2ZERO PROJECT (74829) – TOWARDS ZERO EMISSIONS IN REGIONAL AIRLINE OPERATIONS

D5.1, Modelling and Simulation Summary Report

This report summarises the work in the Modelling and Simulation Work Package, WP 2, of the 2Zero project. It discusses the simulation that was built, how it works, its purpose, how it was used within the project, the results of doing so, and the various lessons learned from the project.

Version: 1.0

Main Authors:

Dr Jason Atkin, University of Nottingham

Dr Daniel Karapetyan, University of Nottingham

Dr Geert De Maere, University of Nottingham

Dr Warren Jackson, University of Nottingham

Dr Yun He, University of Nottingham

Dr Mehdi El Krari, University of Nottingham

Tanja Siggs, University of Nottingham

Primary Point of Contact for this report:

Dr Jason Atkin, jason.atkin@nottingham.ac.uk

This project has received funding from Innovate UK under UKRI's Future Flight Challenge Fund. The grant agreement number is 74829

Table of Contents

This project has received funding from Innovate UK under UKRI's Future Flight Challenge Fund.
The grant agreement number is 74829

This project has received funding from Innovate UK under UKRI’s Future Flight Challenge Fund.
The grant agreement number is 74829

1 Introduction

This report discusses the modelling and simulation work within work package 2 of the 2Zero project. The tasks were to develop a model and simulation which brings together the airspace, airline, and airport constraints, and to use it to evaluate the effects of utilising smaller hybrid-electric aircraft for regional point-to-point networks.

This report summarises the results from a collaboration between different partners. The airline schedule to use was provided by Loganair and considers how to maximise the usage of the aircraft to provide frequent flights to gain passenger demand between airports. Loganair also provided the information for the Twin Otter and how they would utilise it, and a discussion of how they feel airports would need to adapt to make these schedules work with a fast enough turn-around time (some consideration of the effects of these proposed changes, and a comparison with current operations/capabilities is given in the results section). Potential characteristics for the new hybrid-electric aircraft to use were provided by Ampaire, to investigate the use of a hybrid-electric aircraft for these routes and allow a comparison between the Twin Otter and its Hybrid-Electric variant (the Eco Otter), although it should be noted that these are from an 'early conceptual design' and operational studies are expected to tune some of these parameters. The obvious fuel savings from doing this will be observed in the results in Section 6.

For the purpose of this evaluation, the passenger capacity is assumed to be identical for the two aircraft, with a maximum passenger capacity of 19 persons.

This report is structured as follows:

In Section 2, an overview of the simulation is provided, explaining how the initial description was converted into a working system, and how the different elements have been modelled.

In Section 3, the input data is considered, in terms of what data was needed for the evaluation of this future concept for point-to-point regional aircraft usage.

Section 4 provides a relatively high-level overview of the simulation, explaining its different elements and how they are integrated together.

Section 5 provides more information about the algorithms (steps or approaches followed) which are used in each of the elements of the above model.

Section 6 then explains the configurations which were tested and presents the results. In particular it discusses the comparison between two input schedules (manually produced and an automated modification), two aircraft types (comparing Twin Otter vs a hybrid-electric variant), and a number of different assumptions about airport facilities and operation processing times.

This project has received funding from Innovate UK under UKRI's Future Flight Challenge Fund. The grant agreement number is 74829

Section 7, presents an overview of what has been learned from the project, including what the results show, what the simulation could be used to investigate, where we could see this going beyond this project and lessons learned or reinforced by the project which may be of use to others.

Finally, Section 8 presents some final comments and summaries.

This project has received funding from Innovate UK under UKRI's Future Flight Challenge Fund.
The grant agreement number is 74829

2 Problem Model and Simulation

This section talks about the model and the simulation. The model is the set of rules (embedded within Java code in most cases) which define what can and cannot be done, and how to perform calculations. This includes a 'central simulation' element, which takes the input problem, applies the different sub-models to it to perform the simulation, and outputs the results.

A simulation is the process of running the model and applying the different rules to the inputs to create the outputs. This includes:

- taking the input flight schedule, the characteristics of the aircraft, airports, airspace, etc.,
- predicting the kinds of delays that aircraft would have, the fuel burns, energy usages, recharging times, etc.,
- outputting the results in a format which can be interpreted by humans.

In other words, simulating the results is the process of asking the 'central simulation' element to run the various parts of the model for a specific input configuration to determine the expected results.

The model and simulation have been designed to be flexible and extendable for future use. It will probably be noticed in later sections that many features of the model discussed in Sections 3 and 4 are not needed for the evaluations performed within the 2Zero project. In some cases, this is because the final selection was not known when the model was designed (e.g., the model can use recharging policies that are much more complex than the ones utilised in the evaluation, and for most of the project it was expected that partial recharging in-flight would be possible), while in other cases, additional options are provided for completeness and flexibility for later extension/expansion, or for use with other airlines, airports or aircraft types. For example, there is no specific fuel type defined, except in terms of capacity, refuelling rates, usage rates according to speeds, and fuel costs. This allows different fuel types to be evaluated using the model. At present one fuel type and one (electrical) energy type are considered, but the model's implementation means that it could be used to evaluate other options, if desired.

The model has 6 main elements: the main simulation (the element which joins the other elements), demand prediction, flight allocation, airline cost model (includes rostering costs), airport model, and airspace model. Figure 2.1, on the next page, shows how these are integrated.

Further details of what each model achieves can be found in Section 4, describing the simulation. Details of how each element achieves its aims can be found in Section 5, providing an overview of the algorithms.

This project has received funding from Innovate UK under UKRI's Future Flight Challenge Fund.
The grant agreement number is 74829

Figure 2.1: Summary of links between elements of the simulation model

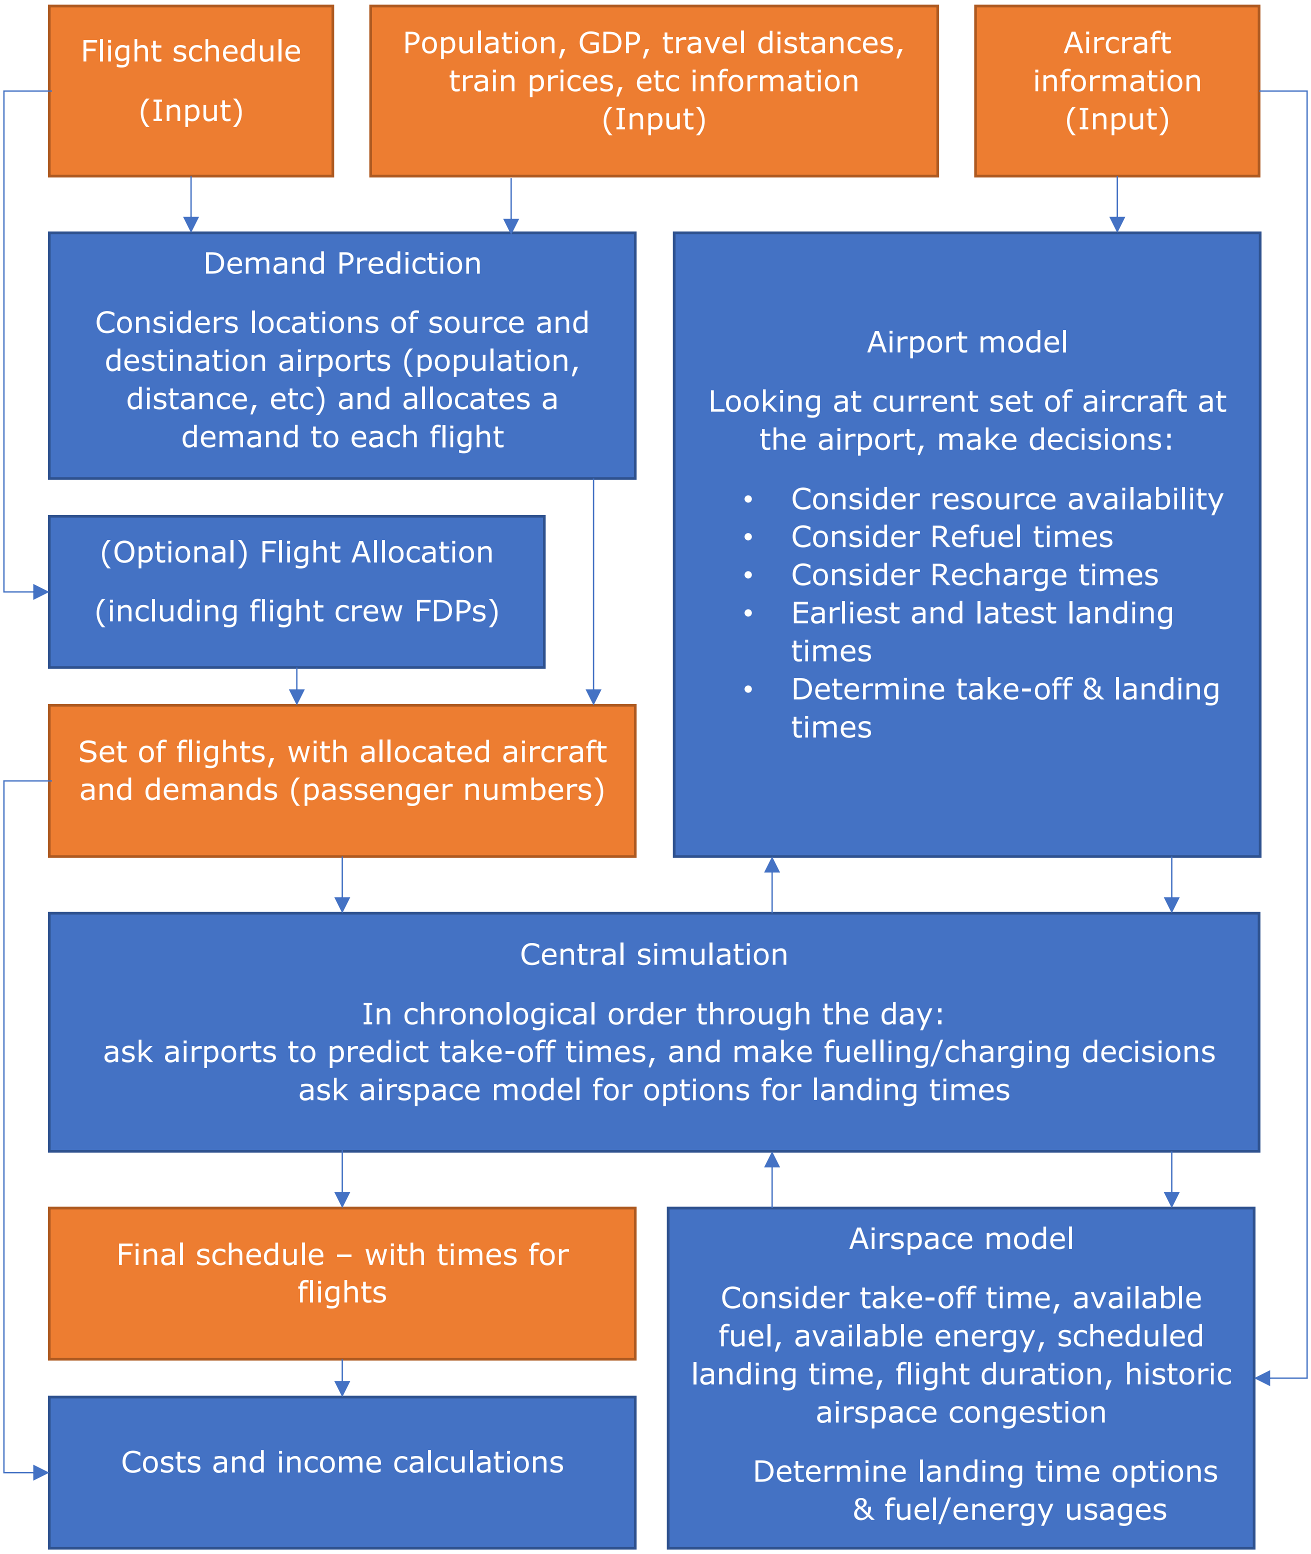

This project has received funding from Innovate UK under UKRI’s Future Flight Challenge Fund.
The grant agreement number is 74829

The initial proposal covered 4 different areas, but this had to expand to cover the requirements of the project, as data which we had expected to be available turned out not to be (e.g. demand data), and decisions which we expected to be needed (e.g. roster patterns) were considered unnecessary and impractical to perform at the planning/evaluation stage.

In this report we consider the problem in terms of the elements of Figure 2.1 (on the previous page), since that is how the model and algorithms have been built, but we first consider how these elements cover the four areas or the original plan.

Airline Model: This was planned to cover the aims and constraints of the airline to ensure that the simulation works and is acceptable from their point of view. The airline has been the main influence upon the model design, and the constraints can be observed in the flight schedule itself (provided by the airline), the cost model (the costs for the airline – including costs propagated by airports – are likely to be the main driver on acceptance of approaches), and the ability to fly the schedule (covered in the flight allocation/verification block).

Airport model: the airport model is a dedicated element in the simulation. It considers that each aircraft must perform a set of operations (e.g. landing, taxi-in, unloading, refueling, rechargeing, loading, etc.), each of which will utilise one or more resources. These resources may or may not be limited and may or may not have costs associated with them. It is often possible to perform operations in parallel (e.g. refueling while cleaning the aircraft, or refueling and rechargeing at the same time), while others have strict precedence constraints (e.g. unloading passengers before refueling, for safety reasons, and not loading new passengers until refueling has finished). Given the set of aircraft on the ground at a given point in time, the airport model is responsible for predicting the landing and take-off times (assuming a single runway) for the aircraft, as well as determining the amount of fuel or recharging that is required (depending upon predictions of fuel/energy consumption obtained from the airspace model, and the refueling/recharging policies that apply). This information is then provided to the central simulation.

Rostering: following advice from the airline, the rostering model was built to consider feasibility rather than optimised rosters for individual crew members. Flight schedules were built to fit within feasible Flight Duty Periods (FDPs) to which individual pilots can be assigned, subject to side constraints (e.g., annual leave, training, illness, etc.). In addition, the maximum flying hours per year affect how many of these schedules a pilot could fly per year (what is feasible for a timespan of one week may not be for a year). Given feasible flight schedules and flight duties that ensure that a pilot can return to base each day and does not exceed the maximum allowable flying time, a standard pilots-to-flights ratio was applied, as suggested by the airline. This approach was also applied for maintenance and ground crew. The latter of these elements was particularly interesting as it is possible that the ground operations may need to change

This project has received funding from Innovate UK under UKRI's Future Flight Challenge Fund.
The grant agreement number is 74829

considerably if the best performance is to be achieved from the schedules – reducing ground times to a minimum, despite additional costs which this could incur for airlines (assuming that airport costs could be offset upon the airlines).

Airspace model: the airspace model is responsible for modelling the behaviour of aircraft in the air. The airspace model is contained in its own component in the model and is given the take-off time of the aircraft, its characteristics, and historic congestion levels for the airspace. The model will then output: a set of possible landing times; expected energy and fuel usages to achieve these times (considering the speeds that the aircraft would have to travel at); the fuel and energy usage characteristics; and the likely congestion (introducing delays).

This project has received funding from Innovate UK under UKRI's Future Flight Challenge Fund.
The grant agreement number is 74829

3 Input Data

This section describes the input data that the model requires, and why they are necessary.

3.1 Flight Schedule

The flight schedule consists of a set of flights performed by a set of aircraft within a day of operation. Each flight has an origin and destination airport, a scheduled take-off time (from the origin airport), a landing time (at the destination airport), and a day of operation (Monday to Sunday).

The schedule is constructed such that each aircraft starts and ends the day at the same airport. It also has sufficient turnaround time allocated between consecutive flights carried out by the same aircraft. Additional ground time is included after a given number of hours such that pilot changeover can be accommodated.

3.2 Aircraft data

The aircraft data contains the characteristics of aircraft.

The aircraft model includes information which is required to model the use of the airspace: the cruise speed and maximum speed; the fuel burns and energy requirements for climb, descent, normal cruise speed and maximum speed; the climb and descent rates; and airspeeds during these phases. The airspace model utilises this information (see Sections 4.5 and 5.5) to model the flight, and to predict flight times, battery usage and fuel burn.

Each aircraft model has a specific configuration for the ground operations. The set of necessary operations (boarding, cleaning, refuelling, etc.) for the turnaround, the resources that are required, and the durations and constraints are not necessarily the same for different aircraft.

3.3 Airspace data

The airspace data divides the airspace into longitude/latitude rectangles and consists of a two-dimensional table giving a congestion value for each square. These congestion values are used by the airspace model to increase associated flight times (by increasing the apparent distance that would be flown).

For the simulations described in this report, this data was generated from pre-covid ADS-B data containing counts of aircraft in each rectangle in each minute. Since these are just tables of numbers, any alternative sources of information can be used, and could be modified to consider the effects of predicted air traffic changes, if desired.

This project has received funding from Innovate UK under UKRI's Future Flight Challenge Fund. The grant agreement number is 74829

3.4 Geographical location related data

The geographical location (latitude/longitude) of each full postcode (e.g. AB10 1AB) in the UK is obtained from https://www.freemaptools.com/download-uk-postcode-lat-lng.htm. The latitude/longitude for each postcode district (e.g. AB10) is calculated as the average latitude/longitude of all postcodes inside this district. There are 2977 postcode districts in total.

The list of airports in the schedule is provided by Loganair. The geographical location of each airport included in this project is obtained from https://www.partow.net/miscellaneous/airportdatabase/.

The geographical location of each train station in the UK is from the public database of National Rail.

Population and GDP data per head of each postcode district is obtained from various public sources of the 2011 population census across the UK. UK population data are obtained from:

- http://www.nomisweb.co.uk/census/2011/ks101ew for England and Wales,
- https://www.isdscotland.org/Products-and-Services/GPD-Support/Population/Census/ for Scotland,
- https://www.ninis2.nisra.gov.uk/public/SearchResults.aspx?sk=postcodes for Northern Ireland.

3.5 Transport data

The Open-Source Routing Machine (OSRM) service is used to calculate the duration and distance of road travel between two geographical locations. The road travel price is set to be 45p per mile.

Train travel duration and ticket prices are obtained from National Rail (http://www.railwaycodes.org.uk/stations/stationa.shtm). The train schedules themselves are obtained from https://ojp.nationalrail.co.uk/.

The ticket price for each flight is calculated as follows:

- If the travel time by car is less than 8 hours, the ticket price by air is the same as travelling by car.
- If the travel time by car is more than 8 hours, it is the cost by car plus an additional - configurable - hotel cost (set to £55).

