Developing discrete optimization methods for a holistic airline scheduling algorithm with sustainability as an objective function

Naoise Barry[1] and Charles Stuart[2]
Trinity College Dublin, the University of Dublin, Ireland

This work outlines the investigation of discrete optimization methods employed to solve an airline fleet optimization and scheduling problem for a small network of airports operating B737-8200 MAX aircraft and approximately 5,000 passengers. The model developed in this study uses real-world inputs for technology, operations, policy and strategy to account for the current and future options available to the aviation sector. The model provides insight into the objective sustainability, assessed simultaneously from an economic and environmental perspective, of the decarbonization design space for the purpose of determining optimal pathways. Through the comparative analysis of the three different linear programming, discrete optimization methods in the solving of this type of problem, an understanding of the limitations and requirements of such a model are established. It was found that solutions from the model were highly solver and scenario dependent. Problems of larger size had difficulty obtaining optimal solutions within tolerance limits regardless of the linear programming solver used but resulting schedules did react in a deterministic way to scenario changes indicating viability as an assessment tool in the evaluation of decarbonization pathways.

I. Nomenclature

BTS	=	Bureau of Transportation Statistics
CO_2	=	Carbon Dioxide
CO	=	Combinatorial Optimisation
CSP	=	Constraint Satisfiability Programming
ETS	=	Emissions Trading System
EU	=	European Union
LH_2	=	Liquid Hydrogen
LP	=	Linear Programming
MILP	=	Mixed Integer Linear Programming
IATA	=	International Air Transport Association
ICAO	=	International Civil Aviation Organisation
NLR	=	Royal Netherlands Aerospace Centre
TTW	=	Tank-to-Wake

II. Introduction

The aviation sector must decarbonize by 2050 to meet the resolution passed by members of the International Air Transport Association (IATA) at their 77th IATA Annual General Meeting in 2021 [1]. Similar agreements and commitments by members of the sector exist, such as “Destination 2050”, an industry alliance made roadmap developed by the Royal Netherlands Aerospace Centre (NLR) [2] and the International Civil Aviation Council’s (ICAO) “Safe Skies – Sustainable Future” specialized agency [3]. How this decarbonization of the sector will come to fruition is not yet clear nor is there a widely agreed upon pathway to this net-zero future. Engine manufacturers,

[1] PhD Researcher - Dept. Mechanical, Manufacturing & Biomedical Engineering (nabarry@tcd.ie)
[2] Assistant Professor - Dept. Mechanical, Manufacturing & Biomedical Engineering (stuartch@tcd.ie)

airframers, component manufacturers and other technology stakeholders are targeting more efficient systems that reduce emissions and energy consumption with examples of Pratt & Whitney's GTF Engine which decreases fuel burn up to 20% per flight relative to prior generation engines [4] and Boeing's 777X which will deliver 10% lower fuel use as well as 10% lower operating costs relative to its competition [5]. Members of the fuel and utility suppliers are exploring alternative forms of energy to aid in the decarbonization effort with Airbus' ZEROe project looking to use hydrogen fuel cells to power the aircraft electrically rather than fossil kerosene [6]. Policy makers and regulatory bodies are seeking to establish rules and practices which reduce emissions like the European Union (EU) based, emissions trading system (ETS) where emissions' penalties are applied to those generating emissions [7]. Airlines and aircraft operators are taking steps to mitigate their emissions, with examples like Ryanair's fleet renewal that offer 16% less fuel burn with 4% more guests [8] and easyJet's development of a descent profile optimization software which reduces CO_2 emissions associated with their operation [9]. With all of these developments in technologies, fuels, operations and policies, it is obvious that there is no single solution that will unilaterally result in the decarbonization of this hard to abate sector. As such, studies that can look into the options the sector has and will have in the future, to explore how impactful they will be for various deployment scenarios is a critical undertaking. These types of investigations can provide the much-needed clarity in solution efficacy and aid in the decarbonization decision-making process as a preliminary, exploratory tool.

Discrete optimization methods are well suited to this style of problem: finding the optimal set of decisions that need to be taken from a variety of options with their own associated advantages and disadvantages, such that a goal or target be achieved under logistical and feasibility constraints. Small problems with a limited number of options have trivial solutions with obvious objectively good versus bad choices. However, when these problems increase in size and the complexity of the available options produces confounding variables and counterintuitive relationships, this is when discrete optimization methods prove invaluable. These methods follow an objective evaluation of the various options' pros and cons through mathematical relations and assess the numerous permutations of decisions that can be made to determine solution sets that are feasible, and in some cases, optimal. For the aviation sector, these decision sets can comprise of the best aircraft and energy carriers to develop, implement and scale, where and how often they should fly, incorporating the capacity to determine the impact of out-of-sector investment and market based-measures on the airline business model, among a variety of additional options. Opting for one aircraft or fuel over another at a given instance may have downstream impacts that are unexpected and result in objectively worse outcomes than if other upstream decisions had been taken. Hence, modelling the aviation sector, with its various technology, fuel, operational, policy and strategy options in the framework of a discrete optimization decision-making problem where the blended impact of different decisions is assessed, is crucial for guiding the decarbonization pathway.

Krömer et al. [10] investigated the ability for airline scheduling models to solve for sustainability, assessed as fuel burn minimization. This work used a mixed integer, quadratically constrained problem to demonstrate that hypothetical airlines may reduce the number of flights in their operations, therefore decreasing total fuel burn, without a substantial sacrifice to profitability. This work concluded that approximately 10% of fuel burn can be saved through these discrete optimization methods being employed to solve the planning and scheduling aspect at the expense of between 0.7% and 1.92% operating profit, for their scenarios. Numerous scenarios were run with varying problem sizes. The smallest scenario of five airports had 6,150 optimization variables and 48,399 constraints. The largest scenario investigated was a twelve airport, 43,923 and 339,317 variables and constraints, respectively. Jalalian et al. [11] explored the task of servicing passengers versus reducing CO_2 emissions from an airline's perspective. This work used a bi-objective, mixed-integer non-linear programming model to study the interdependencies between targeting CO_2 emission reductions while maintaining or increasing the service of passengers to determine a shared optima between both objectives. The optimum scenario identified by the model was capable of facilitating reductions in CO_2 emissions of 11% and increased passenger service by 31%, for a real-world case study. This work has two problem scales, small and large. The small-scale problem has between one and two days of operations, between five and seven aircraft and between 15 and 25 flights. The large-scale problems have between two to three days, 20 to 40 aircrafts, and 100 to 300 flights. Parsa et al. [12] investigated the design of airport networks using discrete optimization methods to determine emission mitigation strategies through hub-and-spoke style networks. The study used a multi-objective, mixed-integer linear programming method to solve for optimal hub locations and target reduced emission production from the network as a whole. This work uses a hub-and-spoke network of 25 airports, with up to 360 aircraft of varying types facilitating over 8 million passengers in this multi-objective, network design optimization study.

These studies within the literature demonstrate a trend whereby researchers are exploring the decarbonization of the aviation sector from a more holistic, multi-faceted approach. It reflects the importance of having a well-rounded perspective on what leads to emissions within the sector and what mitigating measures can be developed to help reduce these emissions outside of technological advancement alone. Furthermore, the aggressive decarbonization targets require revolutionary and disruptive changes to be made to the aviation sector, hence models and studies which evaluate the level of impact different implementation strategies and scenarios comprised of various future technologies, energy carriers, policies, etc. are key to determining the optimal decarbonization pathway ahead.

There is a gap in the literature for this study's unique problem formulation – a scheduling problem that has decision variables for aircraft type, route, passenger allocation and fuel at take-off with a holistic sustainability objective function. This type of problem formulation allows the exploration of trade-offs between all options in a given scenario but enacts change and choice on a per aircraft, per flight basis across a fleet. The model developed in this work takes the form of a decision-making algorithm using discrete optimization methods in relation to the aviation sector, specifically in the area of exploring airline future fleet composition and scheduling, with the goal of maximizing sustainability. The model utilizes inputs relevant to the sector such as aircraft technology, airport operations and logistics and the market's policies and regulations to fully characterize the design space from which optimal, sustainable decarbonization pathways may be established. This work focuses on benchmarking different discrete optimization methods in their application to a problem of this nature and ultimately, seeks to establish a methodology for generating a tool that enables objective, scientifically based and data-driven decisions for the decarbonization of the aviation sector.

III.Methodology

A. Context & Background

A high-level overview of the modelling framework is shown in Figure 1. This flowchart comprises all of the fundamental components of the model and is agnostic to the type of discrete optimization method used such that a fair basis for comparison can be used across the selected methods during the benchmarking study. Prior work by Barry et al. [13, 14] established an initial version of this type of model using mixed integer linear programming (MILP) methods from MATLAB's optimization toolbox [15]. The purpose of these studies was to gain an initial understanding of the fundamental requirements for a model of this kind. The MATLAB versions of these models contained the majority of the high-level functionality to be discussed later in this work but were described for a scope-limited, MATLAB based model whereas this study is designed to generate a solver-agnostic, scalable platform. The focus of this work is on the development of the methodology such that an airline's schedule may be modelled under user defined scenarios regarding airport network, fleet of aircraft, passenger traffic and determine the optimal sustainability performance for these inputs, subject to in and out-of-sector constraints. While familiarity with the prior work is not essential, it may help readers with providing additional context and explanation regarding the development of the methods and focus of the research explored in this current study.

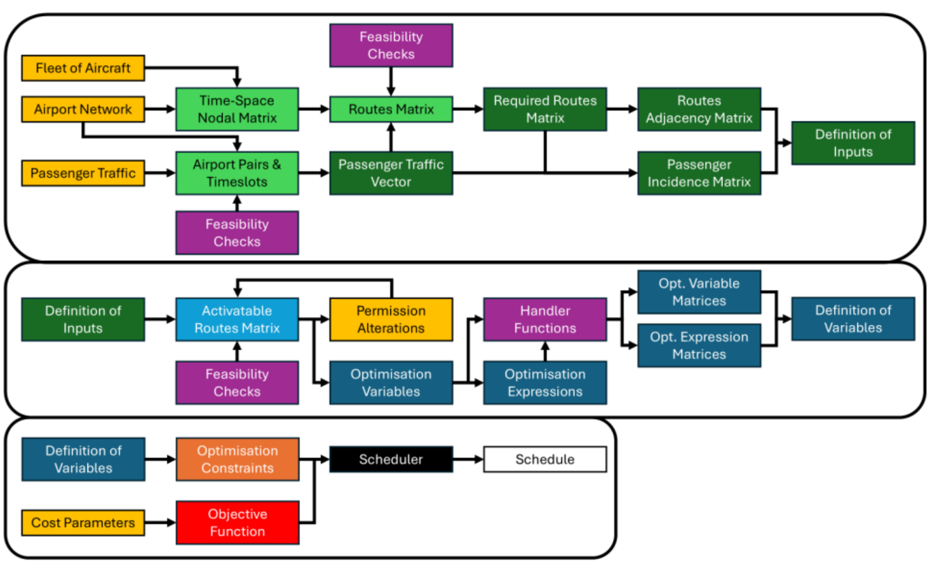

Fig. 1: Fundamental decision-making model overview

The data inputs to this model are the aircraft, airport network and passenger traffic information. These pass through a converter, which handles the manipulation of these raw data into accessible digital versions ready for storage and further action within the model. The aircraft data is converted into what is referred to as a “digital hangar” that acts as a lookup table for all feasible operating points that the aircraft may be expected to complete, i.e every physically permissible combination of passenger number, range and fuel level at take-off for a given aircraft. For each of these operating points, a fuel burn value is calculated and stored which creates a tri-variate contour performance map (graduated by fuel consumption) for the given aircraft where the axes are passenger number, range (ground distance to fly) and fuel level at take-off, respectively as seen in Figure 2.

