

Minimising emissions from flights through realistic wind fields with varying aircraft weights

Article

Published Version

Creative Commons: Attribution 4.0 (CC-BY)

Open Access

Wells, C. A. ORCID: <https://orcid.org/0000-0001-9438-4954>,
Williams, P. D. ORCID: <https://orcid.org/0000-0002-9713-9820>,
Nichols, N. K. ORCID: <https://orcid.org/0000-0003-1133-5220>,
Kalise, D. and Poll, I. (2023) Minimising emissions from flights through realistic wind fields with varying aircraft weights. Transportation Research Part D: Transport and Environment, 117. 103660. ISSN 1361-9209 doi: <https://doi.org/10.1016/j.trd.2023.103660> Available at <https://centaur.reading.ac.uk/110952/>

It is advisable to refer to the publisher's version if you intend to cite from the work. See [Guidance on citing](#).

To link to this article DOI: <http://dx.doi.org/10.1016/j.trd.2023.103660>

Publisher: Elsevier


All outputs in CentAUR are protected by Intellectual Property Rights law, including copyright law. Copyright and IPR is retained by the creators or other copyright holders. Terms and conditions for use of this material are defined in the [End User Agreement](#).

www.reading.ac.uk/centaur

CentAUR

Central Archive at the University of Reading

Reading's research outputs online



Minimising emissions from flights through realistic wind fields with varying aircraft weights

Cathie A. Wells^{a,*}, Paul D. Williams^a, Nancy K. Nichols^{a,b}, Dante Kalise^c, Ian Poll^d

^a University of Reading, UK

^b National Centre for Earth Observation, UK

^c Imperial College London, UK

^d Poll AeroSciences Ltd, UK

ARTICLE INFO

Keywords:

Dynamic programming
Transatlantic aviation routing
Fuel-minimal
Aircraft
Emissions reduction

ABSTRACT

The international aviation community has agreed to advance actions to reduce CO₂ emissions. Adopting more fuel-efficient routes will achieve this goal quickly and economically. Full satellite coverage of transatlantic flight routes is now a reality, allowing us to consider moving from the Organised Track Structure to Trajectory-Based Operations. Here, fuel-optimal trajectories through wind fields from a global atmospheric re-analysis dataset are found using dynamic programming. The control variables of aircraft headings and airspeeds are varied to find free-time, fuel-minimal routes. Aircraft fuel consumption is modelled with a new model-specific fuel-burn function, which incorporates aircraft mass reductions as fuel is burned. From 1 December 2019 to 29 February 2020, fuel use from simulated routes is compared with fuel estimates based on recorded flight data. Results demonstrate that an average fuel reduction of 4.2% is possible without significant changes to flight duration. This equates to a reduction of 16.6 million kg of CO₂ emissions. Therefore, free-time, fuel-minimal routes have the potential to offer substantial fuel and emissions savings.

1. Introduction

Approximately 2.4% of all anthropogenic CO₂ is traceable back to aviation (Graver et al., 2019; Grewe et al., 2019; Lee et al., 2020) and by 2050, emissions are forecast to be three times those recorded in 2015 (ETS, 2021). This means urgent action needs to be taken to reduce aviation emissions. 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). Sustainable Aviation Fuel (SAF) is able to provide an 80% reduction in CO₂ emissions, but the costly nature of the fuel has currently limited it to only 0.05% of the EU market (Uppink, 2021).

An immediate answer to limit the impact of aviation emissions 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 CO₂ emissions (Aireon, 2020; Girardet et al., 2014). This is in contrast to the Organised Track Structure (OTS) across the

* Correspondence to: Department of Mathematics and Statistics, Whiteknights, PO Box 220, Reading RG6 6AX, UK.

E-mail address: cathie.wells@reading.ac.uk (C.A. Wells).

North Atlantic. This daily set of routes, issued by Air Navigation Service Providers, has been necessary for the last 60 years due to limited radar signals over this stretch of water. This paper considers individual flights, as conflict avoidance can be factored in using take-off times and altitude differences once narrower safe flight proximity is fully exploited.

Here we plan trajectories across the North Atlantic, between John F Kennedy Airport in New York (JFK: 40.6° N, 73.8° W) and London Heathrow Airport (LHR: 51.5° N, 0.5° W), considering the problem through the framing of Optimal Control Theory. Dynamic programming is used to ensure a globally optimal solution. By solving the Hamilton–Jacobi–Bellman (HJB) equation, a value function is obtained which allows every point 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 both time-minimal trajectories and fuel-minimal, free-time trajectories can be found.

Trajectory optimisation began with Zermelo discussing a 2D time-minimal flight path in his 1930 lecture (Zermelo, 1930). Since then research has continued to be undertaken by aeronautical engineers, mathematicians, meteorologists and computer scientists. Their interest in optimal paths has included finding time-minimal routes (Kim et al., 2015; Jardin and Bryson, 2012; Nguyen, 2006; Wells et al., 2021), fuel-minimal routes (Wickramasinghe et al., 2012; García-Heras et al., 2014; Rodionova et al., 2014; Ahmed et al., 2021; Wells et al., 2022) and sometimes combining these to optimise for a cost index (Sorensen and Waters, 1981; Burrows, 1983; Villarroel and Rodrigues, 2016; Fan et al., 2020). Here we compare the amount of fuel used for both time-minimal and fuel-minimal trajectories with actual flown routes.

Previously trajectories have often been optimised without accounting for the wind field (Speyer, 1973; Schultz, 1974; Menon, 1989; García-Heras et al., 2014; Ho-Huu et al., 2019; Ahmed et al., 2021) or been found for simplified wind fields to allow for easier mathematical analysis (Franco and Rivas, 2011, 2013; Nguyen, 2006). Instead we use re-analysis 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 enables a useful comparison to be drawn between actual flight data and results from modelled, optimised flights.

