

Conceptual Design of Green Propulsive Systems Using Reinforcement Learning

by

Martijn van Dongeren

to obtain the degree of Master of Science
at the Delft University of Technology,

to be publicly defended on Tuesday, 17th February at 14:00

Student number: 5083486
Project duration: May 2025 – February 2026
Thesis committee: Dr.ir. M.F.M Hoogreef, TU Delft, Chair
Dr.ing. F. Orefice, TU Delft, Supervisor
Dr. R.P. Dwight, TU Delft, External member
Faculty: Faculty of Aerospace Engineering, Delft
MSc Track: Flight Performance & Propulsion



Conceptual Design of Green Propulsive Systems Using Reinforcement Learning

Martijn van Dongeren *

Delft University of Technology, Delft, The Netherlands, 2600AA

Hybrid-electric powertrains offer a solution to significantly reduce aircraft emissions in flight. This study presents a method for automatically generating hybrid-electric architectures and optimizing two different objective functions by evaluating the control parameters of each unique architecture using reinforcement learning. Two ATR 72-600 configurations serve as reference aircraft, and three technology levels are considered. When maximizing the ratio of effective radiative forcing to payload mass, the results indicate that the optimal design is sensitive to both aircraft configuration and technology level; however, architectures fully powered by hydrogen fuel cells are preferred when feasible. When maximizing payload and applying the Flightrpath 2050 sustainability goals as constraints, the optimal architecture shifts to one in which conventional jet fuel and hydrogen are combusted in a gas turbine to power the primary propulsive line, while the majority of the power is delivered by the fuel cells to an auxiliary propulsive line. Compared with a conventional architecture, this design reduces CO₂ and NO_x emissions by up to 74% and 86%, respectively, while reducing payload mass by only 24%.

Nomenclature

Greek symbols

Φ = Supply parameter

φ = Shaft parameter

Acronyms

BAT = Batteries

CG = Center of Gravity

CJF = Conventional Jet Fuel

EI = Emission Index

EM = Electric Machine

ERF = Effective Radiative Forcing

FC = Fuel Cell

GB = Gearbox

HEP = Hybrid-Electric Propulsion

MTOM = Maximum Take Off Mass

MTOW = Maximum Take Off Weight

OEM = Operational Empty Mass

PM = Power Management system

SAC = Soft actor-critic

I. Introduction

As global populations and economies continue to grow, demand for aviation is projected to increase by 3-5 percent per year through 2050 [1]. Air transport, however, contributes to global warming through the emission of carbon dioxide (CO₂), nitrogen oxides (NO_x), water vapor (H₂O), and contrail formation [2]. It is estimated that in 2018, CO₂ emissions from aviation accounted for 1.59% of global effective radiative forcing (ERF) and non-CO₂ aviation emissions for 1.91% [3]. To reduce aviation's impact on anthropogenic climate change, in 2011, the European Union outlined its long-term vision in Flightrpath 2050. The sustainability goals include a 75% reduction in CO₂ per passenger kilometer (pkm), a 90% reduction in NO_x emissions, and a 65% reduction in perceived noise compared to the year 2000 [4]. To achieve these goals, it is necessary to significantly reduce aircraft emissions. One potential solution lies in the development of hybrid-electric powertrains (HEP) [5, 6].

The foundation for developing mathematical HEP models was established by Lorenz et al. [7] and Isikveren et al. [8], who defined various hybridization factors necessary to quantify the power demands of HEP architectures. These factors represent the ratio of power supplied to or delivered by a powertrain component to the total power. Building on this foundation, de Vries et al. [9] reformulated these factors to control parameters in their conceptual powertrain model. Specifically, de Vries et al. defined two types: Supply parameters, which represent the ratio of power from a single energy source to the total generated power, and shaft parameters, which represent the ratio of power in a single

*MSc. Student, Faculty of Aerospace Engineering, Delft University of Technology

shaft to the total shaft power. Due to this model's simplicity and flexibility, it is well-suited to conceptual HEP designs and was incorporated by Orefice [10] into the conceptual design workflow for hybrid-electric aircraft. Furthermore, the model originally proposed by de Vries et al. was later extended by Borgia [11], adding hydrogen as a third energy source, demonstrating potential to reduce energy consumption and emissions. In addition to developing conceptual HEP models, efforts have focused on developing an optimization framework for complex powertrain architectures. In the work conducted by Bussemaker et al. [12], a multi-objective architecture optimizer systematically generated candidate architectures, and for each architecture, the relevant continuous design variables were sized using gradient-based optimization, with the sizing loop minimizing a weighted objective combining fuel burn, maximum take-off weight (MTOW), and flight time. The results compared full-electric, hybrid-electric, and conventional architectures, without considering hydrogen alternatives.

To address this gap, this work proposes a conceptual design method that couples generative hybrid-electric architecture design with a reinforcement learning (RL) framework. This method is applied to a retrofit study of the ATR 72-600 for both passenger and cargo configurations, utilizing technology projections for the 2030, 2040, and 2050 periods. Consequently, this work proposes the following primary research question:

How do projected technological advancements and aircraft configurations influence the optimal architecture and control parameters that minimize effective radiative forcing (ERF) per passenger for a regional turboprop aircraft?

To answer this, the following sub-questions are addressed:

- Which hybrid-electric powertrain architectures and control parameters are identified by a reinforcement learning algorithm as optimal for minimizing ERF per passenger?
- What specific thresholds in hydrogen storage density and fuel cell efficiency must be reached to enable a fully hydrogen-electric regional flight profile?
- How do optimal architectures and control parameters shift when the objective is changed to maximize payload capacity, while meeting the Flightpath 2050 environmental goals?

A flowchart of the methodology used to answer these questions is shown in Fig. 1. First, a HEP architecture is stochastically generated and passed to the optimizer, which sets the control parameters for each flight phase. After a single flight is completed, the fuel, component, and payload weights are determined, along with the flight's climate impact. Then, the design's performance is evaluated and passed to the optimizer to improve its decision-making regarding the control parameters. This process is repeated until a specified training threshold is reached, after which the optimizer determines the optimal control parameters for each HEP architecture. The powertrain model used in this study is introduced in section II and the optimization framework is discussed in section III. Furthermore, the reference design and projected component characteristics are presented in section IV. Finally, the study's results are presented in section V.

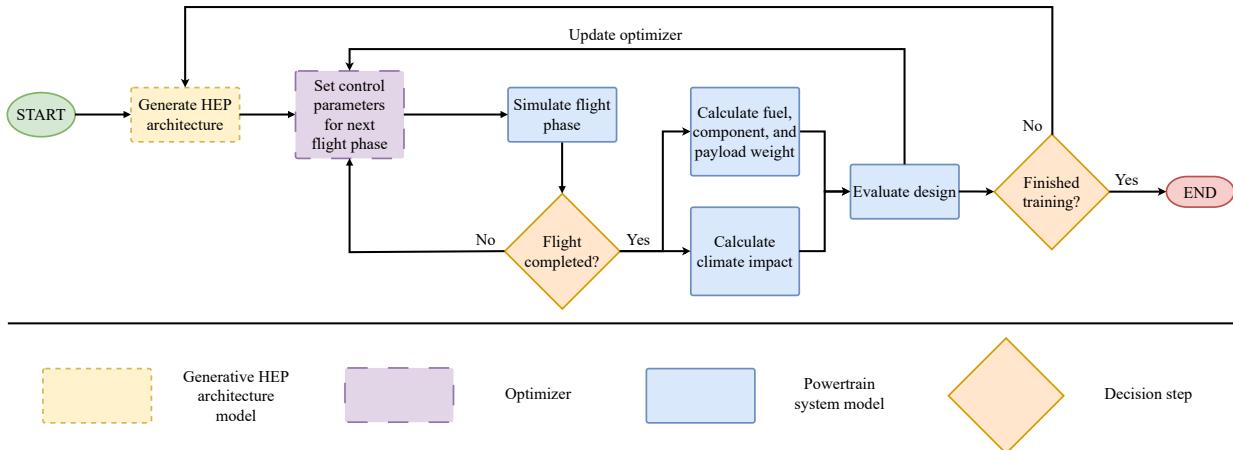


Fig. 1 Schematic of the proposed methodology.

II. Powertrain System Modeling

This section introduces a method for describing a simplified HEP architecture within a flexible mathematical framework that captures interactions among powertrain elements. Then, the flight mechanics equations used to estimate the required propulsive power for a nominal mission are presented. Finally, the metric ERF is used as the primary tool to assess the climate impact of each design.

A. Powertrain Metamodel

In this work, the propulsion system is modeled using simplified representations of its elements, including energy sources, components, and power paths connecting them. The considered elements are similar to those proposed by de Vries et al. [9], and include conventional jet fuel (CJF), a gas turbine (GT), batteries (BAT), a gearbox (GB), a power management system (PM), and various electric motors (EM) and propellers (P). Furthermore, the model incorporates hydrogen (H_2) and fuel cells (FC), aligning with the modifications introduced by Borgia [11]. Of these elements, CJF, H_2 , and BAT are classified as energy sources, and the remaining six elements are classified as components. The elements are linked by mechanical or electrical connections, collectively referred to as power paths. Borgia identified 28 distinct architectures using this set of elements. However, by relaxing specific constraints, such as the requirement that two different energy sources must be indirectly connected, 48 unique architectures are formulated. This is expanded to 120 architectures when increasing the maximum number of auxiliary propulsive lines from one to three on a single wing. The architecture in which all types of energy sources, components, and propulsive lines are present is shown in Fig. 2. The remaining 119 architectures are therefore limit cases of this architecture. A symmetry condition is applied to each architecture, whereby the single-wing configuration is mirrored to represent the aircraft's full powertrain.

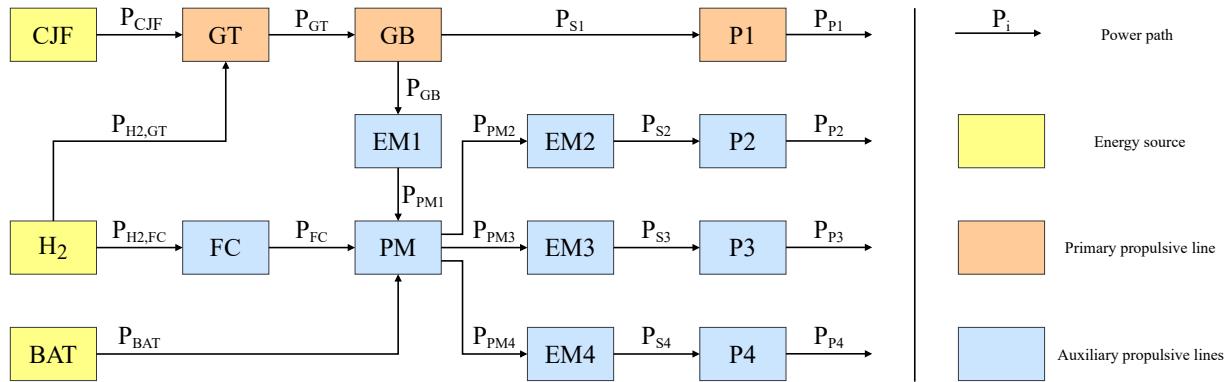


Fig. 2 Architecture consisting of the maximum number of energy sources, components, and propulsive lines.

The architecture shown in Fig. 2 has a total of nineteen different power paths; therefore, nineteen equations are necessary to solve this system. Using the conservation of energy, Eq. 1, an equation is formulated for all thirteen components (not including the energy sources) [9]. An additional equation is formulated by summing the power of each propeller and equating it to the total propulsive power required, see Eq. 2 [9].

$$\eta_{component} \cdot \sum P_{in} = \sum P_{out} \quad (1)$$

$$P_p = 2 \cdot \sum_{n=1}^{N_{Prop}} P_{P_n} \quad (2)$$

The remaining six equations are formulated as control parameters, which are either supply or shaft parameters and are bounded within the domain $[0,1]$ [9]. The supply parameters are the ratios between the power delivered by a single energy source and the total power delivered by all energy sources combined. For a system with N energy sources, there will always be $N - 1$ supply parameters. The architecture shown in Fig. 2 will therefore have three supply parameters, which are defined by Eq. 3, Eq. 4, and Eq. 5.

$$\Phi_{H2,GT} = \frac{P_{H2, GT}}{P_{CJF} + P_{H2, GT} + P_{H2, FC} + P_{BAT}} \quad (3) \quad \Phi_{H2, FC} = \frac{P_{H2, FC}}{P_{CJF} + P_{H2, GT} + P_{H2, FC} + P_{BAT}} \quad (4)$$

$$\Phi_{BAT} = \frac{P_{BAT}}{P_{CJF} + P_{H2, GT} + P_{H2, FC} + P_{BAT}} \quad (5)$$

Finally, the shaft parameters are the ratios between the power delivered by a single shaft and the total shaft power. For the previously discussed architecture, three shaft parameters are required, which are defined in Eq. 6, Eq. 7, and Eq. 8.

$$\varphi_{S2} = \frac{P_{S2}}{P_{S1} + P_{S2} + P_{S3} + P_{S4}} \quad (6) \quad \varphi_{S3} = \frac{P_{S3}}{P_{S1} + P_{S2} + P_{S3} + P_{S4}} \quad (7) \quad \varphi_{S4} = \frac{P_{S4}}{P_{S1} + P_{S2} + P_{S3} + P_{S4}} \quad (8)$$

While the number of supply parameters is directly related to the number of energy sources, this is not the case for the shaft parameters. If in the architecture shown in Fig. 2, the component EM1, were to be removed, as shown in Fig. 3, two power paths disappear, and therefore two equations must also be removed as to not overconstrain the system. Due to the removal of EM1, Eq. 1 is used one time less, and one less shaft parameter is required. The reduction in shaft parameters aligns with the architecture's physical constraints. Since P_{S1} is now dependent only on the summation of P_{CJF} and $P_{H2,GT}$, it is no longer a degree of freedom.

Each unique architecture will have a distinct system of equations that depends on the elements present and their connections. Eq. 9 is the system for the architecture in Fig. 2, and to solve for each power path, the matrix is inverted and multiplied by the vector on the right-hand side.

