. & Calosi, P. (2013). Ocean acidification and climate change: advances in ecology and evolution. *Philos Trans R Soc Lond B Biol Sci*, **368**.

Gratton, G., Padhra, A., Rapsomanikis, S., & Williams, P. (2020). The impacts of climate change on Greek airports. *Climatic Change*, **160**(2), 219–231.

Graver, B. & Rutherford, D. (2018). Transatlantic airline fuel efficiency ranking, 2017. `https://theicct.org/publications/transatlantic-airline-fuel-efficiency-ranking-2017`. Accessed: 2019-10-23.

Graver, B., Zhang, K., & Rutherford, D. (2019). CO_2 emissions from commercial aviation, 2018. `https://theicct.org/sites/default/files/publications/ICCT\_CO2-commercl-aviation-2018\_20190918.pdf`. Accessed: 2019-10-14.

Green, J. (2009). The potential for reducing the impact of aviation on climate. *Techn. Analysis Strat. Manag.*, **21**, 39–59.

Grewe, V., Matthes, S., Frömming, C., Brinkop, S., Jöckel, P., Gierens, K., Champougny, T., Fuglestvedt, J., Haslerud, A., Irvine, E., & Shine, K. (2017). Feasibility of climate-optimized air traffic routing for trans-atlantic flights. *Environmental Research Letters*, **12**(3), 034003.

Grewe, V., Matthes, S., & Dahlmann, K. (2019). The contribution of aviation NO x emissions to climate change: are we ignoring methodological flaws? *Environmental Research Letters*, **14**(12), 121003.

Grobler, C., Wolfe, P., Dasadhikari, K., Dedoussi, I., Allroggen, F., Speth, R., Eastham, S., Agarwal, A., Staples, M., Sabnis, J., & Barrett, S. (2019). Marginal climate and air quality costs of aviation emissions. *Environmental Research Letters*, **14**(11), 114031.

Grooters, F. (2008). Aircraft observations. `https://public.wmo.int/en/bulletin/aircraft-observations`. Accessed: 2019-12-04.

Grönstedt, T. (2021). Opportunities for heat management in hydrogen fuelled aircraft.

Harvey, B., Cook, P., Shaffrey, L., & Schiemann, R. (2020). The Response of the Northern Hemisphere Storm Tracks and Jet Streams to Climate Change in the CMIP3, CMIP5, and CMIP6 Climate Models. *Journal of Geophysical Research: Atmospheres*, **125**(23), e2020JD032701.

Hasan, A., Mamun, A., Rahman, S., Malik, K., Al Amran, I., Khondaker, A., Reshi, O., Tiwari, S., & Alismail, F. (2021). Climate change mitigation pathways for the aviation sector. *Sustainability*, **13**(7).

Hashemi, S., Botez, R., & Grigorie, T. (2020). New reliability studies of data-driven aircraft trajectory prediction. *Aerospace*, **7**(10), 145.

HeartAerospace (2022). Heart aerospace will deliver the first es-19 electric airliner certified for commercial flight by 2026. `https://heartaerospace.com`. Accessed: 2022-03-30.

Henderson, R., Martins, J., & Perez, R. (2012). Aircraft conceptual design for optimal environmental performance. *The Aeronautical Journal*, **116**(1175), 1–22.

Henninger, H., von Ellenrieder, K., & Licht, S. (2020). Energy-minimal target retrieval for quadrotor UAVs: trajectory generation and tracking. In *2020 28th Mediterranean Conference on Control and Automation (MED)*, pages 727–732.

Hersbach, H., Bell, B., Berrisford, P., Hirahara, S., Horányi, A., Muñoz-Sabater, J., Nicolas, J., Peubey, C., Radu, R., Schepers, D., Simmons, A., Soci, C., Abdalla, S., Abellan, X., Balsamo, G., Bechtold, P., Biavati, G., Bidlot, J., Bonavita, M., De Chiara, G., Dahlgren, P., Dee, D., Diamantakis, M., Dragani, R., Flemming, J., Forbes, R., Fuentes, M., Geer, A., Haimberger, L., Healy, S., Hogan, R., Hólm, E., Janisková, M., Keeley, S., Laloyaux, P., Lopez, P., Lupu, C., Radnoti, G., de Rosnay, P., Rozum, I., Vamborg, F., Villaume, S., & Thépaut, J. (2020). The era5 global reanalysis. *Quarterly Journal of the Royal Meteorological Society*, **146**(730), 1999–2049.

Ho-Huu, V., Hartjes, S., Visser, H., & Curran, R. (2019). An optimization framework for route design and allocation of aircraft to multiple departure routes. *ArXiv*, **abs/1908.11086**.

ICAO (2016). On Board:A Sustainable Future 2016 Environmental Report. `https://www.icao.int/environmental-protection/Pages/env2016.aspx`. Accessed: 2019-10-30.

ICAO (2019). Destination green: The next chapter. `https://www.icao.int/environmental-protection/Documents/`

EnvironmentalReports/2019/ENVReport2019_pg111-115.pdf. Accessed: 2020-07-15.

ICAO (2020). Carbon emissions calculator. https://www.icao.int/environmental-protection/Carbonoffset/Pages/default.aspx. Accessed: 2020-07-15.

IEA (2022). Sustainable development scenario. https://www.iea.org/reports/world-energy-model/sustainable-development-scenario-sds. Accessed: 2022-03-30.

IPCC (2019). IPCC special report on the ocean and cryosphere in a changing climate. In press.

Irvine, E., Hoskins, B., Shine, K., Lunnon, R., & Froemming, C. (2013). Characterizing North Atlantic weather patterns for climate-optimal aircraft routing. *Meteorol Appl.*, **20**, 80–93.

Irvine, E. A., Shine, K. P., & Stringer, M. A. (2016). What are the implications of climate change for trans-atlantic aircraft routing and flight time? *Transportation Research Part D: Transport and Environment*, **47**, 44–53.

Jackson, A., Newton, B., Hahn, D., & Bussey, A. (2001). Statistical contrail forecasting. *Journal of Applied Meteorology*, **40**(2), 269–279.