The choice of control and state variables also varies widely, as allowing too many states and controls would lead to experiments too computationally demanding to solve. Where research is based on routes that can visit a large number of given grid points along a journey, often the airspeed or Mach number will be kept constant throughout and just the heading angle will be altered to search for a minimum time or fuel route (Cheung, 2018; Murrieta Mendoza et al., 2020; Girardet et al., 2014; Rodionova et al., 2014, 2016). Other research fixes the entire horizontal trajectory of a flight in order to optimise for airspeed (Franco and Rivas, 2011; Villarroel and Rodrigues, 2016; Fan et al., 2020). By using dynamic programming for a fixed-altitude flight, we are able to allow the trajectory to access all points within a horizontal box covering the North Atlantic, whilst finding both the optimal airspeed and heading angle for each time step. Mass is also considered as a state. In this way we account for the change in the fuel burn rate due to the aircraft's mass change as fuel is burned.

Where real wind data are included, previous research has tended to limit flight simulations to a selection of a few particular days (Sridhar et al., 2011; Girardet et al., 2014; Rodionova et al., 2014; Soler et al., 2020; Yamashita et al., 2020). In contrast, here we use data describing the atmospheric conditions for all 91 days of an entire winter season from 1st December, 2019 to 29th February, 2020. Over this time period we compare simulated fuel- and time-minimal trajectories with more than 3000 actual flight routes, flown by ten different models of aircraft, by four different airlines.

Recent research by Yamashita et al. (2020, 2021) has considered “climate-optimal routes” featuring contrail (short for condensation trail) avoidance in order to address some of the less well understood climate warming effects caused by aviation. However, these lead to greater fuel use and therefore, more emissions (Grewe et al., 2019; Rosenow et al., 2018). In addition contrail avoidance requires hourly, fine-resolution forecasts (Jackson et al., 2001; Rosenow et al., 2018) which are not yet readily available. Teoh et al. (2020) also found that just 2.2% of flights that they considered were responsible for 80% of the contrails observed. We have chosen to focus on reducing CO₂ rather than contrails, as the effect this gas has on radiative forcing is known with less than 1% uncertainty (Lee et al., 2020).

An 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 and Schumann, 2021a,b). This aims to give a more transparent method for calculating a value for fuel consumption than has previously been possible when relying on Eurocontrol's Base of Airline Data (BADA), which is a black box method based on proprietary data (Eurocontrol, 2021; Harada et al., 2013). Although overall BADA data was found to have a 5% error over a flight, which is the same accuracy level as the Poll–Schumann method, BADA was found by Mouillet et al. (2019) to underestimate fuel use in some circumstances (Poll and Schumann, 2021b; Harada et al., 2013). It is, however, considered adequate for trajectory research, if not for financially motivated studies (Mouillet et al., 2019). In our research the fuel-burn rate at any point in space can be estimated by considering atmospheric data for this location as well as the aircraft's instantaneous altitude, airspeed and mass. The function also uses aircraft model-specific parameters to ensure a fair comparison between actual flight data and the simulations.

Previous research comparing time-minimal trajectories with flights taken along the OTS showed that there is the potential to reduce air distance flown (the distance an aircraft flies relative to the air around it) which at constant airspeed and altitude leads to reduced emissions (Wells et al., 2021). It has also been shown that varying the airspeed along a flight route when optimising fuel use makes best use of the wind field. This allows additional savings to be made compared to optimisation dependent on heading angle only (Wells et al., 2022). Here these two ideas are extended so that both time-minimal and fuel-minimal cruise-phase routes are considered for fixed-altitude flights with varying heading angle and airspeed. Using a free-time formulation when minimising fuel use makes the results more easily comparable with actual flight data.

This paper is set out in five sections. In Section 2, the different datasets are described. Section 3 describes the use of HJB for both time- and fuel-minimal optimisation. Results showing potential fuel savings across the 91 day trial period are set out in Section 4 and a discussion of the results is provided. Finally, the results are summarised and future research directions are suggested in Section 5.

2. Data sources

To find time- and fuel-optimal routes atmospheric data is needed, as well as information about fuel burn for specific aircraft models. In order to compare the fuel use of actual flights to those simulated, data detailing flight routes is also required.

2.1. Atmospheric data

All wind field and temperature data are obtained from the National Center for Atmospheric Research re-analysis data set (Kalnay et al., 1996). This combines a large number of observations gathered from a variety of sources with a geophysical model, to give a best estimate of the state of the atmosphere at any time in the past. Average daily horizontal wind velocities and temperatures are given every 2.5 degrees of both latitude and longitude, at a range of pressure levels. However, we consider only the cruise phase of the flights, as this covers 92% of the ground distance (Wells et al., 2021) and fix the altitude. This approach is taken to reduce the dimension of the state space to avoid the “curse of dimensionality” (Bellman, 1957) and spiralling computation costs. This means we only use data from the 250 hPa pressure surface, which corresponds approximately to an altitude of 34 000 feet. Simulations have been fixed at this altitude as the jet stream is on average strongest at this level (Mangini et al., 2018). The choice to vary airspeed rather than altitude along a trajectory is further justified by Lovegren and Hansman (2011), with fixed horizontal flight path trajectories showing 2.4% fuel savings possible from airspeed optimisation compared to 1.5% for altitude optimisation. Lunnnon and Mirza (2007) demonstrated that long-haul journeys are largely insensitive to the resolution of atmospheric data, so the 2.5° grid is sufficient, with linear interpolation applied to obtain meridional (northward) winds, zonal (eastward) winds and temperatures in between the grid points. The wind field is dominated by the slowly evolving jet stream and so here we take a daily average for atmospheric data, Mangini et al. (2018) having shown that the wind field does not alter significantly at this timescale.