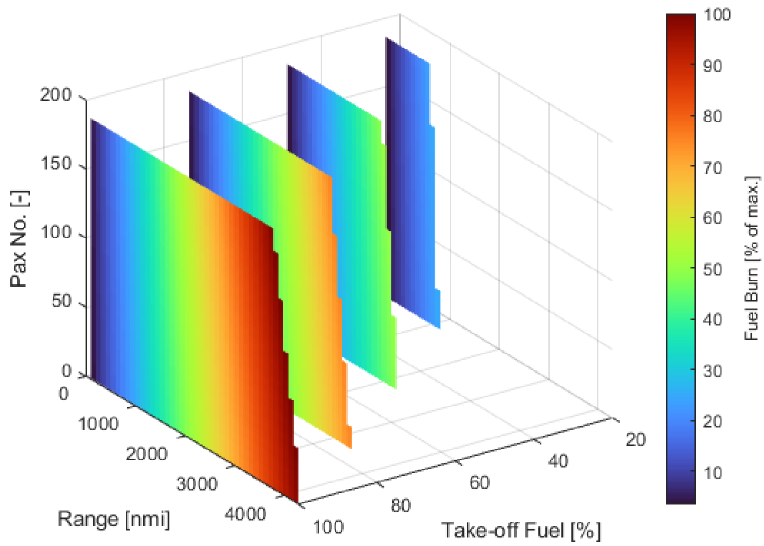

Fig. 2: Sample performance map within digital hangar functionality

The method for determining fuel burn values is not discussed directly in this study as this work focuses on exploring the utilization of different aircraft assuming accurate performance models have previously been established and verified. Therefore, the specific methodology for how these aircraft performance data were generated and validated will not be included here but can be found in Barry et al. [13, 14], which is influenced by the work of Gallagher et al. [16, 17]. This methodology achieved fuel burn predictions with a mean error of 5.87% compared to the real-world flight data of approximately 1706 flights. With this input, the model has access to a lookup table of validated performance data for an aircraft, or variety of aircraft depending on the fleet.

B. Model Configuration

The model developed in this study takes the form of an airline scheduling algorithm. This algorithm receives input information describing the airport network, the fleet of aircraft, the passenger traffic and the cost parameters for a given scenario to be scheduled. These inputs enable the complete definition and configuration of the problem as a discrete optimization, decision-making problem which explores the trade-offs and benefits of various decision sets, i.e different schedules and their respective objective values.

1. Model Configuration: The mathematical representation of the inputs

The input information about the airport network and the fleet of aircraft are used to mathematically define the connectivity of the airport network and routing of flights through it. The airport network information contains details of airport names, codes (IATA), latitude and longitude coordinates and their permissible connections to other airports. The fleet of aircraft information describes what the composition of the fleet is and its capabilities such as maximum ranges, seat counts, fuel burns at each permissible operating point (range, passenger and fuel at take-off allocation) in addition to refueling and turnaround requirements. The mathematical representation of the network is a digraph where nodes represent an airport at a given point in time and the edges connecting nodes represent feasible flights between airports. The digraph’s nodes are initialized in a time-space matrix where each node refers to a distinct position and point in time giving the model both spatial and temporal dimensions. The nodes are referenced according to their index within the time-space, two-dimensional matrix. Firstly, this gives each node a unique identifier (its index value within the matrix) and secondly, to aid in the visualization of the model where columns of nodes all belong to the same airport while increasing row index indicates the advancement of time. For example nodes 1, 200 and 718 refer to airports A at 06:00, B at 10:30 and E at 17:25, respectively. This nodal, time-space matrix representation, N_{TS}, uses a temporal resolution of 5 minutes. Other timesteps, both larger and smaller are feasible, but the total number of nodes depends inversely on this resolution. A larger timestep coarsens the temporal discretization, therefore decreasing the total number of nodes in the digraph which may hinder solution feasibility in some model configurations.

$$N_{TS} = \left[n_{t,s}\right]; t \in T, s \in S \tag{1}$$

The input information for the fleet of aircraft enables the linking of nodes according to the flight time between airports. Knowing the airport coordinates and the aircraft capabilities (maximum range, flight profile and flight times), the flight times for each permissible airport pairing can be determined and used to create a routes matrix, R, that contains all the feasible pairs of nodes from N_{TS}. A feasible pair of nodes creates an edge in the digraph, therefore defining an allowable flight through the network. A feasible nodal pair defines the departure airport and time (given by the first node), the arrival airport and time (second node) and the flight time (difference between the times of the nodes). All feasible nodal pairs must follow spatial and temporal feasibility criteria. The spatial criterion is that the departure airport must be allowed to fly to the arrival airport, i.e the airports are allowed to be connected. The first of the temporal criteria is that the departure time is less than the arrival time, i.e chronological order of take-off and landing is observed. The second temporal criterion is that the difference between the departure and arrival times is equal to the flight time of that specific airport pairing for the given aircraft. These criteria define all feasible edges linking the nodes of the network and are stored in the routes matrix.

The input information for passenger traffic is passed to the model to amend the routes matrix to only consider possible routes that have non-zero passenger traffic. In other words, the passenger traffic information determines which routes available to the model are required since routes that facilitate zero-valued passenger traffic pathways are not required and therefore, may be removed from the analysis to reduce its complexity. The passenger traffic information is passed in a column vector format, $\vec{p}$, where the vector is indexed according to airport pairs and timeslot periods. The passenger traffic vector is a column vector with the number of rows being equal to the number of airport pairs multiplied by the number of timeslots, summarily denoted by m_P. In a network of five airports with no connection restrictions, the number of airport pairs would be 20 (A to B and B to A are considered separate pairs) as it follows the rules for the number of permutations of two items from a set of five items, where order matters.

The timeslots are a segregation of the operating period into distinct windows and are delineated based on the period in which the passengers want to arrive at their destination by. The timeslots used in this analysis segregate the 06:00 to 18:00 operating period into four windows: 06:00-09:00, 09:01-12:00, 12:01-15:00 and 15:01-18:00. Combining the airport pairs with the timeslots, and indexing accordingly, the passenger traffic vector's cell values then define the number of passengers that are to travel between the given airports during a certain time period. Zero-valued cells in this vector represent no passenger traffic between select airports during the prescribed times. The reason the delineation of timeslots is based on this arrival window preference is because it facilitates the model in scheduling flights (activating certain routes) according to when a given route arrives at its destination.

$$R = [n_{dep} \quad n_{arv}]; (n_{dep}, n_{arv}) \in N_{TS} \tag{2}$$

$$\vec{p} = \{p_m\}; m \in m_P \tag{3}$$

After the passenger traffic vector has been loaded and used to amend the routes matrix to only its required constituents (becoming the required routes matrix, R_{req}), the route adjacency, A_r and passenger incidence, I_p matrices can be created. The route adjacency matrix is a square, binary matrix with a side length equal to the number of required routes (n_R). The routes adjacency matrix describes the ability for a preceding route (given by the column index) to be succeeded by another route (row index) where a cell value of 1 permits this chaining of routes, and 0 prohibits it. This adjacency matrix is useful later in the model for verifying the viability of route chaining for a given aircraft. This adjacency matrix is subjected to some general feasibility constraints in its construction. Firstly, a spatial constraint which requires that the succeeding route departs from the same airport as the preceding route arrived to. Secondly, the temporal aspects of the preceding and succeeding routes follow chronological order, i.e the departure time of the second route cannot be before the arrival time of the prior. Lastly, that the minimum turnaround times of an aircraft are observed. This closely resembles the previous chronological criterion but adds in the caveat that the gap between the departure time of the succeeding route and the preceding route's arrival times are no less than the minimum turnaround time. Additional aircraft or specific feasibility criteria may be defined and utilized here to prevent or encourage the chaining of select routes. For example, a maximum dormancy between flights may be used which would act similarly to the minimum turnaround criterion but instead of a minimal, it would be a maximum allowable gap between the landing and following take-off times between routes such that the aircraft is not permitted to dwell unused for a period longer than the maximum dormancy.

$$A_r = \left[a_{i,j}\right];\ a_{i,j} \in [0,1], (i,j) \in n_R \tag{4}$$

The passenger incidence matrix is a rectangular, binary matrix where the number of rows and columns are given by the length of the passenger traffic vector (m_P) and the number of required routes, respectively. When a route travels between a set of airports and arrives during a given timeslot matching that of the passenger traffic vector, the corresponding cell in the passenger incidence matrix is set to 1. This cell is indexed along the columns by the specific route and along the rows by the passenger traffic index associated with the specific airport pair and timeslot window. Hence, the zero-valued cells in this passenger incidence matrix indicate that the route (column index) does not facilitate passengers traveling between a set of airports during a given timeslot (row index). This matrix is useful in assessing the connectivity and solvability of the model based on a set of inputs. If the entire matrix is zero-valued, no routes exist in the model to enable passengers to traverse between airports nor during any timeslot, hence the model would deem this configuration of the network as infeasible since no passengers can travel and no schedule may be created. It can therefore be understood that the less zero-valued cells in this matrix, the more connected the network is. These connections are determined based on the airport and fleet information as well as the passenger traffic inputs, hence altering these may dramatically alter the network and scenario under investigation. Additionally, the model uses this passenger incidence matrix to assess a set of route activations in the scheduler as it compares where and when the model's scheduled passengers are to the passenger traffic vector and determines the overall passenger traffic satisfaction. Ideal schedules would satisfy the passenger traffic completely as this means all passengers (viewed as revenue sources) have been flown thus 100% of the possible revenue has been captured. But this may not always be possible due to a multitude of reasons; insufficient routing or airport connectivity, the fleet size being too small, the cost to operate outweighing the possible revenue capture, among other feasibility and logistical reasons for under-catering. It must be noted that this model prevents over-catering (scheduling more passengers than the passenger traffic vector) elsewise the model would elect to only service the relatively more profitable routes as much as possible and only cater to the less profitable routes when necessary, defeating the purpose of scheduling for a prescribed scenario of airport, fleet and passenger traffic.

