



Simultaneous aircraft allocation and mission optimization using a modular adjoint approach

John T. Hwang¹ *, Satadru Roy² †, Jason Y. Kao¹ ‡,
Joaquim R. R. A. Martins¹ §, and William A. Crossley² ¶

¹*Department of Aerospace Engineering, University of Michigan, Ann Arbor, Michigan, 48109, United States*

²*School of Aeronautics and Astronautics, Purdue University, West Lafayette, Indiana, 47907, United States*

The aircraft design optimization problem is typically formulated to maximize performance at a small number of representative operating conditions. This approach makes simplifying assumptions such as ignoring the climb and descent phases, but they can be avoided by performing simultaneous design-mission-allocation optimization with surrogate models for the aircraft design disciplines. As a first step towards this goal, this paper presents a method for simultaneous allocation-mission optimization. We integrate aerodynamic and propulsion surrogates, a mission analysis tool, and allocation models within a computational framework that automates solving the coupled simulation and computing derivatives using the adjoint method for gradient-based optimization. We solve the mixed-integer allocation-mission optimization problem by using the linear allocation-only optimization to generate a good starting point and applying the branch-and-bound method to find an optimum for the mixed-integer nonlinear allocation-mission problem. The results show that this approach efficiently finds good local optima with, on average, roughly 10 node evaluations in the branch-and-bound method and most of the continuous optimizations converging almost immediately. The results show that adding next-generation aircraft to a fleet yields a 200-400 % profit increase for a 3-route test problem.

I. Introduction

The average rate of fuel burn of commercial aircraft has halved in the last five decades, but the reductions have begun to plateau in recent years [1]. With each generation of aircraft, improvements become more difficult to achieve, necessitating new technologies and revolutionary changes to maintain the high rate of progress from the early days of commercial aviation. There is significant impetus for progress, with air traffic projected to grow faster than efficiency improvements [2] while environmental concerns and fuel costs continue to rise. As a result, there has been much interest in unconventional configurations, and research is being done to study the potential benefits of the blended wing body (BWB), truss-braced wing (TBW), double bubble (D8), and other new configurations.

Multidisciplinary design optimization (MDO) is a useful tool in this area because it enables rapid design space exploration for these new configurations for which prior knowledge and experience are limited. MDO algorithms that focus on the airframe optimize design variables that parametrize the aerodynamic shape and structural sizing by considering their effect on multiple disciplines simultaneously. These algorithms incorporate computational fluid dynamics (CFD) and finite element analysis (FEA) to accurately analyze the aerodynamic and structural performance, respectively, and they use gradient-based optimizers to efficiently solve optimization problems involving hundreds of design variables. The applications of aerodynamic and aerostructural optimization include studies on composite wings [3, 4], the blended wing body configuration [5], and the common research model (CRM) [6, 7], which is based on the Boeing 777. These applications use the multidisciplinary feasible (MDF) approach [8], which fully converges the coupling between disciplines within each optimization iteration. More information on the MDF and other MDO architectures can be found in the survey by Martins and Lambe [9].

CONSIDERING DESIGN, MISSION, AND ALLOCATION SIMULTANEOUSLY. The problem formulation is one aspect that can be improved in the existing aerostructural optimization algorithms. These algorithms provide tools that can capture subtle tradeoffs and reveal useful insights regarding the airframe design problem. They do so through

*Ph.D. Candidate, AIAA Student Member

†Ph.D. Candidate, AIAA Student Member

‡Ph.D. Candidate, AIAA Student Member

§Associate Professor, AIAA Associate Fellow

¶Professor, AIAA Associate Fellow

detailed parametrizations of the aircraft geometry and structure, as well as accurate computational models for the aerodynamic and structural performance. The problem formulation could benefit from applying a similar level of detail and accuracy.

When considering only aerodynamics, the problem is often formulated as a lift-constrained drag minimization, with drag evaluated at one or more lift values at which the aircraft is expected to operate. An improvement upon this is an optimization problem that maximizes range via the Breguet range equation. This not only adds a metric for structural efficiency in the objective, but it also considers the equilibrium equations, taking into account the gradual decrease of weight as the aircraft burns fuel. However, the Breguet range equation assumes that aerodynamic and propulsive efficiency are constant—i.e., lift-to-drag ratio and thrust-specific fuel consumption, respectively, are fixed. This approach has been extended further by using actual mission data of an existing aircraft to determine a set of representative missions [10].

To eliminate the error caused by these assumptions, this paper seeks to analyze the full mission in a particular flight, letting the altitude profile be continuously varied by the optimizer. In this formulation, the net performance over the entire mission is the metric to be optimized, instead of a set of representative conditions. Moreover, the airline-level aircraft allocation problem is also simultaneously considered. A mission analysis assumes a range and a payload; however, an aircraft is flown on a wide variety of routes with different payloads over its life. The aircraft design variables depend on the routes on which the aircraft is flown, and the allocation in turn depends on the aerodynamic and structural performance resulting from the choice of aircraft design variables. This coupling necessitates simultaneous optimization of the aircraft allocation as well. Therefore, the motivation for this paper is to enable MDO algorithms to extend to a simultaneous optimization of the aircraft design, mission, and allocation.

PROPOSED APPROACH. The simultaneous design, mission, and allocation problem has an ambitious scope and significant computational challenges. Namely, the number of high-fidelity evaluations required to solve the problem is

$$\begin{array}{cccccc} \mathcal{O}(100) & \times & \mathcal{O}(10) & \times & \mathcal{O}(100) & \times & \mathcal{O}(100) & = & \mathcal{O}(10^7), \\ \# \text{ mission} & & \# \text{ mission} & & \# \text{ routes} & & \# \text{ optimization} & & \# \text{ aerostructural} \\ \text{points} & & \text{iterations} & & & & \text{iterations} & & \text{evaluations} \end{array} \quad (1)$$

which is too large for the solution of the simultaneous problem to be feasible without modifications. The number of mission points refers to the discretization of the mission profile.