2.2. Fuel burn data

All fuel-burn rate calculations made in this research rely on the formulae given in Poll and Schumann (2021a,b) which culminate in an expression for the fuel burn rate in kg s^{-1} of flight time, with the aircraft and atmospheric conditions in their current state:

$$\frac{dm_f}{dt} = \frac{W \times V_\infty}{\eta_0 \frac{L}{D} \times LCV} \quad (1)$$

where W is the current weight of the aircraft in N, m_f is the mass of the fuel in kg, V_∞ is the airspeed of the aircraft in m s^{-1} , LCV is the lower calorific value of the fuel (here taken as 43 MJ kg^{-1} for kerosene) and $\eta_0 \frac{L}{D}$ gives the maximum value of overall efficiency of the propulsion system multiplied by the aircraft lift to drag ratio.

In order to track mass change across a trajectory, it is important to know the mass of each aircraft at the start of the cruise phase. This is modelled as 97.5% of the take-off mass, which in turn can be calculated as:

$$TOM \approx \frac{0.7 \times MZFM + 0.3 \times OEM}{\exp(-(0.014 + 1.015 \times \frac{g \times R_{t(air)}}{\eta_0 \frac{L}{D} \times LCV})) - 0.05} \quad (2)$$

where $R_{t(air)}$ is the air distance flown by the aircraft in m, g is gravitational acceleration (here taken as 9.80665 m s^{-2}), $MZFM$ is the maximum permitted zero fuel mass and OEM is the operational empty mass, both in kg (Poll and Schumann, 2021a). The values of $MZFM$ and OEM can be found in the Aircraft Characteristics documentation provided by both Airbus (2021) and Boeing (2021). As the air distance for a flight is not known a priori, it is necessary to estimate this, using data from previous research (Wells et al., 2022), an approach justified by the low sensitivity of the TOM to $R_{t(air)}$.

Aircraft-specific parameters are provided by Poll and Schumann (2021b). 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 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.

2.3. 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 actual fuel use. 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: position as degrees of longitude and latitude, ground speed in knots and altitude in feet. However, we recalculated ground speeds as those recorded are not always consistent with given positions and times. A list of all flight numbers considered is given in Table 2. However, if data is missing for all or part of a flight, then this is removed from the dataset. The data was downloaded under the terms of a Flightradar24 gold licence and thus cannot be explicitly shared as part of this paper (Flightradar24, 2020).

Table 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

Table 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		

3. Analysis methods

We now describe how the minimum time and fuel-use trajectories are found, using optimal control theory as the basis of our problem formulation and dynamic programming to provide a globally minimal solution trajectory.

3.1. Formulation of optimal control problems

Here we change heading angle, θ and airspeed, V (the control variables) in order to navigate from an initial to a final state, whilst minimising a particular cost associated with the journey. 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.

Given the wide range of magnitudes of the physical quantities involved in the optimisation, it is beneficial to non-dimensionalise the variables. This is achieved by dividing relevant expressions by reference quantities of mass and airspeed. For mass, M_{ref} is the *TOM* of the aircraft and for airspeed, V_{ref} is the standard cruise airspeed of 240 m s⁻¹. This leads to the following definitions for the non-dimensionalised airspeed, V^* , zonal wind speed, u^* , meridional wind speed, v^* , time, t^* and mass, M^* .

$$V^* = \frac{V}{V_{ref}} \quad (3)$$

$$u^* = \frac{u}{V_{ref}} \quad (4)$$

$$v^* = \frac{v}{V_{ref}} \quad (5)$$

$$t^* = \frac{V_{ref} t}{R} \quad (6)$$

$$M^* = \frac{M}{M_{ref}} \quad (7)$$

where R is the approximate radius of the Earth, 6371 000 m.

The aircraft begins its flight at the departure airport and continues until it is within 225 km of the destination airport. Once the route has been found, the first few time intervals are removed 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. 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.

A 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 225 km. The spherical distance is calculated using the Haversine formula:

$$\begin{aligned} a &= \sin^2(\Delta lat/2) + \cos(lat1) \cos(lat2) \sin^2(\Delta lon/2) \\ c &= 2 \operatorname{atan2}(\sqrt{a}, \sqrt{1-a}) \\ d &= Rc \end{aligned} \quad (8)$$

where $lat1$ is the final latitude in the state vector, $lat2$ is the latitude of the destination airport, $\Delta lat = lat2 - lat1$ is the difference in latitudes between these two points, $\Delta lon = lon2 - lon1$ is the difference in their longitudes and R is the radius of the Earth (Veness, 2019). The $\operatorname{atan2}(Y, X)$ function computes $\arctan(Y/X)$, but adds an extra term of $(\pi/2)\operatorname{sign}(Y)(1 - \operatorname{sign}(X))$ to the computation to extend the range of results to cover $[-\pi, \pi]$. It also accepts the case where $a = 1$, returning $\pi/2$.

In order to guarantee a global minimum for time- and fuel-optimal trajectories, dynamic programming is used. By solving the Hamilton–Jacobi–Bellman equation associated with each formulation, via a semi-Lagrangian method (as explained in more detail in Appendices A.1 and A.2), we create a map of the value function, v . 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.

3.2. Time-minimal formulation

For the time-minimal formulation, the cost functional depends on the duration of the flight, expressed as:

$$J(\lambda, \phi, \theta, V^*) = \int_{t_0^*}^{t_f^*} 1 dt^* \quad (9)$$

which is minimised over the control variables, giving:

$$\min_{\theta, V^*} \int_{t_0^*}^{t_f^*} 1 dt^* \quad (10)$$

The non-linear constraint detailed in Section 3.1 is applied, so that the flight will not continue indefinitely, but has a countably infinite duration, $t_f - t_0$. The dynamical system is modelled with two non-dimensional state variables: longitude, λ and latitude, ϕ .

$$\frac{d\lambda}{dt^*} = \frac{1}{\cos\phi} (V^* \cos\theta + u^*) \quad (11)$$

$$\frac{d\phi}{dt^*} = V^* \sin\theta + v^* \quad (12)$$

The state variables are given boundary conditions of -80 to 10° for longitude and 30 to 70° for latitude, 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° . This discretisation was chosen as atmospheric data is obtained at this spatial resolution. A finer grid would require the same interpolation that 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 does not form part of the optimisation itself.