Jacobson, M., Wilkerson, J., Naiman, A., & Lele, S. (2013). The effects of aircraft on climate and pollution. part ii: 20-year impacts of exhaust from all commercial aircraft worldwide treated individually at the subgrid scale. *Faraday discussions*, **165**, 369–82.

James, E., Benjamin, S., & Jamison, B. (2020). Commercial-Aircraft-Based Observations for NWP: Global Coverage, Data Impacts, and COVID-19. *Journal of Applied Meteorology and Climatology*, **59**(11), 1809 – 1825.

Jardin, M. & Bryson, A. (2012a). Methods for computing minimum-time paths in strong winds. *Journal of Guidance, Control, and Dynamics*, **35**(1), 165–171.

Jardin, M. & Bryson, A. (2012b). Neighboring optimal aircraft guidance in winds. *18th Applied Aerodynamics Conference*.

Jensen, L., Tran, H., & Hansman, R. (2015). Cruise fuel reduction potential from altitude and speed optimization in global airline operations. *USA/Europe Air Traffic Management Research and Development(ATM) Seminar*, **141**(688), 764–773.

Jiang, S.-Y., Luo, X., & He, L. (2021). Research on Method of Trajectory Prediction in Aircraft Flight Based on Aircraft Performance and Historical Track Data . *Mathematical Problems in Engineering*, **2021**(1), 668213.

John, B. (2020). Stranded on the tarmac? Here's what you need to know. `https://www.airhelp.com/en-gb/blog/tarmac-delay-us/`,. Accessed: 2020-07-13.

Kalnay, E., Kanamitsu, M., Kistler, R., Collins, W., Deaven, D., Gandin, L., Iredell, M., Saha, S., White, G., Woollen, J., Zhu, Y., Chelliah, M., Ebisuzaki, W., Higgins, W., Janowiak, J., Mo, K., Ropelewski, C., Wang, J., Leetmaa, A., Reynolds, R., Jenne, R., & Joseph, D. (1996). The NCEP/NCAR 40-year reanalysis project. *Bulletin of the American Meteorological Society*, **77**(3), 437–472.

Karnauskas, K., Donnelly, J., Barkley, H., & Martin, J. (2015). Coupling between air travel and climate. *Nature Climate Change*, **5**(12), 1068–1073.

Kent, T. & Richards, A. (2015). Analytic approach to optimal routing for commercial formation flight. *Journal of Guidance, Control, and Dynamics*, **38**(10), 1872–1884.

Khardi, S. (2012). Aircraft flight path optimization the hamilton-jacobi-bellman considerations. *Applied Mathematical Sciences*, **6**(25), 1221–1249.

Khodayari, A., Olsen, S., & Wuebbles, D. (2014). Evaluation of aviation nox-induced radiative forcings for 2005 and 2050. *Atmospheric Environment*, **91**, 95–103.

Kim, H., Webster, P., & Curry, J. (2012). Seasonal prediction skill of ecmwf system 4 and ncep cfsv2 retrospective forecast for the northern hemisphere winter. *Clim Dyn*, **39**, 2957–2973.

Kim, J.-H., Chan, W., Sridhar, B., & Sharman, R. (2015). Combined Winds and Turbulence Prediction System for Automated Air-Traffic Management Applications. *Appl. Meteor. Climatol*, **54**, 766–784.

Kim, J.-H., Chan, W., Sridhar, B., Sharman, R., Williams, P., & Strahan, M. (2016). Impact of the North Atlantic Oscillation on Transatlantic Flight Routes and Clear-Air Turbulence. *Journal of Applied Meteorology and Climatology*, **55**(3), 763–771.

Kim, J.-H., Kim, D., Lee, D., Chun, H., Sharman, R., Williams, P., & Kim, Y. (2020). Impact of climate variabilities on trans-oceanic flight times and emissions during strong nao and enso phases. *Environmental Research Letters*, **15**(10), 105017.

Kirk, D. (1970). *Optimal Control Theory: An Introduction*. Dover Publications Ltd, New York.

Krozel, J., Klimenko, V., & Sharman, R. (2011). Analysis of clear-air turbulence avoidance maneuvers. *Air Traffic Control Quarterly*, **19**, 147–168.

Kärcher, B. (2018). Formation and radiative forcing of contrail cirrus. *Nat Commun*, **9**(1824).

Lane, M. (2001). Trouble in the air. `http://news.bbc.co.uk/1/hi/in\_depth/uk/2001\\trouble\_in\_the\_air/1449129.stm`. Accessed: 2019-09-28.

Lee, D., Fahey, D., Skowron, A., Allen, M., Burkhardt, U., Chen, Q., Doherty, S., Freeman, S., Forster, P., Fuglestvedt, J.and Gettelman, A., De León, R., Lim, L., Lund, M., Millar, R., Owen, B., Penner, J., Pitari, G., Prather, M., Sausen, R., & Wilcox, L. (2020). The contribution of global aviation to anthropogenic climate forcing for 2000 to 2018. *Atmospheric Environment*, (1352-2310), 117834.

Lee, E. & Markus, L. (1967). *Foundations of Optimal Control Theory*. John Wiley and Sons Ltd, New York.

Lee, H., Olsen, S., Wuebbles, D., & Youn, D. (2013). Impact of aircraft emissions on the air quality near the ground. *Atmos. Chem. Phys.*, **13**, 5505–5522.

Lee, S., Williams, P., & Frame, T. (2019). Increased shear in the North Atlantic upper-level jet stream over the past four decades. *Nature*, **572**(7771), 639–642.

Lester, P. (1994). *Turbulence: A new perspective for pilots*. Jeppesen Sanderson Training Systems.

Levi-Civita, T. (1931). Über Zermelo's Luftfahrtproblem. *ZAMM - Journal of Applied Mathematics and Mechanics / Zeitschrift für Angewandte Mathematik und Mechanik*, **11**(4), 314–322.

Lin, Z., Duan, J., Li, S., Ma, H., & Yin, Y. (2018). ICAO Long-Term Traffic Forecasts: Passengers and Freight. `https://www.icao.int/sustainability/documents/ltf\_charts-results\_2018edition.pdf`. Accessed: 2020-02-16.