The proposed solution is to use a surrogate model for the aerostructural performance that is re-trained every optimization iteration. The outputs of aerostructural analysis depend on the aircraft design variables and the operating conditions. Capturing the dependence on the design variables requires a high-fidelity model because the shape and sizing variables number in the hundreds or thousands in state-of-the-art MDO algorithms. Simplified, low-order models do not capture the relevant physics in sufficient detail—e.g., a vortex lattice code cannot distinguish airfoil thickness changes—and surrogate models with hundreds or thousands of input variables are too inefficient. In contrast, there are only a small number of operating conditions that are inputs to the aerodynamic and structural analyses—angle of attack, Reynolds number, Mach number, and trim deflection. Thus, the proposed solution is to generate the sample points for the surrogate model once in each optimization iteration, and evaluate the surrogate model for each route, each mission iteration, and each mission point. With this solution, the number of high-fidelity evaluations is

$$\begin{array}{cccc} \mathcal{O}(10) & \times & \mathcal{O}(100) & = & \mathcal{O}(10^3), \\ \# \text{ sample} & & \# \text{ optimization} & & \# \text{ aerostructural} \\ \text{points} & & \text{iterations} & & \text{evaluations} \end{array} \quad (2)$$

which is tractable. Note: though the emphasis has been on airframe design, a propulsion surrogate is also necessary for mission analysis.

OBJECTIVE. Even with this improvement, there are many additional challenges that must be overcome. A computational framework is needed for integrating the aerostructural surrogate, propulsion surrogate, aerostructural analysis tool, mission analysis tool, and allocation models, which are all shown in Fig. 1. Many of these components have multiple instances—e.g., due to multiple missions—and some should be run in parallel, so the resulting complexity must be managed by automating parallel data passing between components. Another critical requirement is that gradients must be efficiently computed using the adjoint method to enable gradient-based optimization. All of these requirements are addressed by using a recently-developed computational framework for large-scale optimization [11] that has been used to solve a 7-discipline satellite optimization problem with more than 25,000 design variables [12].

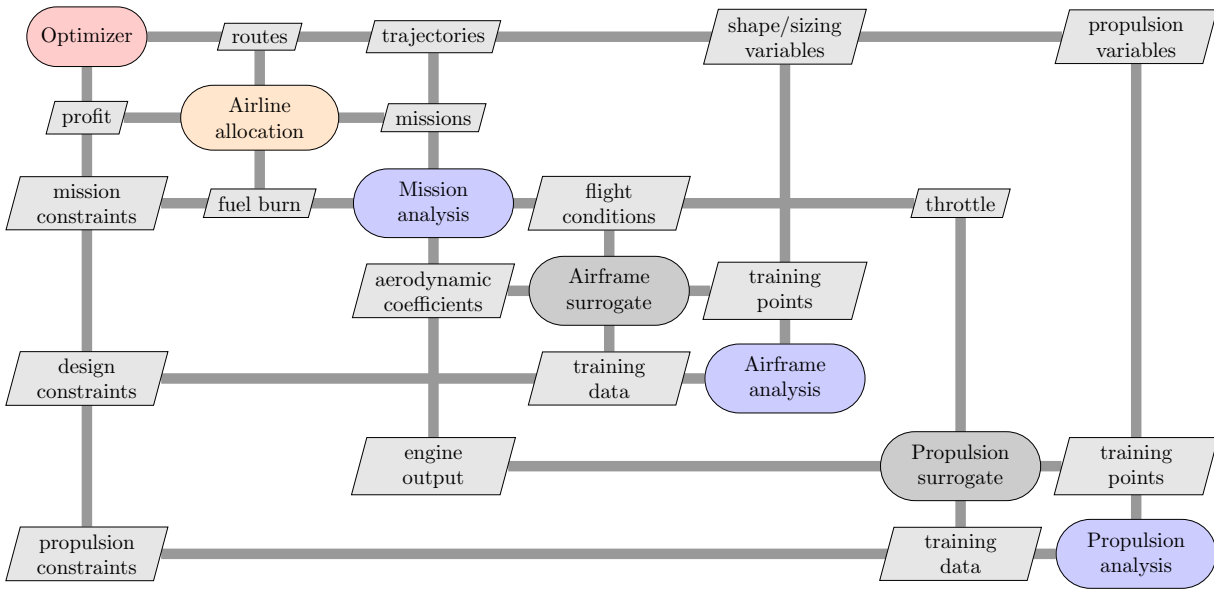


Figure 1: Extended design structure matrix [13] of the components necessary for simultaneous design-mission-allocation optimization.

The objective of this paper is to perform simultaneous allocation-mission optimization with aerodynamic and propulsion surrogates. This problem represents a necessary intermediate step before solving the full design-mission-allocation problem, and one that requires solutions to many of the key challenges for the full problem.

The paper is organized as follows. First, we survey the aircraft allocation and mission optimization literatures in Sec. II. Next, we describe the methodology for the framework, the mission analysis, the allocation problem, and the combined problem in Sec. III. In Sec. IV, we describe the results for a suite of allocation-mission optimization problems.

II. Literature survey

In this section, we survey the literatures in the aircraft allocation and mission optimization fields. In both areas, there are unique requirements introduced by our objective of combining allocation and mission optimization into one problem. On the allocation side, the key difference is that the simultaneous problem is necessarily nonlinear and each evaluation of the objective and constraints is significantly more expensive. On the mission side, modularity is a key requirement for the simultaneous problem because multiple mission instances are combined within a single allocation problem.

A. Aircraft allocation

One of the first applications of linear programming was the airline allocation problem, first published in 1956 [14], less than 10 years after the development of the simplex method in 1947 [15]. The original problem aimed to find the optimal allocation of aircraft to routes given uncertain demand [14], assuming a fixed cost incurred by flying a given route with a given type of aircraft.

Beginning in 1966 [16], researchers began applying optimization to flight schedule design, and later, the fleet assignment problem (FAP)—assigning aircraft to scheduled flights. In 1989, Abara showed that it is possible to use linear programming to solve FAPs of representative sizes [17]. Recent research has focused on improvements upon these formulations with for instance improved revenue models [18, 19], homogeneity of aircraft types at airports [20] or on legs [21], and simultaneous optimization of scheduling and fleet assignment [22]. More information can be found in survey articles by Clarke and Smith [23] and Barnhart et al. [24].

The fruits of research in this field have been used by airlines to improve their operations and increase profit. For instance, algorithms for solving fleet assignment problems have been used by American Airlines [17], Delta Airlines [25], and US Airways [26].

For the current problem, the motivation for considering the allocation problem differs from that of the airlines. The objective is to solve the allocation problem as a tool to evaluate potential new aircraft designs by quantifying their airline-level benefits. We want to be able to simultaneously optimize the variables parametrizing the mission profile and the design of the aircraft, so simplifying assumptions are used. Specifically, we use a variant of the original allocation problem proposed in 1956 [14], and the considerations for scheduling are ignored. The current allocation formulation maximizes the profit across the entire network of the representative airline subject to operational and demand constraints, following the approach of [27, 28, 29, 30].