We use two control variables to minimise the flight duration. The first is the heading angle, θ which varies from 0 to 360° in 2° intervals and is measured anti-clockwise from due East. We also allow airspeed, V , to vary within the acceptable flying range of 200 to 250 m s^{-1} , in steps of 2 m s^{-1} . In practice, for a time-minimal optimisation, the optimal airspeed obtained is always the maximum value of this range.

The corresponding value function for this problem can be given as:

$$v(\mathbf{x}) = \inf_{\alpha \in A} J_{\mathbf{x}}(\alpha) \quad (13)$$

where α is the vector of control variables, A is the set of all possible admissible control variables and \mathbf{x} is the vector of state variables. We introduce the notation $J_{\mathbf{x}}$ to make explicit the dependence of the cost function to the initial condition for departure, \mathbf{x} . The value function is the unique solution of the HJB equation:

$$v(\mathbf{x}) + \sup_{\alpha \in A} \{-Dv \cdot f(\mathbf{x}, \alpha) - 1\} = 0 \quad \mathbf{x} \in \mathbb{R}^2 \setminus C \quad (14)$$

$$v(\mathbf{x}) = 0 \quad \forall \mathbf{x} \in C \quad (15)$$

where C is the target set for the trajectory, f is the dynamical system and Dv is the directional derivative of the value function. The optimal feedback control is found as:

$$\alpha^* = \arg \min_{\alpha \in A} \{-Dv \cdot f(\mathbf{x}, \alpha) - 1\} \quad (16)$$

To solve the HJB equation we apply a semi-Lagrangian Scheme, which is described fully in Appendix A.1.

3.3. Fuel-minimal formulation

The system when minimising fuel is modelled with three state variables, longitude, λ , latitude, ϕ and mass, M , which is bounded by the *TOM* and the *OEM*. These mass boundaries vary with aircraft model. The aim of the optimisation is to minimise a cost based on fuel burned during the journey:

$$\min J(\lambda, \phi, M^*, V^*, \theta) = \int_{t_0^*}^{t_f^*} g^*(\lambda, \phi, M^*, V^*) dt^* \quad (17)$$

where the fuel burn rate, g , is non-dimensionalised as:

$$g^* = \frac{Rg}{M_{ref} V_{ref}} \quad (18)$$

A non-linear constraint is again applied as discussed in Section 3.1, leading to a fuel-minimal, but free-time formulation.

The state mesh is split into cuboids that are:

$$\delta\lambda \times \delta\phi \times \delta M^* = 2.5^\circ \times 2.5^\circ \times \frac{3333 \text{ kg}}{M_{ref}} \quad (19)$$

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:

$$\frac{d\lambda}{dt^*} = \frac{1}{\cos\theta} (V^* \cos\theta + u^*) \quad (20)$$

$$\frac{d\phi}{dt^*} = V^* \sin\theta + v^* \quad (21)$$

$$\frac{dM^*}{dt^*} = -g^* \quad (22)$$

The value function for this continuous formulation, following (Cristiani and Martinon, 2009), is:

$$v(\mathbf{x}) - \inf_{\alpha \in \mathcal{A}} \{g(\mathbf{x}, \alpha) + f(\mathbf{x}, \alpha) \cdot Dv(\mathbf{x}) - (g^*(\mathbf{x}, \alpha) - 1)v(\mathbf{x})\} = 0 \quad \mathbf{x} \in \mathbb{R}^n \setminus C \quad (23)$$

$$v(\mathbf{x}) = 0 \quad \mathbf{x} \in C \quad (24)$$

where C , f and Dv are defined as in Section 3.2. The optimal feedback control for the fuel-minimal formulation is found as:

$$\alpha^* = \arg \min_{\alpha \in \mathcal{A}} \{g(\mathbf{x}, \alpha) + f(\mathbf{x}, \alpha) \cdot Dv(\mathbf{x}) - (g^*(\mathbf{x}, \alpha) - 1)v(\mathbf{x})\} \quad (25)$$

To solve the HJB equation we apply a semi-Lagrangian Scheme, which is described fully in Appendix A.2.

This formulation has not before been applied to fuel minimisation for flights through realistic wind fields. The continuous dynamical system is only discretised to allow an approximate solution to be found using the semi-Lagrangian scheme, unlike other methods where a grid system is used throughout to recover trajectory waypoints (Cheung, 2018; Rodionova et al., 2014, 2016).

3.4. Methods to recover simulated flight fuel use

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).

The value map is found for a particular day and model of aircraft, by updating the values until the norm of the differences in values between one value map and the next falls below a tolerance, set at one fifth of the product of the mesh dimensions. Now the controls and states giving the optimal trajectory can be retrieved by starting at the departure airport and locating the combination of controls which minimises the value function at the end of each timestep. This method continues, stepping forwards in time, until the point reached falls within 225 km of the destination airport. The initial few timesteps are removed so that the cruise phase considered leaves from the last waypoint that is still within 225 km of the departure airport.

In this way the fuel use is found immediately for the fuel-minimal route as it will just be the difference between the first and last fuel states. However, for the time-minimal formulation, the mass is determined once the route is known. In this case fuel use is calculated for each time interval from the relevant atmospheric data from the recorded position and airspeed. The sum of these values is the cruise phase total fuel use.