$$I_p = \left[I_{i,j}\right];\ I_{i,j} \in [0,1], i \in m_p, j \in n_R \tag{5}$$

2. Model Configuration: The definition of the optimization variables

With the input information of the airport network, the fleet of aircraft and the passenger traffic configured within the model as the required routes matrix, the routes adjacency matrix and the passenger incidence matrix, the model creates optimization variables reflecting the choices available to the scheduler. There are five types of optimization variables used in this model, namely: the route, passenger, fuel at take-off, dummy fuel and dummy time allocations. These optimization variables are defined on a per flight, per aircraft basis to ensure that each aircrafts' individual restrictions and capabilities are reflected in and throughout the scheduled operating period. Firstly, an activatable routes matrix, M_k is created for each aircraft (where the total number of aircraft is n_k). This is a binary matrix with the number of columns and rows equal to the maximum number of journeys (n_J) and the number of required routes, respectively. The first column corresponds to activatable routes (cell value of 1) for the given aircraft's first journey, the second column indicates the activatable second flights and so on. These activatable routes matrices per aircraft are stored in a larger, three-dimensional matrix of the same row and column size but indexed in the third dimension by the number of aircrafts. The purpose is to record the permissions of a given aircraft, i.e how this specific aircraft may traverse through the network, what routes it may activate and on which leg of its flight plan. Initially, these permissions are based on the required routes matrix, the routes adjacency matrix and the fleet capabilities each aircraft belongs to. A user may want to specify alterations to the permissions or restrictions for a given aircraft or subset of aircraft, noting that these must still satisfy the previous feasibility criteria regarding spatial and temporal route chaining. These restrictions and changes to the permissions of an aircraft may take the form of specific starting airports, maximum allowable number of journeys, prohibited airport visits, restricted range or seat counts, or other aircraft and or flight permission-alterations as prescribed by the user. These permission alterations, after being mathematically defined, change the values in the activatable routes matrix per aircraft.

For example, the first column of cells for an aircrafts' activatable routes matrix may all be set to 1, initially. This means that this aircraft can fly from anywhere in the network along any of the routes as its first flight. Using the routes adjacency matrix, the second column of activatable routes can be found by multiplying the first column by the adjacency matrix and recording all non-zero-valued cells' indices, i.e any of the preceding routes (first flight) which have a permissible succeeding one (second flight). These recorded indices are assigned 1 in the activatable routes matrix's second column and make up the possible follow-up routes to the previous set of possible routes. This is repeated over the remaining successive sets of columns (remaining legs in the flight plan) to complete the activatable routes matrix's definition for the aircraft. As an example of when the user wants to alter the permissions of this aircraft further; perhaps the user wants an aircraft to start its operations from a specific airport, its first flight must be before a given time, and it may only visit a certain set of airports. These restrictions would deactivate some of the previously permissible routes (change the cell value from 1 to 0). The model would interpret the user specified alterations to the permissions and construct the activatable routes matrix accordingly. The first column would now only have cell values of 1 on routes that departed the correct airport, before the correct time and that arrive at the permissible set of airports only. The second column would change accordingly too as the adjacency matrix is again used but this time the number of activatable preceding routes has reduced therefore the succeeding routes (the second column of permissible routes) will be reduced to reflect the updated restrictions also. This will cascade down through the successive routes to redefine the permissions of this specific aircraft's maneuverability through the network and operating period. It follows that restrictions reduce the number of choices afforded to a given aircraft and conversely, the lack of additional restrictions outside of logical feasibility gives an aircraft the maximum freedom.

$$M_k = \left[m_{i,j}\right]_k;\ m_{i,j} \in [0,1], i \in n_R, j \in n_J, k \in n_k \tag{6}$$

The previously mentioned optimization variables of route (X_R), passenger (X_P) and fuel at take-off (X_{FTO}) allocations are defined on a per flight, per aircraft basis as defined by the activatable routes matrices per aircraft. For all cells in these matrices valued at 1, an optimization variable is defined. For convenience, each optimization variable is named according to the route index, journey (first, second, or so column) index and aircraft index of the activatable routes matrix alongside whether it is a route, passenger or fuel at take-off allocation. The routes optimization variables are binary and indicate whether a route has been activated or not. The passenger optimization variables are defined as integers, bounded between 0 and the maximum seat number per aircraft ($p_{max,k}$) and indicate the number of passengers travelling on the selected route. The fuel at take-off optimization variables are defined as continuous, bounded between 0 and the maximum fuel capacity ($f_{max,k}$) of the aircraft and indicate the fuel onboard at take-off. Handler functionality is utilized to ensure the size of vectors and matrices created with the optimization variables staying consistent with the routes adjacency and passenger incidence matrices such that aircraft with lesser active routes maintain compatibility with the various other matrices and vectors in the model. These handler functions create a three-dimensional matrix for each of these optimization variable groups and initialize cells as zero-valued everywhere. These handler matrices are sized as the required number of routes, maximum number of journeys and the number of aircraft along the row, column and third dimensions, respectively. These handler matrices are updated to store the recently created optimization variables based on the activatable routes matrices per aircraft to only store an optimization variable where an activatable routes matrix value is 1, i.e a possible choice in route, passenger and fuel at take-off exists for the aircraft.

$$X_R = \left[x_{i,j,k}^{r}\right];\ x_{i,j,k}^{r} \in [0,1], i \in n_R, j \in n_J, k \in n_k \tag{7}$$

$$X_P = \left[x_{i,j,k}^{p}\right];\ x_{i,j,k}^{p} \in \left[0, p_{max,k}\right], i \in n_R, j \in n_J, k \in n_k \tag{8}$$

$$X_{FTO} = \left[x_{i,j,k}^{fto}\right];\ x_{i,j,k}^{fto} \in \left[0, f_{max,k}\right], i \in n_R, j \in n_J, k \in n_k \tag{9}$$

The other optimization variables are the dummy fuel (X_{DF}) and dummy time (X_{DT}) and are defined for each aircraft. These optimization variables serve the purpose of maintaining feasibility in the later constraints and logical handling of decision-making within the model. They are two-dimensional matrices sized according to the number of aircraft and maximum number of journeys. They are both defined as continuous, with bounds from 0 to the maximum fuel capacity per aircraft for the dummy fuel and 0 to the finish time (T_E) of the operating period for the dummy time optimization variables. This group of optimization variables are named according to the aircraft and what journey index they represent alongside whether they belong to the dummy fuel or time variable. Hence, a dummy fuel and time optimization variable exists for every journey per aircraft for every aircraft in the model. These dummy variables remain inactive, set to 0 throughout, unless some of the aircraft are under-utilized in the schedule. This under-utilization refers to optimization variables for route, passenger and take-off fuel being set to 0 for all route options of a journey or journeys for one or more aircraft, i.e a set of possible flights that could have been scheduled to one or more aircraft, but were never flown as the model deemed it was preferential to not use them.

$$X_{DF} = \left[x_{j,k}^{df}\right];\ x_{j,k}^{df} \in \left[0, f_{max,k}\right], j \in n_J, k \in n_k \tag{10}$$

$$X_{DT} = \left[x_{j,k}^{dt}\right];\ x_{j,k}^{dt} \in \left[0, T_E\right], j \in n_J, k \in n_k \tag{11}$$

3. Model Configuration: The creation of the optimization constraints

Following the definition of the optimization variables for the given scenario of inputs, optimization constraints (referred to as "Opt. C." in the coming equations) are required to enforce feasibility and logistical behavior of the model. The first optimization constraint enforces only single route selection being made per journey per aircraft. The second optimization constraint enforces successive route feasibility using a handler function referred to as the route paths matrix (X_{paths}). The route paths matrix is the same size as the activatable routes matrix and uses the previously decided routes to determine which successive routes may be selected, functioning very similar to the activatable routes matrix. However, the route paths matrix uses the live optimization variable decisions within the scheduler of the model to update (activate and deactivate) permissions of succeeding routes as the model determines a schedule.

$$Opt.C.\#1 \qquad \sum_{i=1}^{n_R} X_R \leq 1 \tag{12}$$

$$Opt.C.\#2 \qquad X_{paths} \geq X_R \tag{13}$$

The third optimization constraint enforces that passenger allocations only occur on activated route allocations, i.e passenger counts are only put onto routes that have an active flight. The fourth optimization constraint enforces no "dead-heads" which is a colloquial term in aviation for zero-passenger flights. This constraint ensures that a flight must have a passenger allocation. Dead-heads may be useful in some scenarios as it allows aircraft in different parts of the network to be relocated to facilitate passenger traffic elsewhere, this constraint can be switched off to permit this functionality if desired. The fifth optimization constraint is similar to the third but instead of passengers, it enforces that take-off fuel is only allocated to active flights.

$$Opt.C.\#3 \qquad X_P \leq p_{max} \cdot X_R \tag{14}$$

$$Opt.C.\#4 \qquad X_P \geq X_R \tag{15}$$

$$Opt.C.\#5 \qquad X_{FTO} \leq f_{max} \cdot X_R \tag{16}$$

The sixth and seventh optimization constraints enforce that the dummy fuel and time variables may only be activated on inactive flights per aircraft, i.e if an aircraft has 4 of its 6 maximum flights active, the dummy fuel and time variables can only be non-zero for the 5th and 6th flights as these journeys are unused by the scheduler. The eighth and ninth optimization constraints enforce minimum and maximum fuel at landing (X_{FL}) relative to the fuel burn (X_{FB}), respectively. The minimum constraint mimics safety practices such that an aircraft is not permitted to land with 0% of fuel left on board (empty tanks), some safety-reserve fuel (λ_{min}) must be allocated that remains after the aircraft has landed. The maximum constraint is to prevent "tankering" behavior which is the practice of carrying excessive fuel for a given flight to avoid the need to refuel at the destination. While this may save time or money at turnaround, the additional fuel increases the fuel consumption per flight (heavier flight), hence producing more emissions. This behavior is restricted in Europe with the ReFuelEU's regulations penalizing airlines for producing unnecessary emissions due to tankering practices [18]. The maximum fuel constraint enforcing tankering limits (λ_{max}), like the constraint about dead-heads, can be neglected by the model if desired by the user for a given scenario.