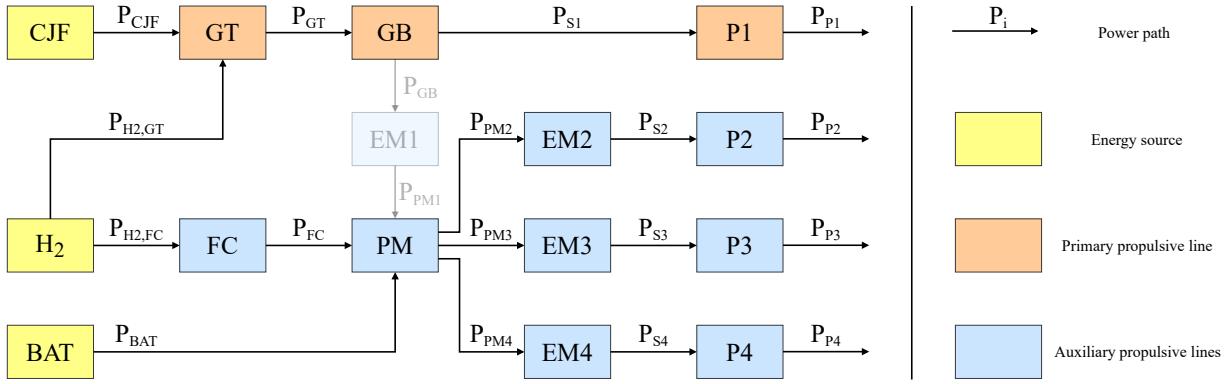


Fig. 3 Architecture where the electric motor between the gearbox and power management system has been removed.

To calculate the weight and volume of each element, the system of equations is solved to determine the power through each power path, which are then doubled to account for the symmetry condition. A distinction is made between energy source and component elements in an architecture. The weight and volume of an energy source are determined by integrating its delivered power over time and dividing by its respective gravimetric and volumetric energy density. While component weight and volumes are calculated by dividing the largest recorded delivered power of each component by its respective gravimetric and volumetric power density. The electric motor and power management components are the two exceptions: instead of the maximum output power, the maximum sum of input power is divided by the density.

For a retrofit application, the total powertrain mass is added to the aircraft's operational empty mass, from which the mass of the original propulsion system has already been subtracted. The remaining available mass is then allocated to the payload, while keeping the maximum take-off mass (MTOM) constant. The fuels, their storage systems, components, and payload must then be distributed throughout the aircraft to ensure the center of gravity (CG) falls within the aircraft's feasible margin to maintain adequate control and stability.

Sizing the powertrain using this method requires that the values of each powerpath be nonnegative; therefore, feasible solutions for each system must satisfy three constraints. First, because both the supply and control parameters express a ratio of a single power path relative to the total supplied or shaft power, the sum of both the supply and shaft parameters cannot exceed one. Secondly, if, for example, in Fig. 2, $\Phi_{BAT} > 0.5$ and $\varphi_{S2} + \varphi_{S3} + \varphi_{S4} < 0.5$, the battery is the main power supplier and the primary shaft delivers the most power to a propulsive element. Power must now flow from the power management system to the gearbox, contrary to the direction of the arrows. As a result, the sign of these power paths becomes negative, which may result in singularities [10, 11]. To avoid this issue, when the value of a power path is negative, its direction is reversed, the system of equations is reconstructed, and the system is solved again. Finally, in the case where the component EM1 is removed, the power through P_{S1} becomes solely dependent on the power paths P_{CJF} and $P_{H2,GT}$, and therefore, one less shaft parameter has to be defined. Although the shaft parameter of this power path is not explicitly defined, its sum with the remaining shaft parameters must not exceed 1; otherwise, the remaining shaft will begin harvesting energy. This is undesirable for certification reasons; therefore, if the sum exceeds 1, the values of the remaining shaft parameters are reduced until the energy-harvesting condition is mitigated.

B. Flight Mechanics for Power Estimation

To solve a system of equations describing one of the possible architectures, it is necessary to define the required propulsive power P_p . In this work, a nominal mission is analyzed that consists of take-off, climb, cruise, and descent, with the assumption that a single value can approximate the propulsive power for each phase. During take-off, the power is estimated by multiplying the normal take-off shaft power by the number of engines and propeller efficiency. For the latter three flight phases, steady state flight is assumed, and Eq. 10, Eq. 11, and Eq. 12 are used to calculate P_p in each phase.

$$P_a = \text{ROC} \cdot W + \frac{1}{2} \cdot C_D \cdot \rho \cdot V^3 \cdot S \quad (10) \quad P_a = P_r = D \cdot V = L \cdot \frac{D}{L} \cdot V = \frac{W}{\frac{D}{L}} \cdot V \quad (11)$$

$$P_a = -\text{ROD} \cdot W + \frac{1}{2} \cdot C_D \cdot \rho \cdot V^3 \cdot S \quad (12)$$

During the climb and descent phases, the average phase-specific power is assumed to be equal to the average of the power determined at the beginning and the end of the phase. The aircraft mass at the beginning of climb, the top of climb, and the end of descent is approximated using the fuel fraction method. The Breguet range equation, Eq. 13 [13], is applied to calculate the fuel consumed during cruise.

$$C_L = \frac{W}{0.5 \cdot \rho \cdot V^2 \cdot S} \quad (14)$$

Here, t denotes the time spent in the cruise phase, and TSFC denotes the thrust-specific fuel consumption. Furthermore, the lift coefficient for each phase is calculated using Eq. 14, and, given the aircraft's drag polar, the drag coefficient can be determined. This method does not aim to predict propulsive power with high precision; rather, it provides a realistic estimate for conceptual analysis.

C. Climate Impact Modeling

The metric chosen to assess the climate impact of each species emitted is effective radiative forcing (ERF). ERF quantifies how species affect the Earth's radiation imbalance at the top of the atmosphere, thereby contributing to climate change [14]. It is preferred over traditional radiative forcing (RF) because it better captures the influence of anthropogenic aerosols and accounts for rapid adjustments in Earth's surface and stratospheric conditions. However, it does not account for ocean and sea-ice responses to the imbalance, which occur over much longer timescales. A different metric that accounts for both short-lived and long-lived forcing components is GWP*, which quantifies warming over multiple decades [3]. However, this metric must be evaluated over a chosen timescale, introducing a temporal bias. Additionally, it should not be applied in scenarios where fuel usage deviates significantly from current trends, such as the potential widespread implementation of HEP systems. Therefore, ERF remains the preferred metric for providing insights into the direct effect of each emitted species on climate forcing and, in turn, on eventual temperature changes.

Lee et al. [3] have calculated the ERF of aviation emission species from kerosene based on the normalized values of ERF per unit emission or distance, specifically for contrail cirrus. These values are tabulated in Table 1 under the ERF column. Positive numbers contribute to net global warming, while negative numbers contribute to net global cooling. While sulfate has a cooling effect, it is associated with the formation of particulate matter, which may pose adverse health effects when exposed to humans [15]. However, this research will only consider the climate effects of aviation emissions. Furthermore, CO, UHC, and soot are omitted as these have a negligible impact on global warming [16]. Lee et al. [3] based their estimations for the annual global emissions of each species on the product of its emission index and fuel burn for each year, where the latter is derived by dividing the estimated annual CO_2 emissions by an emission index of 3.16 [17]. Since global fuel consumption in 2018 was not yet known when Lee conducted his analysis, it was estimated by extrapolating from 2016 data, using an annual growth rate of 3.3%. Global fuel consumption in 2005, 2011, and 2018 is estimated at 229, 242, and 309 million tons, respectively. Furthermore, the ERF per kg of specie is determined by dividing the ERF of each species by its annual emissions, while the value for contrail cirrus in Table 1 is taken directly from Lee et al. [3], as it is based on distance rather than fuel usage. The average is computed across the three years, and the corresponding values are reported in the ERF per kg of specie column of Table 1. When hydrogen is used as fuel, the values for NO_x and contrail cirrus are scaled by factors of 0.76 and 0.31, respectively [18, 19]. While CO_2 emissions are directly related to fuel use and their effects on global warming are independent of location or altitude, the emission indices of other species may vary with the specific aircraft and mission analyzed. Therefore, the ERF per kg of fuel cannot be generalized; it must be determined for each case.

Table 1 ERF of aviation species emissions [3]

	ERF, mW/m^2			ERF per kg of specie, $10^{-6} mW/m^2/kg$		
	2005	2011	2018	CJF	H_2 combustion	H_2 fuel cell
CO_2	25.0	29.0	34.3	0.0359	-	-
H_2O	1.4	1.5	2.0	0.00509	0.0509	0.0509
NO_x	12.9	13.6	17.5	3.86	2.93	-
Sulfate	-5.3	-5.6	-7.4	-19.5	-	-
ERF, mW/m^2			ERF per km traveled, $10^{-6} mW/m^2/km$			
Contrail cirrus	34.8	44.1	57.4	0.000936	0.000290	0.000290

III. Optimization Framework

Optimizing the distribution of power of a simplified hybrid-electric powertrain requires a control strategy that can adapt to various architectures and flight conditions. Consequently, an algorithm must find the optimal set of control parameters for each case that maximizes a certain objective. This section discusses the method for solving the problem at hand and presents the chosen algorithm to address it.

A. Problem Formulation

The power distribution of a hybrid-electric powertrain represents a complex sequential decision-making problem in which optimal control parameters depend on the specific architecture and the current flight phase. Reinforcement learning (RL) is an optimization method that maximizes an objective function by selecting the optimal control parameters. RL is a machine learning approach where an agent learns through trial and error to maximize a specific objective. Other machine learning paradigms, such as supervised and unsupervised learning, aim to make accurate predictions based on labeled data and to uncover hidden patterns or structures in unlabeled data, respectively. However, in uncharted scenarios, where no correct or representative example exists, reinforcement learning offers a distinct advantage by learning from its interactions and experiences [20]. RL problems can be formulated as a Markov decision process (MDP), a mathematical framework where actions (\mathbf{a}_t) influence both immediate rewards and future states (\mathbf{s}_t) [21]. For this application, an architecture is generated stochastically at the beginning of each episode (a flight). Then, for each flight phase (a step), the agent observes the aircraft's status (the state) and outputs the control parameters (the action) to maximize the cumulative reward. The agent must learn to trade off between payload maximization, fuel usage, and climate impact through direct interaction with the powertrain metamodel. To accommodate varying architectures within a single fixed-size interface, the action and state spaces are structured as follows: The Action Space (\mathcal{A}) contains an array of six variables corresponding to the maximum number of supply (3) and shaft (3) parameters. The State Space (\mathcal{S}) consists of a vector providing context to the agent on which it can base its decision of choosing the values for the action space, and the State Space contains the following elements:

- **System Matrix:** A powertrain matrix describing the elements and their connections in a HEP, as shown in Eq. 9. The system matrix has a fixed size of (19×19) representing the maximum possible system dimensions. For limit cases, the powertrain matrix is embedded in the top-left corner, with the remaining elements padded with zeros.
- **System Mask:** Also a (19×19) matrix that contains the active elements of the system matrix. The elements that contain architectural data are set to 1, while the remaining elements are padded with 0.
- **Action Mask:** An array of the same size as the maximum number of control parameters (6). However, most architectures will have fewer control parameters, and the agent uses this array to determine which parameters are relevant. The active control parameters for each architecture will contain a one, while the remaining elements are padded with zeros.
- **Flight Phase Indicator:** An array containing normalized values that contain information regarding the current flight phase, the normalized required power, and the normalized duration of the current flight phase. The first instance is the current flight phase, where 0.2 corresponds to take-off, 0.4 to climb, 0.6 to cruise, 0.8 to descent, and 1 to a completed flight.
- **Violation Array:** The supply and shaft parameters are both constrained such that the sum of both sets may not exceed one. The agent must learn this constraint from the violation array, which contains the absolute difference between each element of the proposed action and the corresponding element of the scaled action. Giving a small penalty for infeasible actions will discourage the agent from selecting such combinations of control parameters.

A simplified example of the state-space vector is shown in Eq. 15, where the size of the System Matrix and Mask is reduced from (19×19) to (3×3) for readability. In this example, the matrix resulting from the architecture contains the elements A, B, C, D , and their locations are denoted by ones in the System Mask. Furthermore, it is assumed that two control parameters are necessary to solve the system. The aircraft is in the cruise phase, and both control parameters were initially set 0.1 too high resulting in a small penalty, as shown in the violation array.

$$\mathcal{S} = \left[\underbrace{\begin{bmatrix} A & B & 0 \\ C & D & 0 \\ 0 & 0 & 0 \end{bmatrix}}_{\text{System Matrix}}, \underbrace{\begin{bmatrix} 1 & 1 & 0 \\ 1 & 1 & 0 \\ 0 & 0 & 0 \end{bmatrix}}_{\text{System Mask}}, \underbrace{\begin{bmatrix} 1 & 1 & 0 & 0 & 0 & 0 \end{bmatrix}}_{\text{Action Mask}}, \underbrace{\begin{bmatrix} 0.6 & \frac{P_{p,\text{current phase}}}{P_{p,\text{max}}} & \frac{t_{\text{current phase}}}{t_{\text{total}}} \end{bmatrix}}_{\text{Flight Phase Indicator}}, \underbrace{\begin{bmatrix} 0.1 & 0.1 & 0 & 0 & 0 & 0 \end{bmatrix}}_{\text{Violation Array}} \right] \quad (15)$$

B. Objective Functions

In conventional RL algorithms, after an action has been completed, i.e., a step has been taken, the new state is determined, and an immediate reward is given. However, for this application, the reward is only calculated at the end of the flight, as powertrain mass, and therefore the payload mass, depends on the fuel usage throughout the entire flight and the

maximum power through each path. As discussed in section I, two different objective functions are defined. The first aims to minimize the ERF per kg of passenger from a single flight, and the second aims to maximize the payload while meeting the Flightpath 2050 goals.