Lin, Z., Duan, J., Li, S., Ma, H., & Yin, Y. (2020). Continuous-time finite-horizon adp for automated vehicle controller design with high efficiency. https://arxiv.org/abs/2007.02070. Accessed: 2022-03-24.

Linde, P. (2020). Multi-functional materials towards environmentally friendly aviation.

Lindsey, R. (2021). Climate change: Atmospheric carbon dioxide. https://www.climate.gov/news-features/understanding-climate/climate-change-atmospheric-carbon-dioxide. Accessed: 2021-09-23.

Lovegren, J. & Hansman, R. (2011). Estimation of potential aircraft fuel burn reduction in cruise via speed and altitude optimization strategies.

Lowe, J., Murphy, J., Harris, G., Sexton, D., Kendon, E., Bett, P., Clark, R., Eagle, K., Fosser, G., Fung, F., McDonald, R., McInnes, R., McSweeney, C., Mitchell, J., Rostron, J., Thornton, H., Tucker, S., & Yamazaki, K. (2018). UKCP18 Land Projections: Science Report. https://www.metoffice.gov.uk/pub/data/weather/uk/ukcp18/science-reports/UKCP18-Land-report.pdf. Accessed: 2019-09-25.

Lunnon, R. & Mirza, A. (2007). Benefits of promulgating higher resolution wind data for airline route planning. *Meteorological Applications*, **14**, 253–261.

Macki, J. & Strauss, A. (1982). *Introduction to Optimal Control Theory*. Springer Verlag, New York.

Mangini, F., Irvine, E., Shine, K., & Stringer, M. (2018). The dependence of minimum-time routes over the North Atlantic on cruise altitude. *Meteorol Appl.*, **25**, 655–664.

Manzini, E., Karpechko, A., & Kornblueh, L. (2018). Nonlinear response of the stratosphere and the north atlantic-european climate to global warming. *Geophysical Research Letters*, **45**(9), 4255–4263.

Matthes, S., Lührs, B., Dahlmann, K., Grewe, V., Linke, F., Yin, F., Klingaman, E., & Shine, K. (2020). Climate-Optimized Trajectories and Robust Mitigation Potential: Flying ATM4E. *Aerospace*, **7**(11).

Matthes, S., Lim, L., Burkhardt, U., Dahlmann, K., Dietmüller, S., Grewe, V., Haslerud, A., Hendricks, J., Owen, B., Pitari, G., Righi, M., & Skowron, A. (2021). Mitigation of Non-CO2 Aviation's Climate Impact by Changing Cruise Altitudes. *Aerospace*, **8**(2).

Mazareanu, E. (2020). Uk airlines: Seat occupancy on international and domestic flights 2002-2018. https://www.statista.com/statistics/304056/international-and-domestic-flight-seat-occupancy-in-the-uk/. Accessed: 2020-09-04.

McSweeney, C. & Bett, P. (2020). Ukcp european circulation indices: Jet stream position and strength. UKCPFactsheet,MetOffice. Accessed: 2020-09-27.

Menon, P. (1989). Study of aircraft cruise. *Journal of Guidance, Control, and Dynamics*, **12**(5), 631–639.

METOffice (2020). Weather and climate science services in a changing world. https://www.metoffice.gov.uk/research/approach/research-and-innovation-strategy. Accessed: 2020-04-29.

Molloy, J. (2020). GHG report: Basis of reporting for NATS' airspace, energy and environmental performance data 2019-20. https://www.nats.aero/wp-content/uploads/2020/11/GHG-19-20.pdf. Accessed: 2021-01-17.

Monbiot, G. (2007). *Heat: How We Can Stop the Planet Burning*. Allen Lane.

Moninger, W., Mamroch, R., & Pauley, P. (2003). Automated Meteorological Reports from Commercial Aircraft. *Bulletin of American Meteorological Society*, **84**, 203 – 216.

Mouillet, V., Nuic, A., Casado, E., & Leones, J. (2019). Evaluation of the applicability of a modern aircraft performance model to trajectory optimisation. https://www.eurocontrol.int/sites/default/files/2019-03/applicability-modern-apm-to-traj-opt.pdf. Accessed: 2020-10-08.

Murrieta Mendoza, A., Hamy, A., & Botez, R. (2016). Lateral reference trajectory algorithm using ant colony optimization.

Murrieta Mendoza, A., Romain, C., & Botez, R. (2020). 3D Cruise Trajectory Optimization Inspired by a Shortest Path Algorithm. *Aerospace*, **7**, 99.

Mårtensson, H. (2021). Harmonic forcing from distortion in a boundary layer ingesting fan. *Aerospace*, **8**(3).

Naresh Kumar, G., Ikram, M., Sarkar, A., & Talole, S. (2018). Hypersonic flight vehicle trajectory optimization using pattern search algorithm. *Optimization and Engineering*, **19**(1), 125–161.

NATS (2007). Use of data link for operational improvements. `https://www.icao.int/EURNAT/OtherMeetingsSeminarsandWorkshops-/NATPBCSWORKSHOP/1stNATPERFORMANCEBASEDCOMMUNICATIONS-ANDSURVEILLANCEWORKSHOP/PPT07-FANSPRESENTATIONFORRCPWORKSHOP.pdf`. Accessed: 2021-04-19.

NATS (2019). Aireon system goes live–trial operations begin over the North Atlantic marking new chapter in aviation history. `https://www.nats.aero/news/aireon-system-goes-live/`. Accessed: 2020-02-15.

NCAR (2022). Cmip analysis platform. `https://www2.cisl.ucar.edu/computing-data/data/cmip-analysis-platform`. Accessed: 2022-03-20.

Ng, H., Sridhar, B., Grabbe, S., & Chen, N. (2011). Cross-polar aircraft trajectory optimization and the potential climate impact. In *2011 IEEE/AIAA 30th Digital Avionics Systems Conference*, pages 1–19.

