MULTI-OBJECTIVE AIRCRAFT OPTIMIZATION FOR MINIMUM COST AND EMISSIONS OVER SPECIFIC ROUTE NETWORKS

Geoffrey C. Bower, Ilan M. Kroo
Stanford University, Department of Aeronautics and Astronautics

Keywords: *Multi-Objective Optimization, Aircraft Design, Environmental Impact*

Abstract

Historically, maximizing profits for corporate shareholders has been the primary goal for aircraft designers. Due to climate change concerns, environmental performance is quickly becoming a major design focus. A methodology to design one or more aircraft to satisfy demand on a given route network is described. The objectives are to minimize direct operating costs, CO_2 emissions and NO_X emissions. A hierarchical decomposition is used with discipline-specific optimization algorithms. A modified version of the NSGA-II multi-objective genetic algorithm is implemented in the system level aircraft design subspace. The CPLEX integer solver is used in the fleet assignment subspace. Results are presented for a test problem that involves designing a single aisle commercial aircraft for a route network consisting of 4 cities and 8 route segments. A Pareto front is found and the characteristics of the optimal aircraft designs and fleet assignment solutions are examined. The results show a definitive trades between operating costs and NO_X emissions and CO_2 emissions and NO_X emissions. Operating costs and CO_2 emissions are strongly related, with little if any trade off evident. All designs show benefits from using higher bypass ratio engines and thicker, higher aspect ratio wings than today's single aisle commercial aircraft.

1 Introduction

The environmental impact of commercial aviation is quickly gaining importance in the aircraft design community as well as with policy makers around the world [1, 2, 3]. As of 1992, aviation accounted for roughly 2% of the total global anthropogenic carbon emissions and 3.5% of the total anthropogenic radiative forcing. As air traffic is expected to grow at approximately 5% per year for the foreseeable future, these fractions are projected to increase to about 3% and 5%, respectively, of the global totals by 2050 [4]. Quantifying the environmental impact of an aircraft design is difficult due to difficulties in modeling the temporal and spatial distribution of emissions, the complex chemical interactions, and uncertainties in global atmospheric simulations [5, 6].

The primary constituents that have an influence on climate change are CO_2, NO_X, H_2O, SO_X and soot [7]. NO_X affects climate through feedbacks to O_3 and CH_4 concentrations. Water and particulate emissions influence the formation of persistent contrails and cirrus clouds. Sulfate emissions have a relatively minor impact on climate change compared to the other exhaust constituents, but contribute to acid rain. Noise generated during takeoff and landing is also a concern to aircraft designers and those living and working near airports.

The true external costs of various pollutants and noise are not currently well understood, but some attempts have been made at quantifying it for specific locales [2]. As such, there is no simple global objective function weighting the costs of CO_2 emissions, NO_X emissions, contrail formation, noise and other emissions. Past studies have shown that improving environmental performance of aircraft often results in higher operating costs and/or reduced performance as measured in traditional parameters such as cruise speed. For this reason, aircraft design for reduced environmental impact lends itself towards finding an optimal set of solutions [8]. A Pareto front of optimal solutions can be generated using a multi-objective optimization algorithm that allows system level trades between the various objectives to be analyzed before a final design is selected.

Ideally, aircraft designers choose an objective function directly related to profit potential for the aircraft manufacturer such as return on investment (ROI) or net present value (NPV) [9]. Historically, surrogates which are easier to compute have been used, such as direct operating costs and maximum takeoff weight [10]. The value of the selected objective is then minimized for a single design mission (a specific payload and range) or a weighted combination of a small set of design missions.

When assessing the environmental impact of aviation, the air transport system as a whole needs to be considered since overall emissions are what impact climate change. Factors such as aircraft utilization and routing become very important in terms of total emissions and operating costs. This is a change from the historical aircraft design philosophy which focuses on minimizing costs on a small set of missions. There has been some recent work on the design of aircraft for single hub route networks, but it did not incorporate environmental performance as an objective or constraint [11].

An approach is presented in this paper that allows the aircraft designer to optimize a set of one or more aircraft at a conceptual design level for a specific route network. The objectives of interest include the economic and environmental performance of a fleet over the route network. A description of the problem and some solution techniques are described in section 2. A small scale test problem is presented in Section 3 and results are shown in Section 4. Conclusions and future work are discussed in Section 5.

2 Problem Formulation

The primary purpose of the commercial aviation industry today is to move people and goods around the world both quickly and economically. An aircraft manufacturer is in business to make money for its share holders by creating products and services to meet the demand of airlines. Governments around the world are also stake holders in the industry to ensure public safety through regulations and to manage the flow of traffic through their airspace. In the past, some of these government regulations have taken the form of limits on noise and emissions around airports [12, 13]. With expanded interest in anthropogenic climate change today, there is discussion of expanding emissions regulations to encompass the complete flight regime or to enforce an economic penalty for emissions [14]. The true external cost of these emissions is not well understood at this time; however, it is thought to be significant [2]. Therefore, it is fortuitous for aircraft designers to consider emissions performance during the conceptual design process to understand and convey to policy makers the tradeoffs between economic performance and environmental performance.

There are many choices to be made by the aircraft designer to meet market demand. The primary choices are the size and performance of the aircraft themselves. The route network over which the aircraft operate also has a large impact on the economic performance and total emissions for given market demand, but this is largely dictated to the aircraft designer through airline demand. In this study, we choose to consider both the route network and aircraft design simultaneously. By considering both facets it is thought that significant gains can be made in environmental performance of the fleet with a small effect on economics. To make the problem tractable, we will assume that there is a known demand for all of the city pair segments in a specific route network. This problem is posed as a multi- objective optimization problem to design one or more aircraft types to satisfy passenger demand while minimizing both operating costs and emissions.

2.1 Objective Functions

This section describes the choice of objective functions considered during the multi-objective optimization for both economic and environmental performance.

2.1.1 Operating Costs

The first objective is related to the economic performance of the aircraft. There have been many economic performance metrics and surrogates used in the past: maximum takeoff weight, direct operating costs, total operating costs, ROI, and NPV to name a few. For this problem, we are considering the performance over an airline route network, so using an objective which directly relates to airline costs is logical. We choose the direct operating cost plus interest ($DOC+I$) method described in Ref. [10]. It is both easy to calculate and adequately reflects the cost to operate an aircraft for an airline on various routes.

The $DOC+I$ method breaks down the operating costs into ten components: Flight Crew, Cabin Crew, Landing Fees, Navigation Fees, Airframe Maintenance, Engine Maintenance, Fuel, Aircraft and Spares Depreciation, Insurance, and Interest. To enable the calculation of costs over the complete route network, these costs were further decomposed into their dependence on block time, flight time, flight cycles, fuel used and number of aircraft. The cost quantities related to the aircraft are the cost per block hour ($\$/bhr$), cost per flight hour ($\$/fhr$), cost per flight cycle ($\$/flight$), cost per year ($\$/yr$) and fuel cost ($\$/lb$). To calculate the total $DOC+I$ the following quantities are also required relating to the operation of the aircraft over the route network: total flight hours (fhr), total block hours (bhr), number of flight cycles ($flights$), number of aircraft (aircraft) and total fuel burn ($fuel$) for the network. The final formula for calculating $DOC+I$ is given in Eq. (1).

2.1.2 Emissions

As discussed previously we are interested in aircraft emissions that affect climate change. The most influential emissions are CO_2, NO_X and contrail formation. For this study we are only considering CO_2 and NO_X emissions as objectives since they have a significant impact on climate change and are simple to model. CO_2 emissions are directly proportional to the total fuel burn. NO_X emissions are related to the design of the combustor, the fuel burn, and details of the engine cycle, primarily the overall pressure ratio. Other pollutants for which there are adequate predictive models could be included as objectives as well. Modeling contrail formation is also possible, but requires detailed knowledge of local atmospheric conditions along the various flight segments [15]. This type of modeling is beyond the scope of the current study.

Multiple event noise metrics estimate the impact of a series of noise events. An accurate noise footprint along with the land use around airports is required to measure the actual effect of noise on the population around a given airport. This type of modeling is also beyond the scope of this conceptual level design study. It is recognized that noise is an important metric for aircraft design today and often drives design decisions [16]. For this reason, noise will be treated as a constraint in this study by requiring the aircraft to satisfy the ICAO stage 4 noise regulations.

2.2 All-at-Once Formulation

Details of the various analyses and optimization algorithms will be described in subsequent sections. Definitions of some important terms follow. A segment refers to a flight path in the route network connecting two airports. A route refers to a series of segments taken by a passenger from their origin to destination or a series of segments flown by a single aircraft during a day. When referring to "routes" we will mean aircraft routes from now on unless otherwise noted. Parameters and variables with subscript i refer to segments and those with subscript j refer to routes. Subscript k refers to the aircraft type if there is more than one being designed.

An all-at-once approach to this multi-objective optimization problem is outlined in Eq.

(2-4).

The objectives are $DOC+I$ for a day, total CO_2 emissions for a day and total NO_X emissions for a day. The design variables include aircraft and engine parameters for each aircraft type, operating conditions for each aircraft type, the number of aircraft of each type flying each route and the average number of passengers on each aircraft type for each segment. Here $AircraftVariables$, $EngineVariables$ and $OperatingConditions$ refer to sets of variables that can be used to specify the aircraft, engine, and operating conditions, respectively.