3.6 Airport data

A set of airports is used during the simulation (and present in the flight schedules). Each airport is identified by a 4-letter ICAO code and a 3-letter IATA

This project has received funding from Innovate UK under UKRI's Future Flight Challenge Fund. The grant agreement number is 74829

code. This identifies their locations, which are used by the demand and airspace models.

Airports are classified as small, medium, or large. The size determines the time an aircraft spends on the airport’s taxiways and the number of resources that are available (ground staff, fuel trucks, etc.). These are used when limited ground resources are modelled.

This project has received funding from Innovate UK under UKRI’s Future Flight Challenge Fund.
The grant agreement number is 74829

4 System Description

This section explains the purpose of each element of the model, the inputs, and outputs. Section 5 explains how each element achieves its goals by summarising the algorithms involved.

4.1 Demand Prediction

4.1.1 Inputs:

Inputs include:

- the network of airports to be served, where each airport is identified by a 4-letter ICAO code with a geographical location defined by its latitude and longitude.
- the flight schedule specifying the departure airport, departure time, arrival airport, and arrival time for each flight.
- a list of all postcode districts (e.g. NG1) within the UK and their central geographical points defined in terms of latitude and longitude.
- the maximum flying distance for the fleet type used.
- minimum layover times to ensure feasible transits.
- transportation data, as described in Section 3.5.
- time-based demand distributions for each day of the week.
- priorities for converting a passenger flight to freight.

4.1.2 Outputs:

The type of each flight (passenger/freight) and the expected number of passengers per flight (if it is a passenger flight).

4.1.3 Purpose/aims:

The goal of the demand model is to estimate the total number of passengers on board of each flight, which is used to estimate the total revenue for the airline. The demand prediction uses a gravity model that considers various factors, including population density, the GDP for each geographical location, and the duration and costs for other modes of transport (e.g. by car or by train). It uses a Mixed Integer Linear Programming (MILP) model to balance air passenger flow in the flight network and maximise the expected profit. It converts passenger flights to freight-services if the expected demand is too low (using time-based demand distributions for each day of week and a time-dependent priority to serve these demands provided by the airline).

This project has received funding from Innovate UK under UKRI's Future Flight Challenge Fund.
The grant agreement number is 74829

4.2 Flight Allocation

4.2.1 Inputs:

The flight schedule provided by the airline.

4.2.2 Outputs:

The aircraft routing, i.e. allocations of aircraft to flights.

4.2.3 Purpose/aims:

Given a set of flights, each with a scheduled departure time at an origin airport and a scheduled arrival time at a destination airport, the flight allocation checks whether the existing schedule is feasible in terms of the maximum flying duty duration (defined by the airline) assuming the minimum turnaround duration at each of the airports. If it is not feasible, the flight allocation proposes a new aircraft allocation using additional aircraft.

The flight allocation uses two mathematical models. The first one is a flow-based model, which considers the number of aircraft in and out of each airport. It ensures the flight connections are on-time and determines the minimum number of aircraft required to operate the schedule with fixed scheduled times of departure and arrival. However, it does not assign a specific aircraft to flights. The second model is based on column-generation and able to re-assign aircraft to flights. It first checks, for each aircraft, whether the existing schedule is feasible in terms of the minimum required turnaround time and the maximum Flight Duty Period (FDP). If feasible, the assignment is kept unchanged. If not, flights are combined to form a new schedule that satisfies the minimum turnaround time and maximum Flight Duty Period constraints (without changing the scheduled time of departure and arrival), while utilising the minimum number of aircraft.

4.3 Central Simulation

4.3.1 Inputs:

All inputs go through the central simulation and are passed on to the other elements.

4.3.2 Outputs:

Simulated “actual” times of events, fuel usages, battery usages, departure delays, and arrival delays.

This project has received funding from Innovate UK under UKRI’s Future Flight Challenge Fund.
The grant agreement number is 74829

Purpose/aims: The central simulation integrates the other elements of the system and drives the simulation. The central simulation provides a framework for integrating all models into an all-in-one simulation system.

Managing data: All data required by any parts of the simulation is read by the central simulation and made available to the other models. Which files to read (see Section 0) is determined by the instance/configuration being evaluated. For example, if the current configuration is for evaluating the Twin Otter aircraft, then the system needs to load the aircraft performance data for the Twin Otter, but if we are evaluating a configuration for the Eco Otter, then all relevant data from the Eco Otter configuration must be read. The other parts of the system 'ask' the central simulation for any data they may need.

Integration of sub-models: The other models can be considered a part of the central simulation, performing elements of the task for it.

The central simulation loads and evaluates the flights from the input schedules based on the input data, policies, and aircraft configuration. It then works through the day using a 'rolling horizon' approach, whereby at any time in the simulation there will be some events which have already happened (they are fixed), some which are about to happen (so decisions made at this point will be unchangeable from then onwards, as soon as the events will have 'happened'), and some events which are still to happen in the future – for which any decisions the simulation planned will be flexible and can be changed later as the simulation time approaches the decision point.

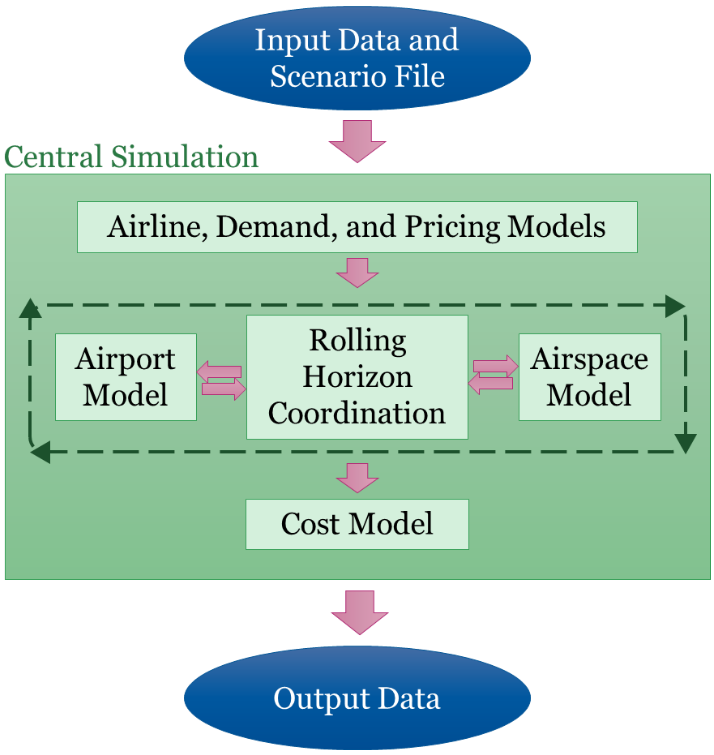

Within any 15-minute time window, each aircraft will either be in the air, heading towards an airport with a planned arrival time, or already at an airport. It is noted that with taxi-in, taxi-out and flight times no aircraft can be at two airports within the same 15-minute period, so each aircraft can be associated with one airport at any point in time (the one it is at or the one it is heading for).

The central simulation will coordinate with the airport model, telling it for each airport which set of aircraft it is responsible for, and coordinating with it to

This project has received funding from Innovate UK under UKRI's Future Flight Challenge Fund. The grant agreement number is [74829](74829)

provide it with information such as predicted arrivals times from information previously provided by the airspace model (for enroute aircraft), flight durations for flights about to take off (asking the airspace model), and recharging/refueling amounts required (according to the current policies).

Decisions from the airport model for each airport will then be used to update the status of each aircraft, after which the time is advanced by 15 minutes and the process is repeated.

4.4 Airport Model

4.4.1 Inputs:

- Runway capacity: limits the number of aircraft movements in a given time window.
- Airport data: defines the taxi duration and the number of available (shared) resources to run the ground operations of different aircraft on stand.
- Aircraft data: sets the regular operations needed for the aircraft before its pushback.
- Airspace data: provides the earliest time of landing for a flight and the amount of fuel and energy (for hybrid models) needed for the next flight.

4.4.2 Outputs:

- Times: provides to the central simulation the departure and take-off times for outbound flights; landing and arrival times for inbound flights.
- Resource usage: returns the number of resources used in each interval.

4.4.3 Purpose/aims:

The airport model simulates ground operations between landing and take-off. It schedules them subject to runway capacity, operational constraints, and resource constraints. The model determines when aircraft can land and take-off and, when the resource constraints are omitted, the minimum number of resources required at the airport.

This project has received funding from Innovate UK under UKRI's Future Flight Challenge Fund.
The grant agreement number is 74829

4.5 Airspace Model

4.5.1 **Inputs:**

- Aircraft details: fuel and energy usage characteristics, cruise and top speed.
- Recharging policy: whether aircraft should try to recharge during flight (if fuel is available, potentially burning extra fuel to recharge batteries), at a specified rate, to the maximum charge.
- Flight details: take-off time for the aircraft, scheduled landing time, starting battery charge and fuel load, optionally a fixed landing time.

4.5.2 **Outputs:**

- The model is used in two different ways: firstly, it can be used to determine the 'ideal' landing time for an aircraft, along with the required fuel and energy usage; secondly, it can be used to determine the fuel and battery energy usage necessary to achieve a specified landing time.

4.5.3 Purpose/aims:

In general, the first mode, to predict a landing time, is utilised by the simulation first, to calculate an aircraft's landing time. It is also used by the airport model to predict the amount of energy and fuel an aircraft will require, so that the aircraft is appropriately prepared for the journey. The refueling/recharging policies will determine how this information is used – e.g. whether refueling just enough (plus contingency fuel) or whether recharging fully.

The airspace model has a concept of an 'ideal landing time', which it will calculate for a flight based upon the flight schedule, aircraft characteristics, and take-off time. Several assumptions underpin the model:

- It is assumed that aircraft will fly at cruise speed unless there is a necessity to fly faster. Normally, the ideal landing time will be the time at which it would land at the destination if it flew at its cruise speed for the entire flight.
- If an aircraft flying at its cruise speed would not reach its destination airport in time for the scheduled landing time (e.g., the input schedule has a very tight flight time, or the aircraft took off later than planned and there is a need to make up time), then the flight will speed up (up to its maximum speed, no more) to try to reach the destination by the scheduled time. The ideal landing time will then be updated to be the new landing time for this speed-up. The available fuel and energy may further limit the maximum speed at which the aircraft can fly – depending upon the aircraft design, speeding up beyond cruise speed will require either more energy per mile or more fuel per mile, or both. A calculation has to

This project has received funding from Innovate UK under UKRI's Future Flight Challenge Fund.
The grant agreement number is 74829

take place to work out what speed could be achieved while still ensuring required reserves of fuel or energy, and if the reserves limit this maximum speed then the adjusted maximum speed will be used in the calculation above and below when working out the ideal landing time. It is assumed that a flight should not be in the air for more than 50 minutes (for reasons of passenger comfort, e.g. there are no refreshment facilities on board). A 50-minute flight time is used in the same way as the scheduled landing time – if an aircraft would be in the air for more than 50 minutes, it attempts to speed up (up to the maximum speed, no more) to reduce the flight time.

- If both the 50-minute flight time and the scheduled time cannot be achieved at the normal cruise speed, then the aircraft will attempt to speed up to achieve both – which means that the speed-up will be required to reach the earliest of these two times.

It should be noted that a speedup from 160 knots to 175 knots makes only a few minutes difference in a one hour flight, so it is highly likely that an aircraft which cannot reach its desired landing time at normal speed will still not be able to do so at its maximum speed either. That is, a speed-up is often to try to reduce the lateness rather than to actually be able to achieve the desired landing time.

A flight is assumed to have three components: the climb, the cruise and the descent. Each has different fuel and energy usage characteristics in the aircraft data. The climb phase happens first. Its duration is determined by the climb rate and the cruise altitude of the flight. The descent phase happens at the end of the flight, and its duration is determined by the descent rate and the starting (cruise) altitude of the flight. Horizontal flight distances for the climb and descent phase are determined, and the cruise phase is assumed to cover the remaining distance from source to destination which is not covered by the climb and descent phases. It is assumed that only the cruise phase has a variable speed to calculate, and that the airspeed for climb and descent is fixed. The distance for the cruise phase is used, along with the calculated speed for the cruise (considering the assumptions above for whether the normal cruise speed can be used or whether the aircraft must speed up) to determine the duration of the cruise. The durations of the three phases are then summed to determine the duration of the flight, from which the landing time can be determined.

Some airspace is congested. The airspace in the model is divided into latitude/longitude rectangles, and an assumption is made that these squares will be roughly rectangular when converted to real positions. It is noted that this is not quite correct, but analysis shows that the difference from this assumption in calculated journey times and distances using the algorithm in Section 5.5.2 is less than 0.1% for UK-related latitudes, with the small rectangles used in our experiments. Congestion values are calculated in advance for each of these 'rectangles' (see Input Data in Section 3.3). An assumption is made that a flight

This project has received funding from Innovate UK under UKRI's Future Flight Challenge Fund.
The grant agreement number is 74829

will fly from source to destination airport in a straight line, and the proportion of the flight which is in each rectangle is determined. The average congestion value for the flight is determined from the congestion values from the rectangles and the proportion of the flight which was spent in the rectangle. The congestion value is then used to increase the apparent distance of the flight, which increases the flight duration for any given flight speed.

In a second mode, the exact landing time is given, instead of asking the model to determine an ideal landing time. In this case, the calculations are all the same, except that the total flight time is known as an input, so the time of the cruise phase can be calculated, thus the cruise speed can be calculated, and the energy and fuel usages can be determined. Apart from not needing to calculate the cruise time/landing time, the calculations are the same as for the other mode. The ability to do this means that the airport model can ask the airspace model for an expected arrival time, fuel usage and energy usage, but can also ask for the effects of landing at other times (in terms of energy usage, fuel burn and cruise speeds). This gives it the potential to find a more economical flight time, for instance not increasing the speed to meet a tight schedule if resources at the airport would not be available anyway.

4.6 Costs and Income Calculations

4.6.1 Inputs:

- Aircraft costs: lease, parking charges, insurance, mechanical and other maintenances, navigation database, charts &manuals, flight data monitoring, de-icing. Aircraft costs given per month, year, or longer.
- Pilots costs: their salary, training, travel and other expenses. They are given per month or per year.
- Engineering costs: A-check consumables, avionics, flight deck monitoring, props, maintenance checks, engine performance, rotable. These costs are provided and calculated per hour of flight. Other engineering costs (battery maintenance or replacement, wheels, tyres, and brakes) are given per cycle.
- Airport costs: taxes, landing fees, ground handling fees, on-stand fees, terminal navigation fees, delay costs. They are counted per flight. The airport charges for the fuel (per litre) and power used (per kWh) to charge the batteries.
- Passenger costs: booking fees and insurance. Both costs are defined per passenger.
- Other costs given per month: customer services, operations department, EU261.
- The expected income per cargo flight.

This project has received funding from Innovate UK under UKRI's Future Flight Challenge Fund. The grant agreement number is 74829

4.6.2 Outputs:

- Expenditure: The above-mentioned input costs are used to calculate the total cost per year for each category and the total cost to run the flight schedule over 365 days.
- Income: Ticket revenues and cargo revenues. The first one is given by the simulation, while the second is calculated based on the number of freight flights and the average income per cargo flight.
- Profit: The difference between the total revenues and total costs.

4.6.3 Purpose/aims:

All cost and income data are important to assess viability and cost-effectiveness of running a flight schedule in a specific configuration.

For a better readability, the data are listed in a spreadsheet, where the first sheet provides the global costs/revenues. Then each sheet provides the detailed costs to allow investigating which costs are the more impactful on profit. The spreadsheet has the input costs/income as static data and uses some output files from the simulation to feed the sheets with the necessary data to perform the calculations.

This project has received funding from Innovate UK under UKRI's Future Flight Challenge Fund.
The grant agreement number is 74829

5 Algorithms

Sections 3 and 4 explained the problems that each element of the simulation is handling, and the inputs and outputs. This section summarises the algorithms that are used to do this, showing how decisions are made and how the inputs are utilised to produce the outputs.

5.1 Demand Prediction

Demand predictions are obtained using a gravity model with Mixed Integer Linear Programming (MILP, solved using CPLEX from IBM). Between each pair of postcode districts, i and j, the demand model provides an estimation of the total number of passengers that travel between i and j per day. This demand is then split between competing modes of transport. The demand between each pair of postcodes is mapped onto the airport network in the flight schedule by putting the demand onto the flight sector(s) that serve those postcode pairs in the quickest time. The demand for each day is distributed into periods inside the day using distribution factors. Finally, the demand for each flight is calculated using a MILP model and flights with low demand are converted to freight.

5.1.1 Total demand estimation for a best-case travel scenario

The passenger volume V_{ij} between two postcode districts i and j is estimated using the following formula:

$$\boldsymbol{V_{ij}} = \boldsymbol{k} \cdot \frac{\left(\boldsymbol{P_i P_j}\right)^{\boldsymbol{\alpha}} \quad \boldsymbol{\theta}}{\boldsymbol{d_{ij}^{\beta} t_{ij}^{\gamma} c_{ij}^{\eta}}}.$$

where k is a constant, P_i and P_j are the population sizes, and G_i and G_j are the GDPs per head for i and j.

To have an estimation of the total number of passengers between postcode districts i and j under the best scenario, we use:

- d_{ij} the shortest distance;
- t_{ij} the shortest travel time;
- c_{ij} the cheapest cost of travel.

In addition to the values above, there are parameters to calibrate the contribution of different factors to the demand. More precisely, α is a constant for the population size; β is a constant for the distance between locations; γ is a constant which changes the influence of the travel time; η is a constant for the

This project has received funding from Innovate UK under UKRI's Future Flight Challenge Fund.
The grant agreement number is 74829

cost of travel; and θ is a constant for the influence brought by the GDP for the given areas.

5.1.2 Demand calibration

The values of $\alpha, \beta, \gamma, \eta, \theta$ and k are calibrated automatically using the irace package[1] in *R*. Given a set of target demand values, the irace package iteratively test various combinations for their values, within a certain range, and outputs the combination for which the calculated demand is closest to the target values.

In our experiments, the target demand values are the historic demand data provided by Loganair for a set of representative origin/destination pairs. After an initial manual test, the range for α, β, η was set to be between 0 and 1, and for γ, θ, k to be between 1 to 2. The final values calibrated by the irace package are 0.3768 for α, 0.4502 for β, 0.0882 for η, 1.0859 for γ, 1.4558 for θ and 1.8323 for k.