1. Minimization of ERF per passenger

This objective function aims to maximize the ratio provided in Eq. 16a. However, depending on the CG location and payload mass, a different function is used.

$$r_{\text{unscaled}} = \begin{cases} \frac{m_{\text{payload}}}{\text{ERF} \cdot 10^6} & \text{if } m_{\text{payload}} > 0 \text{ and CG in feasible range} \end{cases} \quad (16a)$$

$$r_{\text{unscaled}} = \begin{cases} -\frac{|\text{CG}|}{\text{CG}_{\text{aft}}} & \text{if } m_{\text{payload}} > 0 \text{ and CG outside feasible range} \end{cases} \quad (16b)$$

$$r_{\text{unscaled}} = m_{\text{payload}} \quad \text{if } m_{\text{payload}} < 0 \text{ and CG outside feasible range} \quad (16c)$$

Here, m_{payload} is the payload mass, CG is the location of the center of gravity as a percentage of the mean aerodynamic chord (MAC) of the aircraft with the current HEP design, and CG_{aft} is the aft-most feasible CG location as a percentage of the MAC of the reference aircraft. The reward is then scaled using Eq. 17 to ensure it is bounded in the domain [-10, 10].

$$r_{\text{scaled}} = \begin{cases} \frac{2}{\pi} \cdot 10 \cdot \arctan\left(\frac{r_{\text{unscaled}} - b}{b \cdot 10}\right) & \text{if Eq. 16a used} \end{cases} \quad (17a)$$

$$r_{\text{scaled}} = r_{\text{unscaled}} \quad \text{if Eq. 16b used} \quad (17b)$$

$$r_{\text{scaled}} = \frac{2}{\pi} \cdot 10 \cdot \arctan\left(\frac{r_{\text{unscaled}} - b}{b \cdot 1000}\right) - 5 \quad \text{if Eq. 16c used} \quad (17c)$$

Here b is the unscaled reward that is awarded to the conventional design with 2030 technology projections. Furthermore, to ensure the resulting scaled rewards of each function have distinct regions, 5 is subtracted in Eq. 17c. As a result, the rewards from Eq. 17a are in the domain [-1,10], Eq. 17b are in the domain [-5,-1], and Eq. 17c are in the domain [-5,-10]. Finally, although the majority of the reward is received at the end of the flight, a small penalty, namely the sum of the elements in the violation array, is applied after each flight phase to encourage the agent to choose feasible actions.

2. Payload Maximization

When aiming to meet the Flightpath 2050 goals, the objective function changes, potentially altering the optimal design. As noise is not considered in this work, the Flightpath 2050 goals include a 75% reduction in CO_2 emissions per pkm and a 90% reduction in NO_x emissions relative to the year 2000, based on 2050 technology predictions. To account for these new goals, the objective function is reformulated as presented in Eq. 18.

$$r_{\text{unscaled}} = \begin{cases} 10 \cdot \frac{m_{\text{payload}}}{m_{\text{payload, max}}} & \text{if } m_{\text{payload}} > 0, \text{ CG in feasible range and FP 2050 goals are met} \end{cases} \quad (18a)$$

$$r_{\text{unscaled}} = \Delta\text{CO}_2 + \Delta\text{NO}_x \quad \text{if } m_{\text{payload}} > 0, \text{ CG in feasible range and FP 2050 goals are not met} \quad (18b)$$

$$r_{\text{unscaled}} = -\frac{|\text{CG}|}{\text{CG}_{\text{aft}}} \quad \text{if } m_{\text{payload}} > 0 \text{ and CG outside feasible range} \quad (18c)$$

$$r_{\text{unscaled}} = m_{\text{payload}} \quad \text{if } m_{\text{payload}} < 0 \text{ and CG outside feasible range} \quad (18d)$$

Whereas in Eq. 16, the rewards of designs with high emissions were reduced by dividing the payload mass by the resulting ERF, now the payload mass is divided by the maximum achievable payload mass with 2050 technology projections and multiplied by 10 to ensure it is bounded in the domain [0,10]. If the goals are not met, Eq. 18b is used, where ΔCO_2 and ΔNO_x are formulated in Eq. 19 and Eq. 20, respectively. For designs with significant CO_2 and NO_x emissions during a flight, these values are negative, but become increasingly positive as emissions decrease. This serves

as a constraint on the algorithm, as its reward will be significantly reduced if a design fails to meet the Flightpath 2050 goals.

$$\Delta CO_2 = \min\left(\frac{CO_{2,2000}}{m_{payload, 2000}} \cdot 0.25 - \frac{CO_2 \text{ emissions}}{m_{payload}}, 0\right) \quad (19)$$

$$\Delta NO_x = \min\left(\frac{NO_{x, 2000} \cdot 0.1 - NO_x \text{ emissions}}{m_{payload}}, 0\right) \quad (20)$$

Here, $CO_{2,2000}$, $NO_{x,2000}$, and $m_{payload, 2000}$ are the CO_2 , NO_x emissions and the payload mass calculated for the year 2000 with a conventional powertrain design and assuming a turboprop efficiency of 25%. Furthermore, to ensure that the unscaled reward of each function in Eq. 18 remains in its own distinct region, the rewards are scaled with Eq. 21. If the design is feasible, but the goals are not met, Eq. 21b, the reward is scaled through an exponential function. These rewards may be positive for low values of $r_{unscaled}$, but quickly decrease when the emission goals are further exceeded. The scaling for Eq. 21c and Eq. 21d remains similar to Eq. 17c and Eq. 17d.

$$r_{scaled} = \begin{cases} r_{unscaled} & \text{if Eq. 18a used} \\ (10 \cdot \frac{m_{payload}}{m_{payload, max}} - 1) \cdot e^{10 \cdot r_{unscaled}} & \text{if Eq. 18b used} \\ r_{unscaled} & \text{if Eq. 18c used} \\ \frac{2}{\pi} \cdot 10 \cdot \arctan\left(\frac{r_{unscaled} - b}{b \cdot 1000}\right) - 5 & \text{if Eq. 18d used} \end{cases} \quad (21)$$

C. Reinforcement Learning Algorithm

In the RL algorithm, the agent aims to learn a specific policy $\pi(a_t | s_t)$ that helps it choose an action that maximizes cumulative reward. This policy can either learn from data generated by its current decision-making strategy (on-policy) or from data generated by all its past policies (off-policy). While on-policy agents are more stable, they are less sample-efficient than off-policy agents. Furthermore, the policy bases its choice of actions on the value function that estimates the future rewards based on the state, the state-value function $V(s_t)$, or both the state and the action, the action-value function $Q(s_t, a_t)$ [21]. Finally, the two types of RL methods are model-free methods, which rely on trial-and-error learning, and model-based methods, which use a model to predict future rewards and state transitions. Model-based methods are more sample-efficient, but typically require more training time and have lower asymptotic performance than model-free methods. Model-free methods are preferred when the model is very complex, while model-based methods are advantageous when the model is easier to learn than the policy or when only limited interactions with the environment are possible [21].

The metamodel described in subsection II.A requires, at maximum, six different control parameters, which are continuous values bounded in the domain $[0,1]$. Because there is no constraint on the number of interactions and higher asymptotic performance is desired, a model-free method is more appropriate for this application, specifically the soft actor-critic (SAC) algorithm [22, 23]. This algorithm uses a critic, a state-value function, that evaluates previous actions taken in specific states and the expected reward from the next state. The evaluation is then used to improve the policy, which guides the actor's future actions. To incorporate large, continuous domains, the critic and actor are modeled as neural networks and trained with stochastic gradient descent [23]. Furthermore, SAC is a soft algorithm that maximizes expected reward while also maximizing entropy, thereby encouraging greater exploration. The balance between expected reward and entropy depends on the entropy coefficient α , which had to be set to a specific value in the first version of the SAC algorithm, developed by Haarnoja et al. [22]. However, setting this hyperparameter is not trivial and requires tuning. In the second version, also developed by Haarnoja et al. [23], the algorithm automatically adjusts the coefficient to explore more in uncertain regions but is more deterministic when the optimal action is already clear. Therefore, the second version is used for this application. Although the algorithm is off-policy, it has been demonstrated to be stable and more sample-efficient than other algorithms. When applied to various challenging multi-dimensional continuous control tasks, it consistently outperformed other algorithms, with increasing margins for tasks with more dimensions [23]. Therefore, if the architectures and, thereby, the number of control parameters are expanded in future

research, it is expected that SAC will still outperform other algorithms.

A simplified overview of the SAC algorithm is provided in Fig. 4. First 100,000 steps (25,000 flights) are taken without any training to fill the replay buffer. Then, training is initiated, and after a specified interval of steps, n gradient steps are performed, each updating the neural network's weights n times. Depending on the batch size, a specified number of transitions are sampled from the replay buffer to update the critic. This update is based not only on the expected reward for a given state-action pair but also on the entropy associated with the transition. Then, using the updated critic, the policy is improved, which in turn updates the critic target that is used in the next gradient step. The improved policy also constrains the entropy coefficient, which determines the importance of randomness relative to reward when updating the critic in the next gradient step. This process is repeated for each gradient and training step. The hyperparameters used in this work are similar to those used by Haarnoja et al. [23], except for the ones listed in Table 2.

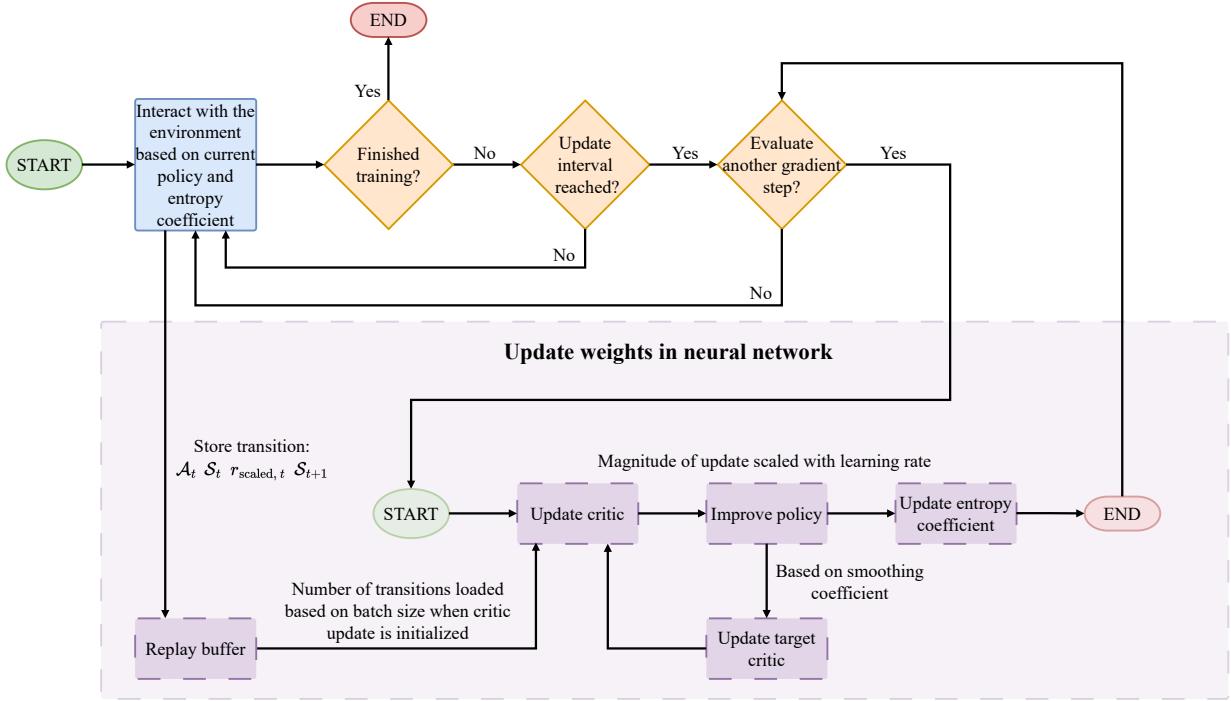


Fig. 4 Schematic of how the SAC model is updated and interacts with the environment.

D. Optimization Verification

While the SAC algorithm does not only aim to maximize its reward, but also entropy by acting as randomly as possible, it does not perform an exhaustive search of the design space. Therefore, the suggested control parameters for each unique architecture are not guaranteed to be global optima. SAC utilizes a stochastic actor to improve stability and exploration in continuous action spaces [23]. However, in design spaces containing steep negative reward gradients near optima, the actor will sample actions near the optima and at the bottom of these gradients. This lowers the expected reward of this region, and the policy may shift to safer areas with higher average returns, potentially moving away from a narrow, sharp optimum.

To identify these narrow, sharp optima and evaluate the performance of the SAC algorithm, the covariance matrix adaptation evolution strategy (CMA-ES) is utilized. This algorithm performs well in black-box continuous optimization problems with non-convex functions [24]. It has been demonstrated to effectively evaluate a large portion of the design space and to outperform other stochastic solvers [25]. CMA-ES optimizes an objective function by stochastically sampling candidate solutions from a multivariate normal distribution. Each sample is evaluated, and the distribution parameters are updated to produce more promising solutions. This process is then repeated for a specific number of generations. The control parameters for each architecture are optimized for 1000 generations using CMA-ES to verify the SAC algorithm's results.

Table 2 Description of relevant changed hyperparameters in the SAC algorithm.

Hyperparameter	Description	Default value	Set value	Motivation for set value
Training frequency	Interval between successive updates. The default is once after every step.	(1, 'step')	(1, 'episode')	Update once after every episode synchronizes updates with the majority of the reward granted at the end of each episode.
Gradient steps	Number of gradient steps that are taken during each update.	1	16	More gradient steps increases sample efficiency by maximizing learning per environmental interaction.
Buffer size	Capacity of the number of transitions that can be used for sampling.	1,000,000	100,000	Decreased buffer size prioritizes more recent data over 'noisy' data from exploration.
Batch size	Collected transitions from the replay buffer for each gradient update.	256	2048	Higher sample size increases computational time, but increases stability.
Learning rate	Step size taken by the optimizer to update the network weights.	0.0003	0.0001	Smaller steps result in higher stability and prevents the policy from collapsing.
Smoothing coefficient	Rate of which the target network weights are updated.	0.005	0.01	Higher value accelerates the integration of the policy into the value estimates.

IV. Case Study: Regional Turboprop Retrofit

The previous two sections discussed a general framework for optimizing the control of hybrid-electric powertrains with a specific objective. This section discusses the chosen reference aircraft to demonstrate the method's application and the projected characteristics of hybrid-electric powertrain elements for three different years.