B. Mission optimization

The prior art in mission profile optimization divides broadly into two categories, direct and indirect [31]. The direct approach first discretizes the equilibrium equations and applies the optimality conditions on the discretized equations. The indirect approach takes the reverse order—it first differentiates the equilibrium equations to derive continuous optimality conditions and then discretizes as the second step.

Only the direct approach is reviewed as it is the one used in this chapter. Betts and Huffman performed trajectory optimization for a hypersonic re-entry vehicle using the direct approach [32]. They used a custom active-set gradient-based optimization algorithm with a sparse finite differencing method. Betts and Cramer applied a similar methodology to the trajectory optimization of commercial aircraft using sparse finite differencing and gradient-based optimization, and B-splines interpolants were used as the aerodynamic model [33]. Park and Clarke solved the trajectory optimization problem using the pseudospectral method, which discretizes state and control variables at points determined by a Gaussian quadrature [34].

An application of mission profile optimization is continuous descent approach (CDA), also known as optimized descent profile (ODP). CDA is one component of the next generation air transportation system (NextGen) currently being developed and implemented by the federal aviation administration (FAA). In addition to providing noise reductions, continuous descents yield fuel burn improvements compared to step descents. In the past decade, continuous and idle descent tests have been designed or implemented at Louisville [35], Denver [36], Los Angeles [37], and Malta [38] international airports.

Based on the current interest in CDA, the approach taken in this paper is to simultaneously optimize the mission profiles for each route while solving the allocation problem. The expected result is a rapid and smooth climb to a cruise-climb stage and then a continuous descent near idle, with the cruise altitudes determined by the optimizer. The numerical approach is similar to that of Betts and Cramer [33], with 2 key improvements. First, the derivatives needed for optimization are computed using the framework as opposed to finite differences for improved accuracy and efficiency. Second, the mission analysis is implemented in a modular way in the framework, facilitating integration with the allocation components, and in the future, with the aircraft design components.

III. Methodology

A. Computational Modeling Framework

THEORY. The theoretical foundation for the computational framework is the representation of the multidisciplinary analysis (MDA) problem as a nonlinear system of equations [11]. A vector that solves this nonlinear system corresponds to the solution of the MDA problem and the process by which the MDA problem is solved can be interpreted as one of the known methods for solving a nonlinear system.

The starting point for the theory is the concatenation of all design, state, intermediate, input, and output variables into a single vector of variables which is the vector of unknowns in the nonlinear system. This includes parameters and variables that are the outputs of the objective function and constraint functions. Each variable must have a residual function of the same size because for every scalar variable, there must be a corresponding equation that constrains its value.

To achieve this, each variable is classified as an explicit, implicit, or independent variable. Let us assume there are n variables in total when all the variables are concatenated into one list. We also assume that each variable is itself a vector. For an explicit variable v_i the residual is $R_i = v_i - V_i(v_1, \dots, v_{i-1}, v_{i+1}, \dots, v_n) = 0$ where V_i is the function that computes v_i , for an implicit variable the residual is simply $R_i = \mathcal{R}_i(v_1, \dots, v_n) = 0$ since v_i already has an associated residual function, \mathcal{R}_i , and for an independent variable the residual is $R_i = v_i - v_i^*$ where v_i^* is the value at which we evaluate the computational model, since this is an input variable.

If we define $R = (R_1, \dots, R_n)$ and $u = (v_1, \dots, v_n)$, we can construct a nonlinear system $R(u) = 0$ as the mathematical representation of the MDA problem. Since the Jacobian $\partial R / \partial u$ is invertible in all practical problems,

the inverse function theorem gives the following relationship for the Jacobian at a converged point [39]:

$$\frac{\partial R}{\partial u} \frac{du}{dr} = \mathcal{I} = \frac{\partial R}{\partial u}^T \frac{du}{dr}^T. \quad (3)$$

Equation (3) unifies all known discrete methods for computing derivatives with the left and right hand sides corresponding to the direct and adjoint methods, respectively, and the forward and reverse modes of algorithmic differentiation. The significance of Eq. (3) is that it allows the framework to automatically compute total derivatives simply by solving this linear system. The user is required to specify partial derivatives of all outputs with respect to all inputs for each discipline or block of code, and from this information, the framework can compute the total derivatives simply by solving the linear system in Eq. (3). Depending on whether the disciplines are coupled and have implicit state variables, the process of solving this linear system reduces to the adjoint method, chain rule, or one of many other known methods.

IMPLEMENTATION The framework implementation has several functions. At its core, it is a piece of software that allows the user to define the components or disciplines of the problem in a modular way. The framework automates transfer of data between components, including parallel data communication when the source and destination of the data are on different processors, and provides the infrastructure for solving the nonlinear system and Eq. (3).

The framework uses object-oriented programming, with instances of a System class as the basic building blocks. In general, the nonlinear system $R(u) = 0$ is hierarchically decomposed, where the top-level system owns the full set of variables u and the set of variables is partitioned at each subsequent level. Figure 2 shows the hierarchy tree that results from this decomposition in which systems contain other systems and each system is uniquely defined by the set of variables it contains.

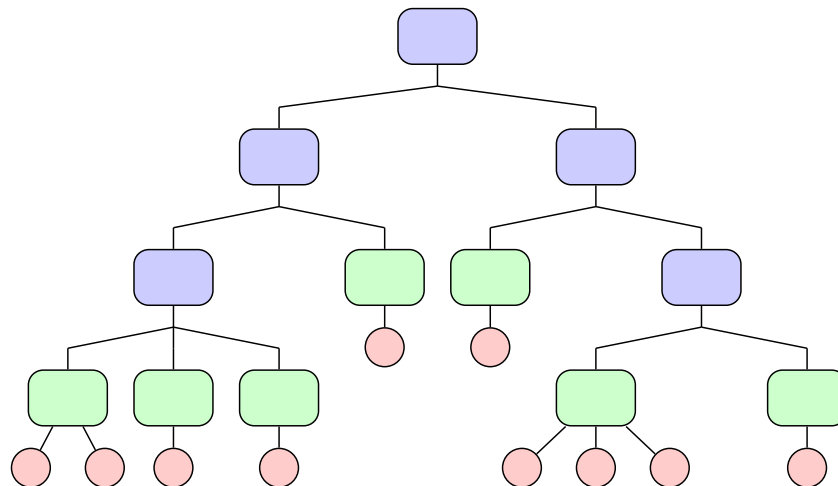


Figure 2: Hierarchy tree for a notional problem. The circles are variables and the rounded rectangles are systems with those in green considered the leaves of the tree (if the circle nodes are ignored).