Constraints sets include aircraft design constraints and fleet assignment constraints. Typical aircraft design constraints include takeoff field length and 2nd segment climb gradient, among others. The fleet assignment constraints include constraints related to route compatibility, demand, airport capacity and continuity of aircraft from day to day. Route compatibility determines if a given design can fly the specified route. Continuity of aircraft from day to day requires that the same number of airplanes of each type start/end the day at each airport so that the schedule can be repeated.

The required analyses include calculating the various cost components, fuel burn, and NO_X emissions for each aircraft type on each segment with the average passenger load specified by $pax_{i,k}$. Fleet assignment parameters also need to be computed along with all of the constraint values.

Immediately, this formulation should appear intractable. It is a constrained, multi-objective, mixed-integer, non-linear global optimization problem. The only technique that might be able to solve the problem as formulated is a multi-objective genetic algorithm (MOGA). However, the large number of integer fleet assignment variables and constraints indicate that finding a feasible (let alone optimal) solution to the complete problem will be difficult. Another formulation should be sought to take advantage of different optimization algorithms to more efficiently solve different parts of the problem.

2.3 Hierarchical Formulation

Integer programming methods are well suited to solving the fleet assignment portion of the problem. MOGAs are one of the best classes of methods for solving multi-objective, non-linear global optimization problems. Multi-objective algorithms which generate a single objective by weighting all of the objectives can also be used. By decomposing the problem into subspaces with each having its own optimization algorithm and analyses, the total computational time can be reduced compared to the all-at-once formulation. However, compatibility between the shared variables and targets in the various subspaces must be maintained, increasing the problem size. In addition, for each design at analyzed at the system level a complete subspace optimization must be performed.

2.3.1 Initial Formulation

A hierarchical optimization methodology was derived to take advantage of our knowledge of the problem. The optimization problem formulation for the system level aircraft design subspace is presented in Eq. (5-7). The fleet assignment subspace optimization problem follows in Eq. (8).

The system level aircraft design subspace passes shared variables and targets to the fleet assignment subspace. These include $OperatingConditions_k$, $flights_{i,k}$, $aircraft_k$, $pax_{i,k}$, bhr_k, fhr_k and one aircraft design variable, the $capacity_k$. The fleet assignment subspace returns the value of its objective function, J_{fleet}. The primed values in the fleet assignment subspace the are the fleet assignment subspace instances of the shared variables or values computed in the subspace.

There are still a number of issues with this formulation. Requiring the average number of passengers on each segment as a shared variable potentially doubles the system level aircraft design subspace variables in the limit of a large route network. It also makes the integer programming problem non-linear, which is beyond the capability of the available solver as will be discussed later. The inclusion of the aircraft capacity and operating conditions in the fleet assignment subspace also make the integer programming problem non-linear. In the system level aircraft design subspace the analyses include calculating the various cost components, fuel burn and NO_X emissions for each aircraft type on each segment. Aircraft design related constraints are also analyzed in this subspace. The fleet assignment subspace tries to match target values for the number of flights on each segment, total flight hours, and total block hours as well as the target values for the shared variables mentioned previously. The fleet assignment local variables are the number of aircraft flying each route and the local copies of the shared variables.

2.3.2 Simplified Formulation

In order to use a linear integer programming algorithm all of the non-linearities in the subspace objectives and constraints must be eliminated. The operating conditions impact which routes are compatible with a given aircraft design and the capacity makes the passenger demand constraint non-linear. In addition, we would like to eliminate the variables indicating the average number of passengers on each flight segment in order to simplify the system level aircraft design subspace optimization problem.

The subspace degrees of freedom from the shared variables are only necessary if the subspace problem is infeasible. By fixing the operating conditions and capacity in the fleet assignment subspace we can make the objective and constraints linear. We are able to fix these variables in the subspace because we can ensure fleet assignment feasibility by merely adding more aircraft.

In order to do this some assumptions must be made. If the fuel burn, emissions and airport compatibility are calculated at maximum payload then the pax_i variables can be eliminated. This assumption is equivalent to assuming a load factor of one. It should be a reasonable assumption since the optimal fleet assignment solutions should tend towards high load factors to minimize both operating costs and emissions. However, the calculated total costs and emissions will be slightly pessimistic. Post optimality analysis can establish the sensitivity of this assumption and estimate a correction based on the actual load factors. For this study no correction was made and the pessimistic solutions were used.

The system level aircraft design subspace optimization problem incorporating these changes is shown in Eq. (9-11). The fleet assignment subspace optimization problem written in its simplest form is given by Eq. (12).

The system level aircraft design subspace only passes targets for the number of flights of each airplane type on each segment and the total number of aircraft of each type to the fleet assignment subspace. The fleet assignment subspace returns the value of its objective function, J_{fleet}. The primed values in the fleet assignment subspace the are values computed in the subspace that try to match the targets. The fhr and bhr compatibilities in the objective function were eliminated since they only depend on $flights_{i,k}$ for a given capacity and cruise speed, both of which are now fixed in the fleet assignment subspace.

2.4 System Level Aircraft Design Subspace

This section details the system level aircraft design subspace analysis methods. A preliminary design tool, PASS (Program for Aircraft Synthesis Studies), is used for all aircraft design analyses [17]. Given a set of aircraft design variables PASS calculates aircraft performance over a specified mission profile. The modularity allows for models of differing fidelity to be incorporated depending on the application. A new propulsion model was added for this study to add degrees of freedom to the engine design to more accurately reflect the impact of engine design variables on emissions.

The engine model is based on the twin spool turbofan analysis methods described in Appendix G of Ref. [18]. This method includes the effects of turbine blade cooling, power extraction, and bleed air. The engine is analyzed on-design at the cruise conditions to obtain specific performance. It is then analyzed off-design at sea level static conditions and sized to meet takeoff thrust requirements. Thereafter, the thrust and specific fuel consumption can be determined at any operating condition. An improved weight model for the engine is also included [19].

The engine carbon dioxide emissions are directly related to fuel burn. For each kilogram of jet fuel that is consumed, 3.13 kg of CO_2 are released into the atmosphere. An empirical model for the NO_X emission index based on curve fits of existing engine data was used. The primary drivers of NO_X emissions are the details of the combustor design and the temperature and pressure at the burner entrance. The temperature and pressure are primarily a function of the overall pressure ratio (OPR) [3]. Simple curve fits of emission index ($EINO_X$) vs. OPR were established from data in the ICAO emissions databank for modern turbofan engines [20]. The engines included in the fits are the CFM56-5Bx/2P, CFM56-7B2x/2, CF34-8C, GE90-110B, GE90-115B, the Rolls Royce Trent series, PW4164, PW4168A, PW6122A, and PW6124A. All of these engines use some form of a low emissions combustor. While future engines may have lower emissions indices than today's engines, the trends with overall pressure ratio should be consistent. Figure 1 shows the fits of $EINO_X$ vs. OPR for the four certification points. The fit is decent for takeoff and climb conditions, but poor for approach and idle. Fortunately, the approach and idle conditions have the smallest indices and only represent a small portion of the total flight time. Emissions indices are not available for cruise conditions since cruise is not a certification point. The cruise emission index is taken to be $(0.45)^{0.7} = 0.73$ of the sea level static value of the climb emission index based on recommendations in Ref. [21]. In the future, an improved cruise NO_X model should be incorporated.

2.5 Fleet Assignment Subspace

The fleet assignment subspace attempts to match the target values of the parameters passed from the system level aircraft design subspace while satisfying local constraints. These parameters are the total number of flights on each segment and the total number of aircraft required to fly the network. The variables in the fleet assignment subspace are the number of aircraft flying each possible route. A system level aircraft design subspace analysis routine determines which routes are compatible with the specific aircraft design being passed to the fleet assignment subspace based on the calculated cruise speed of the design.

This method for the fleet assignment does not deal with the explicit scheduling of each aircraft throughout a day as this would dramatically increase the size and complexity of the problem; it only determines what aircraft routes should be flown during the day. There is a minimum time required to complete each route, and any excess time can be used to vary the start time, end time and/or layover times.

As an example, assume there are 12 useable hours in a day, from 6 am to 6 pm. An aircraft is to fly a route from city A to city B and then back to city A during a day with each flight lasting 4 hours. If one hour is required for a layover, then the aircraft requires a minimum of 9 hours to complete its route. The first flight could leave at 6 am and the second at 11 am ending the day at 3 pm, or the first flight could leave at 9 am and the second at 2 pm ending the day at 6 pm, or the first flight could leave at 6 am and the second flight at 2 pm ending the day at 6 pm. The last scenario has an extended 4 hour layover in the middle of the day whereas the first two have a 3 hr buffer at the beginning or end of the day. Therefore there is some flexibility in the actual scheduling. Any scenario which has the sum of the buffer times at the beginning and end of the day and any additional layover time totaling less than 3 hours would be feasible. The described method captures the important information for calculating operating costs and emissions without explicitly scheduling the flights.

2.5.1 AMPL

The fleet assignment optimization problem was set up using AMPL (A Modeling Language for Mathematical Programming) [22]. The AMPL modeling language is convenient for formulating problems of variable size that are based on sets and indexing. It is primarily used to setup linear, integer and quadratic programming problems. There are a number of available solvers that are tailored to various problems types. First, a model file is defined including names for all of the sets, variables, parameters, objectives and constraints. Then to solve a specific instance of a problem, a data file containing values for the sets and parameters is generated and a solver is specified. The output values of interest can be sent to a text file.