In this work, two different configurations of the ATR 72-600 are analyzed: A passenger aircraft and a cargo aircraft. In the case of a passenger aircraft that uses hydrogen as fuel, its storage system will be stored in the aft bulkhead for safety reasons. This may, if significant in size, result in the CG shifting too far aft; therefore, the placement of the remaining components must aim to counter this shift, if possible. However, for a cargo aircraft, it is assumed that the hydrogen storage system need not be located in the aft bulkhead, and the components can be arranged more easily to comply with the CG margin. Therefore, the CG analysis is not done for the cargo aircraft case. The reason for analyzing both cases is to highlight potential shortcomings of conventional aircraft design (tube-and-wing) and its limitations when transitioning to sustainable fuels. Furthermore, while a generated architecture can contain up to three auxiliary propulsive lines on a single wing, the effects of distributed propulsion are not considered. Therefore, it is assumed that changing the number of propellers on a wing does not affect the aerodynamic performance.

Additionally, because the hybrid-electric architecture is mirrored across both wings, each side will have its own propulsion system independent of the other. Therefore, the certification requirements for one-engine-inoperative conditions, specifically CS 25.147 and CS 25.149 under the CS-25 certification specifications [26], need not be re-evaluated. If the propulsion system of one wing completely fails, the system will still provide adequate power, similar to the one-engine-inoperative condition of the original ATR 72-600.

A. Reference Aircraft and Mission

The ATR 72-600 has a maximum take-off mass of 23 tonnes and can transport 7400 kg of payload, using 2000 kg of CJF, on a nominal mission. Its two engines weigh 1000 kg in total, and the fuel system mass is calculated by averaging

the methods of Raymer [13] and Torenbeek [27], yielding 57 kg. Subtracting the mass of the engines and fuel system, the operational empty mass (OEM) without the propulsion system, of 12,543 kg, is obtained.

Furthermore, the maximum landing weight is 22,350 kg; therefore, at least 650 kg of fuel must be used if the aircraft takes off at maximum weight. However, in hybrid-electric powertrains, gravimetric fuel consumption during nominal flight is lower because the battery mass is nearly constant and hydrogen has a high gravimetric energy density. To ensure a safe landing, three different design considerations may be implemented. First, to account for higher landing loads, the landing gear can be resized, and the resulting mass increase is deducted from the payload mass. Additionally, the landing distance must be re-evaluated, and the increased distance would impose stricter requirements on available landing locations. As this option requires retrofitting beyond the propulsion system and limits the aircraft's operational envelope, it is not implemented. Secondly, a constraint can be implemented, requiring at least 650 kg of fuel to be consumed during a nominal flight. This forces the reinforcement learning algorithm to explore balancing the use of conventional and sustainable fuels and, if successful, verifies its ability to find designs in more complex landscapes. However, applying this constraint will result in designs with higher fuel consumption, favoring less efficient powertrains. Therefore, the final option of limiting the maximum take-off weight by removing payload is implemented, such that the landing weight requirement is not exceeded.

1. Power Estimation

The normal take-off shaft power of each engine used for the ATR 72-600 is 1,846 kW [28]. Assuming a propeller efficiency of 80%, the power output is 2.95 MW. Furthermore, the rate of climb (ROC) is estimated at 1350 feet/min and 300 feet/min at the beginning and top of climb, respectively, while the ROD is constant at 1500 feet/min [29]. The velocity during climb and descent is calculated from the constant indicated airspeed of 170 and 220 kts, and the velocity during cruise is also constant at 270 kts [30]. Finally, the drag coefficient is determined from the ATR 72 drag polar, constructed by Vecchia [31]. The relevant parameters to calculate the required power are reported in Table 3, and the required power itself is listed in Table 4.

Table 3 Flight phase parameters

Flight phase	Mass, kg	ρ , kg/m ³	V, m/s	C_L	C_D
Beginning of climb	22770	1.225	87.46	0.7814	0.05003
Top of climb	22428	0.5686	128.4	0.7697	0.04944
Cruise	21638	0.5686	138.9	0.6342	0.04336
Beginning of descent	20847	0.5686	166.1	0.4272	0.03635
End of descent	20743	1.225	113.2	0.4251	0.03629

Based on data from the ATR 72-600 brochure [30], the total flight time for a 740 nm mission is estimated at 173.5 minutes. By subtracting the estimated durations for the take-off, climb, and descent phases, summarized in Table 4, the remaining time is attributed to cruise.

Table 4 Power required and time of each flight phase.

Flight phase	P_p	Time in phase	P_p full-electric flight
Take-off	2.95 MW	0.5 min	2.95 MW
Climb	2.47 MW	21 min	2.50 MW
Cruise	2.02 MW	146 min	2.09 MW
Descent	0.874 MW	6 min	0.788 MW

The fuel fraction method, used to estimate the weight during each phase of flight, accounts for weight reductions due to fuel consumption. However, the reduction in weight is smaller when hydrogen or batteries are used as energy sources than with conventional jet fuel. Therefore, the calculated powers in Table 4 may be underestimated during climb and cruise, and overestimated during descent. To quantify the maximum value of this error, while assuming full-electric flight, where the mass remains constant at 23,000 kg, the powers are recalculated and listed in the rightmost column of

Table 4. Ultimately, this results in a 3.1% increase in the total energy required. Because this represents the extreme case, full-electric flight without adhering to the maximum landing weight constraint, actual deviations are expected to be smaller across the design space. This suggests that the effect of weight variation, within the bounds of the reference case, on propulsive power is minimal. Therefore, it is assumed that the powers calculated using the fuel fraction method apply to all designs.

2. Stability and Controllability Verification

To achieve adequate controllability and stability, the centre of gravity (CG) of the aircraft must be within 10% and 39% of the mean aerodynamic chord (MAC) [32].

A simplified side view of the aircraft is presented in Fig. 5, where the datum line is defined 2.362 meters in front of the nose of the aircraft. The fuel, engines, and payload are assumed to be a point mass, and their distances with respect to the datum line are taken from the Weight and Balance manual [32] and reported in Table 5. Using a value for the center of gravity of 25% of the MAC [32], the arm of the OEM without the propulsion system is calculated.

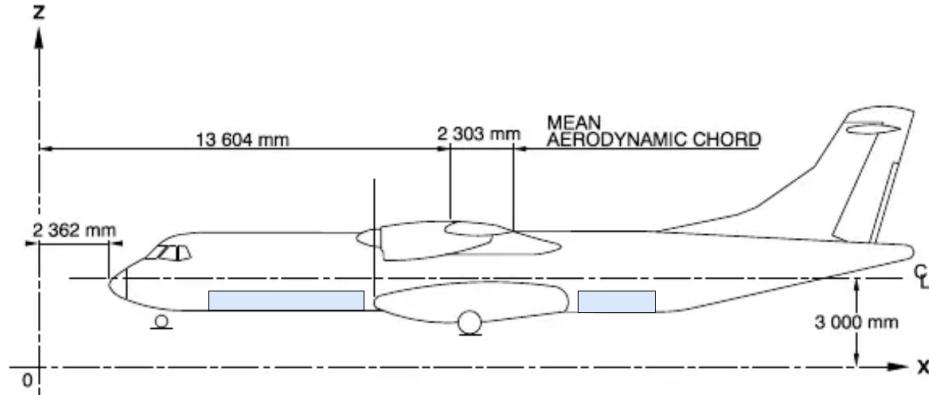


Fig. 5 Side view of the ATR 72-600 [32].

Table 5 Component mass and distance to the datum line of the original ATR 72-600

Component	Mass, kg	Arm, m
Fuel	2000	14.455
Fuel system	57	14.43
Engines	1000	12.7
Payload	7400	14.755
OEM without propulsion system	12,543	13.91

Components such as engines, electric machines, and propellers are all placed on the leading edge (LE) of the wing and will have the same arm as the original engines. Furthermore, CJF and its storage system remain inside the wing, and their CG will move slightly forward depending on the amount of fuel, as reported in the manual [32]. As previously discussed, the hydrogen and its storage system are located in the aft bulkhead for safety reasons, and the CG is slightly dependent on their respective volumes. The batteries, fuel cells, and power management systems can be stored in the compartments located in front or behind the landing gear, in the shaded regions shown in Fig. 5, or in the wing if space is available. In the first iteration, the batteries, fuel cells, and power management system will be loaded from front to back to counter the CG shift caused by the placement of the hydrogen storage system. If this results in the CG being too far forward, the components will be loaded from the rear to the front. The full logic of loading these three types of components is explained in Fig. 6. Additionally, due to a potentially heavier propulsion system, passenger capacity may decrease, leaving more cabin space available for passengers. In the front-loading analysis, the passengers are distributed over the cabin with a seat pitch of 29 inches, similar to the original layout. However, in the aft-loading analysis, the seat pitch is increased to distribute passengers along the entire length of the cabin. The range of distances from the mean and

the respective locations of each component are listed in Table 6, as well as the locations of the compartments used for component storage.

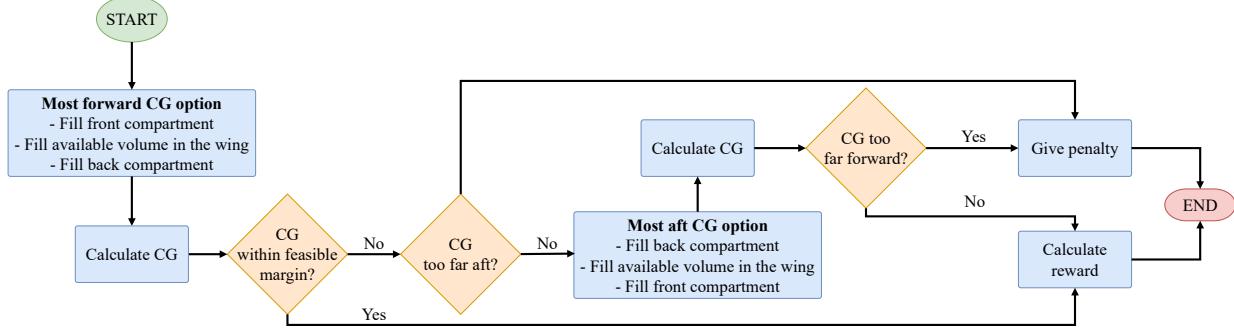


Fig. 6 Placement logic of the batteries, fuel cells, and power management system.

Table 6 Distance to the datum line and location of hybrid-electric aircraft components.

Component or compartment	Arm, m	Location
OEM without propulsion system	13.91	-
Engines and EM's	13.1	LE
CFJ and storage system	14.5-14.6	Wing
H ₂ and storage system	20.9-23.7	Aft bulkhead
Payload	8.49-14.8	Fuselage
Front compartment	5.7-11	Fuselage
Wing	14.7	Wing
Back compartment	18-21	Fuselage

3. Climate Impact Analysis

Lee et al. [3] estimate the effective radiative forcing (ERF) of water vapor only at stratospheric altitudes; below these altitudes, water vapor does not significantly contribute to global warming. As the ATR 72-600 cruises at an altitude of 7.3 km, water vapor emissions are not accounted for throughout the flight. A similar conclusion is drawn for contrail cirrus, which form when humid exhaust plumes from gas turbines mix with cold ambient air. Persistent contrail formation typically occurs between 8-13 km [33-35], below this range, the ambient temperature is not sufficiently cold. While no study has directly compared persistent contrail formation between turbofans and turboprops, it is estimated that the former is responsible for 99.4% of all persistent contrail formation and the latter only for 0.36% [36].

Furthermore, as the throttle setting significantly affects NO_x emissions [37], a statistical model developed by Filippone et al. [38] is used to estimate a more accurate emission index based on the gas turbine output power for each flight phase, rather than the maximum power. This model uses the throttle setting and pressure ratio (OPR) of a turboprop to estimate the NO_x emission index in g/kg, accounting for pressure, temperature, and humidity, as tabulated in Table 7. The OPR is assumed to have a constant value of 14 for all flight phases, which is within the range of the PW127 engine suggested by Yuksel et al. [39].

Finally, the effective radiative forcing per kg of fuel burnt for each species considered is listed in Table 8. Because the NO_x emission index depends on throttle setting, a range based on the minimum and maximum values is provided for each fuel type. It is assumed that water vapor emissions and persistent contrail formation have negligible effects on warming during this specific mission profile; therefore, operating the fuel cell is climate-neutral.

Table 7 Ambient conditions in each flight phase.

Flight phase	Average altitude, m	Temperature, K	Pressure, Pa	RH [40]
Take-off	0	288.15	101325	90%
Climb	3658	264.37	64437.5	45%
Cruise	7315	240.60	39272.1	20%
Descent	3658	264.37	64437.5	45%

Table 8 Effective radiative forcing for each modeled species and fuel type.

	ERF per kg of fuel used, 10^{-6} mW/m ² /kg	
	CJF	H ₂ combustion
CO ₂	0.1134	-
Sulfate	-0.0234	-
NO _x	0.0233 - 0.0645	0.0177 - 0.0490

B. Technology Projection Levels

For component sizing, it is necessary to define the energy or power density of each component. In this work, future technology projection levels for 2030, 2040, and 2050 are presented and summarized in Table 9.

CJF has a gravimetric and volumetric energy density of 43.2 MJ/kg [41] and 33.91 MJ/L [32], respectively, and it is assumed that these will not increase in the future. When the storage system's weight is also accounted for, the fuel's gravimetric density decreases to 42 MJ/kg. Furthermore, the gravimetric energy density of hydrogen itself is a significant 120 MJ/kg; its storage system is of significant weight. For liquefied storage, the gravimetric and volumetric energy densities are estimated at 9 MJ/kg [42] and 6.4 MJ/L [43], respectively, for 2030. Cry-compressed technology is not yet deemed feasible for 2030 due to safety regulations, but can be considered from 2040 onward. It is assumed that, due to future research and development, its energy densities increase by 20% and 30% in 2040 and 2050, respectively, compared to 2030, reaching 12 MJ/kg and 6 MJ/L [42]. Furthermore, Tiede et al. [44] predicted, based on historical data and the practical limitations of each battery type, the state-of-the-art gravimetric energy density for 2030, 2040, and 2050 under conservative, nominal, and aggressive technology-advancement rates. The expected gravimetric energy density and efficiency of the nominal cases are tabulated in Table 9. de Vries et al. [45] assume a charge and discharge rate of 1.2C, meaning the power density, in W/kg, is 1.2 times larger than the energy density, in Wh/kg. The same rate is used in Table 9 to estimate the power density for each year. It is estimated that by 2035, lithium-ion batteries will achieve volumetric energy densities between 600-800 Wh/L, lithium-sulfur batteries between 300-350 Wh/L, and lithium-oxygen batteries between 1000-1600 Wh/L [46]. However, the uncertainty for the latter two is deemed high, whereas that for the lithium-ion type is low. Therefore, an energy density of 600 Wh/L, the conservative side of the 2035 projection, is assumed to be realized by 2030. The ratio between the gravimetric and volumetric energy densities in 2030 is then used to estimate the densities in 2040 and 2050, which are tabulated in Table 9.