$$Opt.C.\#6 \qquad f_{max} \cdot \left(1 - \sum_{i=1}^{n_R} X_R\right) \geq X_{DF} \tag{17}$$

$$Opt.C.\#7 \qquad T_E \cdot \left(1 - \sum_{i=1}^{n_R} X_R\right) \geq X_{DT} \tag{18}$$

$$Opt.C.\#8 \qquad X_{FL} \geq \lambda_{min} \cdot X_{FB} \tag{19}$$

$$Opt.C.\#9 \qquad X_{FL} \leq \lambda_{max} \cdot X_{FB} \tag{20}$$

The tenth and eleventh optimization constraints enforce refueling feasibility and the logical handling of fuel levels between flights for each aircraft. Handler functionality is created within the model by defining optimization expressions for fuel burn, fuel at landing and refueling amounts (X_{FR}) per flight, per aircraft. These are optimization expressions rather than variables as they are not directly varied by the model, instead they are formulae and expressions which contain optimization variables (fuel at take-off and dummy fuel). The fuel burn expression is a function of route, passenger count and fuel at take-off. Knowing the aircraft type and route, the range and flight profile are given. These, alongside a value for passenger count and fuel at take-off, define the operating point, i.e the mission details for an aircraft's flight. This aircraft operating point determines the associated, accurate fuel burn value using the previously discussed "digital hangar" functionality from Section III-A, which the fuel burn expression is set to. The landing fuel expression is equal to the fuel at take-off variables less the fuel burn per flight, per aircraft. The refuel expression is slightly more complicated, the first index per aircraft of this expression, i.e the first flight of each aircrafts' refuel amount is set as the first take-off fuel value. This instructs the model to interpret the first take-off fuel variables' allocations as the first refuel amount which is realistic since the day of operations are only beginning so fuel must be added prior to the schedule of operations. The succeeding refuel expression values are equal to the take-off fuel of the succeeding flight less the landing fuel of the preceding flight plus the dummy fuel amount for that flight per aircraft.

$$X_{FB} = func(type_k, X_R, X_P, X_{FTO}) \tag{21}$$

$$X_{FL,j=1} = X_{FTO,j=1} - X_{FB,j=1} \tag{22}$$

$$X_{FR,j=2} = X_{FTO,j=2} - X_{FL,j=1} + X_{DF,j=2} \tag{23}$$

When all flights are utilized by an aircraft, the dummy fuel variable remains zero throughout and is therefore neglected by the model. However, when there are unused flights by an aircraft, the take-off fuel variable will be zero for that flight and aircraft since no active routes have been called but the dummy fuel variable will be non-zero to ensure a feasible, non-negative refuel expression. This behavior adheres to the fifth and sixth optimization constraints, where the take-off fuel must be zero for inactive flights and the dummy fuel may only be activated on these unused flights. Hence, the tenth optimization constraint enforces that the logical, successive fueling of the aircraft by relating the landing, dummy and take-off fuels per flight per aircraft. And the eleventh optimization constraint enforces the activation of the dummy fuel variable to maintain the refuel expressions feasibility for all aircraft.

$$Opt.C.\#10 \qquad X_{FL} - X_{DF} \leq X_{FTO} \tag{24}$$

$$Opt.C.\#11 \qquad X_{FL} \geq X_{DF} \tag{25}$$

The twelfth through fifteenth optimization constraints enforce temporal logic within the model. Again, handler functions are defined as optimization expressions, namely the earliest possible departure time (X_{EDT}), actual departure time (X_{ADT}), actual arrival time (X_{AAT}) and refuel time (X_{FRT}) per flight, per aircraft. These time-based optimization expressions form matrices of the same size as the dummy time optimization variable, i.e two-dimensional, number of aircraft (rows) by maximum number of journeys (columns). The earliest possible departure time's first index, the first flight per aircraft, is set as the start time of the operating period by default, which in this model is 06:00. The succeeding values are determined by what routes and how much refuel time is required per flight. The actual departure and arrival time expressions are calculated from the required routes matrix and the route optimization variable. If the model selects a route for a flight of an aircraft, the selected route's departure and arrival times are given by the corresponding index in the required routes matrix hence the actual departure and arrival times are simply these times. The actual departure expression requires the addition of the dummy time variable, similar to how the fueling feasibility is maintained with the dummy fuel variable. For aircraft that utilize all of their flights, the dummy time values remain zero throughout. However, when an aircraft does not use all possible flights, the dummy fuel variable activates on the inactive flight or flights to adhere to the governing optimization constraints. The refuel time expression is simply the refuel amount (from the refuel expression discussed earlier) divided by the refuel rate for a given aircraft (σ_k) and is only necessary for consideration if an aircraft uses asynchronous passenger on/off-boarding and refueling. Hence, if an aircraft allows simultaneous refueling and passenger transfer, the refuel time becomes irrelevant as the turnaround time is fixed by the passenger on/off-boarding (assuming this time is greater than the maximum possible refuel time). The earliest possible departure time expression therefore uses the information regarding the previous flights' actual arrival times plus the minimum turnaround times plus the refuel time (if necessary) to determine when the flight is eligible to depart next.

The twelfth optimization constraint, like the eleventh but now for timing logic, enforces that the activation of the dummy time variable to maintain the chronological departure and arrival expressions' feasibility regardless of the number of flights per aircraft. The thirteenth optimization constraint enforces that the actual departure times be greater than or equal to the earliest possible departure times per flight, per aircraft. The fourteenth optimization constraint enforces a maximum lag time (τ_{lag}) between earliest possible and actual departure times. This can be thought of as the maximum allowable slack per aircraft in the schedule, i.e how long between when the aircraft can feasibly depart versus when it actually departs. The user may define this maximum lag to be restrictive or flexible, i.e a very short or long time, respectively. The fifteenth optimization constraint enforces that the first refueling time is feasible for the initial departure time per aircraft, i.e the fuel at take-off for the first flight of the day is capable of being pumped into the tank ahead of the first departure given the aircraft type, refuel rate and operating window.

$$Opt.C.\#12 \qquad X_{EDT} \geq X_{DT} \tag{26}$$

$$Opt.C.\#13 \qquad X_{EDT} \leq X_{ADT} \tag{27}$$

$$Opt.C.\#14 \qquad X_{ADT} - X_{EDT} \leq \tau_{lag} \tag{28}$$

$$Opt.C.\#15 \qquad T_E + \frac{X_{FRT,j=1}}{\sigma_k} \leq X_{ADT,j=1} + 24hrs \tag{29}$$

The sixteenth and final optimization constraint enforces the scheduled passenger allocations to be less than or equal to the passenger traffic vector input. The passenger traffic vector acts as the upper limit of passengers allowed to travel between airport pairs during the timeslots. This constraint prevents additional, uncalled-for passengers being allocated to favorable routes, thus steering the scheduler's allocation of flights away from the prescribed network and passenger inputs. Instead, the schedule may allocate passengers and flights that can, at most, 100% satisfy the passenger traffic index.

$$Opt.C.\#16 \qquad I_p * \left(\sum_{j,k=1,1}^{n_J,n_k} X_P\right) \leq \vec{p} \tag{30}$$

4. Model Configuration: The construction of the objective function

After the optimization constraints are defined, the objective function can be constructed to direct the scheduler's decision making. The model uses a single objective function defined as a holistic, hybrid economic and environmental sustainability metric that ranks the optimality of a given schedule succinctly with units of currency. The input information of the previously mentioned cost parameters now come into consideration within the model. The cost parameters are the user defined objective values associated with each decision that can be made in the model, i.e the cost of each choice. These cost parameters are defined in Table 1 and aid in the construction of the objective function. The objective function is designed to maximize the hybrid sustainability of the schedule. This hybrid sustainability is expressed solely in terms of currency by measuring traditional economic profit (all revenue less all costs) plus the environmental performance (amount of emissions) converted into units of currency via emissions' cost conversion metrics.

Table 1: Cost Parameters

Name	*Symbol and Size*	*Units*
Base Passenger Revenues	$\phi_{p,b,k}; (1 \times n_k)$	€/pax
Distance Passenger Revenues	$\phi_{p,d,k}; (1 \times n_k)$	€/km
Landing Costs	$\phi_{l,k}; (n_R \times n_k)$	€
Handling Costs	$\phi_{h,k}; (n_R \times n_k)$	€
Crew Costs	$\phi_{c,k}; (1 \times n_k)$	€
Refuel Costs	$\phi_{FR,k}; (n_R \times n_k)$	€
Emissions' Costs	$\phi_{CO2,k}; (1 \times n_k)$	€/kg of CO_2

The only emission considered in this analysis at present is carbon-dioxide (CO_2) and is measured on a tank-to-wake (TTW) scale. This TTW scale only considers the CO_2 emissions associated with consuming the fuel in-flight and ejecting CO_2 from the engines into the aircrafts' wakes, it does not consider the embedded CO_2 emissions produced in the generation and transport of the fuel to the aircrafts' tanks. By knowing the type of fuel used on an aircraft, the amount of CO_2 produced from the consumption of that fuel can be inferred using stoichiometry, 1kg of Jet A-1 produces 3.156kg of CO_2 [19]. This mass-equivalence is used to multiply the total fuel consumed to get the total CO_2 emissions which can be converted to CO_2 emissions' costs via the EU ETS penalty of €80/1000kg for CO_2 [7]. The objective function therefore captures the impact on economic behavior and the environmental performance of the schedule and outputs a single metric (γ) with units of currency where higher objective values indicate stronger performing schedules. The objective function is defined as the revenue from passengers (α_p) less the costs of landing (β_l), handling (β_h), crew (β_c), refuel (β_{FR}) and emissions (β_{CO2}). Passenger revenue depends on the aircraft and route selected, where longer routes have a distance weighted increase in pricing (monies paid by passengers to the airline). Landing and handling costs depend on the airport and type of aircraft. Crew costs depend on the type of aircraft. Fuel costs depend on the type of aircraft and airport. Emissions costs depend on the type of aircraft and the emissions' cost conversion metric. It is these cost parameters that act as the weightings against the decision variables in the objective function. It can therefore be understood that tuning these cost parameters between runs of the same scenario can produce substantially different solutions. This tuning is controlled by the user in order to create a preset scenario for a certain technology's cost, energy carrier's cost, emissions' penalty level, or other decarbonization lever/policy/etc.

$$\alpha_p = \phi_{p,b,k} \cdot \left(\sum_{i,j,k=1,1,1}^{n_R,n_J,n_k} X_P\right) + \phi_{p,d,k} \cdot \left(\sum_{i,j,k=1,1,1}^{n_R,n_J,n_k} X_R\right) \tag{31}$$

$$\beta_l = \phi_{l,k} \cdot \left(\sum_{i,j,k=1,1,1}^{n_R,n_J,n_k} X_R\right); \beta_h = \phi_{h,k} \cdot \left(\sum_{i,j,k=1,1,1}^{n_R,n_J,n_k} X_R\right); \beta_c = \phi_{c,k} \cdot \left(\sum_{i,j,k=1,1,1}^{n_R,n_J,n_k} X_R\right) \tag{32}$$

$$\beta_{FR} = \phi_{FR,k} \cdot \left(\sum_{i,j,k=1,1,1}^{n_R,n_J,n_k} X_{FR}\right); \beta_{CO2} = \phi_{CO2,k} \cdot 3.156 \cdot \left(\sum_{i,j,k=1,1,1}^{n_R,n_J,n_k} X_{FR}\right) \tag{33}$$

$$Obj. \qquad \gamma = \alpha_p - (\beta_l + \beta_h + \beta_c + \beta_{FR} + \beta_{CO2}) \tag{34}$$

IV. Results & Discussion

A benchmarking version of the model has been established to test and verify the functionality discussed in Section III using Python as the development environment and the discrete optimization suite known as PuLP [20]. This suite is strictly for linear programming (LP) style problems and uses the "PULP_CBC_CMD" solver by default but can be configured to use a range of alternative open-source and commercial solvers. Two alternative, open-source solvers were tested inside the PuLP configuration, namely "Solving Constraint Integer Programs" (SCIP) [21] and HiGHs [22]. Across each solver within PuLP, the airport network and passenger traffic inputs remain the same. The fleet of aircraft input to the model are varied between scenarios for demonstrative purposes of functionality such as the asynchronous turnarounds for given aircraft, altering aircraft specific permissions when it comes to route chaining.