2.5.2 Calculations

The fleet assignment subspace has one simple analysis associated with it. Before the optimization begins, a set of all the feasible routes is generated. The fastest cruise speed achievable during the optimization (corresponding to the maximum cruise Mach number and the minimum cruise altitude) is calculated. Then, based on the available time in the day for operations and an average layover time between flights it is possible to calculate whether or not a given route can be flown. Once the set of all possible routes is determined, the minimum cruise speed required to complete each route is calculated and saved in an array. Other parameters that are derived from the set of possible routes are the start city, end city, and the number of times each segment is flown on each route. These parameters are saved and added to the data file.

During the optimization run, each aircraft design can be checked for compatibility with each route. If the cruise speed is greater than the minimum cruise speed required, then the route is feasible. Other feasibility checks can be added when multiple aircraft are being designed, such as takeoff field length and maximum range.

2.5.3 Variables, Objectives, and Constraints

The local variables for the fleet assignment problem are the number of aircraft of each type flying each route, $flights_{j,k}$. The objective function for the fleet assignment subspace is to minimize the absolute value of the difference between the target values and local values for the total number of flights on each segment and the total number of aircraft. Slack variables are required to linearize the absolute values in the objective function.

The first set of constraints requires the available capacity to exceed the demand on each segment. Aircraft continuity for each day is also imposed. A further basing constraint can be imposed by specifying a maximum number of aircraft of each type allowed to start/end the day at each airport. The final constraint that can be imposed places an upper bound on the maximum number of aircraft that are allowed to fly a given route during a day. The final formulation of the fleet assignment subspace problem is given by Eq. (13-18).

2.6 Optimization Methods

This section describes the optimization algorithms used for each subspace. The system level aircraft design subspace contains multiple objectives, is non-linear, non-smooth and may contain local minima. The multi-objective algorithms considered for this included weighted sum methods and MOGAs. The weighted sum methods typically utilize gradient based methods and require multiple optimizations to be run to generate a Pareto front. The algorithm choice was narrowed to a MOGA for this study since genetic algorithms (GAs) typically perform better on global optimization problems and because they solve for the Pareto front in a single optimization run. A genetic algorithm was also selected since it is able to better handle discontinuities in the design space and integer variables [23]. The CPLEX integer programming method was selected for efficiently solving the fleet assignment subspace problem [24].

2.6.1 Multi-Objective Genetic Algorithms

There have been a large number of MOGAs developed over the years, starting with the Vector Evaluated Genetic Algorithm (VEGA) in 1984. Two of the more recent algorithms that were considered for this study were the Non-Dominated Sorting Genetic Algorithm II (NSGA-II) and the Strength Pareto Evolutionary Algorithm 2 (SPEA2) [25, 26]. Both of these algorithms are elitist and improve upon the previous generation of MOGAs by eliminating explicit niching parameters to achieve an even spread of solutions over the final Pareto front.

Ultimately, NSGA-II was selected since it required fewer function evaluations on a number of constrained test problems, at the expense of slightly worse performance in evenly spreading designs over the final Pareto front. This algorithm is also easier to modify for handling constraints and controlled elitism based on suggestions in Ref. [23]. The algorithm selects solutions through a domination sorting procedure. A solution dominates another solution if the first one is no worse than the second in all objectives and better than the second in at least one objective. The set of non-dominated solutions is the Pareto front of a given population. These solutions are assigned rank 1. Ignoring this set of solutions, the next non-dominated set is determined and given rank 2. This procedure is repeated until all solutions are ranked. The fitness assignment value of a design is equal to its rank.

NSGA-II uses binary tournament selection to select the parent population for the next generation. In each tournament, the solution with the lower rank is selected. If two designs are of the same rank, then the solution with fewer nearby solutions is selected. This algorithm uses a crowding distance estimate in the objective space that is equal the average side length of the cuboid formed by its nearest neighbors in each objective. Two point crossover and bitwise mutation are then used to generate the child population from the parent population. The child population is then combined with the parent population from the previous generation before the non-dominated sorting procedure is repeated to generate the parent population for the next generation. The algorithm is elitist since it preserves the parent population and compares it to the child population for each successive generation, so the best designs are not lost if the children do not perform better than the parents.

As initially formulated, NSGA-II does not handle constraints. There are a number of methods for handling constraints with GAs that were considered. The first method is to either "repair" designs which violate constraints so that they become feasible, or to eliminate them from the population and replace them with a new random design. This method was not considered due to the large number of constraints in an aircraft design problem and the low probability of having randomly generated feasible designs during the initial generations. The next method considered adds the constraint violations to the objective functions with penalty parameters. This method was not selected because it requires hand tuning to select the penalty parameters for each constraint and objective function pair. If the penalty parameters are not large enough there is the possibility of converging to infeasible solutions.

The final method is the constrained tournament selection method [23]. This method modifies the non-dominated sorting procedure. When comparing two feasible designs, the normal domination comparison operator is used. If a feasible design is compared to an infeasible design, the feasible one is dominant. If two infeasible designs are compared, then the design with the smaller total normalized constraint violation is dominant. It should be noted that there are some flaws with this method. For example, a good design that violates the takeoff constraint by one foot will be ranked worse than a design with twice the cost and emissions. There is on going work by the authors and others on how to better handle constraints in population based methods. The 1 constrained tournament selection method was selected based on its performance on some constrained test problems. The constraints are normalized using Eq. (19).

The final modification made to the NSGA-II algorithm was the addition of controlled elitism [23]. Controlled elitism attempts to maintain a geometric distribution in the number of designs with respect to rank. This is beneficial in higher dimensional objective spaces since it maintains diversity in the population and discourages premature convergence to a deceptive Pareto front.

The parameters which need to be selected for NSGA-II include the population size, maximum number of generations, mutation rate, crossover rate and the controlled elitism parameter. Also, the variables can be represented as either binary strings or real-valued variables. A suggested population size is at least twenty times the number of design variables, although bigger is always better. Convergence criteria for MOGAs are difficult to specify, so the algorithms are typically run until little change in the non-dominated front is seen. For most problems 500 to 1000 generations are sufficient. The bitwise mutation rate for NSGA-II is typically higher than for single objective GAs and is usually in the range of 0.05 to 0.1. The crossover probability used is usually between 0.9 and 1.0. A controlled elitism parameter value of 0.65 was selected based on suggestions in Ref. [23].

2.6.2 Integer Programming Solvers

The fleet assignment subspace optimization problem is solved using the CPLEX integer solver in AMPL. The solver can be used for solving mixed-integer linear programming problems of the type outlined for the fleet assignment subspace. It solves the problem using a branch and cut algorithm. See chapter 7 of Ref. [24] for details.

3 Test Problem Description

In order to test the proposed solution methodology for the problem formulated in Section 2, a test problem was designed. The test problem involves designing a single aircraft type to satisfy demand on a small route network consisting of 4 cities and 8 segments. Figure 2 shows the route network with the range and passenger demand on each segment. For a network this small, it is easy to identify good fleet assignment solutions for each capacity in order to verify that the optimizer is finding good solutions. The optimal fleet assignment solution should be the same if each of the three objectives were optimized individually for such a simple route network with a single aircraft design.

3.1 System Level Aircraft Design Subspace

The system level aircraft design subspace performs analyses using PASS and optimizes the system using the modified NSGA-II algorithm. The aircraft type is a conventional single-aisle commercial transport with two wing mounted engines in the size category of a Boeing 737 or Airbus A320. The following subsections will discuss selection of the design variables and describe the performance analyses and constraints imposed.

3.1.1 Variables

The variables can be split into four categories: aircraft parameters, engine parameters, operating conditions, and fleet assignment targets. The aircraft parameters are the maximum takeoff weight ($MTOW$), maximum zero fuel weight ($MZFW$), capacity ($capacity$), wing loading (W/S), wing aspect ratio (AR), wing taper ratio (λ), wing quarter chord sweep angle (Λ), wing average thickness to chord ratio (t/c), horizontal tail area (S_H/S_{Ref}) and wing location on the fuselage (x_{wing}/L). The engine parameters include the thrust-to-weight ratio (T/W), overall pressure ratio (OPR), fan pressure ratio (FPR) and turbine entry temperature (TET). The optimum bypass ratio for minimum specific fuel consumption at cruise is computed based on the specified OPR, FPR and TET using the method in Ref. [18]. The two operating conditions that need to be specified are the cruise Mach number (M_{cruise}) and cruise altitude (h_{cruise}). The fleet assignment targets are the required number of flights on each segment ($flights_i$) and the total number of aircraft ($aircraft$). Since the fleet assignment solution requires that the same number of aircraft start and end the day at each airport, there is a limited set of values for $flights_i$

that will satisfy these constraints. The set of feasible values for $flights_i$ is generated during pre-processing and saved in an array. The design variable used is the index of this array rather than the eight $flight_i$ variables. This reduced the number of combinations for the $flights_i$ variables from $9^8 (= 43046721)$ to 2868 and cut the number of bits representing these variables in half.

For the genetic algorithm each variable requires an upper bound, lower bound, and bit sting length. There are a total of eighteen design variables. Table 1 lists the minimum, maximum and number of bits for each of the design variables.

Table 1 Minimum, maximum and number of bits for each design variable.