5.1.3 Calculation of travel time by air using existing flight schedule

To have better estimates of the total travel time between post code districts for an existing flight schedule, we enumerate the shortest paths to go from origin to destination. Each postcode district location i is first mapped to its nearest airport AP_i based on the road travel duration (district to airport), and the potential paths between two airports (direct or via other airports) are computed using a dynamic programming approach. Between each pair of airports for the given flights, we calculate the shortest path in terms of total travel time, subject to a minimum transfer time of 15 minutes, a maximum transfer time of 1 hour and 45 minutes, and at most one transfer per journey. The resulting shortest path is defined as the *flight path*.

The total travel time by air is then calculated as the sum of the travel time by car between the origin postcode district and the nearest airport (based on time), the total travel time by air between the origin and destination airport, and the travel time by car between the destination airport to the destination postcode district.

The demand for air travel on different routes between each pair of postcode districts can then be estimated by applying an attractiveness factor. The latter factor compares the total travel time and total cost by air to those by train and by car.

---

[1] Manuel López-Ibáñez, Jérémie Dubois-Lacoste, Leslie Pérez Cáceres, Mauro Birattari, Thomas Stützle, The irace package: Iterated racing for automatic algorithm configuration, Operations Research Perspectives, Volume 3, 2016, Pages 43-58, ISSN 2214-7160, https://doi.org/10.1016/j.orp.2016.09.002.

This project has received funding from Innovate UK under UKRI's Future Flight Challenge Fund.
The grant agreement number is 74829

5.1.4 Estimation of attractiveness factors for air travel

As the total travel time increases for a journey on a given mode of transport, M, the attractiveness of using this mode decreases in proportion to the travel times for alternative modes of transport. Similarly, as the cost of travel increase, the attractiveness of using a particular mode of transport also decreases in proportion to the costs of the alternative modes. The attractiveness factor for transport mode M between i and j due to total travel time and total cost is then defined as:

$$FT_{ij,M} = \frac{1}{t_{ij,M} c_{ij,M}}$$

The relative attractiveness of air travel based on total travel time and total cost of travel is then calculated as:

$$f_{ij,Air} = \frac{\frac{1}{t_{ij,Air} c_{ij,Air}}}{\frac{1}{t_{ij,Air} c_{ij,Air}} + \frac{1}{t_{ij,Train} c_{ij,Train}} + \frac{1}{t_{ij,Car} c_{ij,Car}}}$$

The resulting attractiveness factor can then be used to estimate the total volume of demand for air travel between each pair of postcode districts i and j.

$$V_{ij,Air} = V_{ij} f_{ij,Air}$$

5.1.5 Demand distribution during the day and priority for converting to freight

According to the airline, the demand distribution for different times of different days of the week follows a distribution that can be modelled as a lookup table. The table divides each day of the week into 5 intervals, each one corresponding to a time window of 4 to 6 hours. The total demand for air travel across a day, as calculated in Section 5.1.3, is distributed into periodic demands based on the demand distribution table.

Each flight path, as identified by the algorithm in section 5.1.2, generates a net income per passenger transported. Flights can be converted to "freight flights" if they are not profitable as "passenger flights".

When converting passenger flights to freight flights, a priority order is used. A flight can only be converted to freight if all flights with lower priorities have been converted to freight.

Let p_l denote the priority order of flight l. In our modelling, the higher the value for p_l, the lower the priority for the flight, and hence, the higher the chance that

This project has received funding from Innovate UK under UKRI's Future Flight Challenge Fund. The grant agreement number is 74829

the flight will be converted to a freight flight. That is, two flights l_1 and l_2, if $p_{l_1} < p_{l_2}$, then l_1 is of higher priority and cannot be converted to freight if l_2 has not been converted to freight.

5.1.6 MILP model and solution

The MILP model is formulated as follows:

- Let K be the set of all demands for all origin/destination airport pairs over all time periods. For each demand $k \in K$ between an origin/destination pair in a given interval, D_k represents the maximum number of passengers to be transported in the network from the origin airport to the destination airport.
- Let FL be the set of all flights in the schedule.
- Let Ω be the set of all flight paths found by the algorithm in Section 5.1.2, and $\Omega_k \subset \Omega$ be the subset of feasible paths for demand k. If there is a direct flight between the origin and destination for demand k, set Ω_k contains only direct flights. Otherwise, set Ω_k contains all valid flight paths with one transfer as defined in Section 5.1.3. Let $\Omega_l \subset \Omega$ be the subset of paths containing flight l.
- Q denotes the capacity on each arc, which represents the total passenger capacity of the aircraft in use.
- Let r_p^ω denote the income per passenger transported along flight path ω, and let r_l^f denote the income for flight l if it is converted to a freight flight.
- For each flight path $\omega \in \Omega$, for each flight $l \in FL$, the binary parameter u_ω^l equals 1 if flight l is included in path ω, 0 otherwise.
- Let x_ω^k be an integer variable for the number of passengers k travelling on path $\omega \in \Omega_k$.
- Let y_l be a binary variable which equals 1 if flight l is used to transport passengers in the solution and 0 otherwise.

The MILP formulation is then given by:

This project has received funding from Innovate UK under UKRI's Future Flight Challenge Fund.
The grant agreement number is 74829

$$\max \sum_{k\in K}\sum_{\omega\in\Omega} r^{p}_{\omega} x^{k}_{\omega} + \sum_{l\in FL} r^{f}_{l}(1-y_l) \quad (1)$$

$$\sum_{\omega\in\Omega_k} x^{k}_{\omega} \le D_k \qquad \forall k\in K \quad (2)$$

$$\sum_{k\in K}\sum_{\omega\in\Omega_l\cap\Omega_k} x^{k}_{\omega} u^{l}_{\omega} \le y_l Q \qquad \forall l\in FL \quad (3)$$

$$y_l \le \sum_{k\in K}\sum_{\omega\in\Omega_l\cap\Omega_k} x^{k}_{\omega} u^{l}_{\omega} \qquad \forall l\in FL \quad (4)$$

$$y_{l_1} \ge y_{l_2} \qquad \forall s\in FS, \forall l_1, l_2 \in \{FL_s \mid l_1 \ne l_2, p_{l_1} \le p_{l_2}\} \quad (5)$$

$$x^{k}_{\omega} \in \mathbb{N} \qquad \forall \omega\in\Omega, \forall k\in K \quad (6)$$

$$y_l \in \{0,1\} \qquad \forall l\in FL \quad (7)$$

Constraints (2) state that total number of passengers transported on all routes cannot not exceed the total demand. Constraints (3) state that total number of passengers transported by each flight is less or equal to the total capacity if the flight is a passenger flight. Constraints (4) link variables x and y, stating that there should be at least one passenger on each passenger flight. Constraints (5) ensure that, if l_1 has a higher priority, it cannot be converted to freight, unless l_2 has been converted to freight, with l_1 and l_2 denoting flights belonging to the same sector.

The objective function (1) maximises the total profit from passengers and freight.

5.1.7 General Algorithm

The following algorithm is used to complete the assignment of passengers to flights.

- Estimate the demand for the whole day for all modes of transport
- Initialise the “flight type changed” boolean (true/false – whether it was changed or not) indicator to be true for every period of the day, to ensure that each flight is considered at least once
- While there are still freight flights or the time limit is not exceeded:
    - For each period of a day:
        - If the flight type has changed according to the boolean indicator (note: initially all are true so this will always happen once for each flight, but if they change again it will happen an additional time for each change)
            - For each origin/destination airport pair
                - Calculate the shortest path as defined in Section 5.1.2

This project has received funding from Innovate UK under UKRI’s Future Flight Challenge Fund.
The grant agreement number is 74829

            - Apply the air attractiveness factor for each demand from the origin airport to the destination airport
            - Apply the demand distribution for the current period
            - Apply the increase factor to scale up the demand if needed
        - Otherwise
            - Get the pre-calculated demand from the cached values
        - Add CPLEX variables (Constraints (6) and (7))
        - Set the maximum air demand cap (Constraints (2)).
        - Set priority constraints (5) for passenger to freight conversion
- Add capacity and linking constraints (3, 4)
- If CPLEX can solve the model
    - Set the actual flight types and update the total number of converted flights
    - If no flight is converted to freight in this iteration
        - Set passenger information for each passenger flight

The final demand allocation is then cached so that the next time the same configuration is encountered, calculations do not have to be repeated unnecessarily.

5.2 Flight Allocation

In the flight allocation, the following constraints must be satisfied.

- Each flight should be associated with one and only one aircraft.
- The origin/destination airport for each aircraft must be coherent.
- The departure/arrival time of each aircraft at each airport must be coherent.

5.2.1 Network Flow Model

The model assumes that all departure and arrival times for all flights are feasible, and that no delays can be added to change the existing schedule.

- Let FL denote the set of flight legs with scheduled time of departure and arrival.
- Let AC be the set of aircraft and AP the set of airports. Each aircraft is denoted by $ac \in AC$ and each airport is denoted by $ap \in AP$.
- For each flight $l \in FL$, let STD_l and STA_l be the scheduled time of departure and arrival, respectively.

This project has received funding from Innovate UK under UKRI's Future Flight Challenge Fund. The grant agreement number is 74829

- Let ap_l^d and ap_l^a be the departure and arrival airport.
- The minimum required turnaround time between the arrival and departure of two consecutive flights carried out by the same aircraft is denoted by T.

The flight allocation problem is modelled as a single commodity flow network and its solution gives a lower bound on the number of aircraft required to operate the schedule under the above assumptions, without considering airport and airspace delays.

The graph $G = (FL^*, A)$ for our network flow model is defined as follows:

- The node set $FL^* = FL \cup_{ap \in AP} \{l_{ap}^+, l_{ap}^-\}$ is composed of the set of all flight legs FL and two dummy nodes l_{ap}^+ and l_{ap}^- for each airport $ap \in AP$.
- There is an arc $(l_1, l_2) \in FL^* \times FL^*$ if any of the following three conditions hold:
  1. both l_1 and l_2 belong to FL^*: the departure airport of flight leg l_2 is the arrival airport of flight leg l_1 ($ap_{l_2}^d = ap_{l_1}^a$), and that the minimum turnaround time between flight legs is respected ($STD_{l_2} \geq STA_{l_1} + T$);
  2. there is an airport $ap_{l_2}^d \in AP$ that l_1 = l + ap and l_2 ∈ F L departs from ap (ap = Apd l_2), representing the fact that an airplane starts the route at airport ap with flight l2;
  3. there is an airport ap ∈ AP where l1 arrives and l_2 = l −ap, representing the fact that an airplane ends the route with flight l_1 at airport ap . For each arc (l_1, l_2) ∈ A, let binary variable xl_1,l_2 equal 1 if the arc (l_1, l_2) is selected in the solution.

Below is a single commodity network flow model which minimizes the total number of aircraft used to operate all the flights on one day once and only once.

$$\min \sum_{l_2 \in FL^*} \sum_{ap \in AP} x_{l_{ap}^+, l_2}$$

$$\sum_{l_2 \in FL} x_{l_{ap}^+, l_2} = \sum_{l_1 \in FL} x_{l_1, l_{ap}^-} \qquad \forall ap \in AP \qquad (9)$$

$$\sum_{l_2 \in FL \cup \{f^-_{ap_a^{l_1}}\}} x_{l_1, l_2} - \sum_{l_2 \in FL \cup \{f^+_{ap_d^{l_1}}\}} x_{l_2, l_1} = 0 \qquad \forall l_1 \in FL \qquad (10)$$

$$\sum_{l_2 \in FL^*} x_{l_1, l_2} = 1 \qquad \forall l_1 \in FL \qquad (11)$$

Constraints (9) ensure that for each airport, during the planning horizon, the total number of aircraft starting and ending their route at the respective airport is the same. Constraints (10) are flow conservation constraints for flight leg nodes. Constraints (11) ensure that each flight is operated once and only once.

This project has received funding from Innovate UK under UKRI's Future Flight Challenge Fund.
The grant agreement number is 74829

The above model is a relaxation of the aircraft assignment problem. If the solution of this model requires the same number of aircraft as in the initial solution, constructed by human schedulers, then the initial solution is optimal. If not, i.e. the solution of this model uses fewer aircraft than in the initial solution, the model is refined to ensure that each route start and ends at the same airport. Finally, if the solution of this model uses more aircraft than in the initial solution, the initial schedule is too tight for the given number of aircraft, and the initial schedule should be refined.

5.2.2 Column generation and branch-and-price

A column generation model is used to generate new aircraft schedules if the initial schedule is infeasible. This can, for instance, be caused by turnaround times that are too tight, or by Flight Duty Periods (FDP) that exceed the airline's pre-defined maximum.

The general column generation idea is illustrated in Figure 5.1 below.

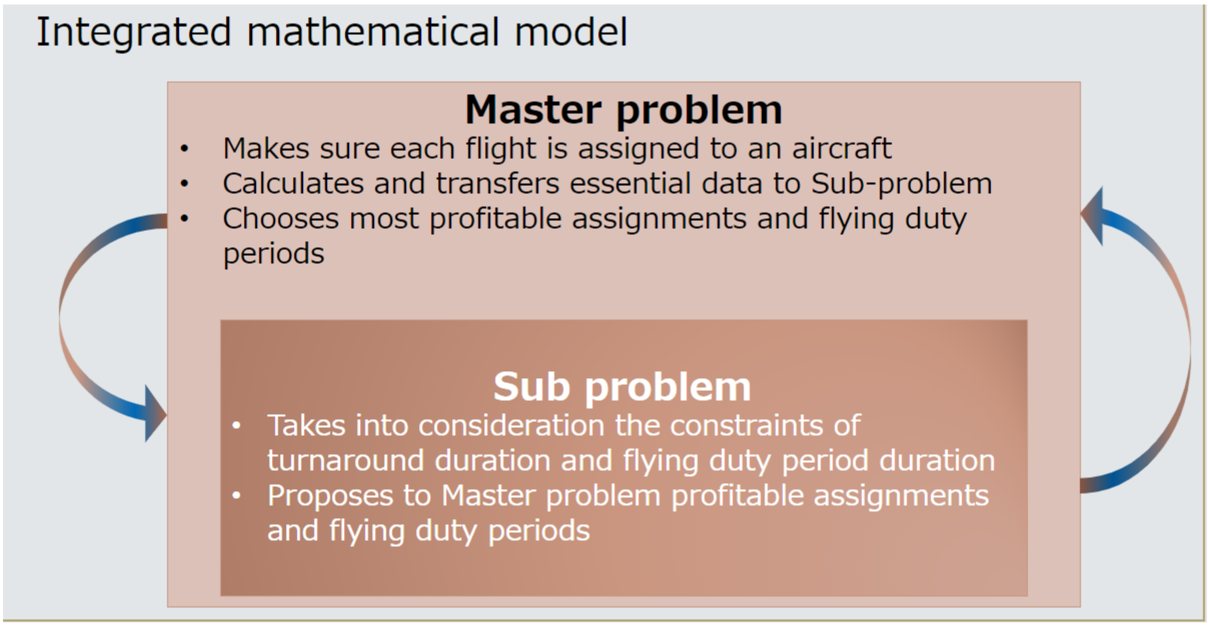

Figure 5.1: The integrated master and sub problems

The following presents the underpinning mathematical model in more detail.

The column generation works with a set of flight paths, in which each path defines a sequence of flights $\{l_1, l_2, \dots, l_n\}$ carried out by a single aircraft and covering one day. In the mathematical model below, constant v_l^p is equal to 1 if flight $l \in FL$ is included in path $p \in P$. A fixed operational cost, c, is incurred for every aircraft that is used. Each flight $l \in FL$ has an operational cost c_l irrespective of the aircraft assigned to it. Binary variable z_p is equal to 1 if flight path $p \in P$ is included in the solution, 0 otherwise.

This project has received funding from Innovate UK under UKRI's Future Flight Challenge Fund. The grant agreement number is 74829

The objective is to minimise the total operational cost for the airline, which is the sum of the operational costs for all flights, augmented by the cost for the number of aircraft used.

$$\min \sum_{p\in P} \sum_{l\in FL} c_l v_l^p z_p + c \sum_{p\in P} z_p$$

$$\sum_{p\in P} z_p v_l^p \geq 1 \qquad \forall l \in FL \qquad (13)$$

$$z_p \in \mathbb{N} \qquad \forall p \in P \qquad (14)$$

Constraints (13) ensure that each flight is covered by at least one flight path. The column generation starts with a linear relaxation of constraints (14) and a limited set of paths included in P. This is commonly known as the Restricted Master Problem (RMP).

It gets $\alpha_l \in R$, the dual variables corresponding to constraints (13). The reduced cost of each path p is:

$$c_p^r = c + \sum_{l\in L} (c_l - \alpha_l) v_l^p$$

The sub-problem is to choose a combination of flights in FL to build cost reducing paths that minimise c_p^r. Each new path p must start and end at the same airport, satisfy minimum turnaround times and, in addition, meet restrictions related to the maximum FDP.

In the sub-problem, v_l^p is a binary decision variable that is equal to 1 if flight l is contained in path p, 0 otherwise. If a path p with a negative reduced cost c_p^r is found, then the path is added to the RMP.

The path-finding sub-problem is modelled as a graph $G = (FL, A)$, in which each node represents a flight leg. There is an arc $(l_1, l_2) \in A$ between two flight legs $l_1, l_2 \in FL$ if, and only if, the departure airport for flight leg l_2 is the same as the arrival airport of flight leg l_1 ($ap_{l_2}^d = ap_{l_1}^a$), and if the minimum turnaround time between flight legs is respected ($STD_{l_2} \geq STD_{l_1} + T$). In addition, restrictions related to the maximum FDP are also validated, and any infeasible paths are not added to the RMP.

5.2.3 Complete algorithms

The main algorithm maintains a set of paths which would end in each flight. After the initialisation, it runs the algorithm to find cost reducing paths (i.e. shorter paths), and extracts and combines paths to form valid ones.

The whole Branch and Price framework works as follows:

This project has received funding from Innovate UK under UKRI's Future Flight Challenge Fund. The grant agreement number is 74829

It maintains a list of active nodes, each node containing a variable with an integer value (0 or 1) and related constraints added to the original model. It also maintains the best LB and UB found, as well as a list of incumbent solutions.

The list of solution is initialised with a dummy node (the root node).

While the list of active nodes is not empty, it builds the model by column, and sets variables to integer values if needed. It then solves the RMP and generates columns depending on reduced cost. It subsequently updates the Lower Bound and sees if a fractional variable in the Linear Programming solution to branch on exists in the solution. If such a variable exists, it creates a new active node with the fractional variable set to 1, and add the new node to the list of active nodes. Otherwise, a new integer solution is found, added to the list of solutions, and the Upper Bound is updated.