The turboprop thermal efficiency is estimated to be in the range of 0.25-0.35 [39, 47, 48]. For this study, it is assumed that efficiency increases from 0.3 in 2030 to 0.35 in 2050. For the gas turbine's continuous power density, a value similar to that of the original ATR 72-600 engines is used [28]. Two types of fuel cells are considered for aviation applications: proton exchange membrane fuel cells (PEMFCs) and solid-oxide fuel cells (SOFCs). However, SOFCs operate most efficiently under steady-state conditions and are therefore not suitable for regional aircraft [49]. For this reason, they are not further considered in this work. Current estimates place power densities of a PEMFC stack at 3.0 kW/kg and 3.4 kW/L [50], which decrease to 1.0-1.5 kW/kg [51] and 0.35 kW/L for the entire system [52]. Therefore, system-level power densities of 1.1 kW/L and 0.35 kW/L, with an efficiency of 55%, are considered reasonable targets for 2030. And it is assumed that power densities will continue to improve to 20% and 30% by 2040 and 2050, respectively, while efficiency will only reach its current upper bound of 60% [49]. Furthermore, permanent magnet synchronous machines (PMSMs) are currently the most widely used electric machines in electric and hybrid electric vehicles. They are considered the most attractive option for aviation, due to their high specific power and efficiency [53-55]. The expected power densities and efficiencies with a normal confidence level for 2030, 2040, and

2050 are reported in Table 9 [53]. The power management system densities are estimated at 30 kW/kg and 70 kW/L in 2030 [49, 56], with an efficiency of 99% [57]. Again, both densities are assumed to increase by 20% and 30% in 2040 and 2050, respectively. Finally, de Vries [57] assumes a 96% efficiency for the gearbox, while its weight and that of the propeller are neglected, which has an efficiency of 80% [58].

Table 9 Overview of the projected technology levels considered in this work.

Year	$\eta_{component}$			Energy density, MJ/kg MJ/L			Power density, kW/kg kW/L		
	2030	2040	2050	2030	2040	2050	2030	2040	2050
CJF Storage	100%	100%	100%	42 34	42 34	42 34	-	-	-
H ₂ Storage	100%	100%	100%	9 6.4	14 7.2	16 7.8	-	-	-
Batteries	89%	90%	90%	1.4 2.2	1.8 3.2	2.2 3.4	0.47	0.61	0.73
GT	30%	33%	35%	-	-	-	3.77	3.77	3.77
FC	55%	58%	60%	-	-	-	1.1 0.35	1.3 0.42	1.4 0.46
EM	97%	98%	98%	-	-	-	13	20	24
PM	99%	99%	99%	-	-	-	30 70	36 84	39 91
GB	96%	96%	96%	-	-	-	-	-	-
Propeller	80%	80%	80%	-	-	-	-	-	-

V. Results

This section discusses the results obtained by training the reinforcement learning model for a certain number of steps on stochastically generated architectures. First, the model's performance when minimizing ERF per kg of payload is discussed, and the proposed designs for various technology levels and aircraft configurations are interpreted and compared with those proposed by the CMA-ES algorithm. The model's performance for the second objective function, maximizing payload while meeting Flightpath 2050, is then evaluated and compared with the CMA-ES algorithm's output.

A. Minimizing ERF per kg of payload

The RL model performance, when trained to minimize ERF per kg of payload, for both passenger and cargo aircraft, is presented in Fig. 7. It is observed that the scaled reward varies with different combinations of energy sources, technology projections, and number of training steps. It is observed that training the model has been successful, as it outperforms random selection of control parameter values (zero training steps) across all architectures and technology projections. Training the model with more steps results in higher rewards. However, this effect generally diminishes after 70,000 or 150,000 steps. Furthermore, in Fig. 7a, the model rarely identifies a significantly improved design relative to the baseline conventional design, which has a scaled reward of zero. The mass breakdown of this design is shown in Fig. 8a. Architectures containing only hydrogen are generally too heavy, or it is not possible to arrange the components such that the CG falls within the feasible margin; the battery-powered architectures are always too heavy. The architecture of the best design identified by the RL algorithm for a passenger aircraft under 2030 technology projections is shown in Fig. 9. The architecture consists of a primary propulsive line powered by CJF combustion, an auxiliary propulsive line powered by hydrogen fuel cells, and an electric motor that connects the two lines. In Fig. 9, the power paths' values for each phase are indicated in black. It shows that supplying fuel cells with hydrogen in each phase reduces ERF by 45.9%, whereas payload decreases by 37.9% relative to the conventional design, as shown in Fig. 8b. However, the CMA-ES algorithm has identified a further-optimized design, indicated in red in Fig. 9. The architecture is similar to that found in RL but now uses hydrogen combustion and eliminates the connection between the gearbox and the power management system, thereby enabling more efficient power distributions. Moreover, using both hydrogen and CJF combustion reduces the ERF resulting from CO₂ emissions, accompanied by a further weight increase of the hydrogen storage system. Additionally, excessive hydrogen use will violate the CG constraint, thereby limiting the amount that can be used. In this design, the CG is located at the very aft of the feasible margin. Overall, the reduction of ERF is greater than that of payload, as shown in Fig. 8c, which indicates a further 42.8% decrease in ERF and a 17.8% decrease

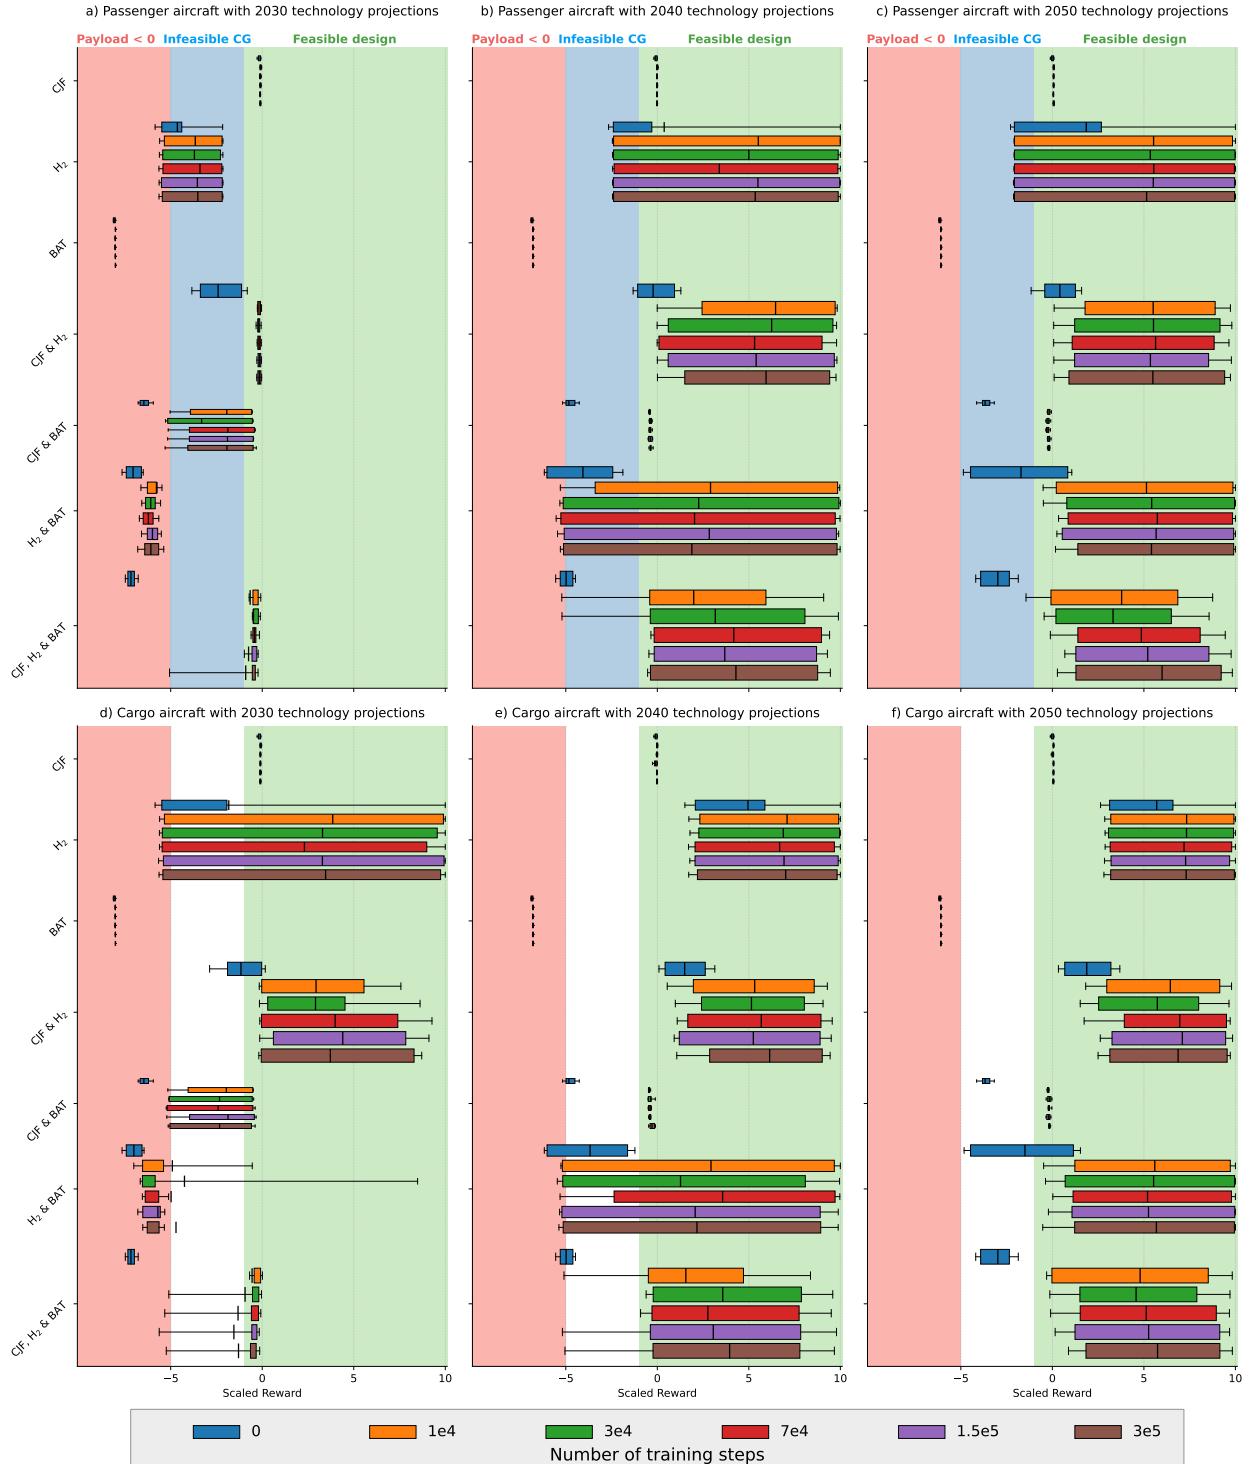


Fig. 7 Optimization results of the first objective function showing the scaled reward distribution for passenger and cargo aircraft architectures, categorized by energy sources under 2030, 2040, and 2050 technology projections. The box plots represent the interquartile range (25th–75th percentiles) with whiskers extending to the 5th and 95th percentiles. The box color indicates the number of training steps; the box height is proportional to the number of unique architectures in that category; and the background shading denotes design feasibility.

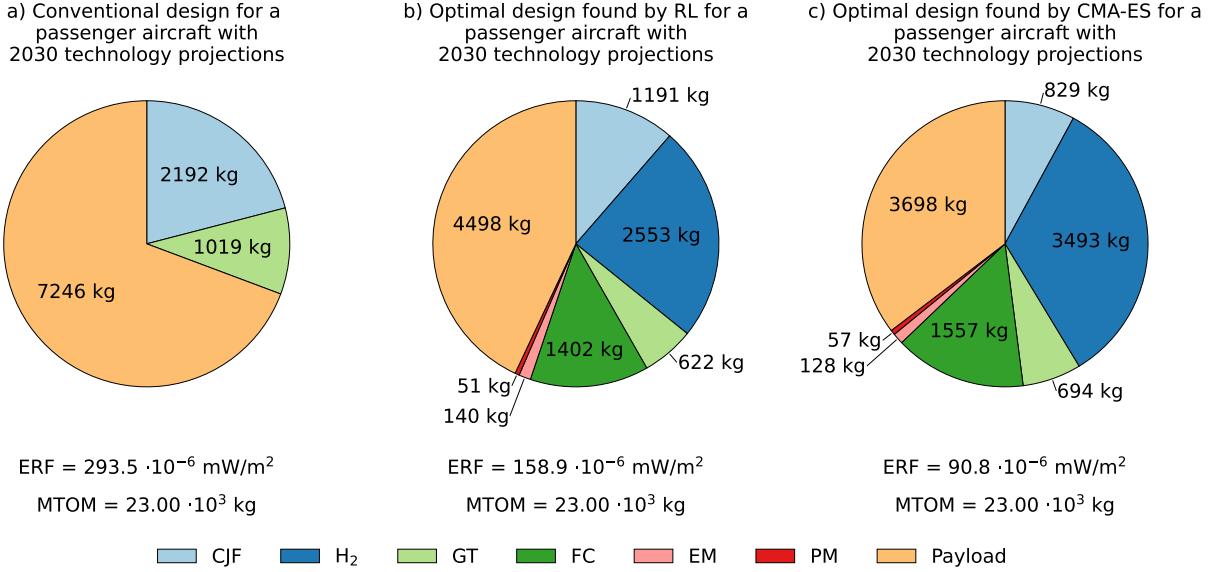


Fig. 8 Mass breakdown of powertrain designs for a passenger aircraft with 2030 technology projections.