Variable	Min	Max	Bits
$MTOW[lb]$	80000	342143	18
$MZFW[lb]$	50000	312143	18
$Capacity$	102	192	4
$W/S[lb/ft^2]$	100	163.5	7
AR	5	20.75	6
λ	0.2	0.51	5
$\Lambda[deg]$	0	42.333	7
t/c	0.07	0.145	4
S_H/S_{ref}	0.1	0.41	5
x_{wing}/L	0.28	0.59	5
T/W	0.2	0.515	6
OPR	15	57	6
FPR	1.6	2.22	5
$TET[R]$	2450	3200	4
M_{cruise}	0.7	0.91	6
$h_{cruise}[ft]$	30000	37000	3
$flights_{index}$	1	2868	12
$aircraft$	5	20	4

3.1.2 Performance Analysis

The performance characteristics of the aircraft design are determined using PASS. The manufacturer's empty weight (MEW) can be computed based on the $MTOW$, $MZFW$ and geometry using the weight model within PASS. To determine the fuel required for each segment, the aircraft is analyzed over the reserve mission shown in Fig. 3. The reserve mission follows the IATA international reserve requirements. For this test problem, the reserve requirements are probably excessive for the ranges specified, but were used regardless. The aircraft is then analyzed over the actual mission range without the reserve requirements to determine the fuel burn and NO_X emissions. For problems with a large number of segments, the fuel burn and NO_X emissions can be analyzed for only a few missions and curve fit to decrease computation time. Noise at the three certification points is estimated using ICAO rules [12]. 2006 is used as the reference year for fuel costs ($1.85/gal) and inflation.

3.1.3 Aircraft Design Constraints

Some constraints are performance and certification driven while others are required to ensure feasible designs based on the selected design variables for the problem. The last set could potentially be eliminated by adding iteration loops within the analysis, however, the optimizer should drive these constraints to near equality, so analysis iteration is not necessarily required.

The balanced field length at takeoff is constrained to be less than 8000 ft and the landing field length less than 7000 ft. The 2nd segment climb gradient with one engine out must exceed 0.024 to meet FAA regulations for a two engined aircraft. The required maximum takeoff weight ($MZFW$+$fuel_{max}$) for all of the missions must not exceed the specified maximum takeoff weight. The required fuel volume for each mission must not exceed the available volume in the main wing box. The required fuel weight must not exceed the maximum fuel weight. The wing must be located along the fuselage such that the main landing gear attachment point lies within the main wing box. The maximum payload weight available ($MZFW$-MEW) must exceed the specified payload weight corresponding to 225 lb/passenger including luggage, plus an additional 15% for freight. Stage 4 noise requirements must be met. The minimum static margin must be at least 0.05. The thrust available during cruise must exceed the drag by at least 14%. The horizontal tail C_L cannot exceed its C_{Lmax} at any operating condition. Similarly, the wing C_L must not exceed its C_{Lmax} at any operating condition. Finally, the wing span is not allowed to exceed 262 ft.

3.2 Fleet Assignment Subspace

For the test problem, there were no constraints on basing or on the number of aircraft allowed on each feasible route. It was assumed that the usable day was 18 hours long (6 am until midnight) and that layovers at each airport were 1 hour long. There are a total of 202 possible aircraft routes that can be flown for this simple route network if the aircraft flies at Mach 0.91 and 30,000 ft, which corresponds to the fastest true airspeed the airplane could fly at.

4 Test Problem Results

Based on the results of many test runs, it was observed that the algorithm had a difficult time finding the optimum aircraft capacity. This occurred because it is much easier to find feasible designs that meet the takeoff and climb constraints for the smaller capacities. The GA would find feasible solutions for these capacities first which would quickly drive the population to capacities less than 150 passengers. It was known this was not correct since better designs could be found in all objectives with a capacity of 168 or 180 passengers through hand tuning.

To overcome this issue, the GA was run for 100 generations for each of the 16 allowed capacities. The 78 best designs were selected from each of these sub-populations and used as seed designs for the final run of the GA. To dramatically speed up computation, a fixed fleet assignment solution was used for these seed generation runs, corresponding to the number of flights required to meet demand for each city pair. The projections of the seed designs into the planes formed by each pair of objectives are shown in Fig. 4. The best performing designs have capacities of 180, 168, 144 and 150 passengers. We expect these designs to proliferate in the final optimization.

The optimization algorithm was run on a 2 GHz laptop PC with 1 GB of RAM for 500 generations with a total computation time of about 38 hours. Eight of these hours were for generating the seed designs. The population size used was 1248. The mutation probability was 0.05, and the crossover probability was 1.0. Each function evaluation took approximately one sixth of a second including the fleet assignment subspace optimization and about 0.014 seconds with a fixed fleet assignment solution during seed generation.

4.1 Final Pareto Front

The final Pareto front resulting from the optimization is shown in Fig. 5. Also included are projections into the planes formed by each set of two objectives of the initial seed designs and the Pareto front at generations 200 and 500. Intermediate generations are not shown for clarity. Between generations 200 and 500 there is little movement in the Pareto front indicating the algorithm has converged. Estimates of the performance of a few current single aisle aircraft produced by Boeing (737-700, -800 and -900) and Airbus (A319, A320 and A321)are also included in the plots for comparison. These aircraft were modeled and analyzed using the same analysis code and objective functions.

Based on these plots we can see that the algorithm progressed to a fairly well-defined Pareto front that forms a curve in the objective space and that it distributed designs well along the front. From the projected views of the final Pareto fronts we see that there are definite trade offs between operating costs and NO_X emissions as well as CO_2 emissions and NO_X emissions. In general, reducing CO_2 emissions also reduces operating costs. The few designs that are at much higher operating cost have a non-optimal number of aircraft for the network, but slightly lower CO_2 and NO_X emissions than some designs on the remainder of the Pareto front. If the number of aircraft were reduced these would slide into the area of the Pareto surface where the majority of the designs reside. For a fixed level of NO_X emissions there is a small trade off between operating costs and CO_2 emissions, however the width of the front in these objectives is small due to the strong relationship between operating costs and fuel burn. With the increasing cost of fuel over the value used in this studies we would expect the width to be even smaller today.

4.2 Aircraft Design Results

In the final Pareto front most of the aircraft designs are very similar. The aircraft parameters, operating conditions and fleet assignment values all converged to small ranges. Almost all of the variation over the Pareto front is due to variations in the engine parameters. Table 2 shows the mean and standard deviation for the variables over the final Pareto front. Figure 6 shows histograms for each of the design variables for the final Pareto front and a top view of the "mean" airplane design. The optimal aircraft capacity turned out to be 168 passengers, one of the values expected from the seed designs.

The variables at or near their bounds are h_{cruise}, TET, λ and t/c. The allowable cruise altitude range was selected somewhat arbitrarily, if the optimization were to be re-run the upper bound would probably be extended to a higher altitude. The turbine entry temperature upper bound was selected at a reasonable turbine material temperature. Today's materials may be able to handle higher temperatures, but reliable data on the turbine entry temperatures of current commercial engines was unavailable. The taper ratio bound was set due to limitations of the analysis code to properly handle constraints that become active as the taper ratio approaches zero. The bound on wing thickness was set at what seemed to be a reasonable value based on root sections of transport wings. The high thickness and higher than expected wing sweep are the result of the analysis code. The trades between sweep and thickness for a given mach number are fairly flat, but higher sweeps tend to perform slightly better leading the optimizer to thicker, high sweep wing designs.

The constraints with the smallest margins (active constraints) for most designs in the final Pareto front were the takeoff field length, second segment climb gradient, maximum takeoff weight, payload weight and either the fuel weight or fuel volume constraints.

Table 2 Mean and standard deviation of variables over the final Pareto front.

Variable	Mean	Standard Deviation
$MTOW\,[lb]$	194690	17.2
$MZFW\,[lb]$	148330	174
$Capacity$	168	0
$W/S\,[lb/ft^2]$	140.1	0.334
AR	12.6	0.228
λ	0.20	0.0
$\Lambda\,[deg]$	35.2	0.5
t/c	0.145	0.0
S_H/S_{ref}	0.27	0.001
x_{wing}/L	0.366	0.01
T/W	0.32	0.002
OPR	23.02	7.98
FPR	1.67	0.021
$TET\,[R]$	3197	12.0
M_{cruise}	0.79	0.004
$h_{cruise}\,[ft]$	36826	383

Figure 7 shows how the engine parameters vary over the Pareto front. As expected the low NO_X designs have lower overall pressure ratios. The lowest CO_2 designs have a high turbine entry temperature, a high overall pressure ratio, and a low fan pressure ratio, which corresponds to a high bypass ratio and high propulsive efficiencies. Almost all of the designs have higher cruise bypass ratios than engines in service today. For reference, the CFM56 series of engines used on the B737 and A320 series of aircraft have bypass ratios of about 6 and overall pressure ratios between 21.6 and 32.8, depending on the model.

4.3 Fleet Assignment Results

The final fleet assignment solution was consistent throughout the final Pareto front. All of the population members converged to having the minimum number of flights per day on each segment to meet demand. Most used seven aircraft to satisfy the required demand. The few outliers at a higher $DOC+I$ in the final Pareto surface had more aircraft than were necessary. Table 3 shows the routes flown by the seven aircraft, their utilization, and the number of aircraft flying each route. The utilization is the sum of the block hours, but not including any layover time between flights. Table 4 shows the number of flights and the average load factors on each segment. Both the load factors and utilizations are high indicating that the fleet assignment solution is optimal for the specific aircraft capacity and cruise speed.

Table 3 Utilization and number of aircraft flying each route.

Route	Utilization [hr]	# of Aircraft
$A.B.C.B.C$	14.56	2
$A.D.A.D$	15.72	1
$C.B.A.D$	11.53	1
$C.B.C.B.A$	14.56	1
$D.A.D.A$	15.72	1
$D.C.D.C.D.A$	12.07	1
Average	14.10	–

Table 4 Number of flights and load factors on each segment.