5.3 Central Simulation

An overview of the simulation was provided in Section 4.3. The central simulation manages the individual flights for the duration of their existence in the simulation:

- At the start of the simulation, the input schedules are read in and flights are created based on these.
- The set of flights are handed to the airline model. This assigns individual aircraft to individual flights based on constraints set out in the instance data. The airline model also estimates passenger demand and assigns passengers to each of the flights, including any transits and the ticket prices paid by passengers.
- After aircraft (physical aircraft) and passengers have been assigned to flights, the simulator communicates with the airport and airspace models to estimate fuel consumption, battery energy consumption, en-route delays, and ground-based delays. This information is used to predict the earliest times that flights can depart, take off, land, and arrive on-stand.
- The whole simulation is broken down into smaller decomposed problems which reduces the overall complexity of the problem to solve. 15-minute time slices, called time-windows, are used to divide the flights into smaller sets for the individual airport models. This eliminates interactions between potentially multiple airports for a single aircraft, which would lead to an intractable problem to solve and would reduce the complexity of the airspace by requiring calculations to be performed for a single take-off time.
- Two planning windows are used (see Figure 5.2):
    - The planning horizon specifies the duration of the look-ahead for which the system will consider the aircraft. Any event which would occur beyond the end of the current planning horizon is ignored in the current iteration.

This project has received funding from Innovate UK under UKRI's Future Flight Challenge Fund.
The grant agreement number is 74829

- The time window specifies the length of time which will be planned.

For example, with a planning horizon of 1 hour and a time window of 15 minutes, an airport will be informed of all aircraft which are at the current airport already or which are in the air and expected to land within the next hour. Decisions made by the airport will only be enacted for the next 15 minutes, and time will be advanced by 15 minutes the next time the airport model solves the problem. In this way the airport model has some visibility of later effects of the decisions it makes at that time, since the other aircraft which may be affected by the planning are visible to the model.

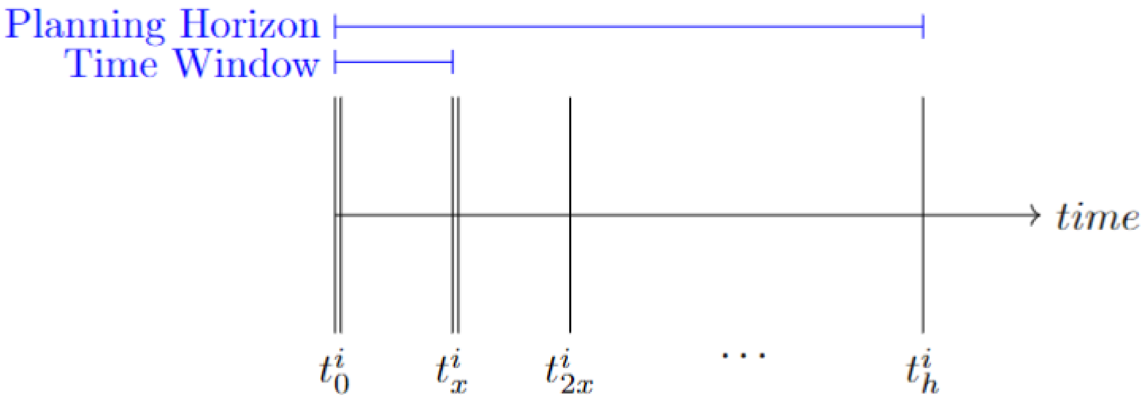

Figure 5.2: Visualisation of the time windows in a planning horizon

- For each new time window, the relevant flights are partitioned based on the specific airport which next needs to handle each of them. The airspace model is used to update fuel and battery usage based on landing times and what the earliest and latest times are that each flight can land based on aircraft performance and energy reserves. These flights are passed to the relevant airport models for each to decide landing times, turnaround operation times, and departure and take-off times.
- At each iteration, the times that are set based on the solution found by the airport model are fixed if they are within the current time-window.
- The simulation coordinates the various models to update any information based on accurate timings (such as fixed ATOT's for aircraft that would have already landed) such that in the next iteration, the models have the most up to date data to solve the sub-problems.
- After all flights have been completely simulated, the central simulation uses the simulated timings to calculate the KPIs and outputs them for evaluation.

This project has received funding from Innovate UK under UKRI's Future Flight Challenge Fund.
The grant agreement number is 74829

5.4 Airport Model

The model schedules runway movements and ground operations for each aircraft. It runs under the rolling horizon and time window defined in the previous section. In each iteration, it takes into consideration all flights arriving or already on ground, but not yet departed, at a particular airport in the planning horizon. However, it only fixes operations scheduled within the current time window, so that later operations can be re-scheduled if necessary. Time is discretised in the model: it uses intervals of half minute, which is the smallest duration any of the ground operations can take to be done.

The airport model uses two sub-models: the first one, *landing model*, schedules the landing operations (landing, taxi-in), the remaining operations at the airport are scheduled by the second sub-model, *post-landing model*.

Any operations scheduled within the current iteration's time window are considered to be fixed. That is, their time can no longer be changed in future iterations.

5.4.1 Runway movements

The airport model considers a single runway for each airfield, used for both landings and take-offs. The runway is assumed to have a capacity of three movements per five minutes time window.

The airport provides to each flight the earliest possible time for either landing or take-off. An inbound flight is kept in the airspace for longer when more time is required while the runway is busy. An outbound flight remains on stand until the runway allows it to take-off.

The taxiways are not modelled but the airport model is considering the taxi-in and taxi-out durations in these operations.

5.4.2 Ground operations

The turnaround starts after aircraft towing and ends before pushback starts.

The model schedules the ground operations as early as possible while making a feasible schedule. This feasibility is determined by the set of constraints associated with each operation:

Timing constraints: different timing constraints exists on the model

- Ground operations cannot be scheduled before the end of the pushback.
- If not yet fixed, an operation cannot start before the first time interval of the rolling horizon.
- The aircraft cannot leave the stand before the scheduled time of the departure.

- Boarding cannot start any earlier than fifteen minutes before the scheduled time of the departure.
- Baggage handling cannot start any earlier than twenty minutes before the scheduled time of the departure.

Precedence constraints: some operations cannot start before the end of others. Boarding for example cannot start before refuelling has finished, which cannot start before de-boarding has finished.

Contiguous constraints: having a gap of time between some operations is not permitted. The whole refuelling process for instance needs to be done continuously to optimise the usage of the fuel truck. The post-refuelling must start immediately after the end of refuelling, which must start once the pre-refuelling finishes.

Asynchronous constraints: this class of constraints was defined specifically for refuelling and recharging. If these two operations cannot be done in parallel (which is the case for our aircraft model due to fire/shock risk), there is no constraint on which one starts first.

Resource constraints: the model ensures that operations have the required resources during the time intervals in which they take place. An operation can be delayed until the number of required resources are available. In each time interval, the sum of all used resources needs to be less than or equal to the number of available resources at the airport.

5.4.3 Landing post-processing phase

Time setting:

After each iteration of the simulation, the airport model sets the starting times for all the operations scheduled before the end of the horizon and fixes them if they fall within the time window. The model will also fix any operation starting after the time window if it has contiguous constraints with a fixed operation. This applies for the post-landing and post-processing phase as well.

The programmed time of the landing is set if the model scheduled it.

Fuel and energy:

The amount of fuel and energy needed for the next flight cannot be decided before aircraft land. Once landed, these amounts can be calculated and thus, the durations for the refuelling and recharging operations can be set.

5.4.4 Post-landing post-processing phase

Time setting:

This project has received funding from Innovate UK under UKRI's Future Flight Challenge Fund.
The grant agreement number is 74829

The scheduled time of departure is set once the pushback operation is scheduled. The same holds for the scheduled take-off time if the model schedules it.

Resource usage:

When resource constraints are considered during the simulation, the model needs to:

- Update the resource usage for the next aircraft in the same iteration.
- Save the resource usage for the fixed operations in each time interval they were it is scheduled in.

5.5 Airspace Model

The algorithms for the airspace model are really deterministic mathematical calculations.

5.5.1 Calculation of flight distance

The Great Circle Distance is used to determine the distance between any two airports. The latitude and longitude of each airport is given in the input data.

5.5.2 Calculation of average congestion value

Considering the airspace as a grid of latitude and longitude points, the latitudes of each horizontal grid rectangle boundary are known and the longitudes of each vertical grid rectangle boundary are known. Therefore, it is trivial to calculate which grid square any latitude/longitude point is in.

A straight line is ‘drawn’ from the latitude/longitude point for the source airport to the latitude/longitude point for the destination airport. It is trivial to work out which horizontal grid boundaries this line crosses, and which vertical boundaries it crosses. Since the start and end points of the line are known, it is trivial to calculate the gradient of the line. Given the starting point, the gradient of the line and any latitude on the line, the longitude can be easily calculated. Similarly for calculating the latitude from a longitude, therefore it is easy to calculate the latitude/longitude points at which the line crosses horizontal or vertical grid rectangle boundaries.

The line from source to destination can therefore be considered to be a series of journeys between different points – each of which is either the source, destination or an intersection point where it changes grid rectangles. Considering each sequential pair of points, the points denote a journey across a specific grid rectangle. The distance across the rectangle is determined from the great circle distance between the two latitude points. The congestion value of each rectangle is known. The two values are multiplied together and the sum of these, divided

This project has received funding from Innovate UK under UKRI’s Future Flight Challenge Fund.
The grant agreement number is 74829

by the sum of the distance values gives the average congestion value for the journey.

5.5.3 Calculation of maximum speed

The normal maximum speed is given as input data, however there is potentially either a fuel burn increase or a stored energy usage cost from increasing the speed beyond the cruise speed – or both. Because of this, once the cruise distance is known the following calculation is performed:

Consider initial fuel and deduct the fuel needed for climb and descent to calculation remaining fuel for the cruise.

Consider initial stored energy and deduct the fuel needed for climb and descent to calculate remaining fuel for the cruise.

Starting at the cruise speed, for each 2% (2% is an arbitrary increment to get enough options) speed increment between the cruise speed and top speed (i.e. top speed is 100%), determine the cruise time (noting that faster speeds reduce cruise time), the fuel burn and the energy usage (at this speed) to complete the cruise. Determine the highest increment for which there is both energy and fuel available to complete the flight. Record this as the maximum speed.

Note: If there is sufficient fuel and battery energy available to speed up then the maximum speed from this calculation will be the same as the normal maximum speed from the input data. If not then the aircraft will not be able to accelerate so much and the calculated maximum speed will be lower.

If the aircraft does not have sufficient fuel or energy to reach the destination using its standard cruise speed then the flight is determined to be infeasible and the schedule is rejected.

It should be noted that minimum energy levels and fuel levels considered by this calculation are assumed to include any contingency fuel or energy, and that this is calculated externally to this calculation.

5.5.4 Calculation of ideal landing time

The ideal landing time is calculated as the earliest of the starting time + 50 minutes and the scheduled landing time. This is the time that the flight would prefer not to be too late for.

5.5.5 Calculation of actual cruise speed

If the aircraft can reach the destination at or before the ideal landing time by flying at its cruise speed then it will do so.

Otherwise, the cruise speed that the aircraft would need to reach the destination by the ideal time is calculated. If this speed is no greater than the maximum

This project has received funding from Innovate UK under UKRI's Future Flight Challenge Fund.
The grant agreement number is 74829

speed calculated from Section 5.5.3 then this speed will be used, to reach the destination at the ideal landing time.

If the maximum speed from Section 5.5.3 is not sufficient to reach the destination by the ideal landing time then the maximum speed is used as the actual speed, to reach the destination as soon as possible.

5.5.6 Calculation of fuel and energy usage for a landing time

The fuel burn and energy usage at the cruise and maximum speeds are available in the input data. It is assumed that the usages for any intermediate speed can be interpolated, noting that this is likely to be slightly pessimistic. Given this assumption, since the speed is known from Section 5.5.5, giving the consumptions per minute (from the interpolation), and the flight duration can be calculated from the take-off time and landing time, giving the number of minutes of flight, the overall fuel and energy consumptions can be calculated by multiplying the two figures.

5.5.7 Consideration of recharging in the air

Recharging and refueling policies are used to determine how to fuel and charge the aircraft. One of these policies allows charging in the air, so that aircraft whose design permits excess energy in the cruise or descent phase to be used to charge the energy storage can be included in the modelling.

If the charging policy indicates that charging in the air is possible then data will be available for the cruise and descent phases, specifying the fuel costs per minute and the energy available per minute in this situation.

Note: It is assumed that a smart management system on the aircraft would manage the charging and later enhancement of this model to consider multiple charging speeds may be needed. This is therefore an approximation to allow aircraft which charge in the air to be modelled, under the assumption than many of these aircraft will actually not be able to charge fully anyway in flight, given the energy needed for the climb and the aim to reduce engine sizes, so it is likely that charging will be maximised through most of the flight in these cases.

The charging model will first determine whether it is more efficient (energy per unit of fuel) to charge in the cruise or descent phase, and utilise that phase by preference, maximising its usage of that phase before using the other phase. As noted, it is likely that a realistic implementation would fully utilise both phases most of the time, but we desired to model an intelligent controller to also handle partial cases.

For each flight phase (in the priority order), the model will consider how much fuel is available (not being used for other phases) and how much charge is needed (between the current charge and maximum it should charge to), and how long this phase is (i.e. maximum charging time in the phase), and will use

This project has received funding from Innovate UK under UKRI's Future Flight Challenge Fund.
The grant agreement number is 74829

this information to determine how many minutes (as a continuous value, not an integer) of charging is desired. This will be assumed to occur, and then the other phase will be considered if there is still excess fuel and energy need.

An important assumption is made that the aircraft design will either be using energy in the cruise phase or will be charging. As such the recharging in the air is applied after the other calculations have been made. This means that energy from charging would NOT be available to increase flight speed if energy is needed to do this.

5.5.8 Special cases

When the final aircraft model was made available for use in the simulation, some of the earlier agreed assumptions were found not to be valid. Two different modifications have therefore been applied for use when needed, and are discussed here.

Firstly, the original models provided had three flight phases, climb, cruise and descent. The latest data has four phases – take-off (up to 500’), climb (the remaining climb distance), cruise and descent. Rather than change all of the model to account for this change, since the take-off is effectively a fixed extra fuel and energy cost (the take-off duration is fixed, as is fuel/energy per minute, therefore excess over normal climb is a constant), two additional values can be set in the airspace model now, specifying excess energy and fuel requirements for take-off. For completeness and to aid future-proofing, two values were specified for landing as well, in case of future needs to have a separate case for this.

Secondly, the energy usage was seen to be low until around 165kts, then to rise steeply, thus a linear cost from 160kts to 175kts was not reasonable. For this reason, the model was given the ability to switch between different (fuel/energy) cost models, one of which is to support a piece-wise linear model rather than linear model, allowing fuel burn and energy usage to be specified for arbitrary speed sample points, and perform a linear interpolation between these sample points. This is done through two functions, to return the fuel and energy usage, respectively, per minute for any speed, so that it is easy to provide arbitrary functions at a later date.

As a minor implementation note: due to the late date at which these needs were identified, incorporation of the changes into the model was performed with a requirement for minimal changes, so both of these need to be turned on manually in the code when the aircraft model needs them – in contrast with most features which are controlled by input data.

This project has received funding from Innovate UK under UKRI’s Future Flight Challenge Fund.
The grant agreement number is 74829

5.6 Costs and Income Calculations

All costs and revenues are calculated based on the input data, detailed previously in this report, and the output of the simulation. While the simulation can run for any given number of days, all cost/revenue calculations are done for one year.

A spreadsheet has been designed to show all the induced costs and generated revenues from the different system components. The file spreadsheet has all the input data embedded and connects to the output of the simulation to get the data needed for global cost and revenue calculation. The main purpose of the spreadsheet is to abstract the analysis of the results from the running of the simulation, so that it is simple to investigate the effects of changing costs parameters without having to re-run the simulation itself.

To give an idea of the detailed information available to the users, the file has nine sheets in total:

1. **"cost_revenue_related_data":** Parameter data used by the other sheets. It contains the aircraft market value (for the lease cost) in USD, then conversion rate to convert costs given in USD to GBP; the total fuel and power provided by the airports and total on-stand duration, which is calculated based on the "airport_logs" sheet. It also contains the number of days in the schedule; total flight hours; total number of aircraft, pilots, passengers, flights, freight flights and ticket revenue are all given by the simulation outputs.

2. **"Costs Revenues":** Shows the global cost related to aircraft, pilots, engineering, airports, passengers and other costs. Each one of these is extracted from its corresponding sheet. This sheet also provides the total revenues which are the sum of the ticket revenues, provided by the simulation, and cargo revenues calculated based on the cargo revenues for a single freight flight and total number of this type of flights provided by the simulation.

3. **"Aircraft"**: Include all the costs related to aircraft used to tun the flight schedule. Each one of the costs is given per aircraft and per period, which could a month, a year or five years. The cost per year per aircraft is firstly calculated then the cost for one year for all the aircraft.

4. **"Pilots"**: Contains costs related to the pilots: their salaries, training, travel expenses and other ones. Each one of these costs is calculated per year and per pilot. Then the total cost per year for all the pilots.

5. **"Engineering"**: Engineering-related costs are given either per flight hour or per cycle. The costs induced by the simulation is then calculated for the

This project has received funding from Innovate UK under UKRI's Future Flight Challenge Fund.
The grant agreement number is 74829

simulation, then for one year, depending on the number of flight hours and flights operated.

6. **"Airports"**: All these costs are given per flight, except the fuel (in Litres) and power (in kWh) sold to the airline. Each one of these costs are calculated for the simulation then for one year.

7. **"Passengers"**: Booking fees and war insurance are the passenger-related costs. They are given per passenger and are calculated for the schedule then for one year.

8. **"others":** Includes customer services, operations department and EU261 cost. These are provided per month and calculated per year.

9. **"airport_logs":** The last sheet is a copy of one of the csv files provided by the airport model. It contains information about arrival and departure delays, on-stand durations then fuel and energy sold for the airline.

This project has received funding from Innovate UK under UKRI's Future Flight Challenge Fund.
The grant agreement number is 74829

6 Evaluations Performed

This section discusses the different configuration options that have been considered, then presents the various results and the conclusions which can be drawn from them.

6.1 Configuration options

The designed simulation model is very configurable, allowing for many different options using different aircraft technologies, under different airline operations, with different airport scenarios and different fuel sources. The simulation has been used to evaluate several different configurations, as shown below.