For demonstrative purposes and to investigate the full breadth of functionality within the model without over-encumbering the solvers with too many variables and constraints, the benchmarking version of this model's scope is kept relatively small. The airport network input is a five airport, fully connected system with an operating window from 06:00 to 18:00, at a resolution of 5 minutes. The fully connected five airport network creates 20 unique spatial, airport pairs. This creates a time-space nodal matrix of 725 unique nodes. Two types of aircraft are used within this study, the Jet A-1 fueled, CFM-LEAP-1B27 powered, 197-seater, B737-8200 MAX aircraft [23] and a liquid hydrogen (LH_2) fueled variant of same. The capabilities and modelling of the LH_2 variant are discussed in detail in work by Gallagher et al. [16, 17] and the connected work by Barry et al. [13, 14]. For this study however, it is sufficient to understand that rigorous fuel consumption modelling has been conducted in previous studies, hence accurate fuel burn approximations are available for use within this analysis. The passenger traffic input is informed by the United States of America's Bureau of Transportation Statistics (BTS) [24] where a two-peak pattern of passenger traffic through the operating period is observed in each airport. Publicly available data such as the BTS passenger traffic is useful for the development of this model type wherein the schedule output depends entirely on the inputs. The overall objective of this model is to create schedules which instruct users on how sustainable given decarbonization scenarios may be (based on their objective functions' scores), hence using real-world data to develop this model in the first instance embeds the required reliable and robust linking of passenger data inputs to the schedule outputs. Within the BTS two-peak data, the first and largest peak is in the morning with a lull in passenger flow through midday, and finally the second peak coming in the evening time. The setup of the model follows the methodology outlined in Section III with regards to defining optimization variables and constraints and the objective function for maximizing hybrid sustainability.

A. Model Functionality Demonstration

The high-level purpose of this model is to generate a flight schedule describing all operations to be undertaken and summarizing all fuel, passenger, aircraft and other pertinent details given a set of inputs and user specified preferences. This schedule is guided by an objective function maximizing the hybrid economic and environmental sustainability, hence schedules that score relatively stronger in this objective can be understood to enact a set of decisions that favor more profitable and "greener" operations. It is this hybridized sustainability metric that relates the scheduling model back to the fundamental research focus, the decarbonization of the aviation sector.

1. Model Functionality Demonstration: Aircraft specific flight plan alterations

Firstly, spatial and temporal chaining of flights for a given aircraft are key factors in the feasibility of a flight plan. Spatially, an aircraft must depart from the arrival airport of the previous flight. Similarly, the chronological order of flights for an aircraft must make sense, succeeding take-offs must happen after preceding landings alongside allotted times for turnarounds being observed. This can be seen in Figure 3 where the extract shows that Aircraft 1 has been scheduled a flight plan of six flights starting in airport A. Its first flight carries 197 passengers ("pax.") from A to B between 06:20 and 07:45. The spatial and temporal logic is achieved in this aircraft's flight plan as one goes through the order of airports visited and the times of take-off and landing between flights.

```
Aircraft 1 Flight Plan & Fuel Details:
Aircraft 1 Flight 1 = A @ 06:20 -> B @ 07:45 w/ 197 pax.
Aircraft 1 Flight 2 = B @ 08:35 -> A @ 10:00 w/ 197 pax.
Aircraft 1 Flight 3 = A @ 11:05 -> D @ 11:55 w/ 197 pax.
Aircraft 1 Flight 4 = D @ 12:50 -> A @ 13:40 w/ 183 pax.
Aircraft 1 Flight 5 = A @ 14:40 -> D @ 15:30 w/ 197 pax.
Aircraft 1 Flight 6 = D @ 16:15 -> A @ 17:05 w/ 91 pax.
```

Fig. 3: Feasible spatial and temporal route chaining

Should the aircraft in the previous extract have its permissions altered (as the user may want the aircraft to only be able to visit specific airports and start somewhere else), its flight plan will be amended when the model is rerun, as seen in Figure 4. In this scenario, the starting airport for Aircraft 1 has been changed from A to B. It has also had its permissible airport connections altered to only include airports A, B and C, removing permissible visits to D and E, a scenario reflective of the limited airport pairs at which new energy carriers may be available as infrastructure is upgraded. This is illustrated in Figure 4 where the starting airport is changed and the flights to airport D and from airport A no longer take place. Instead, there are flights between airports B and C servicing different passengers.

```
Aircraft 1 Flight Plan & Fuel Details:
Aircraft 1 Flight 1 = B @ 07:00 -> A @ 08:25 w/ 197 pax.
Aircraft 1 Flight 2 = A @ 10:00 -> B @ 11:25 w/ 189 pax.
Aircraft 1 Flight 3 = B @ 12:15 -> C @ 12:50 w/ 146 pax.
Aircraft 1 Flight 4 = C @ 13:50 -> B @ 14:25 w/ 197 pax.
Aircraft 1 Flight 5 = B @ 15:45 -> C @ 16:20 w/ 197 pax.
Aircraft 1 Flight 6 = C @ 17:25 -> B @ 18:00 w/ 197 pax.
```

Fig. 4: Altering aircrafts' permissions based on user preferences

Furthermore, perhaps the user would like to simulate an aircraft being taken out of service until a certain point in the day to mimic necessary downtime for scheduled maintenance activities. This would manifest within the scheduling model as an alteration in that aircraft's permissions with no flights allowed to be scheduled until the prescribed maintenance period has concluded. With this reduced availability, the scheduler would have to allocate less flights to the aircraft with all of them taking place after the maintenance window. This can be seen in Figure 5 where Aircraft 1 is not allowed to re-enter service until after 12:00. In Figures 3 and 4, when Aircraft 1 was allowed to start its flight plan before 12:00, it was allocated a full flight plan of six flights, however now it can only manage four flights inside the operating window, reducing its utilization due to required maintenance.

```
Aircraft 1 Flight Plan & Fuel Details:
Aircraft 1 Flight 1 = B @ 12:10 -> C @ 12:45 w/ 146 pax.
Aircraft 1 Flight 2 = C @ 13:55 -> B @ 14:30 w/ 197 pax.
Aircraft 1 Flight 3 = B @ 15:25 -> C @ 16:00 w/ 197 pax.
Aircraft 1 Flight 4 = C @ 17:05 -> B @ 17:40 w/ 197 pax.
```

Fig. 5: Delayed starts and reduced utilisation

This functionality conveys the ability of the model to capture the fundamentals of an airline schedule, allocating where and when flights occur in the network and ordering them in a spatially and temporally feasible manner. The model is capable of considering user specified restrictions and permission alterations such as allowable airport connections or required maintenance windows therefore instilling more realism and application-based operation into the model. This functionality is shown for a single aircraft but works for all aircraft handled within the model.

2. Model Functionality Demonstration: Fuel consumption and feasibility

The fueling of the aircraft is a key component of the model and is connected to all aspects of determining an optimal schedule. Firstly, the fuel required to fly between airports is a function of the aircraft, flight profile, passengers and fuel at take-off (summarized as an operating point). An accurate approximation of the fuel consumed for the two aircraft variants under investigation for all operating points is available to the model. This means that the scheduler knows exactly how much fuel is required to fly between all airports, on either aircraft type, with any permissible combination of passengers and fuel on board. Since the scheduler is trying to maximize the hybrid sustainability metric where fuel costs and emissions reduce this score, minimizing the fuel consumption becomes a proponent of the maximization effort.

Heavier and farther flights burn more fuel, but not all weight and distances can be reduced directly to minimize the fuel consumption per flight. The passengers being served are revenue generators despite their weight contributing to the fuel consumption, hence lowering their number decreases the profitability of the flight and in turn the objective function. The distance between airports cannot be changed as passengers will only generate revenue if they are flying to their intended destination, i.e if the flight does not go between the required airports, no passengers are allocated, and no revenue is generated. This leaves the fuel at take-off parameter as the easiest to directly minimize for lowering fuel consumption since less fuel on board reduces the weight of the aircraft (noting the safety and feasibility requirements of not landing at 0% fuel or having less fuel on board than required, respectively). Conversely, the cost of refueling or the time-taken to turnaround due to refuel rates could be prohibitive for certain aircraft so the scheduler may elect to increase the fuel at take-off of a preceding flight, so the turnaround and take-off of the succeeding flight enables better utilization and a better contribution to the hybridized sustainability metric than alternative fuel level configurations. This is why the selection of route to fly, passenger allocation and fuel at take-off are all optimization variables for each aircraft, they can all be optimized within an aircraft's flight plan such that fuel consumption is minimized as part of maximizing the overall hybrid sustainability for the entire fleet.