Segment	Number of Flights	Load Factor
AB	2	0.685
AD	4	0.978
BA	2	0.673
BC	5	0.851
CB	5	0.856
CD	2	0.872
DA	4	0.939
DC	2	0.884
Average	–	0.862

5 Conclusions and Future Work

Minimizing operating costs has always been a goal for aircraft designers, but environmental performance is quickly becoming a major design focus. In this paper we outline a methodology to design one or more aircraft types to satisfy demand on a given route network to minimize both operating costs and emissions. A hierarchical decomposition is used with discipline specific optimization algorithms. A modified version of the NSGA-II MOGA implementing constrained tournament selection and controlled elitism is used in the system level aircraft design subspace. The CPLEX integer solver is used in the fleet assignment subspace.

Results are presented for a simple test problem with a 4 city and 8 segment route network and a single aircraft design. The aircraft parameters, operating conditions and fleet assignment solutions remained nearly constant over the Pareto Surface, with the surface defined primarily by variations in the engine parameters. The results show a trade between operating costs and NO_X emissions as well as CO_2 and NO_X emissions. Operating costs and CO_2 emissions appear directly related, with little if any trade off evident. All designs show benefits from using higher bypass ratio engines and thicker, higher aspect ratio wings than today's aircraft of the same class.

The methodology presented can be used to solve many problems more difficult than the simple test problem presented. However, reducing the computational time for the optimization is an important consideration. Fortunately, the genetic algorithm is parallelizable, which can greatly reduce computation times if more than one processor is used. The computation time will decrease almost linearly with the increase in the number of processors. This will also allow larger population sizes to be run for more generations, improving Pareto front performance and confidence in the converged designs.

In addition to further study of the simple test problem, possible future work includes a number of modifications and extensions. A problem of immediate interest is the design of two or three aircraft types for more complicated route networks. Current aircraft designs could also be incorporated to see how much improvement new designs will provide in operating costs and emissions. The methodology presented is also well suited to designing a family of aircraft. Fuselage stretches would be easy to incorporate, adding a few variables to the system level aircraft design subspace while giving the fleet assignment subspace more degrees of freedom. Another possibility is including origin-destination passenger demand instead of a fixed demand on each segment. In this case, the fleet assignment subspace would decide which segments should be flown and how passengers should be routed. It would be interesting to see if hub-and-spoke or point-to-point networks evolved to meet the different objectives. Further down the road, probabilistic models with varying demand on the network and the effects of winds and weather would be of interest in developing robust fleet assignment solutions. Another approach that needs to be considered in parallel is to increase the fidelity of many of the analyses. A better engine model and NO_X model are desired as well as including other emissions models such as contrail formation and noise exposure around the airports.

Acknowledgments

The authors acknowledge Steve Altus at Jeppesen for discussions related to fleet assignment, air traffic management and environmental issues. The PASS analysis code was developed and provided by Desktop Aeronautics, Inc. The authors also acknowledge the contributions from colleagues at Stanford University. The first author would also like to thank the National Science Foundation for fellowship funding that made this work possible.

References

[1] Penner, J. et al., *Aviation and the Global Atmosphere*, Cambridge University Press Cambridge, 1999.

[2] Brooker, P., "Civil Aircraft Design Priorities: Air Quality? Climate Change? Noise?" *The Aeronautical Journal*, Vol. 110, No. 1110, 2006, pp. 517–532.

[3] Green, J. et al., "Air Travel-Greener by Design. Mitigating the environmental impact of aviation: opportunities and priorities," *The Aeronautical Journal*, Vol. 109, No. 1099, 2005, pp. 361–418.

[4] Houghton, J., *Climate Change 1995: The Science of Climate Change*, Cambridge University Press, 1996.

[5] Rogers, H., Lee, D., Raper, D., De F. Foster, P., Wilson, C., and Newton, P., "The impacts of aviation on the atmosphere," *Aeronautical Journal*, Vol. 106, No. 1064, 2002, pp. 521–546.

[6] Panel on Atmospheric Effects of Aviation and Board on Atmospheric Sciences & Climate and Commission on Geosciences & Environment & Resources and National Research Council, *Atmospheric Effects of Aviation: A Review of NASA's Subsonic Assessment Project*, 1999.

[7] Sausen, R., Isaksen, I., Grewe, V., Hauglustaine, D., Lee, D., Myhre, G., Köhler, M., Pitari, G., Schumann, U., Stordal, F., et al., "Aviation radiative forcing in 2000: An update on IPCC (1999)," *Meteorologische Zeitschrift*, Vol. 14, No. 4, 2005, pp. 555–561.

[8] Antoine, N. and Kroo, I., "Aircraft Optimization for Minimal Environmental Impact," *Journal of Aircraft*, Vol. 41, No. 4, 2004, pp. 790–797.

[9] Markish, J. and Willcox, K., "Value-Based Multidisciplinary Techniques for Commercial Aircraft System Design." *AIAA Journal*, Vol. 41, No. 10, 2003, pp. 2004–2012.

[10] Liebeck, R. and Center, L. R., *Advanced Subsonic Airplane Design & Economic Studies*, National Aeronautics and Space Administration, Lewis Research Center; National Technical Information Service, distributor, 1995.

[11] Crossley, W. and Mane, M., "System of Systems Inspired Aircraft Sizing Applied to Commercial Aircraft/ Airline Problems," *AIAA 5th Aviation, Technology, Integration, and Operations Conference(ATIO)*, 2005, pp. 1–12.

[12] ICAO, E., "Volume 1–Aircraft Noise, 4th Ed." 2005.

[13] ICAO, E., "Volume 2–Aircraft Engine Emissions, 2nd Ed." 1993.

[14] ICAO, E., "Environmental Report 2007," 2007.

[15] Schumann, U., "Formation, Properties and Climatic Effects of Contrails," *Comptes rendus-Physique*, Vol. 6, No. 4-5, 2005, pp. 549–565.

[16] Norris, G. and Wagner, M., *Airbus A380: Superjumbo of the* 21^{st} *Century*, Zenith Press, Osceola, WI, USA, 2005.

[17] Kroo, I., "An Interactive System for Aircraft Design and Optimization," Aiaa paper 92–1190, AIAA Aerospace Design Conference, Feb. 1992.

[18] Mattingly, J., Heiser, W., and Pratt, D., *Aircraft Engine Design*, Aiaa, 2002.

[19] Torenbeek, E., "Synthesis of Subsonic Aircraft Design," *Student edition, Delft University of Technology/Martinus Nijhoff*, 1982.

[20] ICAO, I., "Engine Exhaust Emissions Data Bank.–2007," Tech. rep., Doc 9646-AN/943. International Civil Aviation Organisation, Montreal, Canada, 2007.

[21] Norman, P., Lister, D., Lecht, M., Madden, P., Park, K., Penanhoat, O., Plaisance, C., and Renger, K., "Development of the technical basis for a New Emissions Parameter covering the whole AIRcraft operation: NEPAIR, Final Technical Report," Tech. rep., CAEP/6-IP/17 Appendix B, 2003.

[22] Fourer, R., Gay, D., and Kernighan, B., *AMPL: A Modeling Language for Mathematical Programming*, Thomson/Brooks/Cole, 2003.

[23] Deb, K., *Multi-Objective Optimization Using Evolutionary Algorithms*, Wiley, 2001.

[24] CPLEX, I., "10.0 User's guide," *ILOG, Mountain View, CA*, 2006.

[25] Deb, K., Agrawal, S., Pratap, A., and Meyarivan, T., "A Fast Elitist Non-Dominated Sorting Genetic Algorithm for Multi-Objective Optimization: NSGA-II," *Proceedings of the Parallel Problem Solving from Nature VI Conference*, 2000, pp. 849–858.

[26] Zitzler, E., Laumanns, M., Thiele, L., et al., "SPEA2: Improving the Strength Pareto Evolutionary Algorithm," *EUROGEN*, 2001, pp. 95–100.

$$DOC + I = (\$/bhr) * bhr + (\$/fhr) * fhr + (\$/flight) * flights + (\$/yr)/365 * aircraft + (\$/lb) * fuel \quad (1)$$

$$\begin{aligned}
&Min: \\
&\quad J_1 = DOC + I = \sum_k [(\$/bhr)_k * bhr_k + (\$/fhr)_k * fhr_k + (\$/flight)_k * flights_k + \\
&\qquad\qquad (\$/yr)_k/365 * aircraft_k + (\$/lb) * \sum_i (fuel_{i,k} * flights_{i,k})] && (2) \\
&\quad J_2 = CO_2 = \sum_k 3.13 \sum_i fuel_{i,k} * flights_{i,k} && (3) \\
&\quad J_3 = NO_X = \sum_k \sum_i (NO_X)_{i,k} * flights_{i,k} && (4) \\
&w.r.t: \quad AircraftVariables_k, EngineVariables_k, OperatingConditions_k, flights_{j,k}, pax_{j,k} \\
&s.t.: \quad AircraftConstraints_k, Capacity_i \geq Demand_i, AirplaneContinuity_k, RouteCompatibility_k
\end{aligned}$$