Ng, H., Sridhar, B., & Grabbe, S. (2014). Optimizing aircraft trajectories with multiple cruise altitudes in the presence of winds. *Journal of Aerospace Information Systems*, **11**(1), 35–47.

Nguyen, N. (2006). Singular arc time-optimal climb trajectory of aircraft in a two-dimensional wind field. volume 7.

Németh, H., Švec, M., & Kandráč, P. (2018). The Influence of Global Climate Change on the European Aviation. *International Journal on Engineering Applications (IREA)*, **6**(6).

OAG (2020). Busiest routes 2019. `https://www.oag.com/reports/busiest-routes-2019`. Accessed: 2020-09-21.

OAG (2022). What is on time performance? `https://www.oag.com/on-time-performance-airlines-airports`. Accessed: 2022-06-03.

Oakley, M. (2019). A century of transatlantic flying. https://www.aerosociety.com/news/a-century-of-transatlantic-flying/. Accessed: 2020-07-29.

Oudar, T., Cattiaux, J., & Douville, H. (2020). Drivers of the northern extratropical eddy-driven jet change in cmip5 and cmip6 models. *Geophysical Research Letters*, **47**(8), e2019GL086695.

Overton, J. (2022). The growth of greenhouse gas emissions from commercial aviation. https://www.eesi.org/papers/view/fact-sheet-the-growth-in-greenhouse-gas-emissions-from-commercial-aviation. Accessed: 2022-07-12.

Partridge, J. (2021). New budget airline will fly from london to new york by 2022. https://www.theguardian.com/business/2021/aug/10/new-budget-norwegian-airline-will-fly-from-london-to-new-york-by-2022. Accessed: 2021-09-23.

Parzani, C. & Puechmorel, S. (2017). On a Hamilton-Jacobi-Bellman approach for coordinated optimal aircraft trajectories planning. *Optimal Control Applications and Methods*, **39**(2), 933–948.

PCMDI (2007). CMIP3 Climate Model Documentation, References, and Links. https://pcmdi.llnl.gov/ipcc/model_documentation/ipcc_model_documentation.html. Accessed: 2021-06-19.

PCMDI (2018). CMIP5 - Coupled Model Intercomparison Project Phase 5 - Overview. https://web.archive.org/web/20180719234013/https://cmip.llnl.gov/cmip5/index.html. Accessed: 2021-06-19.

Peterson, E., Neels, K., Barczi, N., & Graham, T. (2013). The economic cost of airline flight delay. *Journal of Transport Economics and Policy*, **47**(1), 107–121.

Pierson, B. & Ong, S. (1989). Minimum-fuel aircraft transition trajectories. *Mathematical and Computer Modelling*, **12**(8), 925–934.

Poll, D. (2011). A first order method for the determination of the leading mass characteristics of civil transport aircraft. *The Aeronautical Journal (1968)*, **115**(1167), 257–272.

Poll, D. (2018). On the relationship between non-optimum operations and fuel requirement for large civil transport aircraft, with reference to environmental impact and contrail avoidance strategy. *The Aeronautical Journal*, **122**(1258), 1827–1870.

Poll, D. & Schumann, U. (2021a). An estimation method for the fuel burn and other performance characteristics of civil transport aircraft during cruise. Part 2: determining the aircraft's characteristic parameters. *The Aeronautical Journal*, **125**(1284), 296–340.

Poll, D. & Schumann, U. (2021b). An estimation method for the fuel burn and other performance characteristics of civil transport aircraft in the cruise. Part 1: fundamental quantities and governing relations for a general atmosphere. *The Aeronautical Journal*, **125**(1284), 257–295.

Pontryagin, L., Boltyanskii, V.and Gamkrelize, R., & Mishchenko, E. (1962). *The Mathematical Theory of Optimal Processes*. Interscience Publishers John Wiley Sons, Inc., New York.

Poret, M., O'Connell, J., & Warnock-Smith, D. (2015). The economic viability of long-haul low cost operations: Evidence from the transatlantic market. *Journal of Air Transport Management*, **42**, 272–281.

Pouzolz, R., Schmitz, O., & Klingels, H. (2021). Evaluation of the climate impact reduction potential of the water-enhanced turbofan (wet) concept. *Aerospace*, **8**(3).

Quadros, F., Snellen, M., & Dedoussi, I. (2020). Regional sensitivities of air quality and human health impacts to aviation emissions. *Environmental Research Letters*, **15**(10), 105013.

Reed, D. (2018). New York - London Is the World's First Billion-Dollar Airline Route. `https://www.forbes.com/sites/danielreed/2018/07/09/for-real-new-york-london-is-the-worlds-first-billion-dollar-route-for-british-airways` Accessed: 2019-10-03.

Ren, D. & Leslie, L. (2019). Impacts of climate warming on aviation fuel consumption. *Journal of Applied Meteorology and Climatology*, **58**(7), 1593–1602.

Rompokos, P., Rolt, A., Nalianda, D., Isikveren, A., Senné, C., Grönstedt, T., & Abedi, H. (2020). Synergistic technology combinations for future commercial aircraft using liquid hydrogen. volume 3.

Rosenow, J., Fricke, H., Luchkova, T., & Schultz, M. (2018). Minimizing contrail formation by rerouting around dynamic ice-supersaturated regions. *Aeronautics and Aerospace Open Access Journal*, **2**(3), 105–111.

Ross, I. (2015). *A Primer on Pontryagin's Principle in Optimal Control, Second edition.* Collegiate Publishers, San Francisco.

Ross, I. & Karpenko, M. (2012). A review of pseudospectral optimal control: From theory to flight. *Annual Reviews in Control*, **36**(2), 182–197.

Ruban, D.and Yashalova, N. (2020). Ethics prescribed by top world airlines: Empirical evidence. *Climate*, **8**(10), 119.

Rumpfkeil, M. & Zingg, D. (2010). The optimal control of unsteady flows with a discrete adjoint method. *Optimization and Engineering*, **11**(1), 5–22.

Ruosteenoja, K., Vihma, T., & Venäläinen, A. (2019). Projected changes in european and north atlantic seasonal wind climate derived from cmip5 simulations. *Journal of Climate*, **32**(19), 6467 – 6490.