6.1.1 Aircraft configurations

Two different aircraft models have been used for this evaluation, with settings specified by partners – the Twin-Otter of the present, and the Eco Otter as a future hybrid-electric aircraft.

Characteristics for the Twin Otter were provided by Loganair, based upon its operational use in their schedules. Characteristics for a conceptual 'Eco Otter' were provided by Ampaire, based upon the results from their work on hybrid-electric aircraft in other work packages in this project and the requirements of the evaluation. It should be noted that some of these parameters are expected to be adjusted/tuned according to operational studies and that potential customer requirements would also inform these design decisions. In particular, we note that the climb rate provided for the Eco Otter model used in this investigation was significantly lower than for the Twin Otter, however this could be adjusted in designs if operational requirements were shown to need this. We note that the lower climb rate did not cause a problem for these evaluations, and would allow significantly lower engine power/weight/fuel burn.

Considering the characteristics in the air, the hybrid-electric aircraft is designed with a smaller, much more fuel efficient, engine. Fuel burn is therefore significantly lower than for the existing twin otter, however in some cases (particularly for the take-off and climb, as well as for any journeys where a high flight speed is required to achieve the schedule) the electric propulsion is also used.

The passenger capacity is assumed to be identical for the two aircraft, with a maximum passenger capacity of 19 persons.

The climb rate is slightly lower for the hybrid aircraft, resulting in a slightly longer climb phase, which in some cases could increase flight time slightly. However, the top speed is 180kts to compensate (this being very energy hungry if needed), compared with the top practical speed of 175 which is used for the twin otter.

This project has received funding from Innovate UK under UKRI's Future Flight Challenge Fund. The grant agreement number is 74829

The simulation allows various charging policies for the hybrid aircraft, including the originally envisaged idea of charging in the air using excess power from the engines. The policy to charge to full on the ground was used in each case since the charging time was not usually a limiting factor for the ground times and this enables engine sizes to be kept smaller, for lower fuel burns. It will be observed that one of the major differences between the results for the twin otter and Eco Otter models is the large reduction in fuel burn which is predicted.

It was assumed that the aircraft were designed for simplified freight handling, e.g. using Unit Load Devices (ULDs). As such unloading/loading was assumed to be fast – 5 minutes for unloading and another 5 minutes for loading. In contrast, existing handling is more likely to take around 15 minutes for unloading or loading, so 30 minutes in total. Most experiments were executed with a 5 minutes freight loading and 5 minutes freight unloading time, but some were executed with 15 minutes for each operation, to indicate the timings required for current operations without automated handling.

6.1.2 Flight schedule configurations

One of the main purposes of the project is to evaluate the airline's flight schedules to see whether they are profitable under the new regional market, how many delays there would be using automated ground handling technologies and what the costs are using this new type of hybrid-electric aircraft.

The initial schedule ('Loganair' in the experiment section) was thus manually designed by Loganair considering all of the future concepts included in this project. This initial schedule includes high-frequency short-range flights between pairs of airports with a very short turnaround duration on the ground, with coordinated ground times to allow passenger transfers.

Some of the assumptions considered in this schedule would rarely happen within the current technology at the airports – e.g. the freight handling times. It was observed when running the 'Loganair' schedule that a number of flights were late – either due to flight times being slightly longer than expected, or the airport model with various configuration requiring longer ground times and failing to have aircraft take off by the scheduled times. Even small delays then accumulate through the schedule.

To make a comparison with existing scenarios running at the presented airports, an expanded schedule ('Expanded') has been automatically generated based on the 'Loganair' Schedule by increasing the turnaround time between consecutive flights performed by the same aircraft and by removing all flights which cannot be finished on time by the end of the day, while making sure that the FDPs of flight crew are always respected. This was done by splitting the day into three conceptual periods, spreading out the start times of aircraft after 7am by 50% (e.g. an 8am flight becomes 8:30, 9 am becomes 10am, etc, so that flights are still on the ground at the same times for transfer passengers), allocating pilots

This project has received funding from Innovate UK under UKRI's Future Flight Challenge Fund.
The grant agreement number is 74829

to flights so that they start and end the duty period at the same airport (noting that this means that aircraft will do so as well) and moving flights which cross over the boundary between periods into the next period, cancelling flights which cannot occur by the end of the last period. This results in a reduced frequency schedule with much more slack to account for delays, fewer flights per pilot than the original schedule and longer transfer times, but the same connectivity as the original schedule.

The two schedules were evaluated: the first from Loganair, designed with a new type of aircraft in mind which could have a very low ground time, and the second expanded one for comparison.

It should be noted that the expanded schedule has fewer flights per year, fewer flights per pilot, but the same salary costs and time-based costs (e.g. annual maintenance costs) as the original schedule, so should be less cost effective – the revenue dropping by more than the costs. It should also be noted that this was provided for comparison only, to get an idea of the sort of benefits which might be achievable under the current operating processes. The designed system evaluates schedules, it does not produce them – the expanded schedule was therefore produced using a custom algorithm designed for that purpose and applied as a pre-process stage to generate a new schedule from the existing one. The expanded schedule was provided primarily to avoid requiring partners to provide alternative schedules which would work well for each individual configuration, which would result in the evaluation of many different schedules. In this way a comparison can be made against just two schedules. It is expected that a hand-crafted schedule should be able to be more effective, perhaps only adding slack where needed for the desired connectivity of the schedule.

6.1.3 Airport configurations

In line with suggestions from Loganair, the airport models used for the experiments were assumed to have dedicated resources available for each aircraft, given the tight turnaround needed for the schedule provided.

The normal operating model for airport models is to have a duration for each operation, such as unloading passengers, recharging, refuelling, etc, as well as a taxi-in time from the runway and a taxi-out time back to the runway. All of these are parameters which were set for a ‘generic’ airport for this ‘proof of concept’ analysis, but could be set to specific values for airports if needed.

It is known that unloading and loading freight can be time-consuming for the Twin Otter, and was estimated by an airport at 15 minutes per operation. Some comparisons were executed using these values, but the provided schedules were built on the basis that this difficulty was resolved by the time these schedules were to go into operation. The majority of experiments therefore use a default time of 5 minutes for unloading freight and 5 minutes for loading freight, reducing the ground time considerably for freight flights.

This project has received funding from Innovate UK under UKRI’s Future Flight Challenge Fund.
The grant agreement number is 74829

It is assumed that in future these schedules would be able to be performed with much shorter overall ground times, optimising operations to ensure fast turnaround and fast connections. For this reason some experiments were also performed to consider the effects of shorter fixed turnaround times (excluding taxi-in/out times) of 10, 15 and 20 minutes, to see the effects of ground times upon schedules.

6.1.4 Demand modelling and freight flights

As discussed in Sections 4 and 5, a demand model was built to estimate the passenger numbers for flights, for example to consider the population sizes around the source and destination airports with the existing modes of transport. This was designed to be closer to the current passenger numbers for flights, but if this concept takes off it is envisaged that demand would be much higher than the level currently predicted by the model, and it is likely that the parameters used in the demand model for these experiments underestimate the real demand. To observe the effects, alternative values for demands were also considered, increasing the current estimated numbers by 25%, 50%, 100% and 200%. These are labelled 'Demand x 1.25', 'Demand x 1.50' etc in the results.

It should be noted that flights are converted to freight if there is very low demand, with a view that the freight flights will reduce the loss of the airline from running the flights, but not be as profitable as passenger flights with larger passenger numbers, and that higher demand values will mean fewer flights are converted to freight. It should further be noted that this is assumed to be done in the schedule design stage, so the final schedule would show fewer flights available for passengers to book, and having freight space available to book at certain times, rather than flights being changed to freight depending upon number of tickets sold.

It is also considered that as demands become better known over time, the schedule would evolve to take all of the demands into account. To model this kind of behaviour, a number of experiments were also executed to consider the income and costs with different load percentages for passengers (rounded down) – fully loaded (19 passengers per flight, 100%), 90% load average, 75% load average, 50% load average and 25% load average, to see the effect of the average passenger load upon the cost-effectiveness of the flights. We note that this approach assumed that all flights remained as passenger flights (no conversion to freight flights) and that they all ran with that fixed number of passengers (that percentage of maximum, rounded down). In contrast to the demand modelling approach, which will have some routes with large numbers of passengers and some with few, this approach will assume that all routes are equally attractive, so could be considered to model an approach where only routes with a certain attractiveness were actually run by the airline, rather than running all routes and converting some flights on quiet routes to freight (reducing passenger flight frequency).

This project has received funding from Innovate UK under UKRI's Future Flight Challenge Fund.
The grant agreement number is 74829

6.1.5 Confidentiality

We note that some information will be confidential for partners, and all results in this report have been approved by the relevant data owners. For this reason, some breakdown detail may not be available here. Detailed results spreadsheets were generated for each configuration and headline results are shown in the results section.

6.2 Simulation Results

The subsections below discuss the test cases which were considered and the combinations of these which were used for each analysis, along with the results of that analysis.

6.2.1 Test cases and configuration naming

In total, 72 cases were tested in the experiments. These were divided into four sets of 18 experiments, investigating various changes for each of four configurations: either Twin Otter or Eco Otter with either the provided schedule (named 'Loganair') or an expanded schedule (named 'Expanded').

The baseline cases of interest are assumed to be either the normal demand or the full aircraft (100% demand) cases, and ground handling durations from the airport model, with 5 minute freight unloading and loading times. The 18 test configurations then consider variations from these baselines:

- 10 configurations for passenger demands, including:
  - 5 for calculated passenger demands using the demand model: with normal demand, +25%, +50%, +100%, and +200% demands, to get an idea of how benefits change with demand.
  - 5 for load percentages for passengers: with 100%, 90%, 75%, 50%, 25% load
- 2 configurations with 15-minute freight handling durations – one using the demand model and one for 100% demand.
- 6 configurations with fixed turnaround durations, of 10 minutes, 15 minutes or 20 minutes (3 configurations), for each of the normal demand model and the 100% demand case.

Results therefore consider: the effects of the aircraft type; the effects of changing the schedule; the effects of the demand; and the effects of different airport processes – either fixed ground times or increased freight handling times.

All of the results are labelled in the same manner according to four criteria which have been varied across the experiments.

Aircraft: Either **Eco Otter** (hybrid-electric) or **Twin Otter** (chemical engine) to denote the aircraft details which were used to obtain the results.

This project has received funding from Innovate UK under UKRI's Future Flight Challenge Fund.
The grant agreement number is 74829

Demand: One of the 10 different demand levels. Anything labelled '**Normal**' uses the demand modelling system and automatic conversion to freight flights for low demand flights, potentially with an increased (multiplied, e.g. x 1.25 to increase demands by 25%) demand on each route. Anything labelled '**Fixed**' assumes that all aircraft have this percentage occupancy (rounded down to a number of seats).

Schedule: Whether the results are for the provided schedule from '**Loganair'**, or for the automatically '**Expanded'** schedule which was produced to investigate the benefits/effects of longer ground times.

Ground Operations: there are two available models for ground operations, 'Normal' or 'Fixed' duration. Where the Normal operations are used there are two options for freight handling durations, either 5 mins or 15 mins for offloading and loading freight – labelled '**Normal, 5 min**' and '**Normal, 15 min**' respectively in the results. Where the fixed duration model is used, the freight operation durations are irrelevant, but various different durations for the entire ground operations have been investigated, shown as '**Fixed 10 min**', '**Fixed 15 min**', or '**Fixed 20 min**', respectively.

6.2.2 Schedule details

Table 6.1 presents the details of the schedules, showing total number of flights, number of passenger flights, and number of overnight freight flights for each schedule, along with the time of the first departure and the landing time of the last arrival at the end of the day – noting that these will be different aircraft, obviously, so that the day's schedule can be repeated the following day. These characteristics do not depend upon the test case executed.

Table 6.1: Schedule details

Schedule	**Loganair**	**Expanded**
Total Number Flights	2238	1497
# Passenger Flights	1802	1149
# Overnight freight	436	348
Start time first flight	00:00	00:45
End time last flight	24:45	23:20
	Fixed Load 100%	**Fixed Load 100%**
Number passengers	34,238	21,831
Passengers/flight	19	19
Total Ticket Revenue	2,421,969	1,542,692
	Demand Normal	**Demand Normal**
Number passengers	14,970	11,347
Passengers/flight	11.11	12.19
Total Ticket Revenue	£1,583,721	£1,177,021

This project has received funding from Innovate UK under UKRI's Future Flight Challenge Fund. The grant agreement number is 74829

It can be clearly seen from Table 6.1 that the expanded schedule involves the dropping of a considerable number of flights, with a consequent reduction in the maximum passenger numbers who could be accommodated.

Depending upon the number of passengers that are assumed to want to catch each flight, the actual passenger load and the ticket revenue will vary greatly.

The table shows the number of passengers that could be accommodated by each schedule if every flight were full, as well as the ticket revenue that this would bring in for the day of the schedule. It should be noted that it is easy to calculate the revenue for any other number of passengers by scaling this 100% figure. Similarly, the equivalent revenue for an increased ticket price everywhere by 10% is equally simple to calculate.

Finally, the bottom three rows of the table show the results from the demand model. It should be noted that the demand from the demand model is not equal across all flights, and in fact some flights will be converted to fright due to low demands: 455 for the Loganair schedule, and 218 for the expanded schedule, where there are fewer flights already. The results show the total number of passengers predicted by the demand model for the given schedule, which is significantly lower than full flights. It also shows the average number of passengers per passenger flight, and the overall ticket revenue – which is again much lower than with full aircraft, and lower for the expanded schedule than the original Loganair schedule.

6.2.3 Baseline results

The baseline results are shown in the table below. The columns of the table are as follows:

Aircraft: the type of the aircraft used by the airline.

Demand: 'normal' meaning the demand as estimated by the Demand Model; `Fixed 100%' meaning the demand matching the capacity of each flight.

Schedule: 'Loganair' is the schedule provided by Loganair; 'Expanded' is the schedule with the ground times increased to match the modelled ground times (this results in also decreasing the number of flights).

Total cost – the operational costs per year.

Total revenue – the total income through the tickets and freight.

Total profit – the difference between the total revenue and the total cost.

Energy, kWh: the amount of electric energy per day needed to charge the aircraft batteries for the propulsion system, while on the ground. Note that this

This project has received funding from Innovate UK under UKRI's Future Flight Challenge Fund.
The grant agreement number is 74829

excludes any other electrical systems which are not involved in the propulsion system, hence the assumption is a zero-value for this for the twin otter.

Fuel, litres – the amount of fuel used by the network per day.

Table 6.2: Baseline results for costs, revenue, profit, fuel and energy usage

Aircraft	Demand	Schedule	Total cost	Total revenue	Total profit	Energy, kWh	Fuel, litres
Eco Otter	Demand Normal	Expanded	£604 M	£535 M	-£69 M	6,320	171 K
Eco Otter	Fixed Load 100%	Expanded	£622 M	£641 M	£19 M	6,320	171 K
Twin Otter	Demand Normal	Expanded	£640 M	£535 M	-£106 M		334 K
Twin Otter	Fixed Load 100%	Expanded	£659 M	£641 M	-£18 M		334 K
Eco Otter	Demand Normal	Loganair	£636 M	£531 M	-£105 M	13,633	262 K
Eco Otter	Fixed Load 100%	Loganair	£660 M	£711 M	£51 M	13,650	262 K
Twin Otter	Demand Normal	Loganair	£753 M	£607 M	-£146 M		500 K
Twin Otter	Fixed Load 100%	Loganair	£780 M	£812 M	£32 M		500 K

The following results can be observed from the above data:

- The fuel burn of the Eco Otter is significantly lower than that of Twin Otter. For the twin otter this will be the entire power source, whereas for the Eco Otter electric power from a battery can also be used by the hybrid propulsion system whenever high power is needed, allowing for smaller engines and much lower fuel burns.
- The Eco Otter is more profitable, primarily due to the lower fuel burn. Table 6.2 has a £19M-£41M improvement in profit for the Eco Otter compared to the Twin Otter – which in one case (Fixed Load 100%, Expanded) moved it from a £18M loss to a £19M profit. It should be noted that the lower climb rate with the evaluation configuration can slightly increase flight time, however, Eco Otter still remains more profitable than the Twin Otter configuration.
- The 'Normal' demand, i.e., the demand produced by the demand model, is too low to be profitable for these configurations – but please see the consideration of ticket pricing, later, in Sections 6.2.4 and 7.1.10.
- The Expanded schedule has lower delays and is more 'achievable', however when the network is profitable it produces considerable smaller revenue, leading to smaller profits.

6.2.4 Effects of demand

The effects of the demand upon the predicted profit can be seen in Figure 6.1.

This project has received funding from Innovate UK under UKRI's Future Flight Challenge Fund.
The grant agreement number is 74829

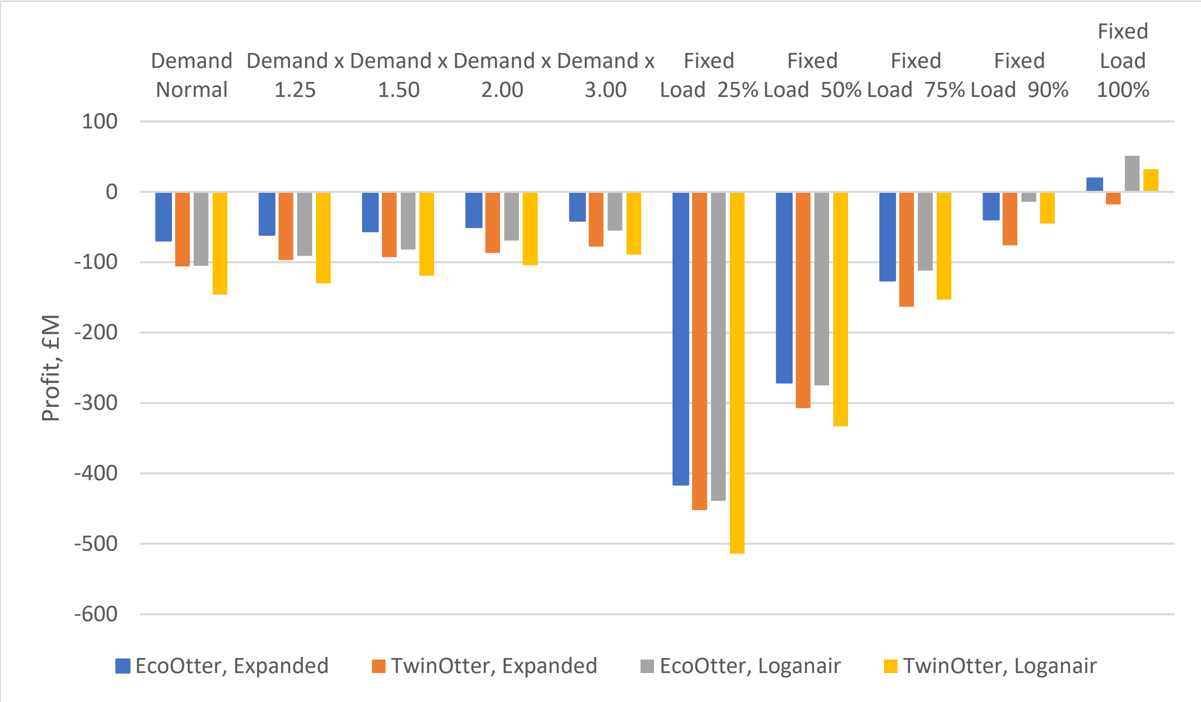

Figure 6.1: Profit for different demands with normal ticket prices

Two types of demand assumptions are considered: predicted demand based on the demand model (labelled 'normal') and fixed demand relative to the capacity of the aircraft/network (labelled 'fixed').
- 'Demand normal' is the default output of the demand model.
- 'Demand x X' is the predicted demand scaled up by the factor X.
- 'Fixed Load X%' creates demand that fills X% of the aircraft capacity in each flight.