The time-minimal code based on the value function update, Eq. (A.1) of Appendix A.1, can be run across the whole state space with a time step of 125 s in a timely manner. However, including the extra state of mass for the fuel-minimal code, based on Eq. (A.4) of Appendix A.2, leads to lengthy optimisation times. 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 Fig. 1. To ensure the accuracy of this method, some examples of fuel-minimal calculations are run for a greatly extended computation time using just the 125 s time step across the whole state space. Results for this method are identical to those for the computationally much faster version.




Fig. 1. 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 is found. The Great Circle (GC) path, shown in magenta, is the shortest distance route around a sphere between the two airports.

3.5. Method to recover fuel use from recorded flight data

Time and horizontal position data from all recorded flights, alongside the re-analysis atmospheric data used for the optimised routes, provide the basis for fuel burn calculations. For a fair comparison, the recorded flight data are assumed to be at a constant altitude of approximately 34 000 feet. Winds and air temperature at a particular position can be found by linear interpolation from the re-analysis data using the 250 hPa pressure level. Once the wind speeds are known, the wind in the direction of the plane is calculated. This 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, which is a valid assumption as these legs do not generally exceed 2 min.

In the discussion section, varying altitudes are included, which entails taking the ICAO definition of the International Standard Atmosphere and interpolating between the altitude levels with a given pressure conversion to estimate the pressure at each point reached. This is paired with the co-ordinates of the point and by 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.

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. Altitude 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).

The fuel burn formula, given in Eq. (1), requires the input of the airspeed, the altitude, the air temperature and the mass of the aircraft. The atmospheric information is calculated at the start of each time step from re-analysis data. Using the product of the fuel burn rate and the step duration, the reduced aircraft mass at the end of a time step is obtained, ready to use as the aircraft mass in the subsequent step's fuel burn calculation. As for the simulated trajectories, only the cruise phase of each flight is considered.




Fig. 2. Plot showing how fuel burn rate varies with altitude for each aircraft model when flight fuel has been reduced to half of its original mass, at a point midway across the North Atlantic. Temperatures at each altitude are taken from a typical day, the 26th January, 2020.

Total trajectory fuel burn is the sum of the fuel use for each time step, which is the difference in the aircraft mass at the start and end of the cruise phase.

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 “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 Fig. 2 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 effect. Whilst the minimum fuel burn for the majority of the aircraft models occurs somewhere in the given cruise range between 34 000 and 40 000 feet, the B764 has a minimum fuel burn rate at a lower altitude and the A35K’s fuel burn rate continues to decrease across this range.


4. 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 altitude of approximately 34 000 ft.


4.1. Comparing fuel and time data: all flights

In total, data was available for 1 567 flights travelling West and 1 547 flights travelling East, between 1 December, 2019 and 29 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. Potential savings are related to the accuracy with which the Aircraft Performance Model (APM) estimates the difference between the fuel flow at different flight conditions and not the absolute accuracy of the fuel-burn method. Since the Poll–Schumann method is physics-based, we expect these relative differences to be accurately predicted (Poll and Schumann, 2021a,b).


In Fig. 3 both violin plots and notched box plots are shown for the fuel use for 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. This is due entirely to the different models of aircraft flown. In Fig. 4, the fuel use for each fuel-minimal trajectory is plotted on a stacked bar chart, showing clearly that the model of aircraft is causing the second mode. In this chart only



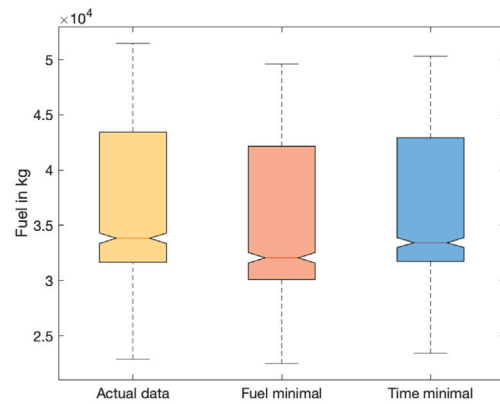
(a) Fuel use West



(b) Fuel use West



(c) Fuel use East



(d) Fuel use East

Fig. 3. 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.

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 4.2 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 boxplots it can be seen that the median of the fuel-minimal routes is significantly lower than 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 is clear from Fig. 3(a).

In Fig. 5 the same charts are plotted for the time of flight for each type of trajectory. 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 shows 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 routes 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.




Fig. 4. 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.

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 Fig. 6. Using the Environmental and Energy Study Institute's estimate that every 1 kg of aviation fuel burned produces 3.16 kg of CO₂, this would amount to a reduction of over 16.6 million kg of CO₂ 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 overall. To check whether airspeed does have an important effect on fuel efficiency, Fig. 7 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 Fig. 7 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 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.


4.2. 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 Fig. 8 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 Fig. 8(c) show the distribution of complete sets of flight data taken across the whole winter period.

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 Figs. 9(a) and 9(c) 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 Figs. 9(b) and 9(d). 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 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




(a) Time of cruise West



(b) Time of cruise West



(c) Time of cruise East



(d) Time of cruise East

Fig. 5. 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.

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 Fig. 10. It can be seen from Figs. 10(a) and 10(c) 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.

4.3. 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. 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 (Student, 1908). Fig. 11 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.




Fig. 6. 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.