In Figure 6, the fueling details of Aircraft 1 are shown where "R", "TO", "B" and "L" stand for the refuel, take-off, burned and landing fuel of the aircraft for the given flight, respectively. One can see that the first refuel amount is equal to the first take-off fuel amount. The landing fuel amounts are equal to the take-off fuel amounts less the burned fuel amounts, which are a function of the aircraft type, route, passengers and take-off fuel. It can be seen that the second and so on, take-off fuel amounts are equal to the previous landing fuel amounts plus the refuel amounts, i.e the third take-off fuel amount is equal to the second's landing fuel amount plus the third's refuel amount. This fueling feasibility was mathematically summarized in Equations 21 to 23 in Section III-B-3. It must be noted that fuel values are expressed in percentage for data sensitivity reasons.

```
Aircraft 1 Flight Plan & Fuel Details:
Aircraft 1 Flight 1 = B @ 06:00 -> A @ 07:25 w/ 197 pax. Fuel Info: R = 26.5222 , TO = 26.5222 , B = 24.1111 , L = 2.4111
Aircraft 1 Flight 2 = A @ 08:15 -> B @ 09:40 w/ 189 pax. Fuel Info: R = 23.1333 , TO = 25.5444 , B = 23.2222 , L = 2.3222
Aircraft 1 Flight 3 = B @ 11:45 -> C @ 12:20 w/ 146 pax. Fuel Info: R = 20.9584 , TO = 23.2806 , B = 21.1642 , L = 2.1164
Aircraft 1 Flight 4 = C @ 13:35 -> B @ 14:10 w/ 197 pax. Fuel Info: R = 28.7000 , TO = 30.8164 , B = 28.0149 , L = 2.8015
Aircraft 1 Flight 5 = B @ 15:15 -> C @ 15:50 w/ 197 pax. Fuel Info: R = 28.0149 , TO = 30.8164 , B = 28.0149 , L = 2.8015
Aircraft 1 Flight 6 = C @ 16:45 -> B @ 17:20 w/ 197 pax. Fuel Info: R = 28.0149 , TO = 30.8164 , B = 28.0149 , L = 2.8015
```

Fig. 6: Fuel details and feasibility within a flight plan

The user may alter the constraints around minimum safety fuel and maximum tankering fuel levels for a given aircraft due to different scenarios of policy or operational preferences. In Figure 6, purely for illustrative purposes, it can be seen that the landing fuel was always 10% of the fuel burn for that flight and the take-off fuel was never more than 20% greater than the fuel burn, this reflects the user specified safety and tankering limits, respectively. In Figure 7, the safety limit is increased so that aircraft cannot land with less than 20% of the fuel burn and this is seen in the increases in fuel values between Figures 6 and 7. The aircraft cannot land with less than 20% of the fuel burn value, i.e must have an additional 20% of fuel burn required for the mission, for safety purposes. This change is reflected with increases in the take-off fuels which in turn increases the fuel burn values for the same route and passenger allocations. Herein lies the need for optimization variables, the scheduler can tune the fuel levels to satisfy the constraints but also maximize the objective function of hybrid sustainability, which as previously stated means fuel consumption is minimized where possible.

```
Aircraft 1 Flight Plan & Fuel Details:
Aircraft 1 Flight 1 = B @ 06:20 -> A @ 07:45 w/ 197 pax. Fuel Info: R = 32.5500 , TO = 32.5500 , B = 27.1250 , L = 5.4250
Aircraft 1 Flight 2 = A @ 09:15 -> B @ 10:40 w/ 189 pax. Fuel Info: R = 25.9250 , TO = 31.3500 , B = 26.1250 , L = 5.2250
Aircraft 1 Flight 3 = B @ 11:55 -> C @ 12:30 w/ 146 pax. Fuel Info: R = 21.3625 , TO = 26.5875 , B = 22.1563 , L = 4.4313
Aircraft 1 Flight 4 = C @ 13:50 -> B @ 14:25 w/ 197 pax. Fuel Info: R = 30.7625 , TO = 35.1938 , B = 29.3281 , L = 5.8656
Aircraft 1 Flight 5 = B @ 15:30 -> C @ 16:05 w/ 197 pax. Fuel Info: R = 29.3281 , TO = 35.1938 , B = 29.3281 , L = 5.8656
Aircraft 1 Flight 6 = C @ 17:25 -> B @ 18:00 w/ 197 pax. Fuel Info: R = 29.3281 , TO = 35.1938 , B = 29.3281 , L = 5.8656
```

Fig. 7: Safety limit changes on landing fuel

The fuel consumption and refueling functionality is demonstrated on a single aircraft basis in this sub-section, but it is replicated at the fleet scale in the model. It is imperative that the logic and mathematical modelling of the fuel consumption and refueling for a fleet being scheduled is reliable and robust such that accurate fuel burns are determined and link in with the required, subsequent refueling and take-off amounts throughout. This model schedules numerous, inter-dependent flights to multiple aircraft to determine the optimal set of decisions that maximize the objective function. The objective function increases with more passengers being allocated to flights as passengers are the revenue source. However, with more passengers and flights comes an increase in the fuel consumption and emissions which have associated costs. Therefore, the scheduler must find the optimal balance between allocating as many passengers as possible (constrained by passenger traffic limits and aircraft availability) while reducing the fuel consumption to keep emissions and fuel costs as low as possible. Also under consideration inside this optimization is the fact that refueling rates and logistical constraints around turnarounds play a role in how many flights a given aircraft can operate in a given window. As such, there are unintuitive secondary and ternary relations between the fuel consumption, the refueling, the turnaround logistics, etc. and the optimal hybrid sustainability of the schedule that can be explored through this model.

3. Model Functionality Demonstration: Turnaround times and other temporal feasibility

The traditional, Jet A-1 fueled B737-8200 MAX is able to complete synchronous turnarounds, i.e transfer passengers on and off the aircraft while simultaneously refueling. The LH_2 variant is setup to have an asynchronous turnaround. In the case of the traditional MAX aircraft, the turnaround time is always 25 minutes (amount of time to on/off-board passengers). Whereas in the case of the LH_2 variant, the turnaround times between flights is 25 minutes plus whatever refuel time is required. The refuel time required is determined by the refuel rate and the amount of refueling to be undertaken. This functionality can be seen in Figures 8 and 9 where the simultaneous versus asynchronous turnarounds based on aircraft type and an example of the turnaround time details of an LH_2 aircraft's flight plan are shown, respectively. The "E Dep.", "A Dep.", "A Arv.", and "R Time" refer to the earliest possible departure time, actual departure time, actual arrival time and required refuel time based on refuel amount and are indexed according to which flight they are in the flight plan, i.e the first column of data belongs to the first flight, and so on.

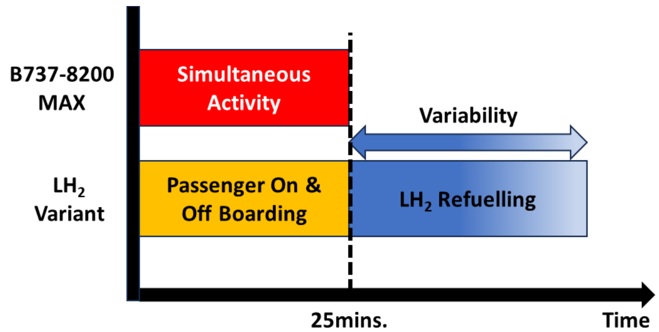

Fig. 8: Simultaneous versus asynchronous turnarounds based on aircraft type

```
Aircraft 1 Additional Time Info:
E Dep.: ['06:00' '08:02' '10:15' '12:59' '15:24' '16:44']
A Dep.: ['06:00' '08:15' '11:45' '14:10' '15:30' '16:50']
A Arv.: ['07:25' '09:40' '12:20' '14:45' '16:05' '17:25']
R Time: ['00:13' '00:12' '00:10' '00:14' '00:14' '00:14']
```

Fig. 9: Asynchronous refuel and turnaround time details

The earliest possible departure time value for the second flight (first value in the second column, Figure 9) is equal to the first actual arrival time (third value in the first column, Figure 9) plus 25 minutes in addition to the second required refuel time (fourth value in the second column, Figure 9). This indicates that the aircraft's second flight cannot be permitted to take-off until the first flight has departed (06:00), landed (07:25), and completed the on and off-boarding of its passengers and refueling (08:02). The second flight departs at 08:15 which is after the earliest possible departure slot, therefore giving the aircraft a slight amount of slack while adhering to the turnaround requirements. The refuel rate is controlled by the refueling infrastructure at airports and the type of technology or energy carrier used. The refuel rate can be set by the user to change between scenarios of faster or slower refueling infrastructure to investigate how turnaround times and the logistics involved therein impact the hybrid sustainability. For the data shown in Figure 10, the refuel rate of the LH_2 aircraft was halved from its value in Figure 9. The drop in refuel rate increases the required refuel time between flights, halving the rate doubles the refuel time in this case as all other parameters remained the same. This can be seen in Figure 10, where the earliest possible departure times for succeeding flights increase proportionally compared to Figure 9 depending on the preceding arrival times and increased refuel times.

```
Aircraft 1 Additional Time Info:
E Dep.: ['06:00' '08:13' '10:26' '13:14' '15:03' '16:43']
A Dep.: ['06:00' '08:15' '11:45' '13:35' '15:15' '16:45']
A Arv.: ['07:25' '09:40' '12:20' '14:10' '15:50' '17:20']
R Time: ['00:27' '00:23' '00:21' '00:29' '00:28' '00:28']
```

Fig. 10: Refuel rate variance and its impact on flight plans

The difference between the earliest possible and actual departure times for a flight is akin to the slack an aircraft has in its flight plan. Larger differences between these times indicates a very flexible schedule, with significant down time between when an aircraft lands and when it is ready to take-off again. Optimization constraint 14 from Section III-B-3 controls how large this gap between actual and earliest possible departure time is allowed to be per flight. In Figure 11, the gap is reduced considerably from its flexible level in Figure 9 to be much more restrictive, therefore enforcing minimal downtime between flights in an aircraft's flight plan. This can be seen between Figures 9 and 11 where the downtime intervals are reduced.

```
Aircraft 1 Additional Time Info:
E Dep.: ['06:00' '08:27' '10:49' '12:25' '14:05' '15:43']
A Dep.: ['06:25' '08:50' '11:10' '12:50' '14:30' '15:50']
A Arv.: ['07:50' '10:15' '11:45' '13:25' '15:05' '16:25']
R Time: ['00:13' '00:12' '00:09' '00:15' '00:15' '00:13']
```

Fig. 11: More restrictive slack in an aircraft's scheduled flight plan

The turnaround of an aircraft is a critical design point which must be minimized for an airline to maximize its utilization and benefit within schedules. Put simply, when aircraft are flying, they are generating income, when they are grounded, they are costing money. The faster, more easily an aircraft can turnaround, the more potentially profitable it is to an airline. Hence, when new or alternative aircraft concepts are being evaluated for their decarbonization potential, understanding the turnaround requirements and limitations during operation is as important as understanding the performance of the aircraft's engine and airframe. While emissions can be tackled on a per flight basis through technology and energy carrier improvements, fleet-wide emissions and economic viability can be assessed by studying how logistical and feasible a conceptual aircraft may be given that it uses a certain type of infrastructure for refueling, or specific energy carrier type or that passenger on/off-boarding must be asynchronous to refueling, etc. Being able to assess a conceptual aircraft beyond its idealized mission to model how it operates in situ and what level of emissions and utilization it sees given various operating scenarios is of paramount importance to objectively evaluating if it may play a role in the decarbonization of the aviation sector. This model enables such analysis and scenario creation. The user has the freedom to restrict or promote certain behaviors, evaluate and compare technologies and other levers on a like-for-like basis and output solutions that reflect the inputs of a given scenario and produce a schedule grounded in logical and deterministic relationships between the variables and design space.