As expected, the demand significantly affects the profitability of the network. The network under the tested configurations is profitable only if utilisation of all the flights is 100% (except for the Twin Otter Expanded schedule option, which was still unprofitable).
The highest profit is achieved by Eco Otter with the Loganair schedule, however this configuration leads to significant delays. A more realistic scenario is Eco Otter with the expanded schedule which is still profitable under 100% utilisation of flights.

The ticket pricing scheme was designed to be 'very competitive' against car or train journeys – matching the prices. However, the results above show that this may be a little hard to justify. To investigate the sensitivity of the profitability to the ticket pricing, calculations were performed with the ticket prices increased by 20%, and the results are shown in the following figure. In this scenario, the Eco

This project has received funding from Innovate UK under UKRI's Future Flight Challenge Fund.
The grant agreement number is 74829

Otter configuration is profitable even under the demand predicted by the demand model (although the twin otter is not). The network is also profitable with 90% utilisation of the capacity.

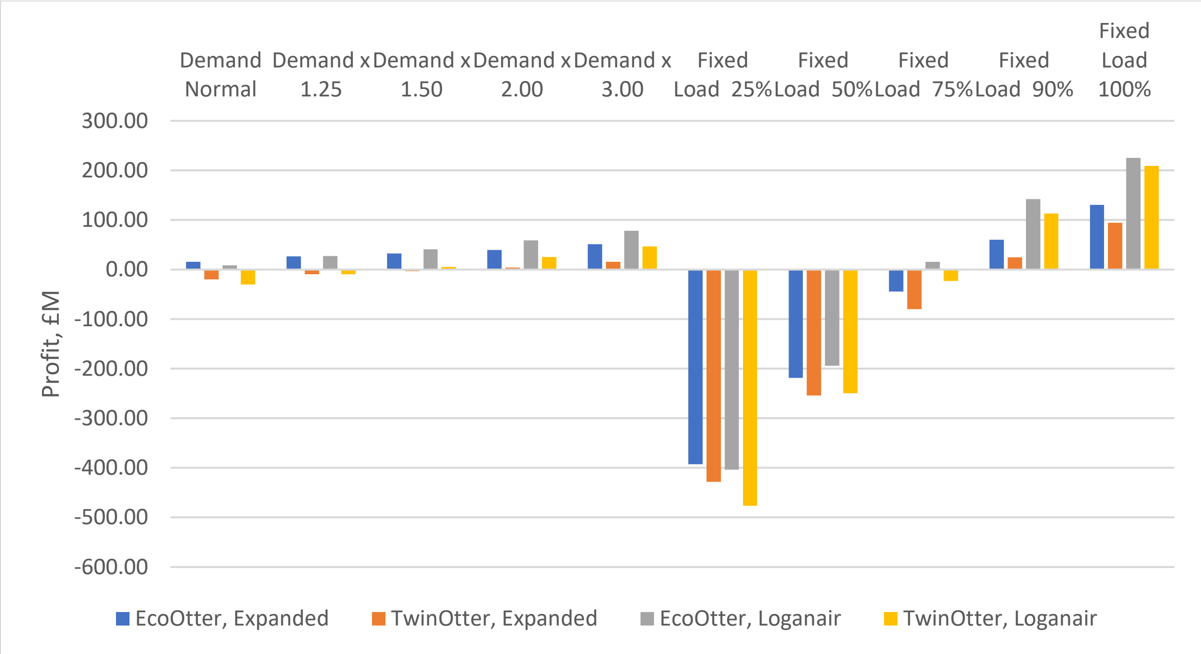

Figure 6.2: Profit for different demands with +20% ticket prices

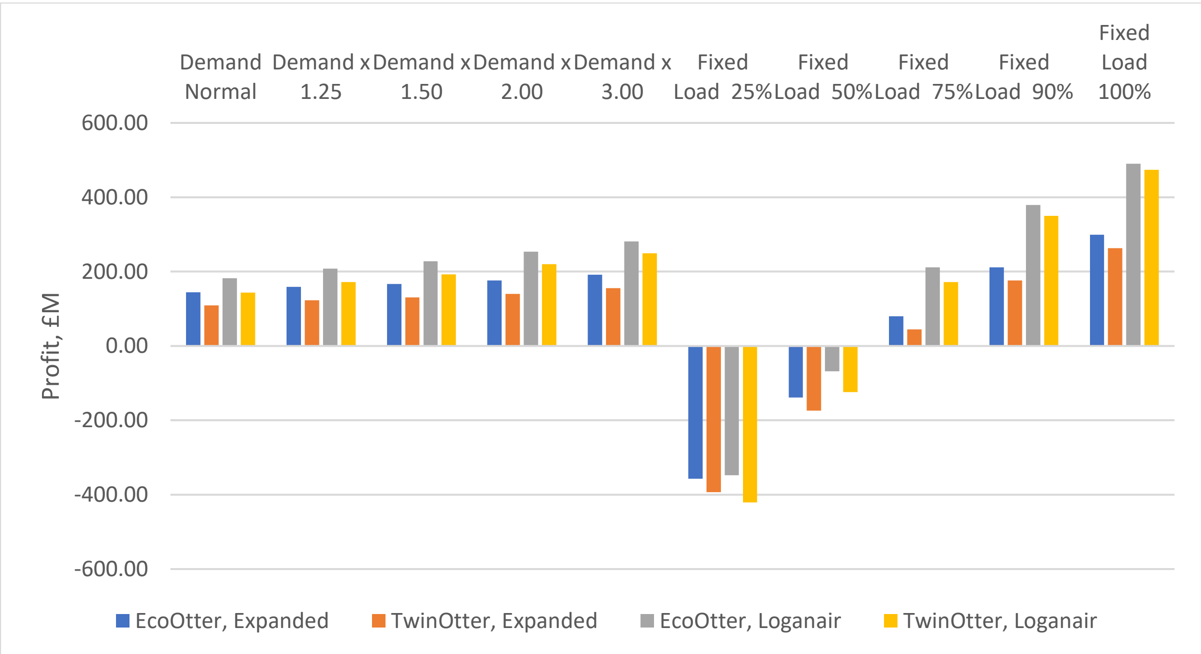

Figure 6.3: Profit for different demands with +50% ticket prices

This project has received funding from Innovate UK under UKRI's Future Flight Challenge Fund. The grant agreement number is 74829

We also calculated profits subject to a 50% increase of the ticket prices and the results are shown in Figure 6.3 on the previous page. It is clear that this kind of ticket price increase greatly improves the profitability of the schedule.

We conclude that to make the system profitable one should either increase the ticket prices or ensure very high demand – noting that these two factors are, or course, likely to be coupled such that increases in prices would likely reduce demand. It should not be unreasonable to expect customers to be willing to pay more if the journey can be made sufficiently fast and easy, but a full analysis of ticket pricing and demand is really needed for any future work.

6.2.5 Effects of ground operation durations

Below we report the effects of the ground time on the lateness of arrivals within the network. Specifically, we report how many flights are expected to have delayed arrivals for each configurations.

The following figures shows the results for the 'Normal' demand, i.e. the demand predicted by the demand model, and then the 100% demand assumption. In each case, 'Fixed X min' means that the turnaround time (excluding taxi time) is fixed to X minutes. 'Normal, X min' refers to the turnaround times estimated by the airport model and each of the freight loading/unloading operations being fixed to X minutes.

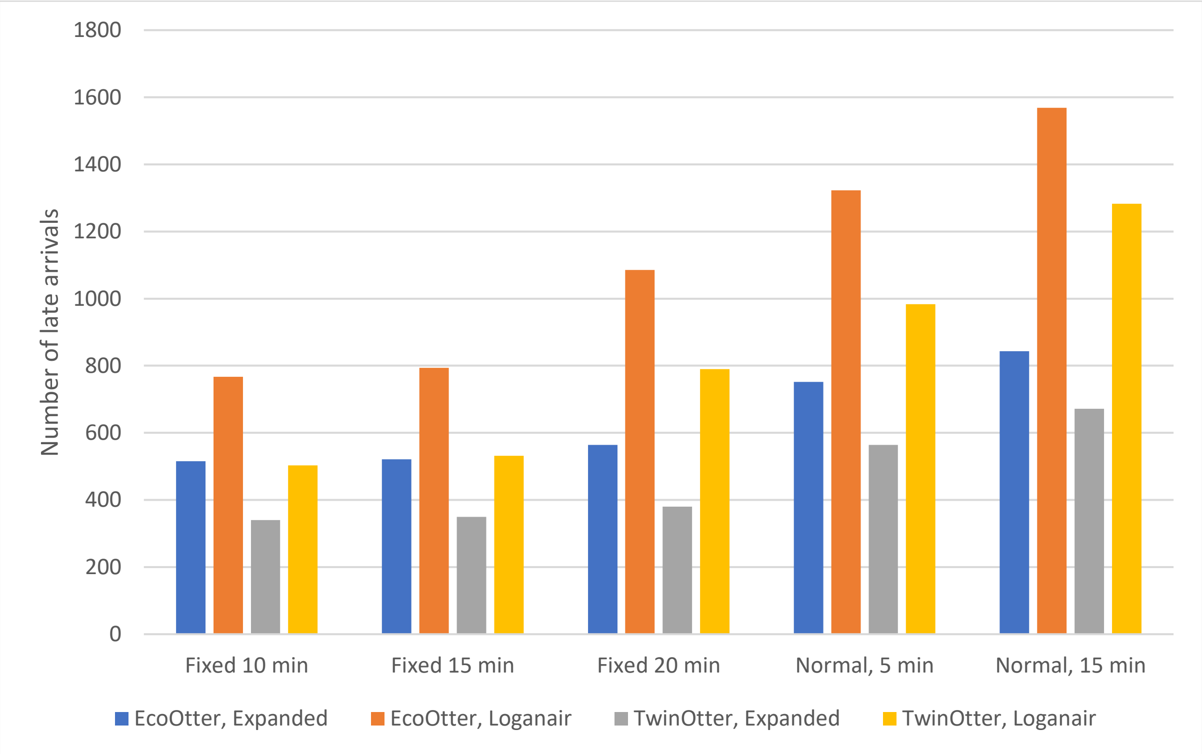

Figure 6.4: Number of late arrivals for ground operation assumptions – normal/predicted demand

This project has received funding from Innovate UK under UKRI's Future Flight Challenge Fund. The grant agreement number is 74829

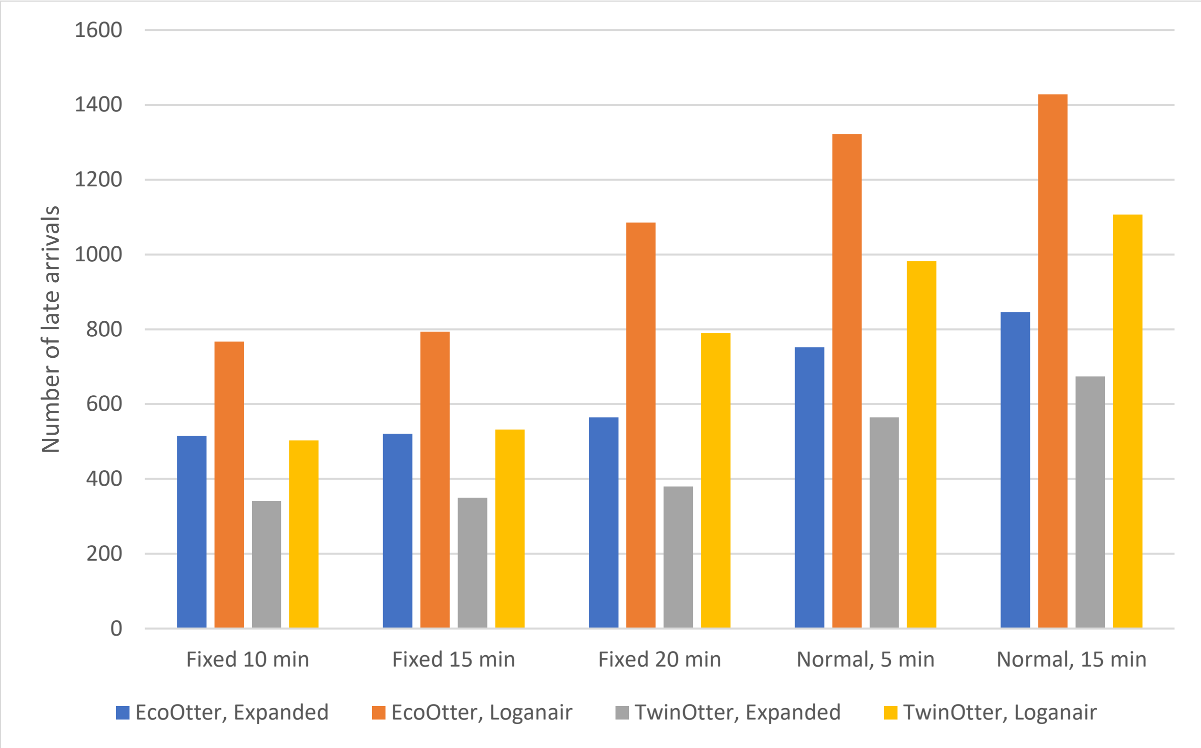

Figure 6.5: Number of late arrivals for ground operation assumptions – fixed 100% demand for all flights

Clearly, the number of delays significantly depends on the turnaround times. Eco Otter has more delays than Twin Otter due to the larger flight times. The Expanded schedule naturally has many fewer delays than the Loganair schedule.

The freight loading/unloading times also have significant effect on the delays, showing the importance of the automated handling.

When considering the results for full aircraft, Fixed 100% load demand. The results are similar, although absolute values are slightly lower in some cases, particularly for the runs with 15-minute freight loading/unloading operations. This is due to the fact that the Fixed 100% load runs assume that no passenger flights are converted into freight flights, whereas the demand model converts some low-passenger flights to freight, which under some configurations can require larger turnaround times.

Detailed simulation results are provided in Table 6.3. This shows that the number of late departures and arrivals ('Late dep' and 'Late arr' columns), the average delays values ('Avg dep delay, min' and 'Avg arr delay, min') and the duration of the day, i.e. the actual time required to complete a day schedule. If the day duration is above 24 hours, further delays will accumulate over multiple days (this accumulation is not a part of our simulation).

This project has received funding from Innovate UK under UKRI's Future Flight Challenge Fund. The grant agreement number is 74829

The ‘day duration’ column shows the reason for the generation of an expanded schedule – delays can accumulate through the schedule, such that some schedules would clearly not be repeatable the next day (where aircraft do not land in time to take off on time the next day) or on the same day (where delays mean that pilots would not actually complete their Flight Duty Period on time, and may, for example, be prevented from making their last flight (back home). It should be noted that it is not a strict cut-off at 24 hours, as some aircraft could be completing flights when others start their flights for the next day, allowing some overlap. Nevertheless, it is easy to see that many of the simulations with the Loganair schedule have durations which would be problematic, and some were tight even with fixed 10 minute ground times. The Expanded schedules fix this issue, however they still include many late departures and arrivals – it’s just that they have capacity to ‘catch up’ at other points in the schedule.

Comparing the number of late arrivals and number of late departures, and the average lateness of arrivals and departures can be interesting. If an aircraft which arrives on time has a long ground delay, it may take off late. Conversely, if the schedule allows enough time on the ground to absorb the late delay, even a late arrival could still take-off on time. Ground delays (more time needed than scheduled) tend to cause more late departures than arrivals, whereas in-flight delays (longer flight time needed than scheduled) can cause later arrivals than departures. These figures show that in general average delays are low, but are better for the expanded schedules.