Each system is responsible for implementing methods that evaluate its owned entries of the R vector and the $\frac{\partial R}{\partial u}$ matrix. Each system can then solve its child systems using nonlinear or linear block Gauss-Seidel or Jacobi, and the child systems may contain their own child systems or solve their variables using a monolithic approach such as Newton's method. Systems can also optionally implement an approximate inverse for its Jacobian if a preconditioner is available. Both the Jacobian and the preconditioner are implemented as linear operators so the framework does not make restrictions on whether they are sparse or dense or have some other special structure, or if only the matrix-vector product operation is available. The framework uses the Portable, Extensible Toolkit for Scientific Computation (PETSc) [40] to handle parallel data communication among systems, and for PETSc's built-in iterative solvers. For more details, we refer the reader to Hwang and Martins [11].

B. The mission problem

The main components of the mission analysis are formed by the equilibrium equations for the aircraft and the ordinary differential equation (ODE) for fuel weight. For the former, the vertical equilibrium determines the required angle of

attack to roughly match lift with weight, the horizontal equilibrium determines the engine throttle setting to roughly match thrust with drag, and the moment equilibrium determines the elevator deflection to trim the aircraft. The equations are

$$L + W \cos \gamma - T \sin \alpha + \frac{W}{g} v^2 \cos \gamma \frac{d\gamma}{dx} = 0 \quad (4)$$

$$T \cos \alpha + D + W \sin \gamma + \frac{W}{g} v \cos \gamma \frac{dv}{dx} = 0 \quad (5)$$

$$M - I \left(\frac{d^2 \gamma}{dx^2} (v \cos \gamma)^2 + \frac{d\gamma}{dx} \frac{dv}{dx} v (\cos \gamma)^2 \right) = 0, \quad (6)$$

where α is the angle of attack and γ is the climb or descent angle.

Assuming quasi-steady flight conditions, the terms containing derivatives with respect to x can be ignored. Dropping these terms and expressing in terms of non-dimensional coefficients yields

$$\frac{1}{2} \rho v^2 S \tilde{C}_L + W \cos \gamma - \frac{1}{2} \rho v^2 S \tilde{C}_T \sin \alpha = 0 \quad (7)$$

$$\frac{1}{2} \rho v^2 S \tilde{C}_T \cos \alpha + \frac{1}{2} \rho v^2 S C_D + W \sin \gamma = 0 \quad (8)$$

$$\tilde{C}_M = 0. \quad (9)$$

The system of equations is completed by adding surrogate models for the aerodynamics and propulsion, which are of the form

$$\tilde{C}_L - C_L(h, M, \alpha, \eta) = 0 \quad (10)$$

$$\tilde{C}_M - C_M(h, M, \alpha, \eta) = 0 \quad (11)$$

$$\tilde{C}_T - C_T(h, M, t) = 0, \quad (12)$$

where the variables with tildes indicate target variables while C_L , C_D , C_M , and C_T are functions representing the aerodynamic or propulsion surrogates. In this problem, the aerodynamic surrogates are tensor-product B-splines interpolants, and the propulsion surrogate is a simple linear function.

The fuel weight is computed by solving the ODE,

$$\dot{W}_f = \frac{\{SFC\} \frac{1}{2} \rho v^2 S C_T}{v \cos \gamma}, \quad (13)$$

where SFC is the thrust-specific fuel consumption, and C_T is the thrust coefficient.

This problem involves multiple disciplines that are coupled. The rate of fuel burn at a given point during the flight depends on the engine output, which is determined by the amount of aerodynamic drag that the engines must overcome. However, the amount of drag depends on how much lift must be produced to counteract the weight of the aircraft, which in turn is affected by the total amount of fuel required, completing the feedback loop.

Figure 3 shows a class diagram containing the components and solvers used in the computational model. The top-level *System* object, *Mission*, contains five systems that have no feedback, so nonlinear and linear block Gauss–Seidel converge in one iteration. One of the five subsystems, *Coupled analysis*, contains the aforementioned feedback loop involving fuel burn, thrust, drag, lift, and weight. For this system, Newton’s method is used to solve the nonlinear system with a few iterations of nonlinear block Gauss–Seidel used as a start-up strategy, the linear systems that arise are solved using fGMRES, and the preconditioner is a few iterations of linear block Gauss–Seidel. *Coupled analysis*, in turn, contains a system that defines implicit state variables, so it also uses a Newton–Krylov solver.

The mission-only optimization is a *nonlinear programming (NLP)* problem with the objective of minimizing fuel burn and is stated below (**NLP-m**). The design variables are the B-spline control points of the parametrized altitude profile with n_{cp} B-spline control points and n_{pt} discretization points. In this paper, we use $n_{cp} = 50$ and $n_{pt} = 250$. The initial and final altitudes are constrained to be zero, and slope constraints are imposed on the altitude profile with a value of 35° for both climbing and descending. The throttle setting is constrained to be between idle (10%) and full (100%). The slope constraints are linearly mapped from the design variables, but the throttle setting is a nonlinear function of the design variables and state variables, so its derivatives are computed using the adjoint method. Since it is costly to solve one adjoint for each of the n_{pt} points, the throttle constraints are aggregated using Kreisselmeier–Steinhauser functions [41] to reduce to a single idle-throttle constraint and a single full-throttle constraint.

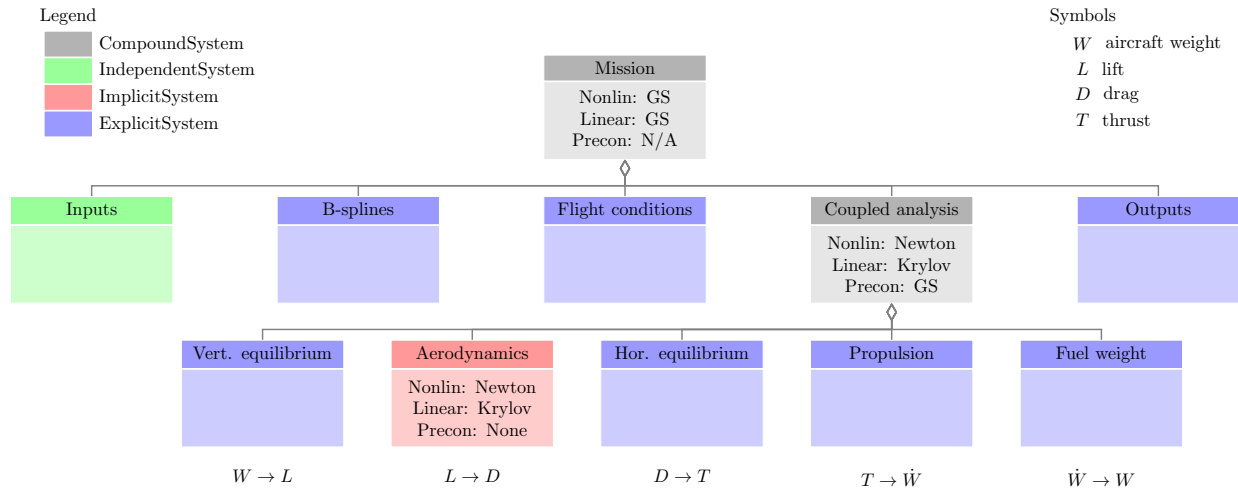


Figure 3: A class diagram showing the *System* instances in the aircraft mission problem.