B. Model Benchmarking

The benchmarking of the model developed in this work took place on a Dell Latitude 5540 laptop operating on Windows 11 (Version 24H2) with a 13th Gen Intel(R) Core(TM) i5-1335U (1.30 GHz) processor and 16GB of RAM. The software used to run the model was Visual Studio Code with the Python Interpreter (Version 3.13.3). The benchmarking of the model investigated solver type (PULP_CBC_CMD, SCIP, HiGHs), problem size and combinations thereof to understand the limitations of different solvers and scenario sizes for this model type. The problem size is defined by the product of the number of optimization variables and optimization constraints. As discussed in Section III-B, these variables and constraints are wholly dependent on the inputs and user preferences supplied to the model.

1. Model Benchmarking: Model outputs versus problem size & solver combinations

The initial benchmarking scenario studies are shown in Table 2. Each column (1, 2 and 3) uses the same solver type whereas each row (A, B and C) increases the problem size through user specified permissions according to the process described in Section III-B. In these scenarios, the user specified permissions being altered were the airports in the network each aircraft could permissibly travel to. Scenarios in row A had the most restriction, i.e fewest permissible airports each aircraft could travel to whereas scenarios in row C had the most flexibility per aircraft. The problem size is the number of optimization variables multiplied by the number of optimization constraints and is seen as such in Table 2 where Scenario 1A has 4,425 variables and 8,110 constraints. The results of the scenarios are shown in Table 3, where the total passengers served, model run time, objective function value and the relative gap to the solver's calculated optimal are displayed per scenario.

Table 2: Model Benchmark Scenarios

	1	*2*	*3*
A	Solver = PULP_CBC_CMD 5 x Airports 5 x B737-8200 MAX Size = 4,425 x 8,110	Solver = SCIP 5 x Airports 5 x B737-8200 MAX Size = 4,425 x 8,110	Solver = HiGHs 5 x Airports 5 x B737-8200 MAX Size = 4,425 x 8,110
B	Solver = PULP_CBC_CMD 5 x Airports 5 x B737-8200 MAX Size = 15,342 x 20,804	Solver = SCIP 5 x Airports 5 x B737-8200 MAX Size = 15,342 x 20,804	Solver = HiGHs 5 x Airports 5 x B737-8200 MAX Size = 15,342 x 20,804
C	Solver = PULP_CBC_CMD 5 x Airports 5 x B737-8200 MAX Size = 32,442 x 41,068	Solver = SCIP 5 x Airports 5 x B737-8200 MAX Size = 32,442 x 41,068	Solver = HiGHs 5 x Airports 5 x B737-8200 MAX Size = 32,442 x 41,068

Table 3: Model Benchmark Results

Scenario	***Parameter***	***1***	***2***	***3***
A	Tot. Pax [-]	3,333	3,196	3,333
	Run Time [sec]	21.29	23.99	30.85
	Obj. Value [€]	32,224.62	30,919.49	32,224.62
	Rel. Gap [%]	0.00%	4.48%	0.00%
B	Tot. Pax [-]	4,198	4,320	4,320
	Run Time [sec]	332.47	173.08	124.19
	Obj. Value [€]	40,494.30	41,722.88	41,649.97
	Rel. Gap [%]	5.00%	1.86%	0.77%
C	Tot. Pax [-]	4,713	4,421	5,471
	Run Time [sec]	1,274.98	1,391.82	1,446.03
	Obj. Value [€]	45,981.16	42,525.67	53,307.80
	Rel. Gap [%]	20.00%	35.35%	4.81%

As expected, when the problem size is increased, the run time of the model increases. An upper limit of five, ten and twenty minutes of solver run time (slightly different from model run time) were set for scenarios belonging to rows A, B and C, respectively. The difference between solver and model run time is due to the fact that the model run time, those values quoted in Table 3, include the reading in and configuration of data, the definition of the model, the solving of the problem (the solver run time) and the displaying of results from the solver. A sufficiently optimal solution was found within the solver time limit for Scenarios 1A, 2A, 3A, 1B, 2B, 3B and 3C. A sufficiently optimal solution was not found within the solver time limits for scenarios 1C and 2C. A sufficiently optimal solution was one that was within 5% of the solver's calculated optimal value for the objective function in each scenario. It can be seen that the "Rel. Gap" values, the relative gap between the final objective function value and the calculated optimal value, were all within the 5% tolerance for scenarios that solved within the time limit. Since these problems contain mixed integer and continuous variables, the ideal, calculated optimal objective function values are not always attainable. Similarly, due to the size of the problems, finding the absolute true optimal may require excessive computational times to determine. The purpose of this model type is as a design space characterization and scenario analysis tool. Having reliably accurate and relatively faster results is of much more value in such an application.

Another trend appeared as the problem size increased, the number of passengers serviced increased for all solvers between problem sizes. The passenger traffic vector input was kept constant for all scenarios so the increase in passenger service was due to the increase in maneuverability of the fleet. As previously stated, the problem sizes increased due to the fact that the aircraft were prohibited from visiting some airports in the smaller sized scenarios but these restrictions were relaxed later in the bigger sized scenarios. By giving the model more choice, it became more complicated (problem size), but it is precisely these additional choices, that unlocked the relatively higher passenger servicing (and higher revenue capture).

Interestingly, the solvers' calculated optimal objective function values were usually different between the solver types for a given problem size. The solvers' calculated optimal values (γ_{opt}) can be determined from Table 3 using the objective function value ("Obj. Value", γ_{res}) and the relative gap (g_{rel}) results for each scenario. The formula for doing this calculation is denoted by Equation 35.

$$\gamma_{opt} = \frac{100\%}{100\% - g_{rel}} * \gamma_{res} \tag{35}$$

The difference between the solvers' optimal values for the same problem was most clear for the scenarios in row B of Table 3. The solvers' calculated optimal values were €42,625.57, €42,513.63 and €41,973.16 for the PULP_CBC_CMD, SCIP and HiGHs solvers, respectively. In this row of scenarios, the problem was identical in terms of inputs, user preferences, time limits and relative gap. The only change across the scenarios was the solver. Conversely, the solvers' optimal values were identical despite different solvers being used between Scenarios 1A and 3A. In both scenarios, the model determined the optimal objective function value of €32,224.62 and catered to 3,333 passengers. It must be noted that these scenarios dealt with the smallest sized problem of 4,425 variables by 8,110 constraints. Hence, the true optimal solution to a problem depends on the solver used by the model with slightly different optimal objective functions being set. This difference is much less pronounced for the smallest problem studied in this work with the optimal values being identical between two of the three solvers, but it is useful to understand that the solution generated from this model may vary, even slightly, due solely to the type of solver used.

As previously mentioned, Scenarios 1C and 2C failed to determine a sufficiently optimal solution (within 5% of the calculated optimal and a solver run time within twenty minutes) while Scenario 3C found a solution catering to 5,471 passengers and generating an objective function score of €53,307.80 with a relative gap of 4.81%. These scenarios all used the largest problem size of 32,442 variables and 41,068 constraints. This re-emphasizes the fact that solution time is problem dependent – the large number of variables and constraints proved too complex to solve to the necessary accuracy within the allotted time for Scenarios 1C and 2C, the default and SCIP solvers, respectively. This again alludes to the fact that the optimal solutions from this model are solver dependent. The HiGHs solver was capable of meeting the optimality criteria for this problem size. It must be noted that the model run time of Scenario 3C was 1,446.30 seconds (24 minutes and 6.3 seconds) which accounts for the setup, running and handling of the outputs of the model. It didn't find a sufficiently optimal solution fast, it took a significant amount of time (boarding on the limit of twenty minutes solver time) but it was the only solver that did find a sufficiently optimal answer for row C in Table 3. The HiGHs solver found the optimal answer for the smallest problem size but was bested in run time by the default solver – 21.29 seconds for PULP_CBC_CMD versus 30.85 seconds for HiGHs. The HiGHs solver was the best solver by all metrics for the medium sized problem of 15,342 variables by 20,804 constraints. This alludes to the fact solution accuracy and solution time are solver dependent, furthermore certain problems and scenarios are better suited to different solver types.

2. Model Benchmarking: Alternative aircraft functionality scenarios versus solvers

Scenarios 1A, 2A and 3A were re-run but with three of the five aircraft replaced with the previously discussed LH_2 variant of the B737-8200 MAX. The B737-8200 MAX aircraft had a synchronous turnaround, whereas the LH_2 variant had an asynchronous turnaround alongside the fact that it used an alternative energy carrier. For these re-run scenarios, henceforth referred to as 1A-LH_2, 2A-LH_2, and 3A-LH_2, all scenario parameters were kept the same as Table 2 except the fleet composition. The results for these scenarios are shown in Table 4.

Table 4: LH_2 Variant Introduction Results

Parameter	*Scenario 1A-LH_2*	*Scenario 2A-LH_2*	*Scenario 3A-LH_2*
Tot. Pax [-]	2,707	2,972	2,972
Run Time [sec]	417.19	316.41	327.86
Obj. Value [€]	26,234.86	28,812.97	28,812.97
Rel. Gap [%]	18.00%	6.81%	6.76%

The purpose of introducing the LH_2 variant aircraft was to assess the solvers' ability to establish a baseline scenario (with Scenarios 1A, 2A and 3A all only using B737-8200 MAX aircraft) then be re-run with an adapted, decarbonization scenario (lower CO_2 emission energy carrier with liquid hydrogen versus Jet A-1). The adapted, decarbonization scenario introduced new functionality to the benchmarking problem such as asynchronous turnaround where the feasibility and logistics of chaining flights per aircraft were altered. These scenarios assessed the ability of the solver to adapt and study the problem for alternative technology concepts while capturing and maintaining the necessary fidelity and functionality already established in the prior scenarios.