Table 6.3: Full results for profitability, number of late departures and arrivals, average departure and arrival delays, and schedule duration, all configurations

Aircraft	Demand	Schedule	Ground time	Profit	Late dep	Late arr	Avg dep delay, min	Avg arr delay, min	Day duration
Eco Otter	Norm	Expanded	Fixed 10,	-£69 M	9	515	0.0	1.2	23:23
Eco Otter	Norm	Expanded	Fixed 15	-£69 M	43	521	0.2	1.4	23:23
Eco Otter	Norm	Expanded	Fixed 20	-£69 M	114	564	0.6	1.7	23:23
Eco Otter	Norm	Expanded	Normal, 5	-£69 M	745	752	2.3	3.2	23:23
Eco Otter	Norm	Expanded	Normal, 15	-£68 M	849	843	5.3	5.7	23:23
Eco Otter	Norm	Loganair	Fixed 10	-£101 M	105	767	0.7	1.9	25:30
Eco Otter	Norm	Loganair	Fixed 15	-£102 M	329	794	2.4	3.6	26:35
Eco Otter	Norm	Loganair	Fixed 20	-£100 M	772	1085	6.8	7.2	27:40
Eco Otter	Norm	Loganair	Normal, 5	-£105 M	1296	1323	18.7	18.9	31:08
Eco Otter	Norm	Loganair	Normal, 15	-£103 M	1549	1569	27.1	26.3	31:43
Twin Otter	Norm	Expanded	Fixed 10	-£105 M	9	340	0.0	0.9	23:21
Twin Otter	Norm	Expanded	Fixed 15	-£105 M	35	350	0.1	0.9	23:21
Twin Otter	Norm	Expanded	Fixed 20	-£105 M	93	380	0.5	1.2	23:21

This project has received funding from Innovate UK under UKRI’s Future Flight Challenge Fund.
The grant agreement number is [74829](74829)

Twin Otter	Norm	Expanded	Normal, 5	-£106 M	733	564	1.6	2.0	23:21
Twin Otter	Norm	Expanded	Normal, 15	-£105 M	845	672	4.6	4.5	23:21
Twin Otter	Norm	Loganair	Fixed 10	-£147 M	105	503	0.6	1.4	25:23
Twin Otter	Norm	Loganair	Fixed 15	-£146 M	263	532	2.2	3.0	26:28
Twin Otter	Norm	Loganair	Fixed 20	-£145 M	659	790	6.0	6.3	27:33
Twin Otter	Norm	Loganair	Normal, 5	-£146 M	1246	983	7.4	7.4	27:49
Twin Otter	Norm	Loganair	Normal, 15	-£143 M	1486	1283	15.9	15.0	28:20
Eco Otter	FL100%	Expanded	Fixed 10	£19 M	9	515	0.0	1.2	23:23
Eco Otter	FL100%	Expanded	Fixed 15	£19 M	43	521	0.2	1.4	23:23
Eco Otter	FL100%	Expanded	Fixed 20	-£7 M	114	564	0.6	1.7	23:23
Eco Otter	FL100%	Expanded	Normal, 5	£19 M	745	752	2.3	3.2	23:23
Eco Otter	FL100%	Expanded	Normal, 15	£20 M	832	846	5.1	5.5	23:23
Eco Otter	FL100%	Loganair	Fixed 10	£90 M	105	767	0.7	1.9	25:30
Eco Otter	FL100%	Loganair	Fixed 15	£82 M	329	794	2.4	3.6	26:35
Eco Otter	FL100%	Loganair	Fixed 20	£78 M	772	1085	6.8	7.2	27:40
Eco Otter	FL100%	Loganair	Normal, 5	£51 M	1297	1322	18.6	18.8	31:08
Eco Otter	FL100%	Loganair	Normal, 15	£52 M	1397	1428	20.9	20.7	31:35
Twin Otter	FL100%	Expanded	Fixed 10	-£17 M	9	340	0.0	0.9	23:21
Twin Otter	FL100%	Expanded	Fixed 15	-£17 M	35	350	0.1	0.9	23:21
Twin Otter	FL100%	Expanded	Fixed 20	-£17 M	93	380	0.5	1.2	23:21
Twin Otter	FL100%	Expanded	Normal, 5	-£18 M	733	564	1.6	2.0	23:21
Twin Otter	FL100%	Expanded	Normal, 15	-£17 M	829	674	4.4	4.4	23:21
Twin Otter	FL100%	Loganair	Fixed 10	£45 M	105	503	0.6	1.4	25:23
Twin Otter	FL100%	Loganair	Fixed 15	£38 M	263	532	2.2	3.0	26:28
Twin Otter	FL100%	Loganair	Fixed 20	£33 M	659	790	6.0	6.3	27:33
Twin Otter	FL100%	Loganair	Normal, 5	£32 M	1246	983	7.4	7.4	27:49
Twin Otter	FL100%	Loganair	Normal, 15	£28 M	1357	1107	10.3	9.9	28:41

This project has received funding from Innovate UK under UKRI's Future Flight Challenge Fund.
The grant agreement number is 74829

7 Conclusions, suggestions and potential extensions

This section first considers the conclusions which arise from the results in the previous section, and their predicted effects upon the potential for utilising hybrid-electric aircraft for regional aircraft schedules, then considers the flexibility of the system and what it can be used to evaluate. Following this, future applications are considered, before summarising the lessons learned and some open questions.

7.1 Conclusions from the results

The part of the 2Zero project considered by this report involved the development of a system to evaluate flight schedules and to use this system to determine the cost-effectiveness of potential new approaches for utilising the hybrid-electric aircraft in the project. By running the different cases considered in Section 6, the developed system has been found to be effective for this purpose, but with some provisos mentioned here – primarily a warning that its accuracy is dependent upon the accuracy of the input data, but secondarily a note that the ideal way to use the system is collaboratively rather than stand-alone – which will be discussed below.

This section considers the lessons learned from the data, experiments and results that were obtained – some of which go beyond the conclusions that can be drawn from the summary data provided in Section 6, and rely upon insights from underlying data breakdowns.

7.1.1 Effects and characteristics of the evaluated schedule

Loganair designed a point-to-point schedule to allow great connectivity across the UK and enable fast transit across the majority of the UK in just two flight legs – taking less than 3 hours in total. This is an impressive task, however it puts various requirements upon the aircraft and airport which may not currently be entirely practical:

1. The aircraft must be capable of an airspeed of at least 175knots, and ideally more, to make some legs of the schedule within the desired transit time. It was observed as well that the climb speed can be more important than expected (the hybrid-electric aircraft has a lower climb speed than the Twin Otter), but turned out not to be problematic in these results. Aircraft manufacturers should be aware that reducing the max speed of aircraft may limit their ability for use in this kind of schedule, not only reducing the length of flight legs (noting the lack of refreshment facilities on aircraft) but also making schedules more affected by wind speed and direction. Climb speeds may also be increasingly important for these kinds

This project has received funding from Innovate UK under UKRI's Future Flight Challenge Fund. The grant agreement number is 74829

of schedules – especially where routes may traverse mountain ranges, as in Scotland.

2. Like many airline schedules, the schedule is designed to facilitate connections with minimal delays for passengers – so that a passenger can make a two-leg journey across most of the country with as much ease as possible. This results in very 'spikey' schedules, with high numbers of aircraft arriving at similar times[2], then taking off at similar times, to maximise the number of passenger connections that can be made between the aircraft on the ground at the same time. However, these aircraft on the ground need the same types of resources, and, of particular interest where charging is concerned, all may need power to charge their internal batteries at the same time. This in turn results in high demands for resources, power and staff at some times, then very low demands at others[3]. This is not ideal for efficient resource usage, or splitting the use of resources between different users, and would be expected to incur additional costs in provision at the airport. The 'spikiness' of schedules will probably be one of the challenges that has to be faced for these kinds of schedules – effectively almost every airport acts as a 'hub' airport, facilitating connections, with all of the consequent problems.

3. The airports need to enable extremely short turn-arounds for aircraft when necessary, to minimise the 'wasted' time for aircraft on the ground and help keep a fast connectivity for passengers making connections. The resource usage problem (of point 2) makes this particularly difficult, since

[2] The spiky nature of the departure schedule is often more obvious than for the arrival schedule, since aircraft schedules refer to departure times, and aircraft which arrive early will still have to wait for these scheduled times to leave. This is both sensible and necessary, since passengers will plan to arrive for their scheduled departure time, so waiting for passengers to arrive will prevent an aircraft taking off early anyway. Thus, there will often be multiple aircraft loading passengers at the same time if their scheduled departure times are similar. In contrast, an aircraft with a shorter flight time than planned will often land early, and may be able to unload, refuel, etc, earlier. This 'wave' structure (where aircraft arrive and depart in waves) for airline schedules is already common at hub airports, where it is similarly used to maximise the number of passenger connections. It is also worth noting that, in this analysis, even where flight durations may be dissimilar, if aircraft have reasonably tight flight times in the schedule, and similar scheduled landing times, the feature of the airspace model to accelerate to get back on schedule where necessary will have the effect of increasing the number of aircraft that arrive at their scheduled arrival time, and hence increase the spikiness.

[3] It should be noted that the contention for resources will depend upon how long resources are needed for. For example, if refuelling takes 3 minutes, two aircraft would require the facilities at the same time if they arrived within 3 minutes of each other, whereas if recharging takes 8 minutes, then they would require them if they arrived within 8 minutes of each other. Also, as noted in [2], this is more commonly a problem with resources needed for departures than arrivals, due to the fact that departures often cannot perform some operations (e.g., passenger loading) early even when ready.

This project has received funding from Innovate UK under UKRI's Future Flight Challenge Fund. The grant agreement number is 74829

the tight turnaround gives no slack for flexibility to 'take it in turns' to use the resources while aircraft are on the ground. Similarly, it may be important to keep taxi times (from runway to stands and back again) low, as this is otherwise wasted ground time. With any move to this kind of high frequency regional flights, increased consideration will need to be taken into how to keep the ground time low at airports.

The results in Section 6 should make clear the importance of understanding these elements of the input schedule, as they greatly affected which constraints the schedule would work under. E.g. that the schedule needed the fast turnarounds on the ground to avoid accumulating large delays became very obvious from the simulation.

7.1.2 Freight Flights and automatic conversion

Overnight flights are used to carry freight in the supplied schedule – ensuring good utilisation of the aircraft which will be available at that time. There is also an aim to run flights as freight when there is expected to be lower passenger demand – keeping the aircraft in use and ideally covering the flight costs, even if not much more than that, but ensuring a regular schedule, relocating aircraft appropriately to enable other (e.g. return flight) passengers to be picked up. This schedule usage assumes that aircraft can be quickly converted to freight from passengers and back again, and that there is a means of quickly loading and unloading freight to ensure fast turnarounds for freight. Neither of these assumptions are probably true for the Twin Otter at the moment. Configurations were tested to considering both current and potential future freight operations – in the future concept this kind of conversion would be fast and easy, whereas in current operations the loading/unloading of freight into these aircraft can be time-consuming.

A default low freight loading and unloading time of 5 minutes was used in the schedule evaluation, since it was observed that making the provided 'future' schedule work really requires this, despite the fact that airport discussions implied that with a Twin Otter at present it would probably take closer to 15 minutes for each operation (30 minutes total for unloading and loading) at present. The results for the comparison between fixed (unrealistic) ground times and those with 5 and 15-minute loading and unloading time (see Tables 6.4 and 6.5 for the effects on the number of late arrivals, or Table 6.6 for further results) show the effects that this ground time and loading/unloading time would have. The expanded schedule mitigates a lot of this, with its longer ground times, as would be expected, but not fully. Therefore, it is clear from the simulation that some kind of fast/automated freight loading/unloading would be beneficial if the schedule is to be used as envisaged, and having other process changes to keep total ground time down would be even more beneficial…

This project has received funding from Innovate UK under UKRI's Future Flight Challenge Fund.
The grant agreement number is 74829

7.1.3 Considering fixed, short ground times

Results showed that the provided schedule was only really feasible using extremely short ground times – much shorter than currently reasonable at airports, and probably shorter even than would be needed with current airport operations and a 5 minute freight handling time. This is not impossible for future airports, and Loganair have some ideas about this already, but would need some work to facilitate, and may perhaps imply that dedicated facilities which could be placed close to the runways, may be a better way forward for these schedules than the positioning of terminals at larger airports. The expanded schedule is more feasible, with longer ground times, but makes less profit, as discussed.

7.1.4 Ground resource usage

Although no such results were presented in this report, the simulation monitors which resources are needed by which aircraft at which time. Due to the nature of the schedule and the tight turnarounds, it was assumed for these investigations that resources such as recharging points, refuelling facilities, etc, were unlimited. This is obviously not the case at real airports, at least for the moment. As previously mentioned, the structure of the flight schedule is such that a high number of aircraft will arrive at an airport at the same time, then depart at the same time, maximising opportunities for passenger connections, but also meaning that the ground resources are all needed by different aircraft at the same time. This is far from ideal for ensuring sufficient resource capacity.

One interesting side-effect of delayed schedules was observed – as aircraft were more delayed, and these delays were not usually evenly distributed between aircraft, so the arrivals naturally became staggered, which staggered the resource usage by aircraft, and in many cases actually reduced the number of simultaneously required resources. Obviously, delays are still not desirable, but this does emphasize the down-side of having waves of arrivals and departures (which maximise connections, as discussed), upon the usage of the airport resources. Perhaps some compromise would be possible if slightly higher ground times could be attained, and this may be worth considering.

7.1.5 Automatically generating an alternative schedule

An alternative for many of the problems has been to utilise a less tight schedule, which was automatically generated in this case. The alternative schedule ‘stretches’ time out, so that the gaps between flights (planned ground times) are greater, so longer timespans are permitted for the flights both in the air and on the ground. This has been observed to be a more feasible schedule, even with current limitations, however it had two major adverse effects: 1) the ticket revenue can be seen in Section 6 to be considerably lower; 2) with fewer flights being delayed, they keep to the schedule, so (in the absence of resource limit constraints in this evaluation), the simultaneous usage of different types of resources was actually observed to be higher, as considered in Section 7.1.4.

This project has received funding from Innovate UK under UKRI’s Future Flight Challenge Fund.
The grant agreement number is 74829

A potentially better alternative to this 'alternative schedule' approach will be discussed in Section 7.2.2.

7.1.6 Effects of airspeed and climb speed

The effects of both airspeed and climb speed were seen in the variations between aircraft results, although these will not be so evident from the top-level results reported in this report. The Eco Otter data allowed for a slightly higher maximum speed than the Twin Otter preferences (180 kts vs 175 kts), but had a slower climb speed (more focus on being eco-friendly and keeping the engine size down). Initially we had expected that this would mean that Eco Otter flights need take no longer than Twin Otter flights, but this was not always the case in this model.

The airspace model works by simplifying the problem to consider a number of stages, using the climb speed to determine the duration of the first/climb stage, the descent speed to determine the duration of the last/descent stage, and allocating the remaining time to the cruise stage. There is an assumption that climb already uses the max power, so that it is the cruise stage where aircraft can speed up to make up time. Slower climb speeds result in longer climb stages, hence comparably less time in the cruise stage to make up for delays. In some cases, this resulted in longer flight durations when it was necessary to speed up, when the climb was a higher proportion of the flight time (higher flight altitudes with lower flight distances).

We believe that lower climb speeds may be a useful means for keeping engine sizes of the future more environmentally friendly, however we highlight that the increase in the proportion of the flight which is a climb may become increasingly important and schedules will need to explicitly consider these effects. There is a clear trade-off here between lower fuel burns and faster flights.

7.1.7 Trade off – airtime vs ground time recharging

With the provided design, flying faster needs considerably more electrical battery energy, which then needs to be replaced through recharging on the ground. Even apart from the additional 'wear' on batteries and other electrical components, this results in additional demands for ground facilities, and for the aircraft to spend longer on the ground charging.

The developed system allows holistic effects of different charging policies to be evaluated – such as to only charge when it will be needed, or at every opportunity, and to only charge with the necessary energy for the next flight or to fully charge each time. In these experiments the aircraft was assumed to be kept fully charged before each flight, allowing maximal flexibility, however in a number of cases the speed-up to 'get to the destination airport on time' was actually inadvisable due to the additional charging time that the aircraft spent on the ground because of this. i.e. the charging time to replace the energy could

This project has received funding from Innovate UK under UKRI's Future Flight Challenge Fund.
The grant agreement number is 74829

exceed the time reduction from using the energy to speed up the flight. More complex policies for when to speed up in flight, which took into account the resulting recharging times, would probably be worth considering in further work, to explicitly capture and utilise this trade-off.

7.1.8 Demand levels and income

It should be clear from the profit data that economically running the schedule with the parameters provided to us for these operations will be demanding, and will require high passenger demands.

Given the configurations used in these evaluations, aircraft need to be virtually fully occupied in order for schedules to make a profit. This seems unlikely, to say the least. However, we note that ticket prices were set to low amounts (comparable to the cheapest of car or train), greatly limiting the ticket revenue. Table 6.2 showed that even a small increase in ticket price, which would seem reasonable given the lower travel times, made the Eco Otter profitable, although not the Twin Otter. We also note that the hybrid-electric aircraft are predicted to have significantly lower maintenance costs, which are not considered here, see Section 7.1.9.

As previously mentioned, full demand modelling is really needed to be able to utilise this model to the fullest extent. This project did not initially include any demand modelling, aiming to use demand figures from a partner, however this is not really realistic since demands would be expected to change for a new aircraft concept, and existing flight schedules (and demand data) cover only a small part of the geography covered by the provided evaluation schedule. Therefore, this turned out to be more problematic than originally envisaged.

The developed demand model is relatively simple, and predicts relatively low demands, having been calibrated against a different type of schedule for existing passenger flights. Because of this, simulation results are presented not only from the developed demand estimation model, but also the results for augmented demands (since predicted demands are probably unrealistically low for a network giving such good coverage and benefits to passengers) and results which consider various fixed demand percentages for all flights.

The importance of these comparisons is, perhaps, hidden by the fact that almost all configurations make a loss under the given cost and ticket price assumptions. For this reason, alternative augmented ticket price effects were also considered and discussed in Section 6.

The conversion of mostly-empty flights to freight would have an effect upon demand, which should be modelled by any enhanced demand model. The current model assumes a certain demand in each time period and a spill-over to move a proportion of the demand to adjacent flights if flights are cancelled. It should be acknowledged that one of the aims of the provided schedule was to

This project has received funding from Innovate UK under UKRI's Future Flight Challenge Fund.
The grant agreement number is 74829

provide frequent flights (around every 75-90 minutes), to encourage passengers to be use the network. Conversion of passenger flights to freight flights has adverse effects upon passenger flight frequency, possibly affecting demand.

It should be noted that the simulation has been designed to allow any alternative demand and costing model to be included – just by setting ticket prices and predicted demands for each flight in the input schedule.

7.1.9 Maintenance and fuel/energy costs

Maintenance costs are a large proportion of overall running costs. Aircraft manufacturers may be able to greatly facilitate adoption of new aircraft types for these kind of regional point-to-point schedules by moving to technologies which have lower maintenance costs. Indeed, maintenance costs has been a good driver for moving to more-electric aircraft at the larger scale.

These experiments assumed that the Eco Otter had the same maintenance costs as the Twin Otter – which is expected to be pessimistic for new hybrid-electric aircraft. Moving to a smaller, simpler engine type, augmented by an electric motor for the high-power flight stages, such as the climb, is expected to reduce maintenance costs longer term, and hence make the flight schedules more cost-effective.

The results for the fuel and energy usage clearly show the financial benefits of the move to a future hybrid engine design, even though they are not a huge percentage of the total costs. These benefits should hopefully help to facilitate any future move towards using this kind of aircraft for regional point-to-point flights.

7.1.10 Ticket Pricing should be reasonable

It should be clear from the demand data that economically running the schedule with the parameters provided to us for these operations will be more than a little demanding. Much of the reason for this is the desire to compete on price with existing transport and their costs. This keeps ticket prices down to what in some cases are very low levels, where competitive train fares may be low. Results for higher ticket prices showed that it became much more feasible to make a profitable schedule in those cases, and that ticket pricing will be extremely important if this sort of regional flight schedule is to be adopted.