Ryley, T., Baumeister, S., & Coulter, L. (2020). Climate change influences on aviation: A literature review. *Transport Policy*, **92**, 55 – 64.

Scaife, A., Arribas, A., Blockley, E., Brookshaw, A., Clark, R., Dunstone, N., Eade, R., Fereday, D., Folland, C., Gordon, M., Hermanson, L., Knight, J., Lea, D., MacLachlan, C., Maidens, A., Martin, M., Peterson, A., Smith, D., Vellinga, M., Wallace, E., Waters, J., & Williams, A. (2014). Skillful long-range prediction of european and north american winters. *Geophysical Research Letters*, **41**(7), 2514–2519.

Schennings, A., Larsson, J., & Markus, R. (2019). Regional sensitivities of air quality and human health impacts to aviation emissions. *Sustainable Environment Research*, **29**(1), 25.

Schultz, R. (1974). Fuel optimality of cruise. *Journal of Aircraft*, **11**(9), 586–587.

Schultz, R. & Zagalsky, N. (1972). Aircraft performance optimization. *Journal of Aircraft*, **9**(2), 108–114.

Schumann, U. (1997). The impact of nitrogen oxides emissions from aircraft upon the atmosphere at flight altitudes—results from the aeronox project. *Atmospheric Environment*, **31**(12), 1723–1733.

Schwab, A., Thomas, A., Bennett, J., Robertson, E., & Cary, S. (2021). Electrification of aircraft: Challenges, barriers, and potential impacts. `https://www.nrel.gov/docs/fy22osti/80220.pdf`. Accessed: 2021-10-22.

Seitz, A. (2012). Advanced methods for propulsion system integration in aircraft conceptual design.

SESAR (2022). Delivering the digital european sky. `https://www.sesarju.eu`. Accessed: 2022-04-11.

Sethi, V., Sun, X., Nalianda, D., Rolt, A., Holborn, P., Wijesinghe, C., Xisto, C., Jonsson, I., Gronstedt, T., Ingram, J., Lundbladh, A., Isikveren, A., Williamson, I., Harrison, T., & Yenokyan, A. (2022). Enabling Cryogenic Hydrogen-Based CO2-Free Air Transport: Meeting the demands of zero carbon aviation. *IEEE Electrification Magazine*, **10**(2), 69–81.

Simpson, I. (2016). Climate change predicted to lengthen transatlantic travel times. *Environmental Research Letters*, **11**(3), 031002.

Singh, V. & Sharma, S. (2015). Fuel consumption optimization in air transport: a review, classification, critique, simple meta-analysis, and future research implications. *European Transport Research Review*, **7**(2), 12.

Soler, M., González-Arribas, D., Sanjurjo-Rivo, M., García-Heras, J., Sacher, D., Gelhardt, U., Lang, J., Hauf, T., & Simarro, J. (2020). Influence of atmospheric uncertainty, convective indicators, and cost-index on the leveled aircraft trajectory optimization problem. *Transportation Research Part C: Emerging Technologies*, **120**, 102784.

Soomer, M. & Franx, G. (2008). Scheduling aircraft landings using airlines' preferences. *European Journal of Operational Research*, **190**(1), 277–291.

Sorensen, J. & Waters, M. (1981). Airborne method to minimize fuel with fixed time-of-arrival constraints. *Journal of Guidance and Control*, **4**(3), 348–349.

Speyer, J. (1973). On the fuel optimality of cruise. *Journal of Aircraft*, **10**(12), 763–765.

Sridhar, B., Ng, H., & Chen, N. (2011). Aircraft Trajectory Optimization and Contrails Avoidance in the Presence of Winds . *J. Guidance, Control, Dyn.*, **34**, 1577–1584.

Sridhar, B., Chen, N., Ng, H., Rodionova, O., Delahaye, D., & Linke, F. (2015). Strategic Planning of Efficient Oceanic Flights.

Steiner, M., Bateman, R., Megenhardt, D., Liu, Y., Xu, M., Pocernich, M., & Krozel, J. (2010). Translation of ensemble weather forecasts into probabilistic air traffic capacity impact. *Air Traffic Control Quarterly*, **18**, 229 – 254.

Storer, L., Williams, P., & Joshi, M. (2017). Global response of clear-air turbulence to climate change. *Geophysical Research Letters*, **44**(19), 9976–9984.

Student, B. (1908). The Probable Error of a Mean. *Biometrika*, **6**(1), 1–25.

Teoh, R., Schumann, U., Majumdar, A., & Stettler, M. (2020). Mitigating the Climate Forcing of Aircraft Contrails by Small-Scale Diversions and Technology Adoption. *Environmental science technology*, **54**(5), 2941–2950.

Timperley, J. (2019). Corsia: The UN's plan to "offset" growth in aviation emissions after 2020. `https://www.carbonbrief.org/corsia-un-plan-to-offset-growth-in-aviation-emissions-after-2020`. Accessed: 2019-02-08.

TSCM (2020). Mach number and airspeed versus altitude. `http://www.tscm.com/mach-as.pdf`. Accessed: 2022-12-10.

Ugray, Z., Lasdon, L., Plummer, J., Glover, F., Kelly, J., & Martí, R. (2007). Scatter Search and Local NLP Solvers: A Multistart Framework for Global Optimization. *INFORMS Journal on Computing*, **19**(3), 328–340.

UKCOP26 (2021). International aviation climate ambition coalition: Cop26 declaration. `https://ukcop26.org/cop-26-declaration-international-aviation-climate-ambition-coalition/`. Accessed: 2021-12-09.

Uppink, L. (2021). The future of flying is closer than ever. `https://www.weforum.org/agenda/2021/07/the-future-of-flying-is-closer-than-ever-sustainable-fuel-is-the-key`. Accessed: 2021-07-18.

Veness, C. (2019). Calculate distance, bearing and more between latitude/longitude points. `https://www.movable-type.co.uk/scripts/latlong.html`. Accessed: 2019-10-27.