Scenarios 1A-LH_2, 2A-LH_2 and 3A-LH_2 had significant changes in their solutions compared to the original problems defined in Scenarios 1A, 2A and 3A. The introduction of the LH_2 variant aircraft saw a reduction of 626 passengers and a deficit of €5,989.76 between Scenarios 1A and 1A-LH_2. This drop in passenger service has an associated drop in revenue capture which directly impacts the objective function. Similarly, the alternative aircraft has different costs compared to the Scenario 1A B737-8200 MAX such as the change in fuel cost and CO_2 emissions' cost (due to LH_2 having no CO_2 emissions on a TTW basis). This reduction in passenger service and objective function values between the original and secondary versions of each scenario was replicated for each solver type, as seen by comparing Table 4 with Table 3, row A. Interestingly, none of the solvers from Scenarios 1A-LH_2, 2A-LH_2 and 3A-LH_2 were capable of determining a sufficiently optimal solution according to the previous criteria of 5% relative gap and a solver run time limit of five minutes. While the problem sizes remained the same, the added complexity of the LH_2 variants' turnaround and refueling logistics alongside the changes to the cost parameters meant that the solvers could not determine a sufficiently optimal solution like they had for Scenarios 1A, 2A and 3A. The SCIP and HiGHs (Scenarios 2A-LH_2 and 3A-LH_2, respectively) solvers both determined identical solutions, as seen by the equal passenger numbers served and objective value metrics. However, the trend of different solvers having different calculated optimal objective function values that was seen previously was demonstrated again even though the difference is very minute. Lastly, the model run times all significantly increased between the original and secondary scenarios, going from an average of 25.38 seconds for Scenarios 1A, 2A and 3A to a mean run time of 353.82 seconds for 1A-LH_2, 2A-LH_2 and 3A-LH_2 conveying the increased complexity the fleet compositional change brought.

V. Conclusions

This work developed a discrete optimization model with the goal of assessing different decarbonization pathways for the aviation sector. To do this, the model must be capable of defining different scenarios to be analyzed and as such take inputs of aircraft data, the network of airports and the passenger traffic information. The model takes the form of a decision-making algorithm that aims to schedule flights, moving passengers through the network, in the most holistically sustainable manner. The holistic sustainability objective combines the measure of the economic and environmental viability of the decisions being made in tandem with one-another such that optimal solutions represent the best, hybrid set of decisions that maximize this metric.

The development of this module saw the definition and creation of optimization variables and constraints which curate the specific scenarios to be studied. Section III details the mathematical procedures and logic used to create these scenarios and embed the feasibility necessary to generate schedules of flights which represent the decisions being made by the model's discrete optimization solver. Section IV-A demonstrated this functionality and the ability of the model to create schedules of flights for aircraft with varying feasibility and logistical requirements and limitations. Multiple test cases were run to convey how the different variables and constraints change within the setup of the model given the inputs and user preferences. Similarly, test cases were run to highlight route chaining, fuel consumption and turnaround functionality.

Section IV-B follows on from the functionality demonstration to assess the solvers used by the model. Three solvers were used in this part of the model development, the default PuLP discrete optimization suite solver, PULP_CBC_CMD, and two open-source solvers, SCIP and HiGHs. Twelve scenarios were run of varying problem size and solver type to investigate the impact these parameters have on the model's ability to generate schedules. It was found that the optimal solution for each scenario was solver dependent but that this was less pronounced for smaller problem sizes. The smallest sized problems solved the fastest and with the highest degree of accuracy (minimal relative gap) whereas the largest sized problems only had one sufficiently optimal solution provided by the HiGHs solver. The final part of Section IV-B saw changes to the aircraft data input from previous scenarios to assess how the model and solvers handled more complex and varied scenarios. In all cases of these final three scenarios, the model's outputs changed to reflect the aircraft changes which was indicative of the model faithfully capturing the necessary functionality. However, the run time and relative gap between the optimal and final objective value were outside the previously used, acceptable bounds for those scenarios indicating that the solvers struggled with the increased complexity despite the problem size remaining the same.

In summary, this model is capable of assessing the scenarios of interest but the solutions it generates are highly solver and scenario dependent. The scenario dependence is beneficial as it assesses the inputs, creates variables and constraints that reflect these inputs and generates a schedule that maximizes the holistic sustainability of that scenario. When changes are made, the schedule is adapted logically and according to an objective function with tangible and perceptible cost parameters such that the user can assess what change had what impact and the degree to which this impact was felt – ideal behavior for an assessment tool especially in the aviation sector where the relationship between changes and decarbonization pathways can be unintuitive and complex. Further work is required to investigate if decoupling the solver dependence is achievable. The solvers used in this study were strictly for linear programming type problems and subject to an objective function where an optimal solution was attempted to be found. Extending the solvers investigated to include constraint satisfaction problem (CSP) solvers rather than direct optimizers will be a necessary step. The model may benefit from establishing feasible regions of the design space before seeking optimal solutions, especially for the more complex and larger problem types where the current model found difficulty determining sufficiently accurate and quick solutions. The scope of the work was kept relatively small for development purposes as these types of combinatorial optimization (CO) problems are susceptible to combinatorial explosion where the number of possibilities becomes unwieldy and unmanageable for discrete optimization solvers. Further research into pre-solving and, or reducing the complexity of these problem types while still affording the necessary freedom and functionality is required to scale this model up to a level where more impactful studies and analyses may be conducted.

Acknowledgements

The authors would like to thank Ryanair for their technical support of this study through the use their data.

References

[1] IATA, "Our commitment to Fly Net Zero by 2050", Press release no. 66: 77th IATA Annual General Meeting in Boston, USA 4th Oct 2021. (Last accessed 16 May 2025) Available from: https://www.iata.org/en/pressroom/pressroom-archive/2021-releases/2021-10-04-03/

[2] Destination 2050, "A route to net zero European aviation''. (Last accessed 16 May 2025) Available from: https://www.destination2050.eu/

[3] International Civil Aviation Council, "Future of Aviation". (Last accessed 16 May 2025) Available from: https://www.icao.int/Meetings/FutureOfAviation/Pages/default.aspx

[4] Pratt & Whitney, "Advancing Technology – GTF Engine", Future of Flight – Advanced Technology. (Last accessed 16 May 2025) Available from: https://www.prattwhitney.com/en/future-of-flight/advancing-technology

[5] Boeing, "Meet the 777X", Boeing – Commercial – 777X. (Last accessed 16 May 2025) Available from: https://www.boeing.com/commercial/777x#overview

[6] Airbus, "ZEROe: a hydrogen-powered, fully electric aircraft", Innovation – Energy Transition – Hydrogen. (Last accessed 16 May 2025) Available from: https://www.airbus.com/en/innovation/energy-transition/hydrogen/zeroe-our-hydrogen-powered-aircraft

[7] European Commission, "EU Emissions Trading System (EU ETS) – What is the EU ETS?". (Last accessed 16 May 2025) Available from: https://climate.ec.europa.eu/eu-action/eu-emissions-trading-system-eu-ets/what-eu-ets_en

[8] Ryanair, "Our Fleet" (Last accessed 16 May 2025). Available from: https://corporate.ryanair.com/about-us/our-fleet/

[9] easyJet, "Our destination is net zero – Our pathway to net zero". (Last accessed 16 may 2025) Available from: https://www.easyjet.com/en/netzero

[10] Krömer, M. M., Topchichvili, D., Schön, C., "Sustainable airline planning and scheduling" Journal of Cleaner Production, Vol. 431, 1 Jan 2024. DOI: https://doi.org/10.1016/j.jclepro.2023.139986

[11] Jalalian, M., Gholami, S., Ramezanian, R., “Analyzing the trade-off between CO_2 emissions and passenger service level in the airline industry: Mathematical modeling and constructive heuristic” Journal of Production, Vol. 206, 1 Jan 2019.
DOI: https://doi.org/10.1016/j.jclepro.2018.09.139

[12] Parsa, M., Nookabadi, A. S., Flapper, S. D., Atan, Z., “Green hub-and-spoke network design for aviation industry” Journal of Cleaner Production, Vol. 229, 20 Aug 2019.
DOI: https://doi.org/10.1016/j.jclepro.2019.04.188

[13] Barry, N., Gallagher, C., Fiztgerald, S., Stuart, C. J., “Optimizing for Sustainability as an Objective Function Within Airline Fleet Scheduling: An Ireland – EU Mobility Case Study” AIAA AVIATION 2024-4166, Session: Airport and Airline Operation II, 27 Jul 2024.
DOI: https://doi.org/10.2514/6.2024-4166

[14] Barry, N., Gallagher, C., Fiztgerald, S., Stuart, C. J., “Characterising the role of fleet renewal on the pathway to 2050: A European airline case study”, 34th Congress of ICAS, Sep 2024. (Last accessed 19 May 2025). Available from: https://www.icas.org/icas_archive/icas2024/data/papers/icas2024_0126_paper.pdf

[15] MathWorks, “MATLAB Optimization Toolbox, Solve linear, quadratic, conic, integer, and nonlinear optimization problems”. (Last accessed 7 Nov 2025). Available from: https://uk.mathworks.com/products/optimization.html

[16] Gallagher, C., Stuart, C., and Spence, S. Validation and Calibration of Conceptual Design Tool SUAVE, AIAA 2023-3225. AIAA AVIATION 2023 Forum. June 2023 DOI: https://doi.org/10.2514/6.2023-3225

[17] Gallagher, C., Stuart, C., and Spence, S. Energy Analysis of Fleet Operations Using Green Liquid Hydrogen and Synthetic Kerosene, AIAA 2023-4301. AIAA AVIATION 2023 Forum. June 2023. DOI: https://doi.org/10.2514/6.2023-4301

[18] European Commission – Mobility and Transport, “ReFuelEU Aviation, Key Facts – Aircraft Operators (Anti-tankering practices, Article 5)”. (Last accessed 7 Nov 2025). Available from: https://transport.ec.europa.eu/transport-modes/air/environment/refueleu-aviation_en

[19] IATA, “IATA Carbon Offset Program, Frequently Asked Questions” (Version 10.2, 19 April 2022). (Last accessed 7 Nov 2025). Available from: https://www.iata.org/contentassets/922ebc4cbcd24c4d9fd55933e7070947/icop_faq_general-for-airline-participants.pdf

[20] Mitchell, S., Kean, A., Mason, A., O’Sullivan, M., Phillips, A., Peschiera, F., “Optimization with PuLP – PuLP is an linear and mixed integer programming modeler written in Python”. (Last accessed 26 Nov 2025) Available from: https://coin-or.github.io/pulp/

[21] Achterberg, T., Berthold, T., Koch, T., Wolter, K., “Integration of AI and OR Techniques in Constraint Programming for Combinatorial Optimization Problems”, CPAIOR 2008, LNCS 5015, Pages 6–20, January 2008. DOI: https://doi.org/10.1007/978-3-540-68155-7_4 (Last accessed 26 Nov 2025). Also available from: https://www.scipopt.org/

[22] Huangfu, Q. and Hall, J. A. J., “Parallelizing the dual revised simplex method”, Mathematical Programming Computation, 10 (1), 119-142, 2018. DOI: 10.1007/s12532-017-0130-5 (Last accessed 26 Nov 2025). Also available from: https://highs.info/

[23] Boeing, “737 MAX – Maximize your potential with the only true single aisle airplane family”. (Last accessed 26 Nov 2025). Available from: https://www.boeing.com/commercial/737max

[24] Bureau of Transportation Statistics USA, “Passengers, All Carriers – All Airports” (Last accessed 26 Nov 2025). Available from: https://www.transtats.bts.gov/Data_Elements.aspx?Data=1