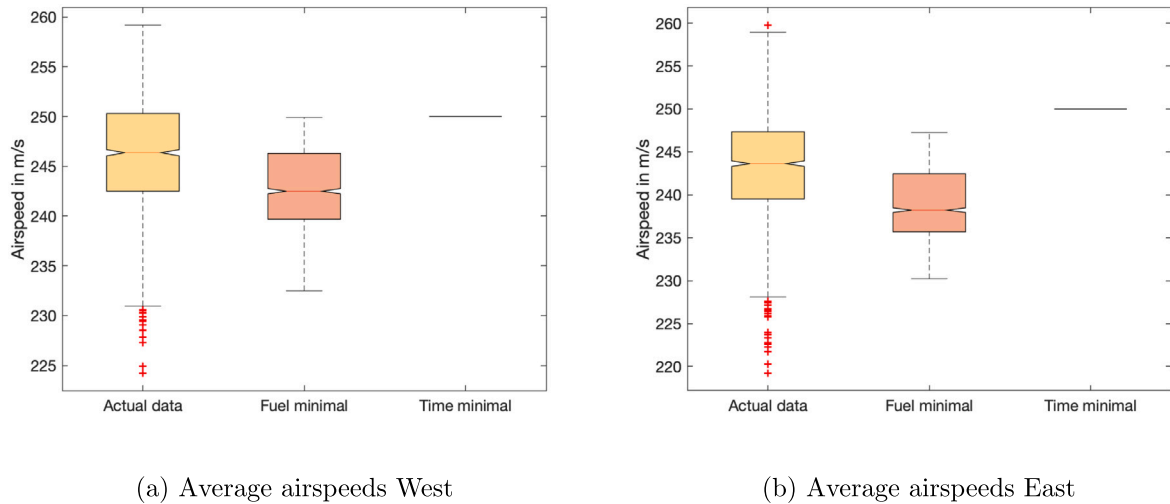


Fig. 7. Box plots showing average airspeed for each aircraft during the cruise phase of flights between LHR and JFK during winter 2019–2020. Average airspeed is considered for actual flights and simulated fuel-minimal trajectories. Time-minimal trajectories have a constant airspeed of 250 m s^{-1} , as this highest airspeed level is chosen as the optimal feedback control in the optimisation.

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 Fig. 2.

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 2.3, including the altitude as part of the optimisation was prohibited by computational restrictions. Although winds at each altitude level are very similar, temperatures are different, so the temperature used in the fuel burn calculations will be taken from the appropriate altitude. By using the horizontal paths found in Section 4 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

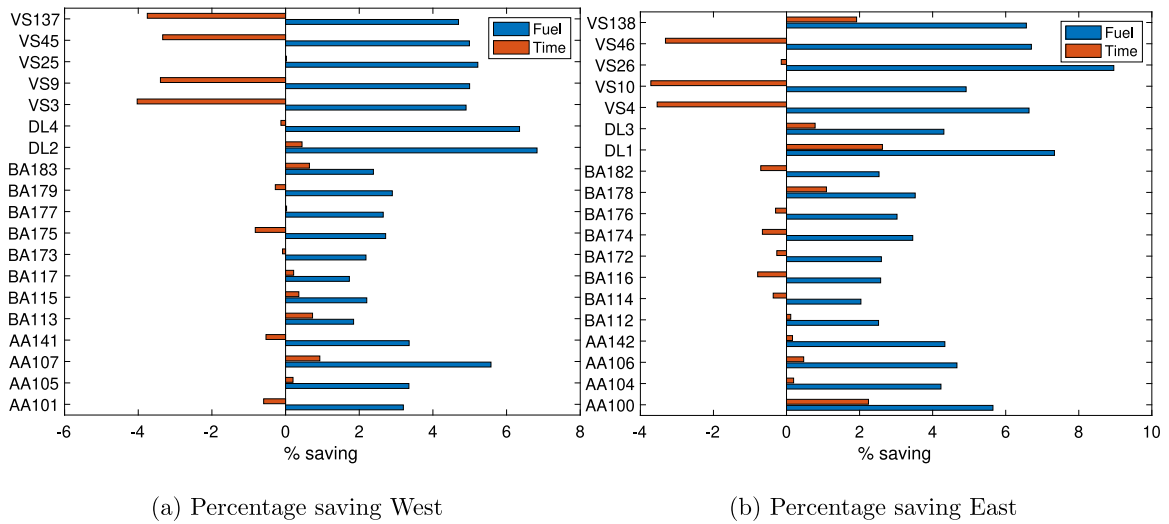


Fig. 8. Bar charts showing percentage savings in time and fuel made by using fuel-minimal trajectories rather than actual routes during the cruise phase of flights between LHR and JFK during winter 2019–2020. The box plot shows the distribution of the number of flights per route code.

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 million kg eastbound. This would reduce CO₂ emissions by 18 million kg across the 2019–2020 winter season.

4.4. Days when 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






Fig. 9. 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.

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 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 Fig. 12,




(a) Absolute average fuel savings West



(b) Proportion per aircraft West



(c) Absolute average fuel savings East



(d) Proportion per aircraft East

Fig. 10. 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.

which shows plots of routes flown by the Boeing 747-436 on the 19th February, 2020. Here the simulation takes a route south of the GC path, whilst the British Airways routes go north of it. Only flights BA113 and BA115 use significantly less fuel (here taken as 0.5%) than the fuel-minimal route. This shows that the optimised route flying South of the GC path is still more fuel efficient than 88% of the chosen flight paths taken by all of the aircraft making the journey from LHR to JFK on the 19th February, 2020.

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. 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.

5. Summary and conclusions

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 Hamilton–Jacobi–Bellman (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




Fig. 11. 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.




Fig. 12. Plots showing routes taken by actual flights and the time- and fuel-minimal routes simulated for the 19th February, 2020.

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 3.9%. Overall this amounts to a reduction of 16.6 million kg of CO₂ 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 trajectory-based operations 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. This study is not intended to provide an operational tool, but does furnish further evidence that North Atlantic routes could be made more fuel-efficient given the increased situational awareness now available.

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. Finally, the method applied in this research could be adapted to consider turbulence and convection avoidance, probabilistic atmospheric data and flight scheduling in future studies.