Vennam, L., Vizuete, W., Talgo, K., Omary, M., Binkowski, F., Xing, J., & Arunachalam, S. (2017). Modeled full-flight aircraft emissions impacts on air quality and their sensitivity to grid resolution. *Journal of Geophysical Research: Atmospheres*, **122**, 13472–13494.

Verhagen, C., Visser, H., & Santos, B. (2018). A decentralized approach to formation flight routing of long haul commercial flights. *Journal of Guidance, Control, and Dynamics*, **233**(8), 2992–3004.

Villarroel, J. & Rodrigues, L. (2016). Optimal control framework for cruise economy mode of Flight Management Systems. *Journal of Guidance, Control and Dynamics*, **39**(5), 1022–1033.

Waltz, R., Morales, J., Nocedal, J., & Orban, D. (2006). An interior algorithm for nonlinear optimization that combines line search and trust region steps. *Mathematical Programming*, **107**(3), 391–408.

Warntz, W. (1961). Transatlantic flights and pressure patterns. *Geographical Review*, **51**(2), 187–212.

Wells, C., Williams, P., Nichols, N., Kalise, D., & Poll, D. (2021). Reducing transatlantic flight emissions by fuel-optimised routing. *Environmental Research Letters*, **16**(2), 025002.

Wells, C., Kalise, D., Nichols, N., Poll, D., & Williams, P. (2022). The role of airspeed variability in fixed-time, fuel-optimal aircraft trajectory planning. *Optimization and Engineering*. https://doi.org/10.1007/s11081-022-09720-9.

Wickramasinghe, N., Harada, A., & Miyazawa, Y. (2012). Flight trajectory optimisation for an efficient transportation system. In *28th Congress of the International Council of the Aeronautical Sciences 2012, ICAS 2012*, volume 6, pages 4399–4410.

Williams, P. (2016). Transatlantic flight times and climate change. *Environmental Research Letters*, **11**(2), 024008.

Williams, P. (2017). Increased light, moderate, and severe clear-air turbulence in response to climate change. *Advances in Atmospheric Sciences*, **34**(5), 576–586.

Williams, P. & Joshi, M. (2013). Intensification of winter transatlantic aviation turbulence in response to climate change. *Nature Climate Change*, **3**(7), 644–648.

WMO (2021). Future of weather and climate forecasting. `https://public.wmo.int/en/resources/library/wmo-open-consultative-platform-white-paper-1-future-of-weather-and-climate`. Accessed: 2021-09-23.

Woollings, T. & Blackburn, M. (2012). The north atlantic jet stream under climate change and its relation to the nao and ea patterns. *J. Climate*, **25**, 886–902.

Xing, J., Mathur, R., Pleim, J., Hogrefe1, C., Wang, J., Gan, C., Sarwar, G., Wong, D., & McKeen, S. (2016). Representing the effects of stratosphere–troposphere exchange on 3-d o3 distributions in chemistry transport models using a potential vorticity-based parameterization. *Atmos. Chem. Phys.*, **16**(1), 10865–10877.

Yamashita, H., Grewe, V., Jöckel, P., Linke, F., Schaefer, M., & Sasaki, D. (2015). Towards climate optimised flight trajectories in a climate model: Airtraf. In *Proceedings of the 11th USA/Europe Air Traffic Management Research and Development Seminar, ATM 2015*. Eurocontrol. 11th USA/Europe Air Traffic Management Research and Development Seminar, ATM 2015 ; Conference date: 23-06-2015 Through 26-06-2015.

Yamashita, H., Yin, F., Volker, G., Jöckel, P., Matthes, S., Kern, B., Dahlmann, K., & Frömming, C. (2020). Newly developed aircraft routing options for air traffic simulation in the chemistry–climate model EMAC 2.53: AirTraf 2.0. *Geoscientific Model Development*.

Yamashita, H., Yin, F., Grewe, V., Jöckel, P., Matthes, S., Kern, B., Dahlmann, K., & Frömming, C. (2021). Analysis of Aircraft Routing Strategies for North Atlantic Flights by Using AirTraf 2.0. *Aerospace*, **8**(2).

Young, J. (2021). Is it time to disband the organised track structure? `https://nats.aero/blog/2021/02/is-it-time-to-disband-the-organised-track-structure/`. Accessed: 2021-02-22.

Zermelo, E. (1930). Über die Navigationin der Luft als Problem der Variationsrechnung. *Jahresbericht der Deutschen Mathematiker-Vereinigung*, **39**, 44–48.

Zermelo, E. (1931). Über das Navigationsproblem bei ruhender oder veränderlicher Windverteilung. *Zeitschrift für Angewandte Mathematik und Mechanik*, **11**(2), 114–124.

Zhang, Z., Yang, R., & Fang, Y. (2018). Lstm network based on on antlion optimization and its application in flight trajectory prediction. In *2018 2nd IEEE Advanced Information Management,Communicates,Electronic and Automation Control Conference (IMCEC)*, pages 1658–1662.

Appendix A

Deriving the Zermelo equations for planar flight

The minimal time problem across a plane, as described by Zermelo in the 1930s (Zermelo, 1930, 1931), can be formulated in terms of optimal control theory (Bryson & Ho, 1975). The Pontryagin Hamiltonian is written to describe the relationship between the state and adjoint equations. The adjoint equations can be derived from this Hamiltonian and they form part of the necessary conditions for a weak relative extremum. The Hamiltonian is minimised with respect to the control variable, according to Pontryagin's Minimum Principle (Pontryagin et al., 1962).

To find the minimum time path through a region of position dependent vector velocity we first need to define parameters, variables, boundary conditions, a dynamical system and a cost functional.

Parameters:

These set out information about the initial and final values of time:

$$t_0 = 0,$$
$$t_f \geq 0.$$

State Variables:

Here the way in which the zonal (eastwards) and meridional (northwards) positions of the aircraft vary with time are considered:

$$\mathbf{x} = \begin{bmatrix} x(t) \\ y(t) \end{bmatrix}. \tag{A.1}$$

Boundary Conditions:

The fixed values for the state variables are given as the airports at either end of the flight.

$$x(0) = x_{dept},$$
$$y(0) = y_{dept},$$
$$x(t_f) = x_{dest},$$
$$y(t_f) = y_{dest},$$

where (x_{dept}, y_{dept}) and (x_{dest}, y_{dest}) are the co-ordinates of the departure and destination airports respectively.