<p style="text-align: center;"> minimize fuel_burn with respect to altitude_k ∈ [0, max_altitude] , 1 ≤ k ≤ n_{cp} subject to altitude₁ = 0 altitude_{n_{cp}} = 0 idle_throttle ≤ throttle (KS) throttle ≤ max_throttle (KS) min_slope ≤ slope_k ≤ max_slope , 1 ≤ k ≤ n_{pt} </p> <p style="text-align: right;">(NLP-m)</p>

C. The allocation problem

The allocation problem mimics the motivation of a profit-seeking airline. Given n_{rt} routes and n_{ac} types of aircraft, the allocation problem we wish to solve in this paper seeks to maximize profit by optimally allocating the available aircraft to the routes with constraints on the number of each aircraft type available as well as the demand for each. In this paper, i is used to index routes, and j is used to index types of aircraft. The design variables are $pax_flt_{i,j}$, which is the number of passenger per flight, and $flt_day_{i,j}$, which is the number of flights per day. Both design variables are arrays over routes and types of aircraft. There are 3 relevant equations for the allocation problem.

The first relevant equation defines the profit:

$$profit = \sum_i^{n_{rt}} \sum_j^{n_{ac}} [price_pax_{i,j} \cdot pax_flt_{i,j} \cdot flt_day_{i,j}] \tag{14}$$

$$- \sum_i^{n_{rt}} \sum_j^{n_{ac}} [(cost_flt_{i,j} + cost_fuel \cdot fuel_flt_{i,j}) \cdot flt_day_{i,j}], \tag{15}$$

where $price_pax_{i,j}$ is the ticket price per flight, $cost_flt_{i,j}$ is the total cost of operating a flight minus fuel, $cost_fuel$ is the cost per unit fuel, and $fuel_flt_{i,j}$ is the total fuel burn on a flight. In some cases, $fuel_flt_{i,j}$ is assumed to be a constant during an allocation optimization, and in others, it is the output of the mission analysis.

The second equation is an inequality that enforces that the total number of passengers flown on a route is less than the demand on that route, and it is given by

$$total_pax_i = \sum_j^{n_{ac}} [pax_flt_{i,j} \cdot flt_day_{i,j}] \leq demand_i , \quad 1 \leq i \leq n_{rt}, \tag{16}$$

where $demand_i$ is the total number of passengers who are looking to fly on a given route. This inequality is enforced for each route i .

The final equation is another inequality that constrains aircraft operation time based on the total number of available aircraft for a given type, and it is given by

$$\text{total_usage}_j = \sum_i^{n_{rt}} [\text{flt_day}_{i,j} \cdot (\text{time_flt}_{i,j}(1 + \text{maint}_j) + \text{turn_flt})] \leq 12\text{hr} \cdot \text{num_ac}_j, \quad 1 \leq j \leq n_{ac}, \quad (17)$$

where $\text{time_flt}_{i,j}$ is the block time for a flight, maint_j is the maintenance time required as a multiple of block time, turn_flt is the turnaround time between flights, and num_ac_j is the number of aircraft available for a given type. As with block fuel, block time is either a constant during optimization, or it is computed by the mission analysis. This inequality is enforced for each aircraft type j .

The allocation-only optimization problem is a *mixed-integer linear programming (MILP)* problem and is stated below (MILP-a). In this case, block fuel and block time are constants, and they are not dynamically re-computed during optimization by running mission analyses. The objective function is profit, and the two integer design variables are passengers per flight and flights per day for each route and type of aircraft. Demand constraints are enforced for each route, and aircraft availability constraints are enforced for each type of aircraft. Here, the demand is assumed symmetric and the model assumes that each aircraft type makes a round trip on each route. This partially addresses the scheduling-related constraints.

<p style="margin: 0;">maximize profit</p> <p style="margin: 0;">with respect to $\text{pax_flt}_{i,j} \in [0, \text{ac_capacity}_j] \cap \mathbb{N}, \quad 1 \leq i \leq n_{rt}, 1 \leq j \leq n_{ac}$</p> <p style="margin: 0;">$\text{flt_day}_{i,j} \in [0, \infty) \cap \mathbb{N}, \quad 1 \leq i \leq n_{rt}, 1 \leq j \leq n_{ac}$</p> <p style="margin: 0;">subject to $\text{total_pax}_i \leq \text{demand}_i, \quad 1 \leq i \leq n_{rt}$</p> <p style="margin: 0;">$\text{total_usage}_j \leq \text{num_ac}_j, \quad 1 \leq j \leq n_{ac}$</p>	(MILP-a)
--	----------

D. The allocation-mission problem

The main optimization problem we wish to solve in this paper is the simultaneous allocation-mission optimization with profit as the objective function. This top-level problem includes the allocation problem as well as the mission problems for each route and each type of aircraft. However, for this paper, mission analyses are only run for the new types of aircraft. The rationale is that in future work with the simultaneous design-mission-allocation optimization, the existing types of aircraft have fixed designs so the block fuel and block time only depend on the routes, which are fixed, and the weight of the passengers, which does not have a large influence. In contrast, the new types of aircraft would be simultaneously designed during the design-mission-allocation optimization, so it is necessary to re-compute the mission analyses dynamically during optimization. In light of this, let us call the number of new types of aircraft n_{nac} .

The simultaneous allocation-mission optimization is a *mixed-integer nonlinear programming (MINLP)* problem and is stated below (MINLP-a-m). To lower the difficulty of the problem, the number of passengers per flight can be relaxed as a continuous variable; however, the same assumption cannot be made for the number of flights per day because the small values signify that relaxing the integer constraints would significantly change the optimization problem.