CRedit authorship contribution statement

Cathie A. Wells: Conceptualization, Methodology, Software, Validation, Formal analysis, Investigation, Resources, Data curation, Writing – original draft, Writing – review & editing, Visualization, Project administration, Funding acquisition. **Paul D. Williams:** Conceptualization, Methodology, Writing – review & editing, Supervision, Project administration. **Nancy K. Nichols:** Conceptualization, Methodology, Writing – review & editing, Supervision, Project administration. **Dante Kalise:** Conceptualization, Methodology, Writing – review & editing, Supervision, Project administration. **Ian Poll:** Conceptualization, Methodology, Writing – review & editing, Supervision, Project administration.

Data availability

Data will be made available on request.

Acknowledgments

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).

Appendix

A.1. Semi-Lagrangian scheme for the minimum time formulation

The value function is found numerically using the method given in [Falcone \(1997\)](#) and [Falcone and Ferretti \(2014\)](#). This semi-Lagrangian scheme uses both a spatial and temporal discretisation to give a value policy update of:

$$[\hat{v}]_i^{m+1} = \min_{a \in \mathcal{A}} \{e^{-\Delta t} I[\hat{v}]_i^m(x_i + \Delta t f(x_i, a))\} + 1 - e^{-\Delta t} \quad (\text{A.1})$$

in which the subscript i refers to a grid position and the superscript m to a fixed point iteration used to approximate the solution of the partial differential equation given in Eqs. (14) and (15). The derivative in the dynamical system is approximated by an Euler scheme, so that $\dot{\mathbf{x}} = f(\mathbf{x}, \alpha)$ is replaced by its discrete approximation:

$$\mathbf{x}_{k+1} = \mathbf{x}_k + \Delta t f(\mathbf{x}, \alpha) \quad (\text{A.2})$$

and $I[\hat{v}]$ is the linear interpolant used to give the value from the value map at this new off-grid position. An algorithm for the method of solution is shown as Algorithm 1.

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

A.2. A semi-Lagrangian scheme for the minimum fuel formulation

In the fuel-minimal formulation we use a semi-Lagrangian Scheme given in [Cristiani and Martinon \(2009\)](#), which incorporates both a spatial and temporal discretisation. The value policy update is found to be:

$$[\hat{v}]_i^{m+1} = \min_{a \in \mathcal{A}} \{I[\hat{v}]_i^m(x_i + \Delta t f(x_i, a)) + \Delta t g(x_i, \alpha)(1 - \hat{v}_i^m(x_i))\} \quad (\text{A.4})$$

Our algorithm for creating the value function map follows Algorithm 1, but uses Eq. (A.4) in place of Eq. (A.3) in step 3.

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 $I[\hat{v}^m]$ (here we use a linear interpolant)
- 3: Solve

$$[\hat{v}]_i^{m+1} = \min_{a \in \mathcal{A}} \{e^{-\Delta t} I[\hat{v}]_i^m(x_i + \Delta t f(x_i, a))\} + 1 - e^{-\Delta t} \quad (\text{A.3})$$

for each combination of control variables and for all gridpoints.

- 4: If a particular set of controls takes the arrival point, $x_i + \Delta t f(x_i, a)$, outside of the specified grid, then set the value for this point to 1.
- 5: Find the minimum $[\hat{v}]_i^{m+1}$ values and the controls used to obtain them and store these.
- 6: Measure $||\hat{v}^{m+1} - \hat{v}^m||$. If it is below a pre-set tolerance, stop, if not go to next step.
- 7: Update so $\hat{v}^m = \hat{v}^{m+1}$
- 8: Reset target values to 0.
- 9: Go back to step 2.