$$\begin{aligned}
&Min: \\
&\quad J_1 = DOC + I = \sum_k [(\$/bhr)_k * bhr_k + (\$/fhr)_k * fhr_k + (\$/flight)_k * flights_k + \\
&\qquad\qquad (\$/yr)_k/365 * aircraft_k + (\$/lb) * \sum_i (fuel_{i,k} * flights_{i,k})] && (5) \\
&\quad J_2 = CO_2 = \sum_k 3.13 \sum_i fuel_{i,k} * flights_{i,k} && (6) \\
&\quad J_3 = NO_X = \sum_k \sum_i (NO_X)_{i,k} * flights_{i,k} && (7) \\
&w.r.t: \quad AircraftVariables_k, EngineVariables_k, OperatingConditions_k, flights_{i,k}, pax_{i,k}, aircraft_k \\
&s.t.: \quad AircraftConstraints_k, J_{fleet} = 0
\end{aligned}$$

$$\begin{aligned}
&Min: J_{fleet} = \sum_k [(OperatingConditions_k - OperatingConditions'_k)^2 + (capacity_k - capacity'_k)^2 + \\
&\qquad (fhr_k - fhr'_k)^2 + (bhr_k - bhr'_k)^2 + (aircraft_k - aircraft'_k)^2 + \sum_i (pax_{i,k} - pax'_{i,k})^2 + \\
&\qquad \sum_i (flights_{i,k} - flights'_{i,k})^2] && (8) \\
&w.r.t: \quad OperatingConditions'_k, capacity'_k, flights_{j,k}, pax'_{i,k} \\
&s.t.: \quad Capacity_i \geq Demand_i, AirplaneContinuity_k, RouteCompatibility_k
\end{aligned}$$

$$
\begin{aligned}
&Min: \\
&\quad J_1 = DOC + I = \sum_k [(\$/bhr)_k * bhr_k + (\$/fhr)_k * fhr_k + (\$/flight)_k * flights_k + \\
&\qquad\qquad (\$/yr)_k/365 * aircraft_k + (\$/lb) * \sum_i (fuel_{i,k} * flights_{i,k})] && (9) \\
&\quad J_2 = CO_2 = \sum_k 3.13 \sum_i fuel_{i,k} * flights_{i,k} && (10) \\
&\quad J_3 = NO_X = \sum_k \sum_i (NO_X)_{i,k} * flights_{i,k} && (11) \\
&w.r.t: \quad AircraftVariables_k, EngineVariables_k, OperatingConditions_k, flights_{i,k}, aircraft_k \\
&s.t.: \quad AircraftConstraints_k, J_{fleet} = 0
\end{aligned}
$$

$$
\begin{aligned}
&Min: \quad J_{fleet} = \sum_k (|aircraft_k - aircraft'_k| + \sum_i |flights_{i,k} - flights'_{i,k}|) && (12) \\
&w.r.t: \quad flights_{j,k} \\
&s.t.: \quad Capacity_i \geq Demand_i, AirplaneContinuity_k, RouteCompatibility_k
\end{aligned}
$$

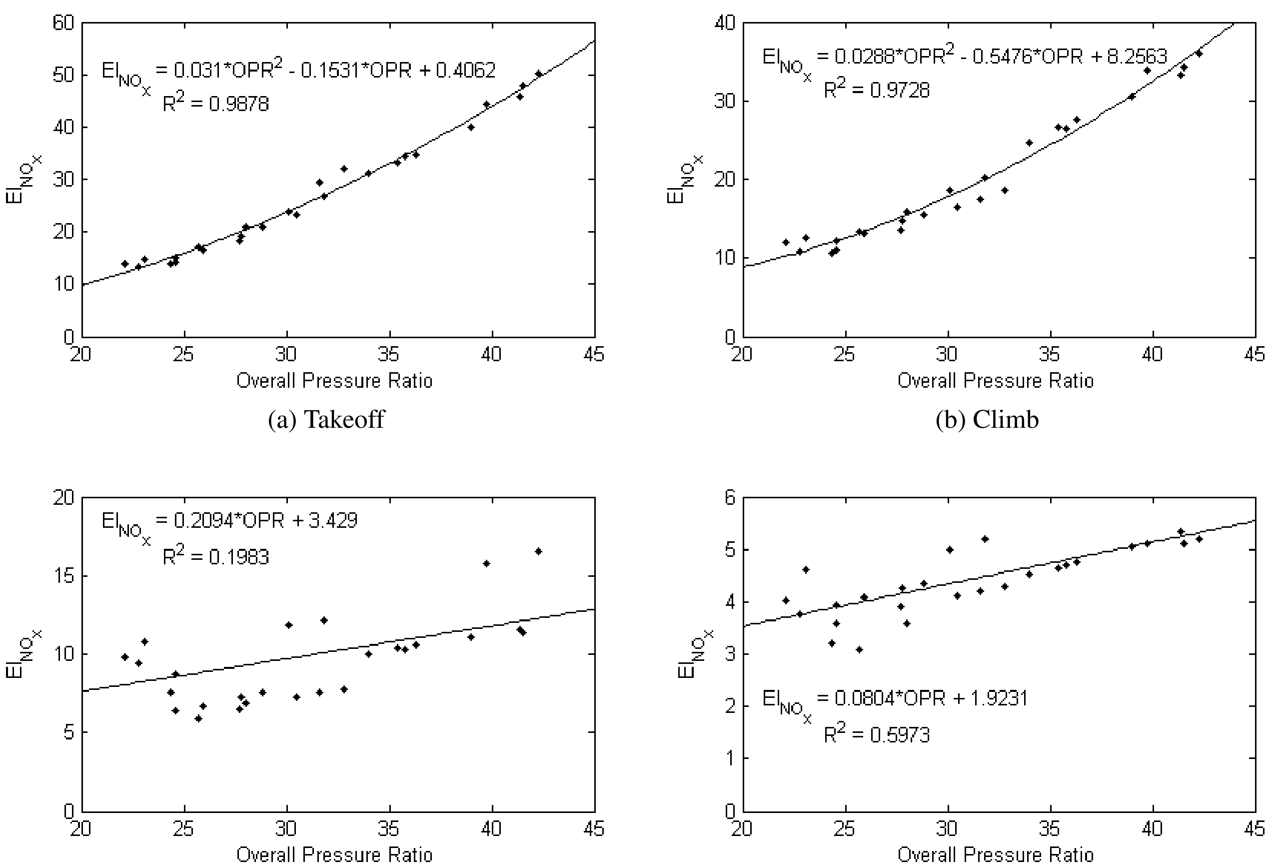

(a) Takeoff

(b) Climb

(c) Approach

(d) Idle

Fig. 1 Fits of NO_X emission index at the four certification points as a function of overall pressure ratio.

$$\begin{aligned}
Sets: \quad & Segment(i), Routes(j), Aircraft(k), City(l) \\
Variables: \quad & flights_{j,k}, aircraftSlack_k, flightsSlack_{i,k} \\
Objective: \quad & J_{fleet} = \sum_k (\sum_i flightsSlack_{i,k} + aircraftSlack_k) && (13) \\
Parameters: \quad & pax_i, capacity_k, flightsTarget_{i,k}, aircraftTarget_k, startCity_{j,l}, \\
& endCity_{j,l}, maxFlights_{j,k}, segmentCount_{i,j} \\
Constraints: \quad & \pm \sum_j flights_{j,k} \mp aircraftTarget_k \leq aircraftSlack_k && (14) \\
& \pm \sum_j segmentCount_{i,j} * flights_{j,k} \mp flightsTarget_{i,k} \leq flightsSlack_{i,k} && (15) \\
& \sum_j startCity_{j,l} * flights_{j,k} = \sum_j endCity_{j,l} * flights_{j,k} && (16) \\
& \sum_j \sum_k segmentCount_{i,j} * capacity_k * flights_{j,k} \geq pax_i && (17) \\
& flights_{j,k} \leq maxFlights_{j,k} && (18)
\end{aligned}$$

$$c(x) \geq b \longrightarrow c(x)/b - 1 \geq 0 \qquad (19)$$

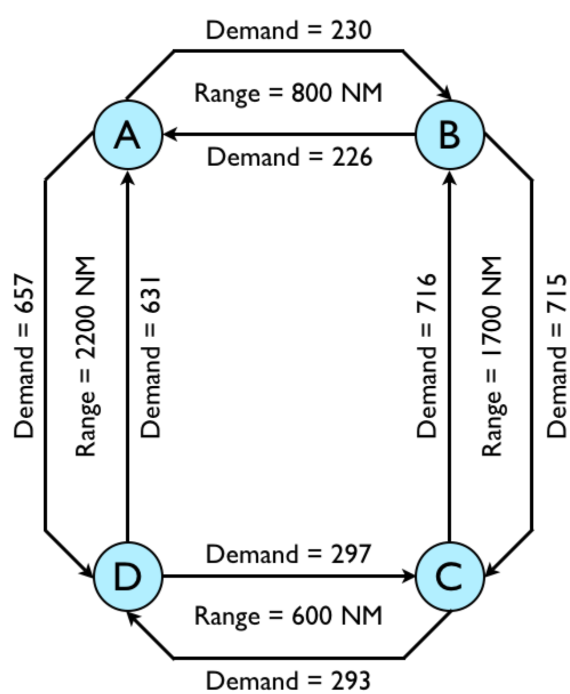

Fig. 2 Route Network for test problem.