7.2 What the system could be used for

The model and simulation that have been developed are powerful and flexible, however, their use is limited by the data which is available. To consider what the model could be used to evaluate, it is useful to consider what data is used by the model – since a common use for such a model is for 'what if' analysis – to consider the effects on different results of changing some data.

The most useful purpose of the model is likely to be to compare two different scenarios and evaluate the differences. E.g., if fuel burn could be reduced to this amount, how much effect would it have upon schedules and costs, or what would be the effect if we lowered the nominal maximum flight speed (and fuel burn) for the execution of our schedules?

To see the variety of things which would be evaluated, we consider in the following subsections the different types of data which are used, and that any of this data could be varied to evaluate the effects of changes. For any evaluation, values are needed for every data item, with multiple values being needed for the data items whose effects are to be evaluated. In addition, one or more evaluation flight schedules will always be needed. The model could then be used to consider the effects of any one of the following areas, or any combination of these areas, resulting in a huge number of potential scenarios which could be considered:

7.2.1 Flight schedule evaluation

Perhaps the most obvious use of a calibrated system is for determining whether flight schedules would work, identifying the 'pinch points' and tight turnarounds, and understanding where delays are likely to occur. Running the simulation results in predicted take-off and landing times for each aircraft in the flight schedule – so delays and problem points can be easily identified.

7.2.2 Interactive flight schedule design

It should be noted that the evaluations here used static schedules – two input schedules were used and the outputs showed how effective they were. This is not the best way to use the system for schedule design, despite being an easy way to evaluate different configurations. An interactive flight schedule design approach would be much more desirable and powerful. To do this, first determine the appropriate settings, and run the prospective schedule through the simulation system. The output will indicate which aircraft will be on time and which will be late (and by how much). Ideally, the input schedule would then be adjusted to take this into account (by hand), to eventually build the best schedule incrementally – basically the input schedule is iteratively refined, by delaying schedule times where flights are predicted to be late, until eventually the schedule is achievable as specified.

This approach works on the assumption that, rather than taking a single pass, observing that certain aircraft will be late, and updating the times for those aircraft, the schedule would be considered holistically and updated in that way. For example, if aircraft A arrives late, is there some way to recover that time, or is the schedule too tight? Can anything else be modified to make up for that? Do

This project has received funding from Innovate UK under UKRI's Future Flight Challenge Fund.
The grant agreement number is 74829

other connecting flights also need to have delayed timings, to allow connections to still take place, etc. The answers to some of these questions may be obvious, but for others they may not – such as whether to accept the problems of delaying other flights (including consequent knock-on delays from these) vs making the connections. This approach also gives the designer an opportunity to go back to the drawing board upon discovering that something just would not really work and is too tight – perhaps changing configurations to accommodate it (e.g. purchasing more recharging facility usage at an airport to avoid delays – see Section 7.2.3).

7.2.3 Airport data

In the supplied evaluations, the airport data was not considered in detail due to the low required ground times, which would require considerable changes to the airports, however the model/simulation is able to consider these elements, whereby each aircraft has to perform a number of operations, each operations has certain resource requirements, and resources are shared between aircraft. The model/simulation could therefore be used for a number of different evaluations beyond those considered here.

At the higher level, the effect of different overall ground times can be considered.

At a more detailed level, the effect of varying the expected durations of different ground operations can be considered – such as was seen in the results from varying the freight unloading and loading operation durations. For example, what if the recharging rate were halved, or doubled?

Unlimited resources were assumed in the experiments performed here (although the number of each which was used at any time is accessible, as is the start/end time of each flight using each resource). If resources are not available, aircraft start to get delayed waiting for their turn at the restricted resource. An important use for the system could therefore be to evaluate the effects of different resource availability at airports – for example, considering how many charging points are needed optimally, and what are the effects of losing one, two, or more? There may, therefore, be a trade-off for airlines from wanting to have no queue/delay vs having to pay for increased resource availability, and airports could start to pass on more of the costs of increasing resource availability if there was a clearer understanding of the longer term as well as immediate costs to airlines of any lack.

A related question could also be considered for questions such as recharge rate – is it better to have fewer but larger/faster recharging points or more but slower ones? Each of these are parameters than can be considered in the model.

This project has received funding from Innovate UK under UKRI’s Future Flight Challenge Fund.
The grant agreement number is 74829

7.2.4 Recharging/refuelling policies

Different recharging or refuelling policies can also be considered within the airport operations, but are highlighted separately as these are really airline decisions about the usage of facilities rather than the presence or effectiveness of facilities.

Important questions could be asked about when aircraft should refuel/recharge, and by how much when they do? This will be extremely important if the costs of operations vary across airports – e.g. fuel is more expensive at some airports than others, and/or recharging facilities are not available or are extremely costly at some airports. It is also quite possible that the policies will vary for the two propulsion methods for hybrid aircraft - carrying unnecessary fuel is obviously an excess weight, but this is not the case for battery charge, and charging points may not be (at least initially) available at all airports, whereas refuelling facilities may be more common. Fuel/energy costs may vary greatly between airports. A schedule may have greater slack ground time at some airports than at others, so charging may be more feasible without causing delays. With so many considerations, being able to evaluate (combinations of) policies could have huge value for operators.

7.2.5 Aircraft and flight model data

The model could help aircraft manufacturers or potential purchasers of aircraft evaluate the effects of design changes. For example, if the climb rate were changed, the descent rate, the cruise speed, the fuel burns/energy usage rates, the fuel/energy capacity etc.

Fuel and energy capacity can also have a large effect upon the flight schedule performance – perhaps implying more frequent recharging/refuelling needs, or limiting the maximum flight speed (to keep usage lower).

The assumption of three flight stages was already challenged late in the project, adding in additional discrete take-off and landing phases on top of the climb and descent phases, and it would be possible to further modify the flight model if desired, perhaps to evaluate the effects of continuous descent vs descending in steps, or to allow optional speed increases for part of the descent. The effects of such changes could then be easily determined by the simulation.

7.2.6 Airspace data and flight model

One important element is the ability to utilise arbitrary congestion data and consider the effects of delays. As airspace gets busier and more congested due to potential unmanned aerial vehicle usage, or even increased regional electric/hybrid-electric flights, congestion may become an increasingly important factor to consider. The congestion information could be useful for determining flight routes, or for choice of airports/airfields to use, ideally to automatically

This project has received funding from Innovate UK under UKRI's Future Flight Challenge Fund.
The grant agreement number is 74829

choose less congested airports without trading off proximity to customers. Realistically though, perhaps its most important usage will be to see the effects upon planned schedules if expected congestion delays increase over time, perhaps as a set of what-if scenarios.

7.2.7 Airline data

Schedules are currently built to consider pilot flight duty periods. Before running the schedule, pilots are allocated to flights and flights are joined together into journeys which start and end at the same place within a flight duty period. Interesting analysis could include the effects of additional constraints on where pilots can go, which connections they can make, or the effects of changes to flight duty period rules.

7.2.8 Propulsion Methods

The developed model is very generic. It supports two fuel types, but does not limit what these are. For each, it has a recharge/refuel rate, a usage in different flight stages at different speeds, and a cost per unit. There is no problem with using alternative fuel sources in the model, and as such the developed model could be used for many other types of aircraft, including:

- Chemical fuel only – either one type or multiple types
- Hybrid chemical/battery
- Battery only
- Hydrogen fuel cell only
- Hybrid hydrogen fuel cell + battery
- Hybrid fuel cell + chemical engine

It is also possible to investigate the trade-offs within fuel types – for example to consider the costs/benefits of different fuels with different costs and burn rates.

In other words, in this evaluation we compared a pure chemical engine aircraft against a hybrid aircraft, primarily due to the interests of project partners and availability of data, however it would be equally possible to compare against an aircraft using hydrogen for fuel, or a hybrid hydrogen+battery aircraft.

7.3 Future work with partners

Given the wide variety of different scenarios considered above, we would be interested in working with future partners with interests in any of these areas, particularly where partners may be using different fuel types, who may be looking for comparisons to see the effects with a specified schedule. We see a number of potential future partners who may find value from using this simulation:

This project has received funding from Innovate UK under UKRI's Future Flight Challenge Fund. The grant agreement number is 74829

Aircraft users or manufacturers: to see the differences that aircraft designs could have in practice, or how to adjust schedules or facility availability to accommodate these changes. As discussed in Section 7.2.5, it would also be possible to consider the effects of altering the flight characteristics of an aircraft to see the effects of reducing engine size, fuel burn, climb rate, etc.

Airports: to see the likely effects of different schedule concepts, the potential benefits of different amounts of facility provision, and potentially to see the costs/value to airlines of provision, perhaps to understand what costs could acceptably be passed on.

Airlines: to evaluate potential schedules, perhaps for potential aircraft types.

In summary, as should be apparent from Section 7.2, the model is abstract enough to be able to cope with a huge variety of configurations, and therefore to have potential value to many potential partners beyond those is this consortium, and we look forward to future collaborations to use and further extend the model and simulation.

As was identified previously, we feel that utilising the system for realistic evaluation of costs will require a good demand prediction system – to estimate the number of passengers who would take any flight. We feel that this is an important area for future work, which could be done via providing predicted demands as input data rather than needing to be integrated into the simulation itself - I.e. for each flight in the schedule, something should predict a number of passengers. We would be very interested in working with someone with interests in this area.

7.4 Lessons learned

In addition to what can be learned from the results (Section 7.1), the observations about the flexibility of the developed system (Section 7.2) and the potential applications beyond this project (Section 7.3), some lessons could be learned from the development process itself – many of which may be obvious in hindsight, but we hope some are of value for others to consider:

Importance of accurate and timely data: the importance of good data in a timely manner was clear to all working on this project. In some cases, data had to be generated (e.g. generating a demand estimation model with partners), in some cases it had to be augmented (e.g. generating an alternative flight schedule with larger gaps to cope with accumulated delays), and in some cases assumptions were made (e.g. that maintenance costs for Eco Otter would match those of Twin Otter, in the absence of any data). Some take-away messages are therefore perhaps: **be ready to generate data** that does not become available, allow time for liaison with partners to more accurately do so; **expect that there**

This project has received funding from Innovate UK under UKRI's Future Flight Challenge Fund.
The grant agreement number is 74829

may be extra data needed that you didn't even think you would need, either because no stakeholder thought it would be needed, or it was expected to be available. The contingency plan of generating the data appears to have worked well in this project, but it involved significant input from partners/experts to get reasonable values for each data item.

The problem of predicting the future: here we wished to look at how a system would be used in future, not now. This meant that various elements had to use 'future' values – estimating what may be achievable in the future, rather than now. This is somewhat in conflict with the aim of getting a 'realistic' simulation, and calibrating it – since good calibration data is by its very nature 'current' rather than future. This is particularly 'interesting' when different partners may have different data or data from different assumptions. For example, the airports in this project could tell us what current operations involve, and how long operations take, whereas airlines may need them to do things 'more efficiently' in future, to make concepts work. Some future discussions between stakeholders should be assumed here. We hope that we achieved the right balance in this work, for example showing results for current ground handling times vs what seems to be hoped for in future. Being able to do this kind of comparison is a valuable feature of the simulation, perhaps allowing stakeholders to make better-informed decisions about where things may need to change.

The difficulties of considering the future and the present: there was another interesting problem in the project related to current or future data – namely the question of whether we model current airport operations, and consider how things could be made to work now, or whether we model future operations, and assume that things will be changed to work. We attempted, somewhat, to do both in this project, to parameterise everything, and to provide comparisons of operations in Section 6 (e.g. fixed ground handling times, or differing freight handling durations). Care has to be taken, however. For example, when considering recharging speeds, with the speed of improvement of battery technology, is the aim to target current speeds, or expected speeds in 5 years? In the end, we built a system which can cope with any of the configurations but provided it the appropriate data for specific configurations. Data management became an increasingly important element of the project – ensuring that the correct data is used for each execution can be non-trivial as the amount of data and number of combinations increase. In this project, a hierarchical data structure was used, whereby you select which of the airport/airspace/aircraft data configurations to use, and each of those configurations is responsible for identifying the actual data to use to achieve that configuration.

Manage the co-development of different parts of the project: here the project was both problematic but interesting because some partners were developing elements at the same time as we developed the simulation which included those elements. For example, aircraft characteristics and battery parameters were never going to be known precisely until quite late in the project. There are common software engineering techniques for managing change within a project, however there are limits to what can be done. In this case we aimed for a modular and parameterised design, so that changes to one element would have minimal effects to other components and that many changes could be accommodated through parameter changes rather than model changes. We aimed for as flexible a design as practical, trying to consider at the start what partners *may* want to do, and ensure that it would be possible – the effects of this being seen in Section 7.2 in the wide variety of things that can now be considered, so the effort was not wasted. In most cases this meant that catering to final configurations from partner requirements was just a case of changing data, however, as mentioned, in at least one case an assumption had to be revisited right at the end of the project and a 'quick fix' had to be added.

It's impossible to predict everything – some changes from increased understanding of partners were not able to be handled by the flexibility we built into the system; for example the late changes to the aircraft model in terms of increased number of flight stages and the changed energy usage characteristics were a consequence of updates from the partner involved and required a number of changes to the part of the simulation which dealt with those factors. In this case, these changes had some fundamental implications upon assumptions and calculations, however we could restrict the changes to within the airspace model component. A decision has to be made in these cases about whether the change is necessary (in this case we decided yes) and how it can be achieved (in this case the priority was upon minimising changes and risks of breakages, but it still required significant re-testing to be performed).

Enable your partners to understand your requirements: again, this is not 'rocket science' but in any collaborative project, consider it part of your role to help your partners to understand what you need. It was much clearer to our partners why we needed certain information when we had a demonstrable system and they could see where data was going. In some cases, understanding why we wanted some information also gave them insights into what we actually meant by some data, so that they could appropriately measure or calculate the requested information. Linked with this, don't underestimate the value of a good visualisation. When evaluating results, we tended to look at the numbers internally within the research group, but it was very clear from discussions with partners, however, that graphically showing them the aircraft flying around the UK, from one airport to the next, was much easier to comprehend than a series of landing and take-off time figures.

Airport operations may need to change and be more streamlined for regional connectivity. It should be apparent from the results, that streamlined ground operations will be key for getting frequent connectivity with the point-to-point aircraft concept. Either consideration of how to do this needs to be taken

This project has received funding from Innovate UK under UKRI's Future Flight Challenge Fund.
The grant agreement number is 74829

or alternative higher ground-time concepts need to be utilised. Importantly, the simulation can be used to evaluate the effects of these choices.

Aircraft may need to change and be more streamlined for regional connectivity. A similar argument applies for aircraft. Under this concept, where flights which are not needed can operate as freight instead, to appropriately relocate aircraft and pilots ready for the next flight, it is clear that fast freight handling will be needed, as will the ability to change quickly from passengers to freight and back again. It is not clear that aircraft designs are currently at that point. If this is not achievable then schedules would need to be modified accordingly – either to fly a mostly empty aircraft anyway (ignoring conversion to freight), or to strategically cancel sets of (rather than individual) flights to reduce costs. The simulation should accurately model the characteristics of the aircraft, and the ability of the simulation here to do so could help in future aircraft design.

7.5 Open questions

We end this section by considering some open questions which we did not cover, but which will need answers.

Will passengers be willing to pay for faster connectivity? There is a huge question mark over the financial viability of the schedule design, however this is primarily because the ticket prices were set to match cheap train tickets and car journeys. It may be more realistic to consider that, since these networks can get passengers to their destinations much faster than alternative means, it should be viable to increase ticket prices, making these operations much more cost effective. With the move towards electric cars with their (currently) lower ranges between recharges, alternative means of transport may become increasingly attractive – especially for businesses.

How much will passengers pay to be green? This is a key question, since it may be feasible, especially if routes could be run entirely electric, that passengers may be willing to pay more for the lower emissions of their journey, as well as for the speed of the journey.

How should demand be estimated? How should ticket prices be set? Demand predictions and ticket pricing are big areas which are of intense interest already to airlines. Due to the very short duration of this project, the developed system evades these questions somewhat by developing a system to which a pricing system and/or demand prediction system could be attached as a part of the input data, and a simple process for each was developed within this project. There are still open questions about how these should be done and we would look forward to working with experts in these areas – within or outside of airlines.

This project has received funding from Innovate UK under UKRI's Future Flight Challenge Fund.
The grant agreement number is 74829

8 Final comments

This project considered producing a model and simulation to evaluate the use of smaller aircraft (particularly hybrid-electric aircraft) to run point-to-point schedules for regional connectivity within the UK.

A computer system (model and simulation) was developed which was able to do this, taking an input schedule, putting it through the various parts of the simulation and getting an output predicting what each aircraft would do, what time it would take off and land for each flight, whether pilots would return to base within their flight duty periods, etc, and providing final values for fuel burn, energy usage, lateness, etc.

A flight schedule was provided for this purpose by a project partner and was evaluated using the simulation. This schedule makes various innovative demands upon different stakeholders within the air transportation system, e.g., requiring very short ground times, efficient airport operations and low taxi times. The results from the system showed the various delays that would be likely to occur if assumptions built into the schedule are not met – indicating that the simulation was able to evaluate the schedules appropriately.

An alternative flight schedule was automatically produced, to provide larger slack in the schedule. This schedule has fewer flights, so may be much less attractive to passengers (schedule frequency), as well as being less profitable, due to having fewer flights/lower ticket income, but removed the issues of short ground times, making the schedule achievable with current ground operations.

The hybrid-electric aircraft were observed to reduce fuel costs considerably, being predicted to improve profitability even without considering any likely reduction in maintenance costs. The environmental benefits may be even more important, of course. It should also be noted, again, however, that the model for the Eco Otter which was used was not necessarily optimal for this, that some parameters are subject to tuning, and that the results and profitability of schedules will depend upon the aircraft characteristics of any final aircraft used to fly these schedules.

In summary, the results of the project imply that a regional point-to-point schedule could work for ensuring good regional connectivity, with the ability to get across most of the UK in only one or two connections – producing an option for extremely fast travel for passengers. Hybrid-electric aircraft would seem to be a more effective means to achieve this than conventional aircraft, perhaps providing a steppingstone to achieving these schedules. Importantly, there are some questions about profitability, given the assumptions made about ticket pricing here, and a consideration of demand modelling and ticket pricing is an important future area for consideration.

This project has received funding from Innovate UK under UKRI's Future Flight Challenge Fund.
The grant agreement number is 74829