<p style="margin: 0;">maximize profit</p> <p style="margin: 0;">with respect to $\text{pax_flt}_{i,j} \in [0, \text{ac_capacity}_j], \quad 1 \leq i \leq n_{rt}, 1 \leq j \leq n_{ac}$</p> <p style="margin: 0;">$\text{flt_day}_{i,j} \in [0, \infty) \cap \mathbb{N}, \quad 1 \leq i \leq n_{rt}, 1 \leq j \leq n_{ac}$</p> <p style="margin: 0;">$\text{altitude}_{i,j,k} \in [0, \text{max_altitude}], \quad 1 \leq i \leq n_{rt}, 1 \leq j \leq n_{nac}$</p> <p style="margin: 0;">$1 \leq k \leq n_{cp}$</p> <p style="margin: 0;">subject to $\text{total_pax}_i \leq \text{demand}_i, \quad 1 \leq i \leq n_{rt}$</p> <p style="margin: 0;">$\text{total_usage}_j \leq \text{num_ac}_j, \quad 1 \leq j \leq n_{ac}$</p> <p style="margin: 0;">$\text{altitude}_{i,j,1} = 0, \quad 1 \leq i \leq n_{rt}, 1 \leq j \leq n_{nac}$</p> <p style="margin: 0;">$\text{altitude}_{i,j,n_{cp}} = 0, \quad 1 \leq i \leq n_{rt}, 1 \leq j \leq n_{nac}$</p> <p style="margin: 0;">$\text{idle_throttle} \leq \text{throttle}_{i,j}(\text{KS}), \quad 1 \leq i \leq n_{rt}, 1 \leq j \leq n_{nac}$</p> <p style="margin: 0;">$\text{throttle}_{i,j} \leq \text{max_throttle}(\text{KS}), \quad 1 \leq i \leq n_{rt}, 1 \leq j \leq n_{nac}$</p> <p style="margin: 0;">$\text{min_slope} \leq \text{slope}_{i,j,k} \leq \text{max_slope}, \quad 1 \leq i \leq n_{rt}, 1 \leq j \leq n_{nac}$</p> <p style="margin: 0;">$1 \leq k \leq n_{pt}$</p>	(MINLP-a-m)
---	-------------

The MINLP problem is solved using the branch-and-bound method [42], which formulates a series of continuous NLP problems that eventually generate an integer solution, and the continuous problem is shown below (NLP-a-m). The continuous NLP problems are solved using SNOPT [43], using a modified version of the pyOpt interface [44]. SNOPT is a gradient-based optimizer that implements the sequential quadratic programming (SQP) method, and it is specifically designed for nonlinear constrained optimization problems that are large-scale and sparse.

	maximize profit with respect to pax_flt _{<i>i,j</i>} ∈ [0, ac_capacity _{<i>j</i>}], flt_day _{<i>i,j</i>} ∈ [0, ∞), altitude _{<i>i,j,k</i>} ∈ [0, max_altitude],	$1 \leq i \leq n_{rt}, 1 \leq j \leq n_{ac}$ $1 \leq i \leq n_{rt}, 1 \leq j \leq n_{ac}$ $1 \leq i \leq n_{rt}, 1 \leq j \leq n_{nac}$ $1 \leq k \leq n_{cp}$ $1 \leq i \leq n_{rt}$ $1 \leq j \leq n_{ac}$ $1 \leq i \leq n_{rt}, 1 \leq j \leq n_{nac}$ $1 \leq i \leq n_{rt}, 1 \leq j \leq n_{nac}$ $1 \leq i \leq n_{rt}, 1 \leq j \leq n_{nac}$ $1 \leq i \leq n_{rt}, 1 \leq j \leq n_{nac}$ $1 \leq i \leq n_{rt}, 1 \leq j \leq n_{nac}$ $1 \leq k \leq n_{pt}$
	subject to total_pax _{<i>i</i>} ≤ demand _{<i>i</i>} , total_usage _{<i>j</i>} ≤ num_ac _{<i>j</i>} , altitude _{<i>i,j,1</i>} = 0, altitude _{<i>i,j,n_{cp}</i>} = 0, idle_throttle ≤ throttle _{<i>i,j</i>} (KS), throttle _{<i>i,j</i>} ≤ max_throttle(KS), min_slope ≤ slope _{<i>i,j,k</i>} ≤ max_slope,	(NLP-a-m)

The overall algorithm that solves (MINLP-a-m) is described in Algorithm 1. The first step is to solve (NLP-m) for each route and each aircraft type using SNOPT in order to determine the block times and block fuels to be used in the next step. We use the *linprog* function in *SciPy*'s *optimize* package via NASA's OpenMDAO framework [45] to solve (MILP-a) to generate a good initial solution for the branch-and-bound algorithm. Next, we begin the branch-and-bound phase by initializing the tree and adding the first node, which is (NLP-a-m) using the initial values for the allocation design variables from the solution to (MILP-a) and the initial values for the mission design variables from the solution to (NLP-m) for each mission. The remainder of the algorithm implements the depth-first branch-and-bound algorithm. An important characteristic of this algorithm is that when branching for a variable x_k , all of the initial values from the parent node are used, except for x_k itself, which is taken to be the same as the upper or lower bound. This allows (NLP-a-m) to be solved in one or a very small number of NLP iterations in many cases, yielding improved efficiency.

Algorithm 1 Solution of the allocation-mission optimization problem, (MINLP-a-m)

```

1: for  $i \leftarrow 1, n_{rt}$  do
2:   for  $j \leftarrow 1, n_{ac}$  do
3:     Solve (NLP-m) to determine block fuel and time for the  $(i, j)$ th mission
4:   end for
5: end for
6: Solve (MILP-a)
7: Create the branch and bound tree
8: Add the first node for an (NLP-a-m) starting from (MILP-a) and (NLP-m) solutions
9: while Nodes exist do
10:  Solve (NLP-a-m) in the deepest node (depth-first strategy)
11:  if Solution is better than the best integer solution then
12:    if Solution is integer then
13:      Store solution—this is the new solution
14:    else
15:       $k \leftarrow$  index of integer variable furthest from an integer
16:      Add a node for an (NLP-a-m) with upper bound and initial value of  $x_k$  changed to  $\text{floor}(x_k)$ 
17:      Add a node for an (NLP-a-m) with lower bound and initial value of  $x_k$  changed to  $\text{ceiling}(x_k)$ 
18:    end if
19:  end if
20:  Delete current node
21: end while

```

IV. Results

This section presents a suite of allocation-mission optimization results obtained using Algorithm 1. We start by describing the routes and the types of aircraft in the problem we solved, and present 4 results. First, we show the predicted increase in profit for an airline if it purchases new aircraft instead of existing aircraft. Second, we show that there is a difference between the results of allocation-only optimization (MILP-a) and allocation-mission optimization (MINLP-a-m). Next, we explore the presence of local optima in (NLP-a-m) and show that (MINLP-a-m) is highly sensitive to the starting point. Finally, we discuss the numerical performance of the allocation-mission analysis and optimization algorithm.