Control Variable:

The heading angle is called the control variable, as by changing this value optimally, time can be minimised:

$$\alpha(t) = \theta.$$

Dynamical System to be controlled:

The dynamical system is based on the motion of an aircraft in the x (zonal) and y (meridional) directions, where V is the aeroplane's constant velocity, u is the speed of the wind in the x direction, v is the speed of the wind in the y direction and θ is the heading angle, measured from due east anticlockwise.

From here the state equations are drawn, which are the airspeed components in each direction added to the corresponding wind speeds:

$$\dot{x} = V\cos\theta + u(x, y), \tag{A.2}$$
$$\dot{y} = V\sin\theta + v(x, y). \tag{A.3}$$

Cost/Payoff functional:

We define optimality as a minimisation of time, giving the cost functional to minimise as:

$$J(\mathbf{x}, \alpha) = \int_{t_0}^{t_f} L(\mathbf{x}, \alpha)\mathrm{d}t = \int_{t_0}^{t_f} 1\mathrm{d}t = t_f, \tag{A.4}$$

where L is the Langrangian or running cost, $\mathbf{x}$ is a vector including both x and y positions, t_0 is initial time and t_f is final time. In this case:

$$L(\mathbf{x}, \alpha) = 1. \tag{A.5}$$

In order to minimise time, two adjoint variables arise from the necessary conditions for the solution to the problem, p_x and p_y. A Hamiltonian relationship exists between the state and adjoint equations, which can be written as:

$$H = p_x(V\cos\theta + u) + p_y(V\sin\theta + v) + 1. \tag{A.6}$$

The partial derivatives of H with respect to each state variable can be found, giving the adjoint equations. We also find the rate of change of the Hamiltonian with respect to the control variable:

$$\dot{p}_x = -\frac{\partial H}{\partial x} = -p_x \frac{\partial u}{\partial x} - p_y \frac{\partial v}{\partial x}, \tag{A.7}$$

$$\dot{p}_y = -\frac{\partial H}{\partial y} = -p_x \frac{\partial u}{\partial y} - p_y \frac{\partial v}{\partial y}, \tag{A.8}$$

$$0 = \frac{\partial H}{\partial \theta} = V(-p_x \sin\theta + p_y \cos\theta). \tag{A.9}$$

Equation (A.9) shows that, for a non-zero V value:

$$\tan\theta = \frac{p_y}{p_x} \implies p_x = \frac{p_y}{\tan\theta} \implies p_y = p_x \tan\theta. \tag{A.10}$$

As H is not explicitly a function of time, it is a constant of integration and can be replaced by zero in order to minimise the cost functional. Setting $H = 0$ and substituting Equation (A.10) into Equation (A.6) we obtain:

$$0 = \frac{p_y}{\tan\theta}(V\cos\theta + u) + p_y(V\sin\theta + v) + 1. \tag{A.11}$$

Equation (A.11) can be rearranged to make p_y the subject of the formula, giving:

$$p_y = \frac{-\sin\theta}{V + u\cos\theta + v\sin\theta}. \tag{A.12}$$

Following the same method, but finding p_x in terms of V, u, v and θ this time gives:

$$p_x = \frac{-\cos\theta}{V + u\cos\theta + v\sin\theta}. \tag{A.13}$$

These can now be substituted back into Equations (A.7) and (A.8) giving:

$$\dot{p}_x = -\frac{\partial H}{\partial x} = \frac{\cos\theta}{V + u\cos\theta + v\sin\theta}\frac{\partial u}{\partial x} + \frac{\sin\theta}{V + u\cos\theta + v\sin\theta}\frac{\partial v}{\partial x}, \tag{A.14}$$

$$\dot{p}_y = -\frac{\partial H}{\partial y} = \frac{\cos\theta}{V + u\cos\theta + v\sin\theta}\frac{\partial u}{\partial y} + \frac{\sin\theta}{V + u\cos\theta + v\sin\theta}\frac{\partial v}{\partial y}. \tag{A.15}$$

The aim here is to find the rate of change of θ so that a path can be plotted, minimising time. In order to do this, we apply the chain rule:

$$\frac{\mathrm{d}\theta}{\mathrm{d}t} = \frac{\mathrm{d}\theta}{\partial p_x} \times \frac{\partial p_x}{\mathrm{d}t}. \tag{A.16}$$

From Equation (A.10) it is clear that:

$$\theta = \arctan\frac{p_y}{p_x}, \tag{A.17}$$

and differentiating with respect to p_x, via the quotient rule gives:

$$\frac{\mathrm{d}\theta}{\partial p_x} = \frac{p_x\dot{p}_y - p_y\dot{p}_x}{p_x^2 + p_y^2} \times \frac{1}{\frac{\partial p_x}{\mathrm{d}t}}, \tag{A.18}$$

$$\frac{\mathrm{d}\theta}{\mathrm{d}t} = \frac{p_x(-p_x\frac{\partial u}{\partial y} - p_y\frac{\partial v}{\partial y}) - p_y(-p_x\frac{\partial u}{\partial x} - p_y\frac{\partial v}{\partial x})}{p_x^2 + p_y^2}. \tag{A.19}$$

To complete the derivation, Equations (A.12) and (A.13) are substituted into Equation (A.19), thus giving:

$$\frac{\mathrm{d}\theta}{\mathrm{d}t} = -\cos^2\theta\frac{\partial u}{\partial y} + \sin^2\theta\frac{\partial v}{\partial x} + \sin\theta\cos\theta(\frac{\partial u}{\partial x} - \frac{\partial v}{\partial y}). \tag{A.20}$$

This completes the working, resulting in the dynamical system shown in Equations (2.1)-(2.3):

$$\frac{\mathrm{d}x}{\mathrm{d}t} = u + V\cos\theta, \tag{A.21}$$

$$\frac{\mathrm{d}y}{\mathrm{d}t} = v + V\sin\theta, \tag{A.22}$$

$$\frac{\mathrm{d}\theta}{\mathrm{d}t} = \sin^2\theta\frac{\partial v}{\partial x} - \cos^2\theta\frac{\partial u}{\partial y} + (\frac{\partial u}{\partial x} - \frac{\partial v}{\partial y})\sin\theta\cos\theta. \tag{A.23}$$

Appendix B

Conformal mapping of the Zermelo equations onto a sphere

A conformal mapping is one which preserves the angles between curves. Latitude and longitude require different factors to move from a planar to a spherical system. In the original system the rate of change of distance in the zonal and meridional directions is used, but in the spherical system the distances are translated into rates of change of longitude and latitude.