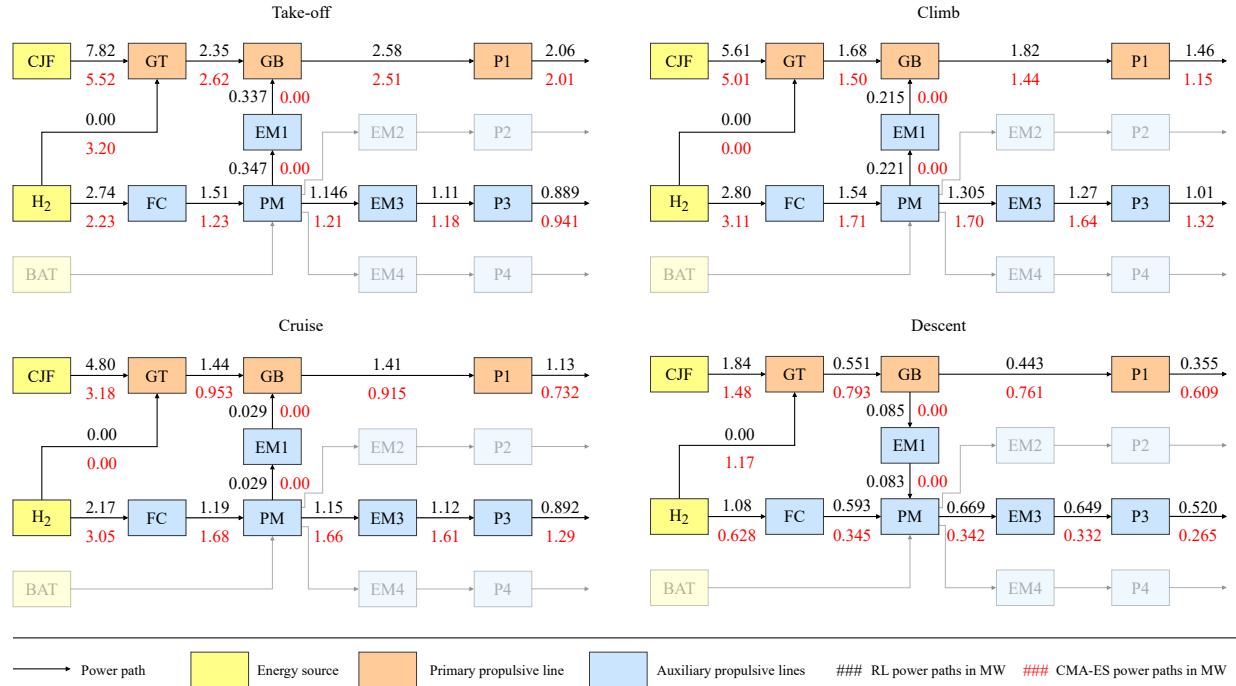


Fig. 9 Optimal architectures and corresponding powers through each power path, in MW, in each phase of flight for a passenger aircraft with 2030 technology projections. The black numbers indicate the optimal designs identified by the RL model, and the red numbers indicate those identified by the CMA-ES algorithm.

in payload relative to the RL design.

With future technological improvements, the model achieves a maximum reward of 10 for architectures containing hydrogen for a passenger aircraft, as shown in Fig. 7b and Fig. 7c. This is primarily driven by improved fuel cell efficiency and higher energy densities of the hydrogen storage system, resulting in lower hydrogen requirements and overall lighter storage systems. While some of these architectures achieve the maximum reward, the rewards are

distributed across a wide range. This is the result of using two different methods to convert hydrogen into power: combustion in a gas turbine or its use in fuel cells. Whereas the former has efficiencies between 30-35% and still emits NO_x , contributing to global warming, the latter's efficiency is between 55-60% and completely climate neutral. This effect is shown in Fig. 10a, Fig. 10b, and Fig. 10c, where instead of differentiating all the architectures by their energy sources, the presence of a fuel cell now separates the two types of combinations. It is clear that, under the 2040 and 2050 technology projections, the maximum reward can only be achieved by using fuel cells. Therefore, the optimal architectures for these years are ones fully powered by hydrogen through fuel cells.

For a cargo aircraft, the location of the hydrogen storage system is not constrained, and therefore, the propulsion system elements can always be arranged in such a way that the CG falls within the feasible margin. While, the CG constraint limits further hydrogen use for passenger aircraft, as already observed in Fig. 7a. For the cargo configuration, the model finds the maximum reward with 2030 technology projections, for architectures fully powered by hydrogen fuel cells, as shown in Fig. 7d. However, as the number of energy sources increases, the model cannot consistently achieve maximum rewards when possible and does not do so across the three energy source architectures. Additionally, the potential of conventional and battery-powered architectures remains unchanged and does not improve relative to the baseline. Furthermore, the performance difference between architectures with or without a fuel cell remains significant, and maximum obtained rewards remain solely obtainable in architectures fully powered by fuel cells, as shown in Fig. 10d, Fig. 10e, and Fig. 10f. Finally, the model performance of the cargo aircraft, as shown in Fig. 7e and Fig. 7f, exhibits trends similar to those of the passenger aircraft. The main difference between these two configurations is that, for a cargo aircraft, the minimum reward obtained by architectures containing hydrogen is shifted from the negative region (infeasible CG design) to the positive region (feasible design).

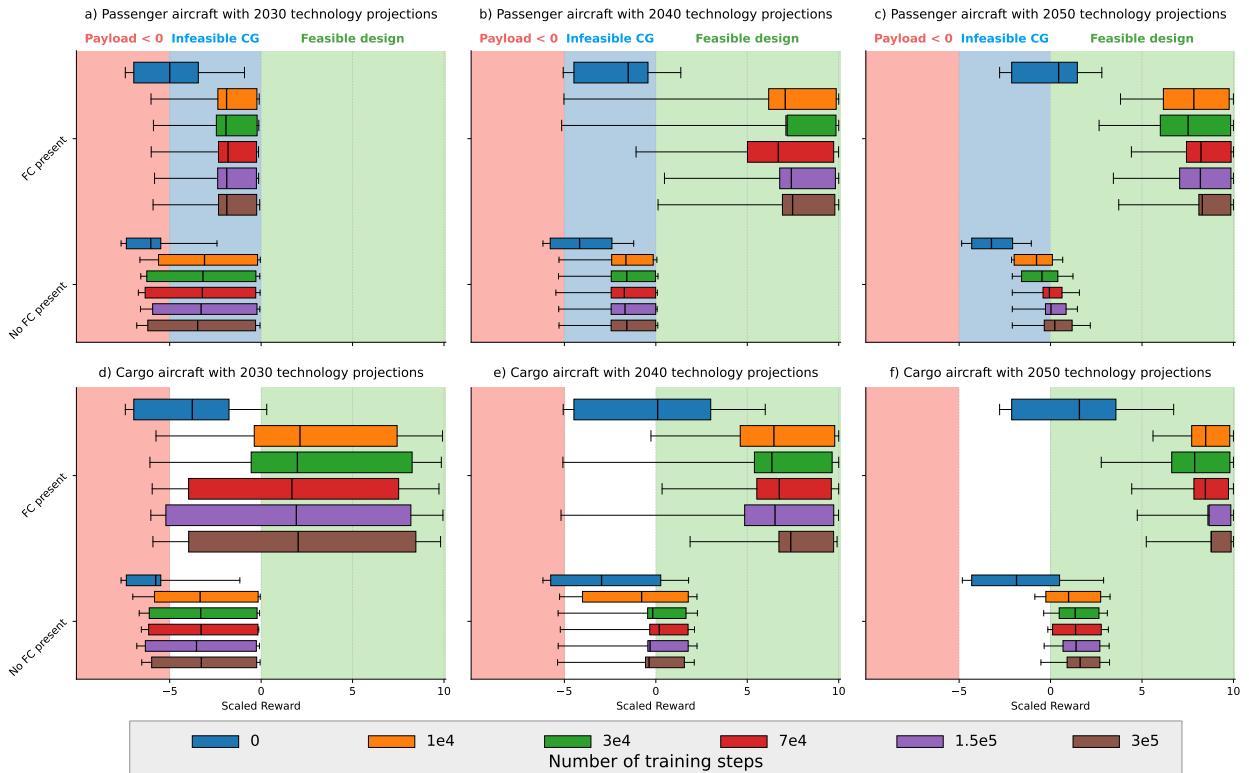


Fig. 10 Optimization results of the first objective function showing the scaled reward distribution for passenger and cargo aircraft architectures containing hydrogen, categorized by the presence of a fuel cell under 2030, 2040, and 2050 technology projections. The box plots represent the interquartile range (25th–75th percentiles) with whiskers extending to the 5th and 95th percentiles. The box color indicates the number of training steps; the box height is proportional to the number of unique architectures in that category; and the background shading denotes design feasibility.

For passenger aircraft under 2030 technology projections, an architecture combining CJF and hydrogen achieves the highest reward. This changes for cargo aircraft or under future technology projections, where the optimal architecture is one fully powered by hydrogen fuel cells, achieving the maximum reward. The mass breakdown of these cases is presented in Fig. 11. Compared to Fig. 8, when switching to architectures powered only by hydrogen fuel cells, the payload decreases significantly due to the weight of the hydrogen storage system and fuel cells. Compared with the optimal design identified by CMA-ES for a passenger aircraft in 2030, the payload decreases by 80.0%, but because only fuel cells are used, the climate impact is zero, and a maximum reward is achieved. However, it is expected that, with increasing technology levels in 2040 and 2050, the payload can increase to 51% and 60% of the conventional design's value in 2030, respectively, while the flight remains climate-neutral. Furthermore, because hydrogen has a high gravimetric density, these designs do not use enough fuel during a nominal flight, resulting in aircraft that exceed the ATR 72-600's maximum landing weight. Therefore, a portion of the payload is removed prior to take-off to ensure that the landing weight requirement is met. This reduction is shown in Fig. 11 by incomplete pies and a decrease in MTOM.

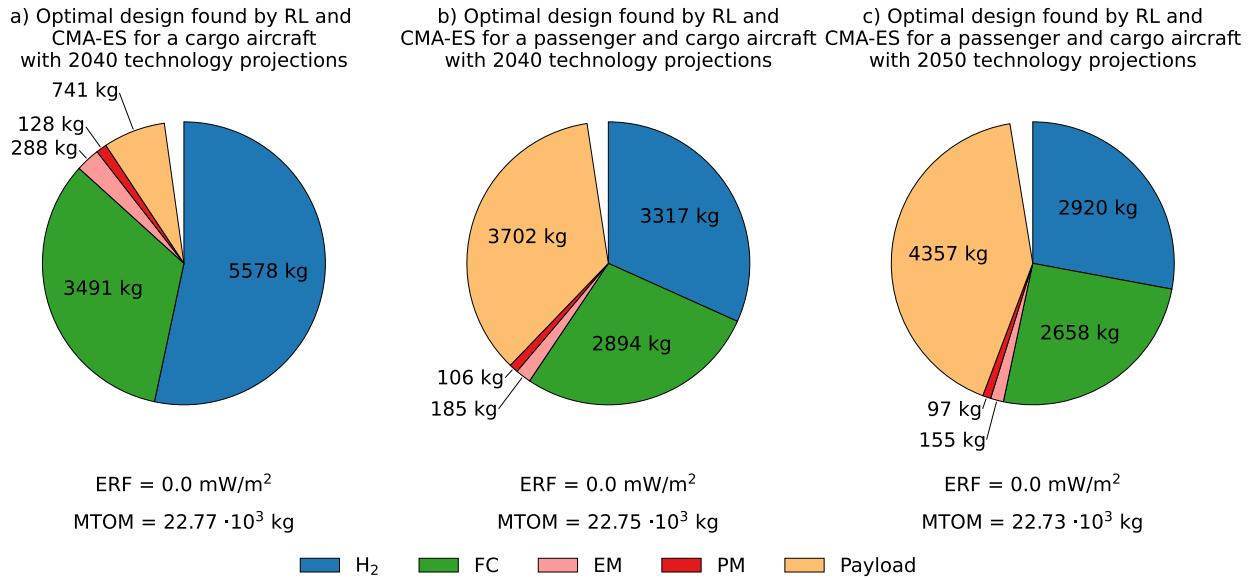


Fig. 11 Mass breakdown of the optimal powertrain designs for the passenger and cargo aircraft with 2040 and 2050 technology projections.

The reinforcement learning model successfully identified the optimal design across all future technology levels for cargo configurations and, for passenger configurations, across 2040 and 2050. Whereas, for a passenger configuration with 2030 technology levels, the optimal design was not achieved. For designs in which the RL model has found an optimum, the objective function is formulated such that the agent is rewarded for increasing hydrogen fuel cell utilization within the architecture and not punished when it is maximized. However, for the 2030 passenger case, a delicate balance is required to maximize the objective. Using excessive hydrogen results in a heavy storage system that shifts the CG too far aft, rendering the design infeasible and penalizing the agent. During training, it was observed that rewards very close to the optimum were achieved, but this policy could never be stabilized. Due to the stochastic nature of SAC, it would sample actions close to the optimum, including designs in which the CG has shifted too far aft. Therefore, the average reward for this policy decreases, and the critic forces the policy to a safer region, away from the optimum.