References

- Ahmed, K., Bousson, K., Coelho, M., 2021. A modified dynamic programming approach for 4D minimum fuel and emissions trajectory optimization. *Aerospace* 8 (5), URL <https://www.mdpi.com/2226-4310/8/5/135>.
- Airbus, 2021. Commercial aircraft. <https://www.airbus.com/aircraft.html>.
- Aireon, 2020. Operations overview. <https://aireon.com/resources/overview-materials/operations-overview/>.
- Bellman, R., 1957. *Dynamic Programming*. Princeton University Press.
- Boeing, 2021. Boeing Commercial Airplanes. <https://www.boeing.com/commercial/>.
- Burrows, J.W., 1983. Fuel-optimal aircraft trajectories with fixed arrival times. *J. Guid. Control Dyn.* 6 (1), 14–19.
- Cheung, J.C.H., 2018. Flight planning: node-based trajectory prediction and turbulence avoidance. *Meteorol. Appl.* 25 (1), 78–85.
- Cristiani, E., Martinon, P., 2009. Initialization of the shooting method via the hamilton-Jacobi-Bellman approach. URL <https://arxiv.org/abs/0910.0521>.
- ETS, E., 2021. Reducing emissions from aviation. https://ec.europa.eu/clima/eu-action/transport-emissions/reducing-emissions-aviation_en.
- Eurocontrol, 2021. Base of aircraft data. <https://www.eurocontrol.int/model/bada>.
- Falcone, M., 1997. Optimal control and viscosity solutions of Hamilton–Jacobi–Bellman equations. In: Bardi, M., Capuzzo-Dolcetta, I. (Eds.), Birkhauser, Boston, pp. 471–504.
- 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, vol. 1, pp. 103–110.
- Flightradar24, A., 2020. Flight data. <https://www.flightradar24.com/data/flights>.
- Franco, A., Rivas, D., 2011. Minimum-cost cruise at constant altitude of commercial aircraft including wind effects. *J. Guid. Control Dyn.* 34 (4), 1253–1260.
- Franco, A., Rivas, D., 2013. Analysis of optimal aircraft cruise with fixed arrival time including wind effects. *Aerosp. Sci. Technol.* 32.
- 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 Eng.* 80, 231–244, 3rd International Symposium on Aircraft Airworthiness (ISAA 2013). URL <https://www.sciencedirect.com/science/article/pii/S187770581401176X>.
- 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.
- Graver, B., Zhang, K., Rutherford, D., 2019. CO₂ emissions from commercial aviation, 2018. https://theicct.org/sites/default/files/publications/ICCT_CO2-commercial-aviation-2018_20190918.pdf.
- Grewe, V., Matthes, S., Dahlmann, K., 2019. The contribution of aviation NO_x emissions to climate change: are we ignoring methodological flaws? *Environ. Res. Lett.* 14 (12), 121003.
- Harada, A., Miyamoto, Y., Miyazawa, Y., Funabiki, K., 2013. Accuracy Evaluation of an Aircraft Performance Model with Airliner Flight Data. *Trans. Jpn. Soc. Aeronaut. Space Sci. Aerosp. Technol. Jpn.* 11, 79–85.
- 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* arXiv:1908.11086.
- Jackson, A., Newton, B., Hahn, D., Bussey, A., 2001. Statistical contrail forecasting. *J. Appl. Meteorol.* 40 (2), 269–279.
- Jardin, M., Bryson, A., 2012. Neighboring optimal aircraft guidance in winds. In: 18th Applied Aerodynamics Conference.
- 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.C., Ropelewski, C., Wang, J., Leetmaa, A., Reynolds, R., Jenne, R., Joseph, D., 1996. The NCEP/NCAR 40-year reanalysis project. *Bull. Am. Meteorol. Soc.* 77 (3), 437–472.
- Kim, J.-H., Chan, W., Sridhar, B., Sharmar, R., 2015. Combined winds and turbulence prediction system for automated air-traffic management applications. *Appl. Meteor. Climatol.* 54, 766–784.
- Lee, D., Fahey, D., Skowron, A., Allen, M., Burkhardt, U., Chen, Q., Doherty, S., Freeman, S., Forster, P., Fuglestedt, 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. *Atmos. Environ.* (1352–2310), 117834.
- Lovegren, J., Hansman, R., 2011. Estimation of potential aircraft fuel burn reduction in cruise via speed and altitude optimization strategies.
- Lunnon, R., Mirza, A., 2007. Benefits of promulgating higher resolution wind data for airline route planning. *Meteorol. Appl.* 14, 253–261.
- 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.
- Menon, P., 1989. Study of aircraft cruise. *J. Guid. Control Dyn.* 12 (5), 631–639.
- 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>.
- Murrieta Mendoza, A., Romain, C., Botez, R., 2020. 3D cruise trajectory optimization inspired by a shortest path algorithm. *Aerospace* 7, 99.

- Nguyen, N., 2006. Singular Arc Time-Optimal Climb Trajectory of Aircraft in a Two-Dimensional Wind Field. vol. 7.
- 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>.
- 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. *Aeronaut. J.* 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. *Aeronaut. J.* 125 (1284), 257–295.
- Rodionova, O., Delahaye, D., Sridhar, B., Ng, H., 2016. Deconflicting wind-optimal aircraft trajectories in North Atlantic oceanic airspace. In: AEGATS '16, Advanced Aircraft Efficiency in a Global Air Transport System. AAF, Paris, France, URL <https://hal-enac.archives-ouvertes.fr/hal-01304633>.
- Rodionova, O., Sbihi, M., Delahaye, D., Mongeau, M., 2014. North atlantic aircraft trajectory optimization. *Intell. Transp. Syst. IEEE Trans.* 15, 2202–2212.
- Rosenow, J., Fricke, H., Luchkova, T., Schultz, M., 2018. Minimizing contrail formation by rerouting around dynamic ice-supersaturated regions. *Aeronaut. Aerosp. Open Access J.* 2 (3), 105–111.
- Schultz, R., 1974. Fuel optimality of cruise. *J. Aircr.* 11 (9), 586–587.
- 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>.
- 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. *Transp. Res. C* 120, 102784.
- Sorensen, J., Waters, M., 1981. Airborne method to minimize fuel with fixed time-of-arrival constraints. *J. Guid. Control* 4 (3), 348–349.
- Speyer, J., 1973. On the fuel optimality of cruise. *J. Aircr.* 10 (12), 763–765.
- Sridhar, B., Ng, H., Chen, N., 2011. Aircraft trajectory optimization and contrails avoidance in the presence of winds. *J. Guid. Control Dyn.* 34, 1577–1584.
- 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. *Environ. Sci. Technol.* 54 (5), 2941–2950.
- UKCOP26, 2021. International aviation climate ambition coalition: COP26 declaration. <https://ukcop26.org/cop-26-declaration-international-aviation-climate-ambition-coalition/>.
- 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>.
- Veness, C., 2019. Calculate distance, bearing and more between latitude/longitude points. <https://www.movable-type.co.uk/scripts/latlong.html>.
- Villarroel, J., Rodrigues, L., 2016. Optimal control framework for cruise economy mode of Flight Management Systems. *J. Guid. Control Dyn.* 39 (5), 1022–1033.
- Wells, C., Kalise, D., Nichols, N., Poll, D., Williams, P., 2022. The role of airspeed variability in fixed-time, fuel-optimal aircraft trajectory planning. *Opt. Eng.* <http://dx.doi.org/10.1007/s11081-022-09720-9>.
- Wells, C., Williams, P., Nichols, N., Kalise, D., Poll, D., 2021. Reducing transatlantic flight emissions by fuel-optimised routing. *Environ. Res. Lett.* 16 (2), 025002.
- 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. vol. 6, pp. 4399–4410.
- 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. 8 (2), URL <https://www.mdpi.com/2226-4310/8/2/33>.
- 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. *Geosci. Model Dev.*
- Zermelo, E., 1930. Über die Navigation in der Luft als Problem der Variationsrechnung. *Jahresber. Dtsch. Math.-Ver.* 39, 44–48.