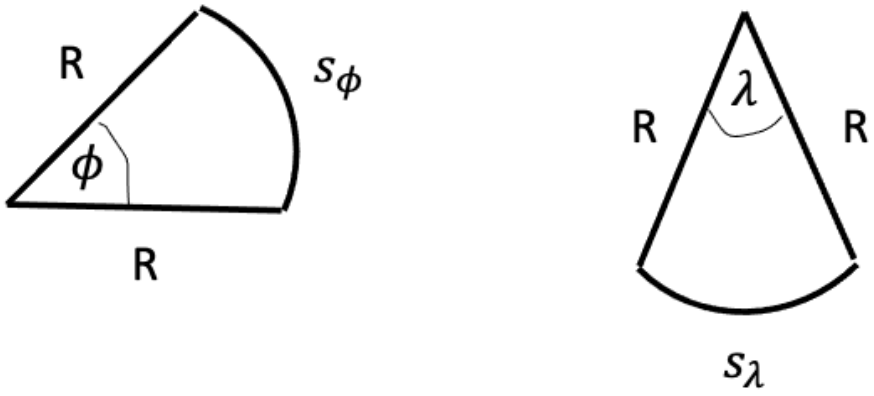

Figure B.1: Sectors showing how the planar distances (s_ϕ and s_λ) relate to the Earth's radius (R) and the change in latitude (ϕ) and longitude (λ).

In Figure B.1 it can be seen that the conformal mapping uses the rule that for an angle in radians, the length of an arc, is found by multiplying the radius by the angle in the sector formed by the arc and two radii.

The distance covered as one radian of latitude is swept through, will be the same at any

point on the sphere, so we can say that:

$$s_\phi = R\phi. \tag{B.1}$$

Rearranging Equation (B.1) to make ϕ the subject of the equation and differentiating with respect to s_ϕ gives:

$$\frac{d\phi}{ds_\phi} = \frac{1}{R}. \tag{B.2}$$

The rate of change of meridional distance with time is:

$$\frac{ds_\phi}{dt} = V\sin\theta + v. \tag{B.3}$$

Using the chain rule:

$$\frac{d\phi}{dt} = \frac{d\phi}{ds_\phi} \times \frac{ds_\phi}{dt} = \frac{V\sin\theta + v}{R}, \tag{B.4}$$

which shows the rate of change of latitude with time.

However, the distance covered as one radian of longitude is swept through, will depend on the latitude of the horizontal sector we are considering. A change in zonal distance corresponding to a change in longitude, is defined to be this change in zonal distance at the equator multiplied by $\cos\phi$. This gives:

$$s_\lambda = R\lambda\cos\phi. \tag{B.5}$$

Rearranging Equation (B.5) to make λ the subject of the equation and differentiating with respect to s_λ gives:

$$\frac{d\lambda}{ds_\lambda} = \frac{1}{R\cos\phi}. \tag{B.6}$$

The rate of change of meridional distance with time is:

$$\frac{ds_\lambda}{dt} = V\cos\theta + u. \tag{B.7}$$

Using the chain rule:

$$\frac{d\lambda}{dt} = \frac{d\lambda}{ds_\lambda} \times \frac{ds_\lambda}{dt} = \frac{V\cos\theta + u}{R\cos\phi}, \tag{B.8}$$

which shows the rate of change of longitude with time. Equations (B.4) and (B.8) are the first two lines of the dynamical system used in Chapters 4-7 of this thesis.

Appendix C

Aircraft model specific parameters

Table C.1: Table showing aircraft model specific parameters as given in Poll & Schumann (2021b). MTOM is maximum take-off mass, OEM is operational empty mass, MZFM is maximum permitted zero fuel mass and MMO is maximum operational Mach number. Other parameter values are discussed in Section 3.4.

Aircraft	Code	MTOM	MZFM	OEM	MMO	τ	$\psi 1$	$\psi 2$	$\psi 3$	$\psi 4$	$\psi 5$	$\psi 6$
777-236(ER)	B772	286900	195000	137050	0.89	0.123	0.211	8.09	0.614	0.811	1.27×10^8	0.632
777-323(ER)	B77W	351530	237683	167829	0.89	0.143	0.219	8.25	0.59	0.811	1.27×10^8	0.774
747-436	B744	396894	246074	178756	0.92	0.107	0.193	7.84	0.621	0.83	1.47×10^8	0.652
767-432(ER)	B764	204116	149685	103872	0.86	0.146	0.182	8.12	0.566	0.772	9.81×10^7	0.748
A330-223	A332	233000	170000	124500	0.86	0.138	0.206	8.17	0.63	0.786	1.13×10^8	0.645
A330-323	A333	233000	175000	127000	0.86	0.142	0.194	8.18	0.612	0.786	1.13×10^8	0.645
A350-1041	A35K	311000	223000	157000	0.89	0.134	0.244	8.09	0.625	0.82	1.31×10^8	0.569
787-9 Dreamliner	B789	254011	181450	128850	0.9	0.149	0.233	8.13	0.595	0.815	1.17×10^8	0.657
A330-343	A333	233000	175000	127000	0.86	0.142	0.194	8.18	0.612	0.786	1.13×10^8	0.645
A340-642	A346	368000	245000	181606	0.86	0.136	0.208	8.26	0.583	0.796	1.26×10^8	0.822