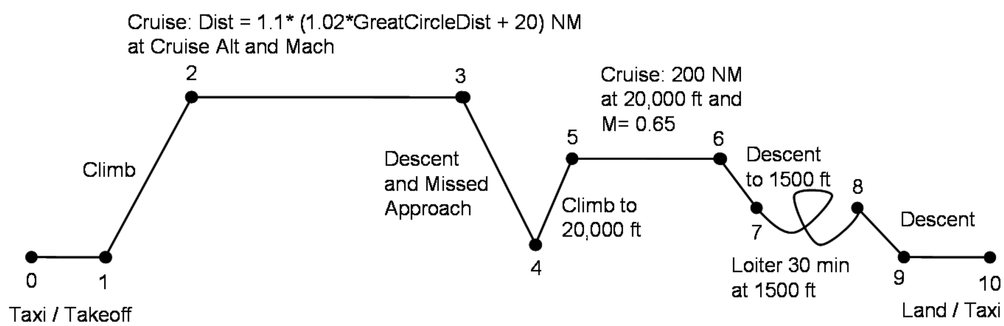

Fig. 3 Reserve mission profile.

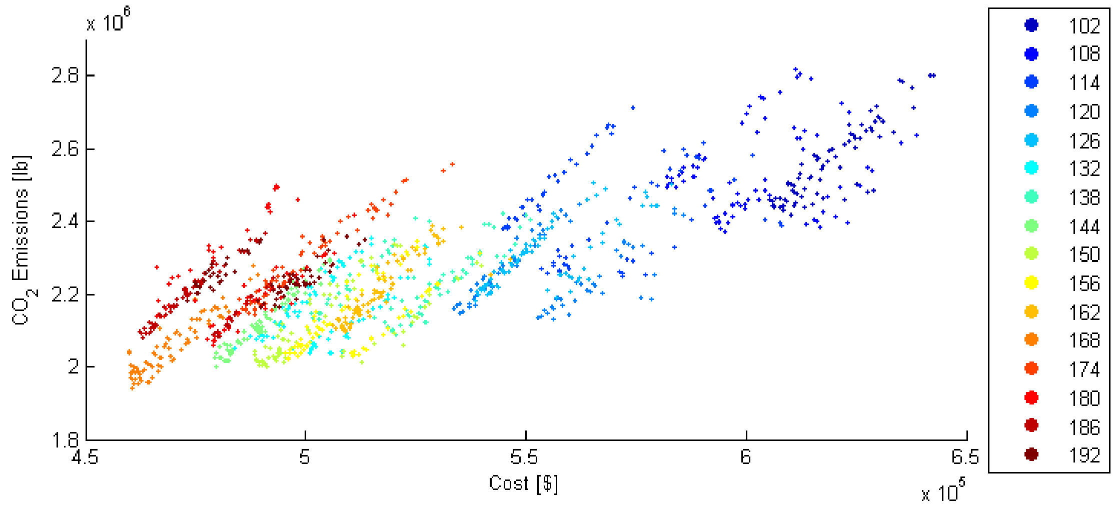

(a) CO_2 vs. $DOC+I$

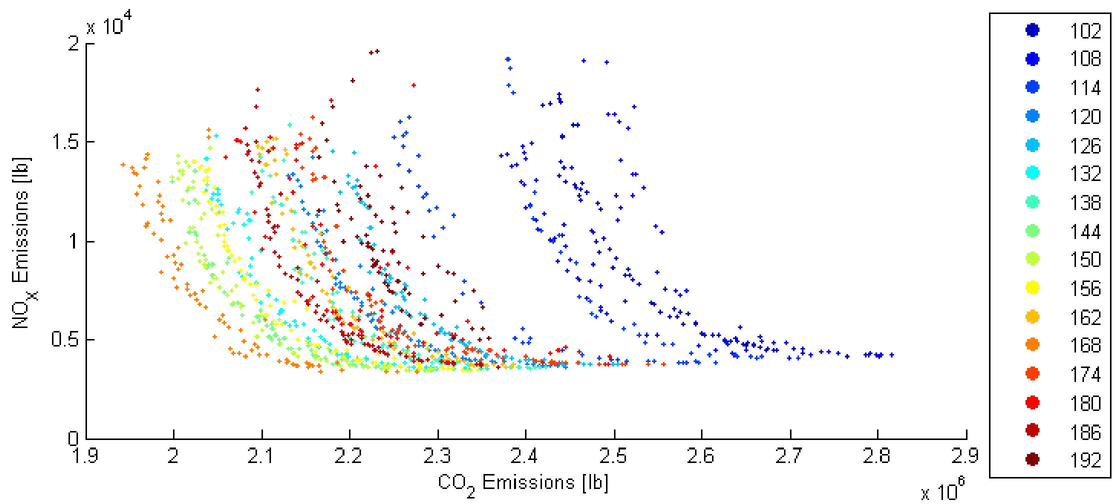

(b) NO_X vs. CO_2

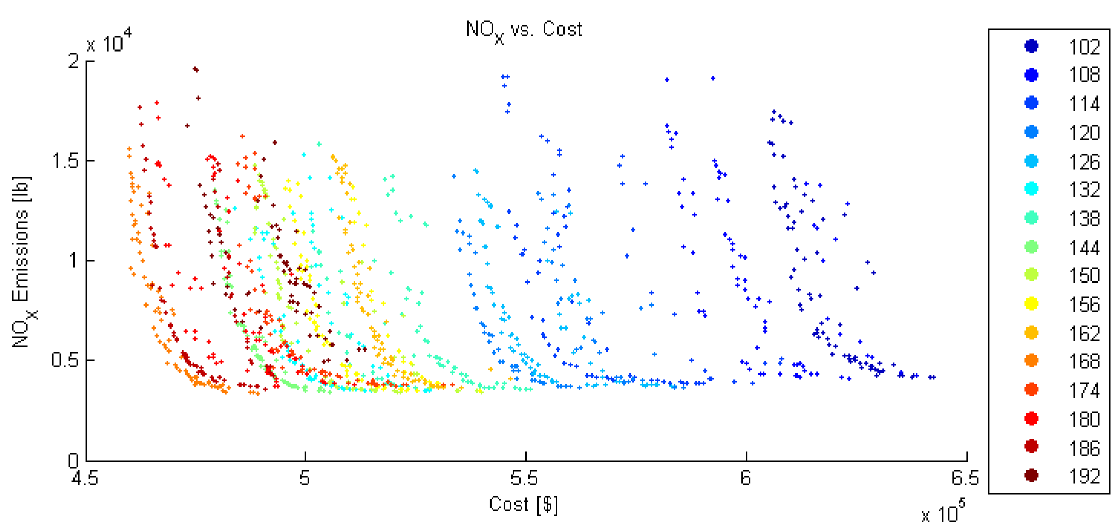

(c) NO_X vs. $DOC+I$

Fig. 4 Seed designs for the final optimization run.