1. Sensitivity Analysis

A very distinctive switch in the optimal architecture and control strategy between 2030 and 2040 for a passenger aircraft is noticed when comparing Fig. 8 and Fig. 11. This is a result of the 2040 technology projections reducing the weight of the hydrogen storage system in the aft bulkhead relative to 2030, thereby reducing the moment arm that initially causes the CG to shift into the infeasible region. This is also highlighted in Fig. 12, where the different design regions are visualized as a function of fuel cell efficiency and gravimetric energy density of the hydrogen storage system. The powertrain design, powered entirely by hydrogen fuel cells with 2030 technology projections (indicated by the plus),

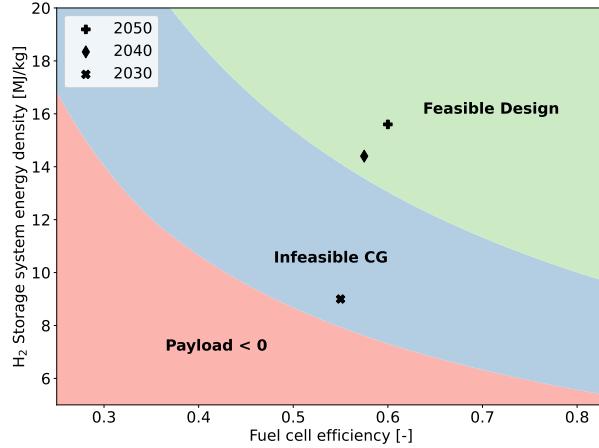


Fig. 12 Design feasibility regions as a function of fuel cell efficiency and H₂ storage system gravimetric energy density, where the background shading denotes design feasibility. The markers indicate projected efficiency and density for 2030 (✖), 2040 (♦), and 2050 (+).

falls within the infeasible CG region and is close to the negative payload region. Here, its capacity to carry payload is severely hindered by the mass of the hydrogen storage system and the fuel cells, as already observed in Fig. 11b. However, unlike powertrains powered solely by batteries, the negative payload region is never reached when hydrogen fuel cells supply all power. Increasing the fuel cell efficiency to its maximum theoretical value of 83% [59] does not allow the design to escape the infeasible CG region. Therefore, it is necessary to improve hydrogen storage system technology to enable a feasible retrofit design for the ATR 72-600. The powertrain designs based on hydrogen fuel cell architectures, with 2040 and 2050 projected technology levels, are indicated by the diamond and cross in Fig. 12. Both designs, due to their proximity to the feasibility boundary, are likely to move into the infeasible CG region if the projected technologies for both years are not realized.

B. Maximizing payload and meeting Flightrpath 2050 goals

Considering the payload maximization objective, the model was again trained for the same number of steps as the first objective function under 2050 technology projections. The results identified a single optimal design, and the corresponding architecture, along with the value of each power path for each phase of flight, is shown in black in Fig. 13. This shows a control strategy similar to that in Fig. 9, in which the gas turbine is used during phases with high propulsive power requirements, whereas in other phases the relative share of fuel cell output power increases. It is also observed that the power delivered by the fuel cell remains relatively constant, minimizing redundant fuel cell weight during phases when it does not deliver its maximum power. The mass breakdown of this design is shown in Fig. 14a, which shows a 26.5% increase in payload compared to the design, which was minimized for ERF per kg of payload, as shown in Fig. 11c. However, the RL model has overshot the Flightrpath 2050 goals, reducing CO₂ emissions per pkm by 83.9% and NO_x emissions by 90.2%, whereas 75% and 90% reductions would have sufficed. While the agent was previously punished for overusing hydrogen fuel cells and the CG not being within the feasible margin, it is now punished for overusing CJF, resulting in noncompliance with the Flightrpath 2050 goals. Therefore, it has opted for a safe design with limited payload to reduce the risk of heavy penalties.

The design identified by the CMA-ES algorithm, which maximizes payload while meeting the Flightrpath 2050 goals, is shown in Fig. 14b. Again, this algorithm has achieved higher rewards than the RL model, demonstrating that a further 4.52% increase in payload is possible. This is achieved by maintaining a constant fuel cell output power, while the remainder of the required power is supplied by hydrogen combustion during take-off and climb and by CJF combustion during cruise. Because NO_x emissions are lower with hydrogen combustion in a gas turbine than with CJF, hydrogen combustion is beneficial for phases with higher propulsive power requirements. Conversely, an alternative control strategy in which CJF is used during take-off and climb instead of hydrogen, and that hydrogen is reallocated to the cruise phase, such that the design is similar to Fig. 14b, results in an increase of NO_x emissions, and noncompliance with the Flightrpath 2050 goals. Optimizing the distribution of power, as shown in red in Fig. 13, reduces CO₂ emissions per pkm by 77.9% and NO_x emissions by 90%, indicating that the NO_x goal limits further increases in payload.

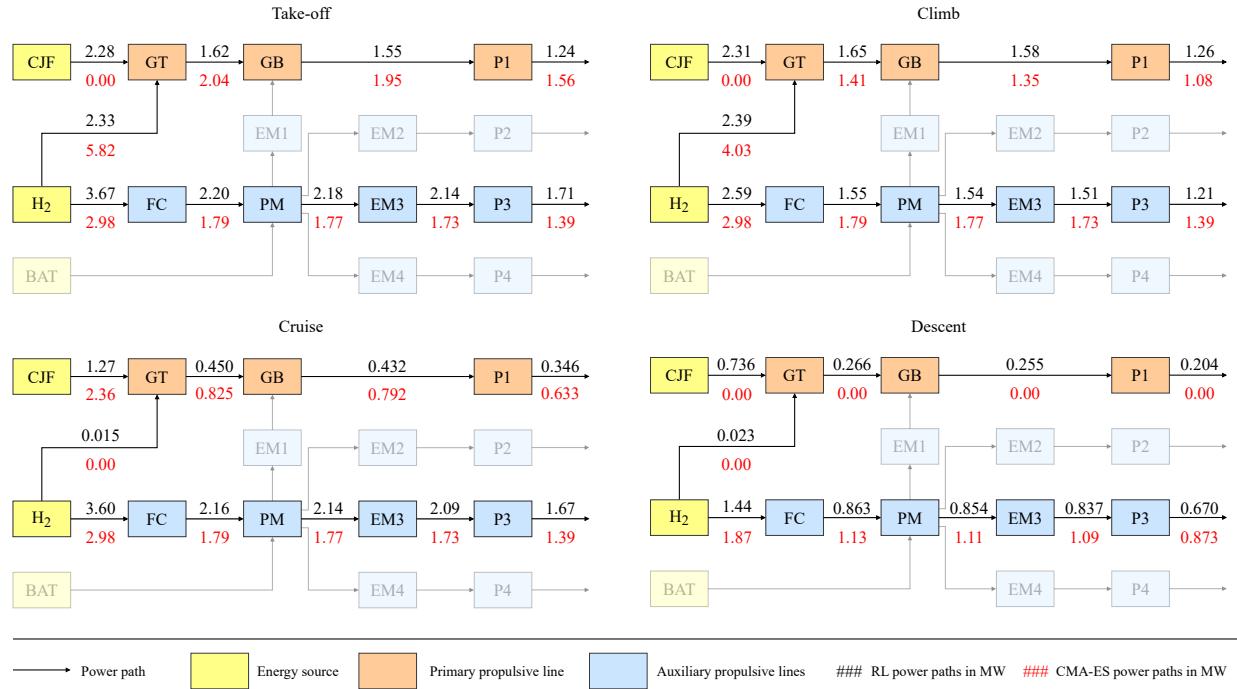


Fig. 13 Optimal architecture and corresponding powers through each powerpath, in MW, in each phase of flight. The black numbers indicate the optimal design found by the RL model and the red numbers the optimal design found by the CMA-ES method.

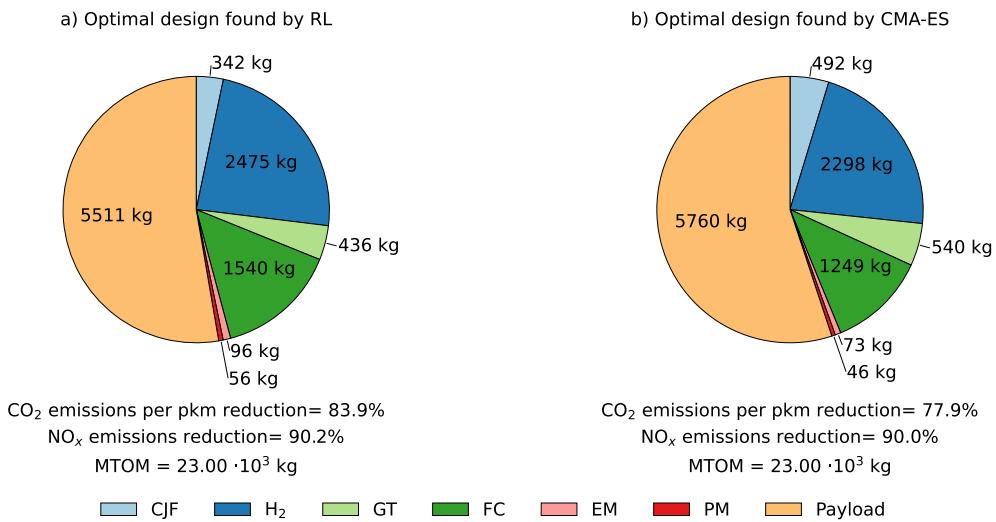


Fig. 14 Mass breakdown of the optimal powertrain designs using RL and CMA-ES that meet the Flighthpath 2050 goals.

VI. Conclusions and Future Outlooks

This work proposes a generative algorithm for creating hybrid-electric powertrain architectures coupled with reinforcement learning to optimize the conceptual design of propulsive systems. This method is applied to retrofit the ATR 72-600 propulsion system, considering two different objective functions. The results indicate that the optimal architecture, when minimizing effective radiative forcing per kg of payload, and the corresponding control strategy, depend on the projected technology level and aircraft configuration. For a passenger aircraft at 2030 technology levels, an architecture that primarily uses conventional jet fuel supplemented with hydrogen is most suitable. Here, the amount of hydrogen is limited by its heavy storage system, which may move the center of gravity too far aft. However, for cargo configurations or future technology levels, the optimal architecture shifts to a fully hydrogen fuel cell electric powertrain. Unless hydrogen storage system density increases significantly, this architecture is not feasible for a passenger aircraft at 2030 technology levels, indicating that the CG margin of tube-and-wing aircraft may limit the adoption of sustainable fuels. Furthermore, the optimal architecture shifts again when the objective is to maximize payload while meeting the Flightpath 2050 sustainability goals. Similarly, in the passenger configuration with 2030 projections, this architecture uses a combination of conventional jet fuel and hydrogen combustion, with the majority of power now delivered by hydrogen fuel cells. The payload mass is allowed to increase by 32% relative to a fully hydrogen fuel cell electric design in 2050, without exceeding the emission thresholds.

While this method provides a quick, straightforward propulsion design strategy for hybrid-electric architectures, several adjustments can further improve it. Firstly, accounting for variable component efficiencies may yield different optima. Specifically for passenger aircraft in 2030, electrically assisting the primary shaft may be beneficial during phases when the gas turbine operates in less efficient regions. Secondly, the effects of distributed propulsion are not accounted for, limiting the benefit of using multiple auxiliary propulsive lines. Additionally, sizing the thermal management system is not included in this conceptual design method, which would increase the drag. In the optimization framework, the reinforcement learning algorithm struggled to produce optimal designs when the designs were near a feasibility constraint. Therefore, reformulating the objective function or further tuning of the hyperparameters may improve learning. Furthermore, this research could be expanded by including sustainable aviation fuel as an energy source, which would require a full life-cycle analysis of all elements. Finally, analyzing other aircraft types, such as turbofans, would also provide valuable insights into how mission characteristics affect the optimal architecture and corresponding control parameters.

References

- [1] Fuel Cells and Hydrogen 2 Joint Undertaking, “Hydrogen-powered aviation – A fact-based study of hydrogen technology, economics, and climate impact by 2050,” Tech. rep., Publications Office, 5 2020. <https://data.europa.eu/doi/10.2843/471510>.
- [2] Lee, D. S., Fahey, D. W., Forster, P. M., Newton, P. J., Wit, R. C. N., Lim, L. L., Owen, B., and Sausen, R., “Aviation and global climate change in the 21st century,” *Atmospheric Environment*, Vol. 43, No. 22, 2009, pp. 3520–2537. <https://doi.org/10.1016/j.atmosenv.2009.04.024>.
- [3] Lee, D. S., Fahey, D. W., Skowron, A., Allen, M. R., Burkhardt, U., Chen, Q., Doherty, S., Freeman, S., Forster, P., Fuglestvedt, J., Gettelman, A., Léon, R. D., Lim, L., Lund, M. T., Millar, R., Owen, B., Penner, J., Pitari, G., Prather, M., Sausen, R., and Wilcox, L. J., “The Contribution of Global Aviation to Anthropogenic Climate Forcing for 2000 to 2018,” *Atmospheric Environment*, Vol. 244, 2021. <https://doi.org/10.1016/j.atmosenv.2020.117834>.
- [4] European Commission and Directorate-General for Research and Innovation and Directorate-General for Mobility and Transport, *Flightpath 2050 – Europe’s vision for aviation – Maintaining global leadership and serving society’s needs*, Publications Office, 2011. <https://data.europa.eu/doi/10.2777/50266>.
- [5] Voskuijl, M., van Bogaert, J., and Rao, A. G., “Analysis and Design of Hybrid Electric Regional Turboprop Aircraft,” *CEAS Aeronautical Journal*, Vol. 9, No. 1, 2017, pp. 15–25. <https://doi.org/10.1007/s13272-017-0272-1>.
- [6] Salem, K. A., Palaia, G., and Quarta, A. A., “Review of Hybrid-Electric Aircraft Technologies and Designs: Critical Analysis and Novel Solutions,” *Progress in Aerospace Sciences*, Vol. 141, 2023. <https://doi.org/10.1016/j.paerosci.2023.100924>.
- [7] Lorenz, L., Seitz, A., Kuhn, H., and Sizmann, A., “Hybrid Power Trains for Future Mobility,” *62. Deutscher Luft- und Raumfahrt Kongress (DLRK)*, 2013, pp. 1–17. <http://dx.doi.org/10.13140/2.1.2741.7925>.
- [8] Isikveren, A., Kaiser, S., Pornet, C., and Vratny, P., “Pre-Design Strategies and Sizing Techniques for Dual-Energy Aircraft,” *Aircraft Engineering and Aerospace Technology*, Vol. 86, No. 6, 2014, pp. 525–542. <https://doi.org/10.1108/aeat-08-2014-0122>.
- [9] de Vries, R., Brown, M. T., and Vos, R., “A Preliminary Sizing Method for Hybrid-Electric Aircraft Including Aero-

Propulsive Interaction Effects,” *2018 Aviation Technology, Integration, and Operations Conference*, 2018, pp. 1–29. <https://doi.org/10.2514/6.2018-4228>.