A. The problem

The results in this paper solve a 3-route problem with ranges of roughly 7000 nmi, 5500 nmi, and 2500 nmi. The routes have been chosen to represent a network with a hub in Newark, New Jersey and the following destinations: Hong Kong; Kuwait City, Kuwait; and Quito, Ecuador. The routes are summarized in Tab. 1 and shown graphically in Fig. 4.

Route	Newark, New Jersey (EWR) Hong Kong (HKG)	Newark, New Jersey (EWR) Kuwait City, Kuwait (KWI)	Newark, New Jersey (EWR) Quito, Ecuador (IQT)
Range [nmi]	6998	5546	2509
Demand	1200	550	700

Table 1: The 3 routes considered in the allocation-mission optimization



Figure 4: Map of the cities and the routes (generated using the Great Circle Mapper: www.gcmap.com). Newark, New Jersey is chosen as the hub.

We consider 6 aircraft types, 4 existing ones and 2 hypothetical next-generation aircraft. The Boeing 737-800 (B737), Boeing 777-200ER (B777), Boeing 747 (B747), and Boeing 787 (B787) have been chosen as the existing aircraft to cover a variety of design ranges and seating capacities. The two new aircraft are a notional advanced conventional design based on the Common Research Model [46] and a blended wing body concept based on Liebeck [47] with a reduced seating capacity. The aircraft types are summarized in Tab. 2 with seating capacities shown after applying an 80 % load factor. Table 2 also shows the 4 allocation-mission optimization problems we solved, representing

Aircraft	Boeing 737-800	Boeing 777-200ER	Boeing 747-400	Boeing 787-8	CRM: advanced conventional	BWB: blended wing body
Category	Existing	Existing	Existing	Existing	New	New
Capacity	122	207	294	200	300	400
Scenario						
S-base	20	24	24	8		
S-CRM	20	24	24		8	
S-BWB	20	24	24			8
S-both	20	20	20		8	8

Table 2: The types of aircraft considered in the allocation-mission optimization. The bottom four lines show the number of each aircraft type available in each of the four scenarios.

different scenarios in which the hypothetical airline chooses to buy different aircraft.

The cost, ticket price and the performance data of the existing aircraft for the different routes in the network are obtained using the simulation tool FLEET [29, 30]. For the new aircraft, the ticket price, no-fuel direct operating cost and the indirect operating cost are also obtained from an equivalent aircraft modeled in FLEET. The current model does not account for airline competition and assumes the ticket prices are fixed across types of aircraft on a given route.

B. Profit increase with new aircraft

Figure 5 shows the profit after optimization for each of the 4 scenarios. The results agree with intuition because the CRM and BWB both represent an improvement over the existing aircraft. The CRM design is based on the B777, but it is assumed to have a larger seating capacity and higher aerodynamic efficiency since it is a next generation aircraft. The B787 has the range of the CRM but a lower seating capacity, while the B747 has the seating capacity but a lower range and lower efficiency as it is a much older design. Thus, the S-CRM scenario provides a 192 % improvement in profit over the baseline S-base, the S-BWB scenario provides a 323 % due to its larger seating capacity, and the S-both scenario provides a further improvement with 414 % compared to the baseline.

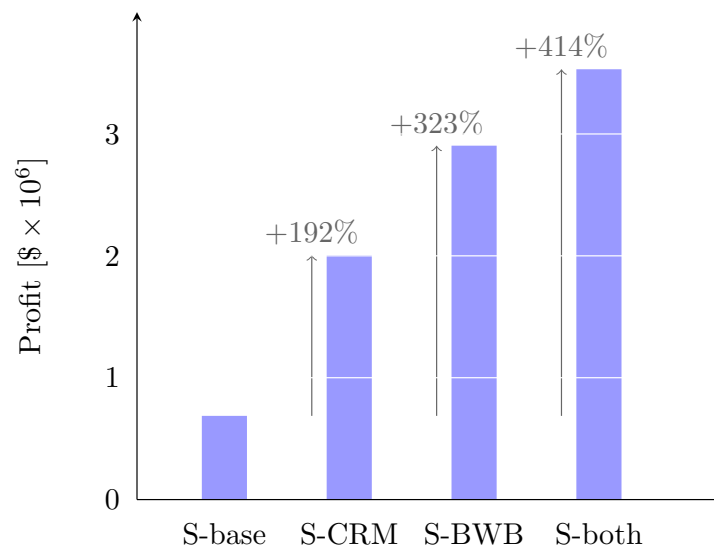


Figure 5: Comparison of profit for the four scenarios at the solution to (MINLP-a-m). The results show an increase in profit when the hypothetical airline purchases the next-generation aircraft.

C. Allocation versus allocation-mission optimization

Figure 6 compares the optimized profit among 4 cases: (MINLP-a-m) without the (MILP-a) initialization, (MILP-a), (MINLP-a-m), and (NLP-a-m). This subsection focuses on (MILP-a) versus (MINLP-a-m)—i.e., allocation-only optimization versus allocation-mission optimization. We see a small increase in profit in all three scenarios with the allocation-mission optimization because it more accurately computes the block fuel and block time for the new aircraft. In particular, the block fuel values used in (MILP-a) assume the aircraft always fly at full capacity—thus, (MILP-a) likely underpredicts profit because the actual fuel burn is lower when the aircraft flies with fewer passengers than its seating capacity. However, this is a minor difference, and we remind the reader that the true motivation for incorporating mission analysis in the allocation problem is to enable design-mission-allocation optimization, and this paper develops many of the methods needed for optimization with all three simultaneously considered.