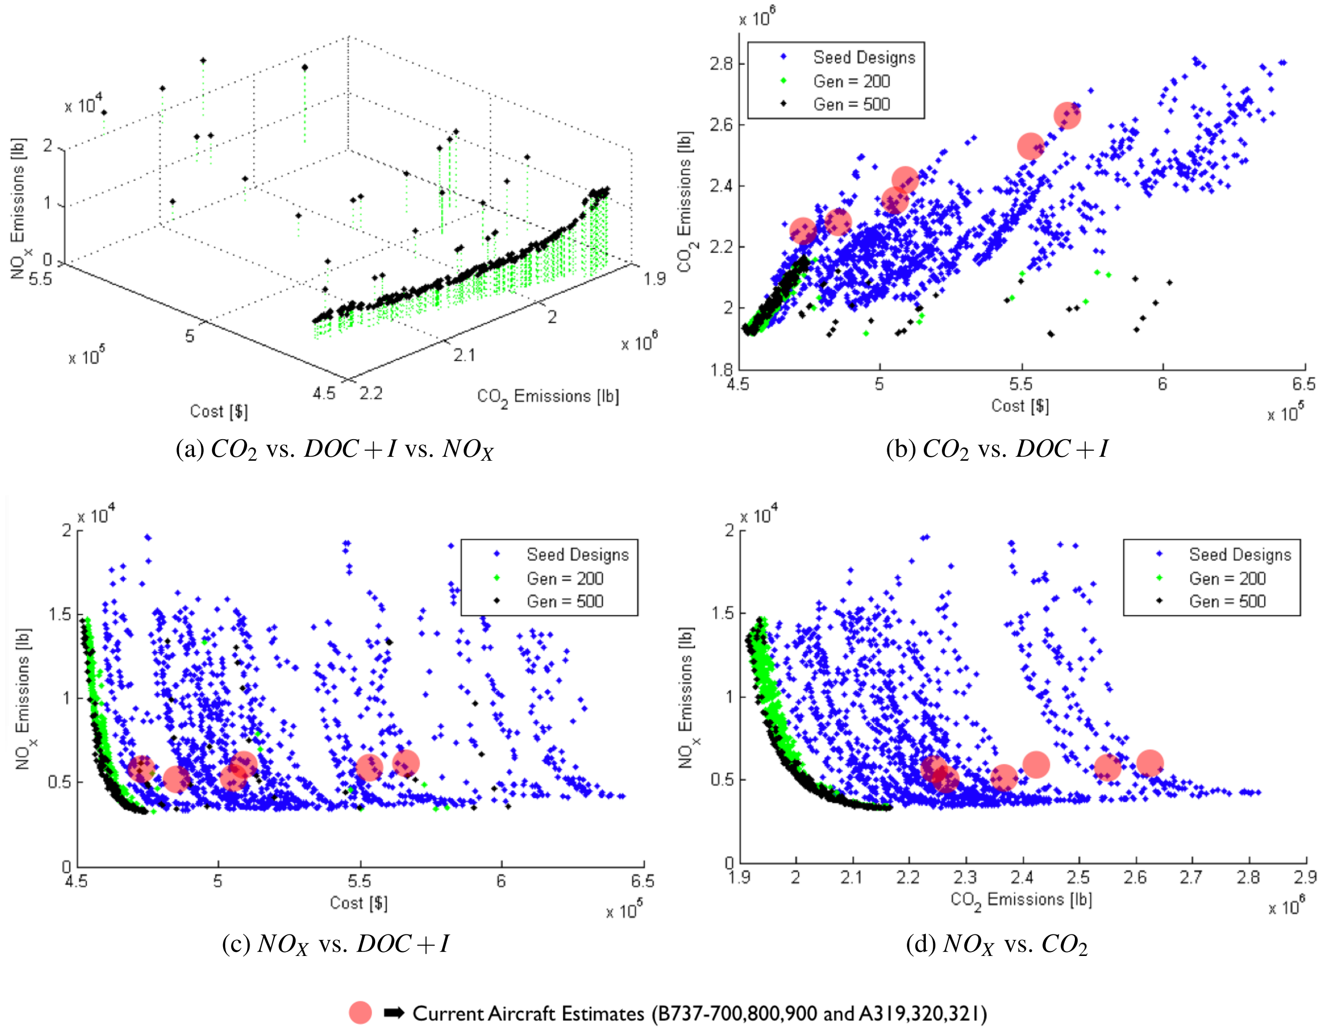

(a) CO_2 vs. $DOC+I$ vs. NO_X

(b) CO_2 vs. $DOC+I$

(c) NO_X vs. $DOC+I$

(d) NO_X vs. CO_2

Fig. 5 The final Pareto front and the projections into the objective planes of the population seeds, the Pareto front at generations 200 and 500, and the estimated performance of current aircraft.

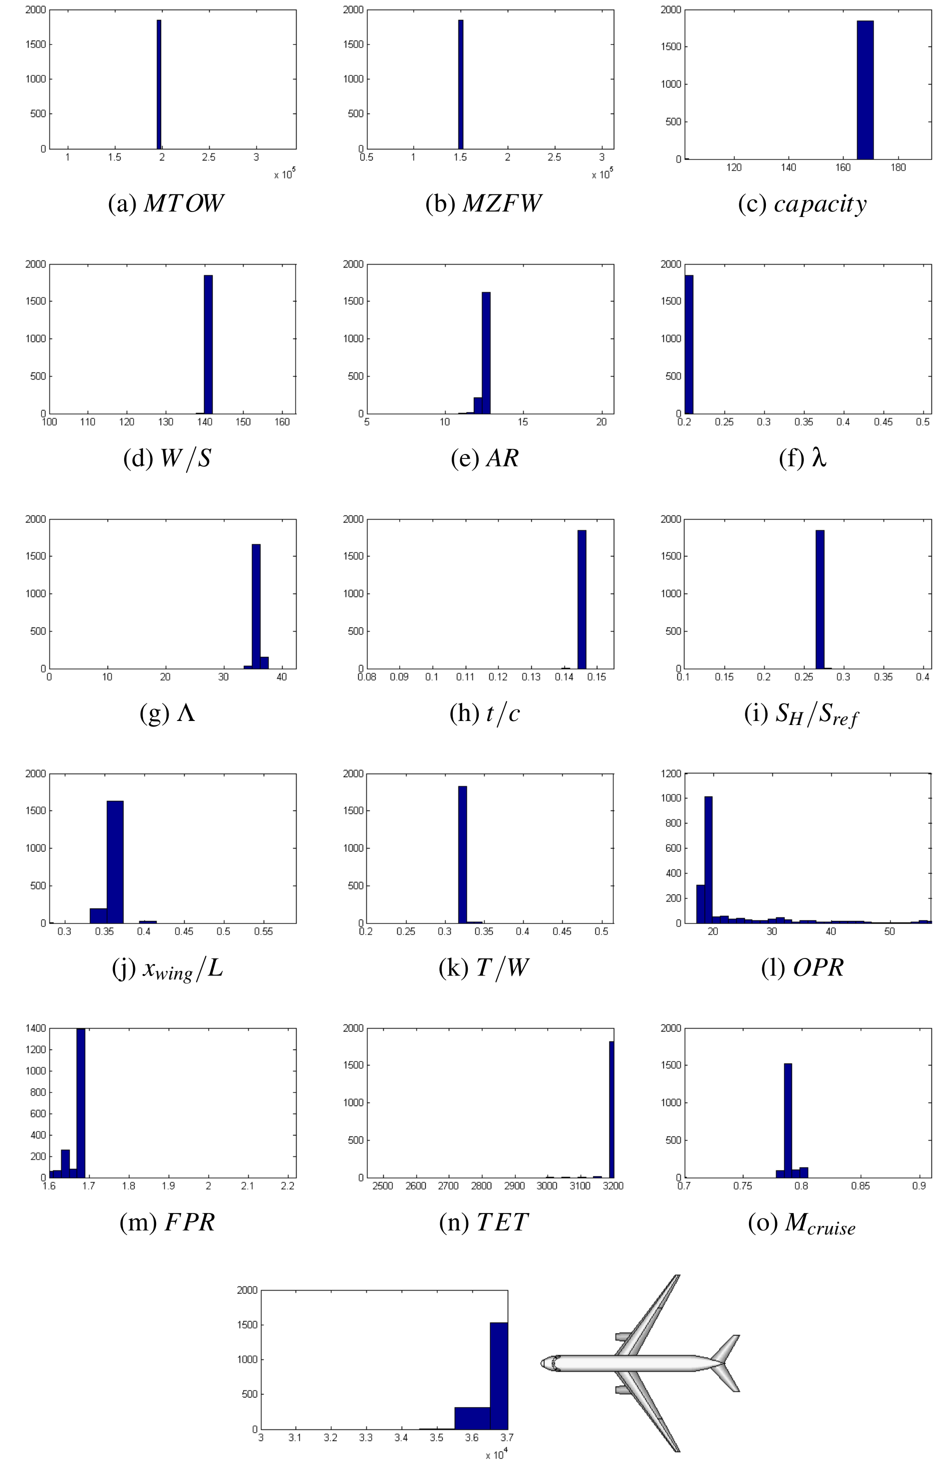

(a) $MTOW$ (b) $MZFW$ (c) $capacity$

(d) W/S (e) AR (f) λ

(g) Λ (h) t/c (i) S_H/S_{ref}

(j) x_{wing}/L (k) T/W (l) OPR

(m) FPR (n) TET (o) M_{cruise}

(p) h_{cruise} (q) Mean Airplane

Fig. 6 Histograms of each airplane design variable and a top view of the mean airplane from the final Pareto front.

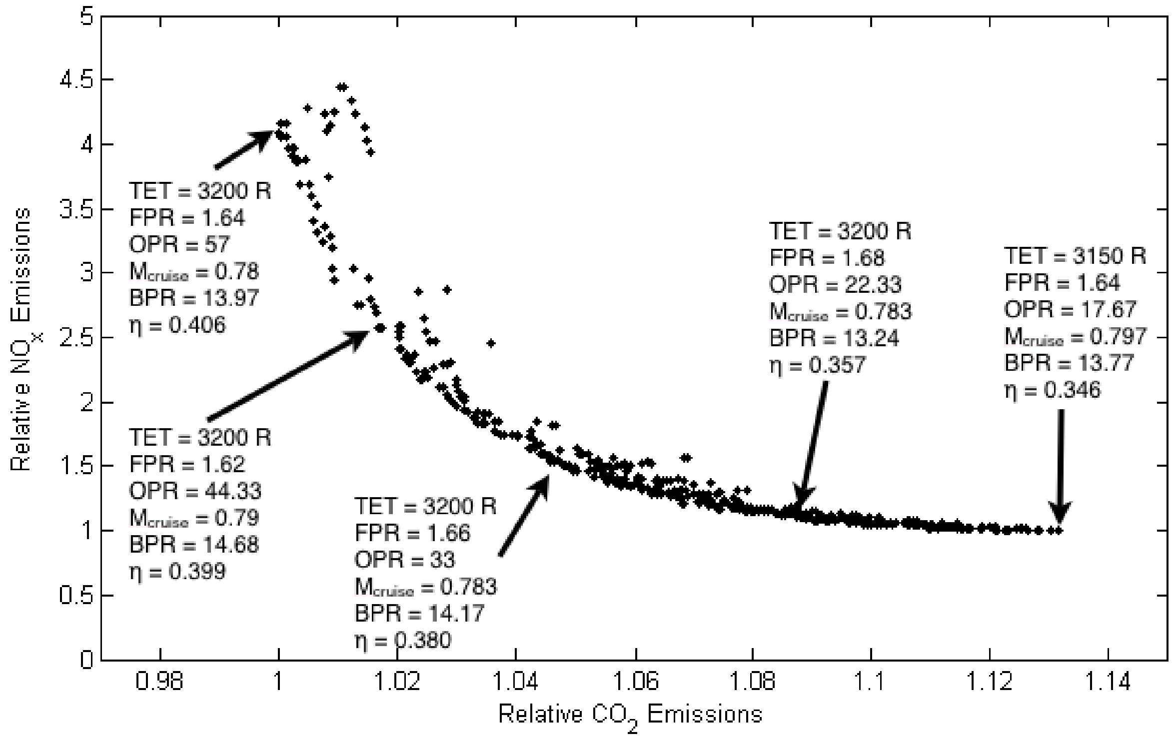

Fig. 7 Engine parameter variation over the final Pareto front.