[10] Orefice, F., “Conceptual Design of Hybrid-Electric Aircraft,” Ph.D. thesis, Università degli Studi di Napoli Federico II, Napoli, Italy, 4 2022. <https://doi.org/10.5281/zenodo.6463971>.

[11] Borgia, A., “Double-Hybrid Powertrain Modelling and Integration in the Conceptual Design Method,” Master’s thesis, Politecnico di Torino, 2025. <http://webthesis.biblio.polito.it/id/eprint/35123>.

[12] Bussemaker, J., Sánchez, R. G., Fouda, M., and Boggero, L., “Function-based Architecture Optimization: An Application to Hybrid-electric Propulsion Systems,” *INCOSE International Symposium*, Vol. 33, No. 1, 2023, pp. 251–272. <https://doi.org/10.1002/iis.2.13020>.

[13] Raymer, D. P., *Aircraft Design: A Conceptual Approach*, sixth Edition, American Institute of Aeronautics and Astronautics, Inc., Reston, VA, 2018. <https://doi.org/10.2514/4.104909>.

[14] Myhre, G., Shindell, D., Bréon, F.-M., Collins, W., Fuglestvedt, J., Huang, J., Koch, D., Lamarque, J.-F., Lee, D., Mendoza, B., Nakajima, T., Robock, A., Stephens, G., Takemura, T., and Zhang, H., *Anthropogenic and natural radiative forcing*, Cambridge University Press, Cambridge, UK, 2013, Chap. 8, pp. 659–740. <https://doi.org/10.1017/CBO9781107415324.018>.

[15] Kapadia, Z. Z., Spracklen, D. V., Arnold, S. R., Borman, D. J., Mann, G. W., Pringle, K. J., Monks, S. A., Reddington, C. L., Benduhn, F., Rap, A., Scott, C. E., Butt, E. W., and Yoshioka, M., “Impacts of Aviation Fuel Sulfur Content on Climate and Human Health,” *Atmospheric Chemistry and Physics*, Vol. 16, No. 16, 2016, pp. 10521–10541. <https://doi.org/10.5194/acp-16-10521-2016>.

[16] Fahey, D. W., and Lee, D. S., “Aviation and Climate Change: A Scientific Perspective,” *Carbon & Climate Law Review*, Vol. 10, No. 2, 2016, pp. 97–104. <https://www.jstor.org/stable/44135212>.

[17] IEA, “CO₂ emissions in aviation in the Net Zero Scenario, 2000–2030,” 2025. URL <https://www.iea.org/data-and-statistics/charts/co2-emissions-in-aviation-in-the-net-zero-scenario-2000-2030>, accessed: 2025-06-23.

[18] Ortúñol, M. S., Yin, F., Rao, A. G., Vos, R., and Proesmans, P.-J., “Climate Assessment of Hydrogen Combustion Aircraft: Towards a Green Aviation Sector,” *AIAA SciTech Forum*, 2023, pp. 1–22. <https://doi.org/10.2514/6.2023-2513>.

[19] Burkhardt, U., Bock, L., and Bier, A., “Mitigating the Contrail Cirrus Climate Impact by Reducing Aircraft Soot Number Emissions,” *npj Climate and Atmospheric Science*, Vol. 1, No. 37, 2018. <https://doi.org/10.1038/s41612-018-0046-4>.

[20] Sutton, R. S., and Barto, A. G., *Reinforcement Learning: An Introduction*, 2nd ed., MIT Press, 2018. <https://mitpress.mit.edu/9780262039246/reinforcement-learning/>.

[21] Lonza, A., *Reinforcement Learning Algorithms with Python*, Packt Publishing, Brimingham, UK, 2019. <https://www.oreilly.com/library/view/reinforcement-learning-algorithms/9781789131116/>.

[22] Haarnoja, T., Zhou, A., Abbeel, P., and Levine, S., “Soft Actor-Critic: Off-Policy Maximum Entropy Deep Reinforcement Learning with a Stochastic Actor,” *ArXiv*, Vol. <https://doi.org/10.48550/arXiv.1801.01290>, 2018.

[23] Haarnoja, T., Zhou, A., Hartikainen, K., Tucker, G., Ha, S., Tan, J., Kumar, V., Zhu, H., Gupta, A., Abbeel, P., and Levine, S., “Soft Actor-Critic Algorithms and Applications,” *ArXiv*, Vol. <https://doi.org/10.48550/arXiv.1812.05905>, 2019.

[24] Nomura, M., and Shibata, M., “cmaes : A Simple yet Practical Python Library for CMA-ES,” *ArXiv*, 2024. <https://doi.org/10.48550/arXiv.2402.01373>.

[25] Rios, L. M., and Sahinidis, N. V., “Derivative-Free Optimization: A Review of Algorithms and Comparison of Software Implementations,” *Journal of Global Optimization*, Vol. 56, No. 3, 2012, pp. 1247–1293. <https://doi.org/10.1007/s10898-012-9951-y>.

[26] European Union Aviation Safety Agency (EASA), “Certification Specifications and Acceptable Means of Compliance for Large Aeroplanes (CS-25), Amendment 28,” Tech. Rep. ED Decision 2023/021/R, European Union Aviation Safety Agency, 2023. <https://www.easa.europa.eu/en/document-library/certification-specifications/cs-25-amendment-28>.

[27] Torenbeek, E., *Synthesis of Subsonic Airplane Design*, Springer Netherlands, Dordrecht, 1982, 1982. <https://doi.org/10.1007/978-94-017-3202-4>.

[28] European Union Aviation Safety Agency, “Type Certificate Data Sheet No. IM.E.041 – PW100 Series Engines,” Tech. rep., EASA, 2023. Retrieved from <https://www.easa.europa.eu/en/document-library/type-certificates/engine-cs-e/easaime041-pratt-and-whitney-canada-pw100-series>, Accessed: 2025-06-16.

[29] DEMOQUAS, “Designing, Manufacturing and Operating Quantification of Uncertainties to increase Aviation Safety,” 2025. URL <https://www.demoquas.eu/>, accessed: 2025-07-10.

[30] ATR, “ATR 72-600,” 2025. URL <https://www.atr-aircraft.com/aircraft-services/aircraft-family/atr-72-600/>, accessed: 2025-06-16.

[31] Vecchia, P. D., “Development of Methodologies for the Aerodynamic Design and Optimization of New Regional Turboprop Aircraft,” Ph.D. thesis, Università degli Studi di Napoli Federico II, Napoli, Italy, 4 2013.

[32] Avions de Transport Régional, *WBM Weight and Balance Manual*, ATR, 2022.

[33] Megill, L., and Grewe, V., “Investigating the limiting aircraft-design-dependent and environmental factors of persistent contrail formation,” *Atmospheric Chemistry and Physics*, Vol. 25, No. 7, 2025, pp. 4131–4149. <https://doi.org/10.5194/acp-25-4131-2025>.

[34] Schumann, U., “Formation, Properties and Climatic Effects of Contrails,” *Comptes Rendus Physique*, Vol. 6, No. 4-5, 2005, pp.

549–565. <https://doi.org/10.1016/j.crhy.2005.05.002>.

[35] Park, J., "Seasonal and Geographical Trends on Contrail Persistent Regions Over CONUS," NASA Technical Memorandum NASA/TM-20240011983, NASA Ames Research Center, Moffett Field, California, 2024.

[36] Schumann, U., Poll, I., Teoh, R., Koelle, R., Spinielli, E., Molloy, J., Koudis, G. S., Baumann, R., Bugliaro, L., Stettler, M., and Voigt, C., "Air Traffic and Contrail Changes over Europe During COVID-19: A Model Study," *Atmospheric Chemistry and Physics*, Vol. 21, No. 10, 2021, pp. 7429–7450. <https://doi.org/https://doi.org/10.5194/acp-21-7429-2021>.

[37] Lefebvre, A. H., and Ballal, D. R., *Gas Turbine Combustion*, CRC Press, 2010. <https://doi.org/10.1201/9781420086058>.

[38] Filippone, A., and Bojdo, N., "Statistical Model for Gas Turbine Engines Exhaust Emissions," *Transportation Research Part D: Transport and Environment*, Vol. 59, 2018, pp. 451–463. <https://doi.org/10.1016/j.trd.2018.01.019>.

[39] Yuksel, O., and Aygün, H., "Comparative Performance Analysis of a Turboprop Engine Used in Regional Aircraft by Considering Design and Flight Conditions," *Aircraft Engineering and Aerospace Technology*, Vol. 97, No. 3, 2025, pp. 345–355. <https://doi.org/10.1108/aeat-07-2024-0194>.

[40] Enteria, N., Santamouris, M., and Eicker, U., *Urban Heat Island (UHI) Mitigation*, Springer Singapore, Singapore, 2021. <https://doi.org/10.1007/978-981-33-4050-3>.

[41] Edwards, T., "Kerosene" Fuels for Aerospace Propulsion - Composition and Properties," *38th AIAA/ASME/SAE/ASEE Joint Propulsion Conference & Exhibit*, 2002, pp. 1–11. <https://doi.org/10.2514/6.2002-3874>.

[42] Massaro, M. C., Biga, R., Koliuchenko, A., Marocco, P., Monteverde, A. H. A., and Santarelli, M., "Potential and Technical Challenges of on-Board Hydrogen Storage Technologies Coupled With Fuel Cell Systems for Aircraft Electrification," *Journal of Power Sources*, Vol. 555, 2023. <https://doi.org/10.1016/j.jpowsour.2022.232397>.

[43] Baroutaji, A., Wilberforce, T., Ramadan, M., and Olabi, A. G., "Comprehensive Investigation on Hydrogen and Fuel Cell Technology in the Aviation and Aerospace Sectors," *Renewable and Sustainable Energy Reviews*, Vol. 106, 2019, pp. 31–40. <https://doi.org/10.1016/j.rser.2019.02.022>.

[44] Tiede, B., O'Meara, C., and Jansen, R., "Battery Key Performance Projections based on Historical Trends and Chemistries," *2022 IEEE Transportation Electrification Conference & Expo (ITEC)*, 2022, pp. 754–759. <https://doi.org/10.1109/ITEC53557.2022.9814008>.

[45] de Vries, R., Wolleswinkel, R. E., Jacobson, D. R., Bonnema, M., and Thiede, S., "Battery Performance metrics for large electric passenger aircraft." *34th Congress of the International Council of the Aeronautical Sciences*, 2024, pp. 1–15. https://www.icas.org/icas_archive/icas2024/data/papers/icas2024_0636_paper.pdf.

[46] Zamboni, J., "A method for the conceptual design of hybrid electric aircraft," Master's thesis, Delft University of Technology, 2018. <http://resolver.tudelft.nl/uuid:7b7dc56b-6647-4cc9-98f6-2ed5d488c759>.

[47] National Academies of Sciences, Engineering, and Medicine, *Commercial Aircraft Propulsion and Energy Systems Research: Reducing Global Carbon Emissions*, The National Academies Press, Washington, D.C., 2016. <https://doi.org/10.17226/23490>.

[48] Dinc, A., and Gharbia, Y., "Exergy Analysis of a Turboprop Engine at Different Flight Altitude and Speeds Using Novel Consideration," *International Journal of Turbo & Jet-Engines*, Vol. 39, No. 4, 2022, pp. 599–604. <https://doi.org/10.1515/tjj-2020-0017>.

[49] Ciliberti, D., Vecchia, P. D., Memmolo, V., Nicolosi, F., Wortmann, G., and Ricci, F., "The Enabling Technologies for a Quasi-Zero Emissions Commuter Aircraft," *Aerospace*, Vol. 9, No. 6, 2022. <https://doi.org/10.3390/aerospace9060319>.

[50] PowerCell Group, "Fuel Cell Stacks," , 2025. <https://powercellgroup.com/fuel-cell-stacks/>, Accessed: 2025-06-19.

[51] Datta, A., "PEM Fuel Cell MODEL for Conceptual Design of Hydrogen eVTOL Aircraft," Tech. Rep. NASA/CR-20210000284, NASA, 1 2021. Available at: <https://ntrs.nasa.gov/citations/20210000284>.

[52] PowerCell Group, "Aviation Systems," , 2025. <https://powercellgroup.com/segments/aviation/>, Accessed: 2025-06-19.

[53] Pastra, C. L., Hall, C., Cinar, G., Gladin, J., and Mavris, D. N., "Specific Power and Efficiency Projections of Electric Machines and Circuit Protection Exploration for Aircraft Applications," *2022 IEEE Transportation Electrification Conference and Expo (ITEC)*, 2022, pp. 766–771. <https://doi.org/10.1109/ITEC53557.2022.9813927>.

[54] Lin, H., Guo, H., and Qian, H., "Design of High-Performance Permanent Magnet Synchronous Motor for Electric Aircraft Propulsion," *21st International Conference on Electrical Machines and Systems (ICEMS)*, 2018, pp. 174–179. <https://doi.org/10.23919/ICEMS.2018.8549030>.

[55] Andersen, H., Chen, Y., Qasim, M. M., Amato, M., Cuadrado, D. G., Otten, D. M., Greitzer, E. M., Perreault, D. J., Kirtley, J. L., Lang, J. H., and Spakovszky, Z. S., "Design and Manufacturing of a High-Specific-Power Electric Machine for Aircraft Propulsion," *AIAA AVIATION 2023 Forum*, 2023, pp. 1–17. <https://doi.org/10.2514/6.2023-4158>.

[56] Yamaguchi, K., Katsura, K., Yamada, T., and Sato, Y., "High Power Density SiC-Based Inverter With a Power Density of 70kW Liter or 50kW Kg," *IEEJ Journal of Industry Applications*, Vol. 8, No. 4, 2019, pp. 694–703. <https://doi.org/10.1541/ieejjia.8.694>.

[57] de Vries, R., "Hybrid-Electric Aircraft with Over-the-Wing Distributed Propulsion: Aerodynamic Performance and Conceptual Design," Ph.D. thesis, Delft University of Technology, Delft, The Netherlands, 2022. <https://doi.org/10.4233/uuid:ef87dc11-e7b2-4726-a41f-28588d64c58d>.

[58] Gudmundsson, S., *General Aviation Aircraft Design*, second edition ed., Elsevier, 2022.

[59] Peng, S., *Hydrogen Fuel Cells*, Springer Nature Singapore, Singapore, 2024. <https://doi.org/10.1007/978-981-96-0157-8>.