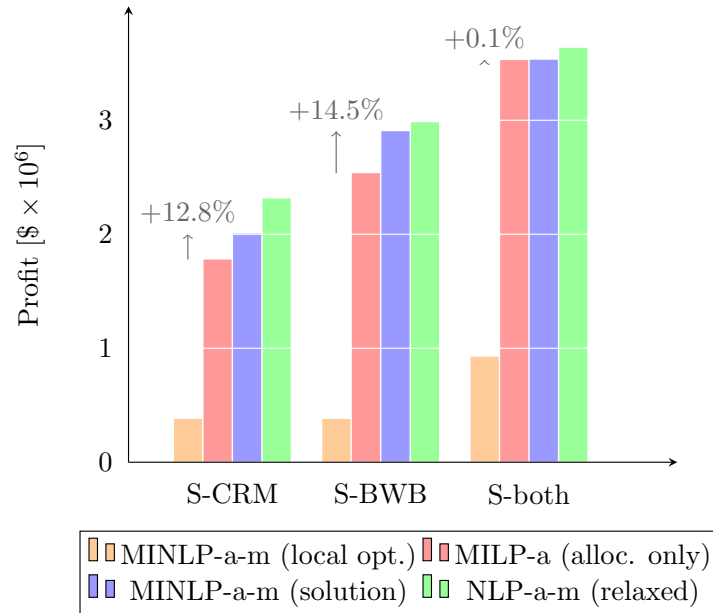


Figure 6: Comparison of optimized profits. The colors represent, in order: (MINLP-a-m) not initialized to the solution to (MILP-a) (orange), (MILP-a) (red), (MINLP-a-m) initialized to the solution to (MILP-a) (blue), and (NLP-a-m) solved in the first node of the branch-and-bound algorithm. The results show clear local minima and an increase in profit in the allocation-mission optimization compared to the allocation-only optimization.

D. Local optima

One significant concern when incorporating mission analysis into the allocation problem is that the problem becomes nonlinear, so there is potential for local optima that represent significantly worse profit than the global optimum. Figure 6 shows that there is indeed a large difference between two solutions to (MINLP-a-m): one in which Algorithm 1 is run as is and one in which Algorithm 1 is run without the (MILP-a) initialization (line 6). In Fig. 6, the orange bars are several times smaller than the blue bars for each scenario. Figure 7 confirms that the integer design variables—flights per day—consistently show large changes between the two optima.

These results highlight two aspects to the importance of the (MILP-a) initialization step (line 6) in Algorithm 1. First, it is not possible to guarantee convergence to a global minimum in nonlinear optimization problems that do not satisfy any convexity conditions, but Algorithm 1 provides a reasonably good local optimum by starting from the global optimum of the linear problem, (MILP-a). The initial (NLP-a-m) and the branch-and-bound method then guarantee that the solution to which it converges is no worse than the initial point provided by the linear optimization.

The second benefit is the tremendous efficiency afforded by this two-step approach. The linear optimization (MILP-a) is inexpensive, so it quickly gets close to the solution of the nonlinear optimization (MINLP-a-m), after which branch-and-bound begins from a good starting point. Fig. 8 shows that the continuous solution of the initial (NLP-a-m) solved in the first node of the branch-and-bound is not significantly different from the solution of (MILP-a),

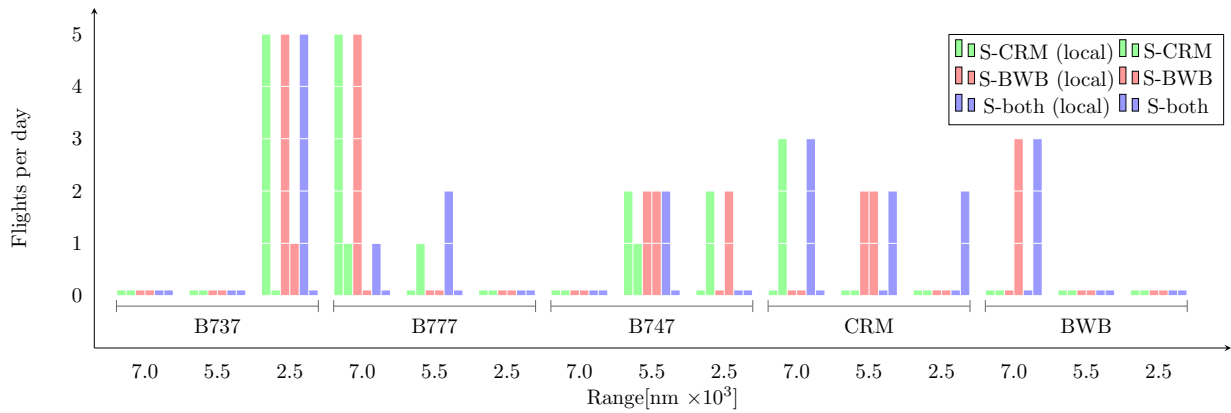


Figure 7: Optimal values of the (round trip) flights per day design variable comparing (MINLP-a-m) with and without (local) initialization to the solution to (MILP-a). The results show two distinct local minima in (MINLP-a-m).

and the solution of (MINLP-a-m) is in turn very similar as well. In fact, the branch-and-bound only forces the integer variables to one of the nearest integers, suggesting that the (NLP-a-m) problem is locally convex in a large enough area containing the nearby integers, but multi-modal over a larger region since at least one other local optimum was found.

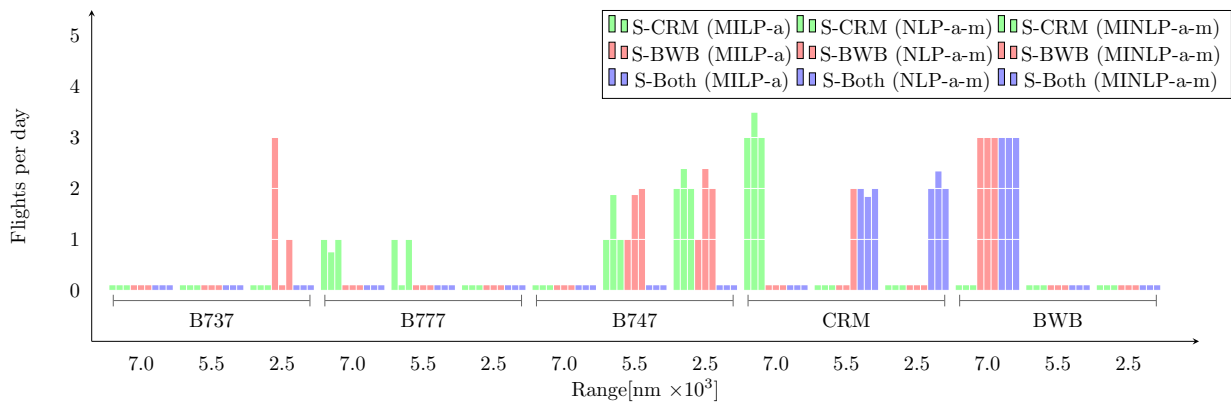


Figure 8: Optimal values of the (round trip) flights per day design variable comparing the allocation-only solution (MILP-a), the relaxed allocation-mission solution (NLP-a-m), and the integer allocation-mission solution (MINLP-a-m). The results show that the relaxed simultaneous solution is close to the allocation-only solution and the integer simultaneous solution is close to the relaxed simultaneous solution.

This observation is significant from a computational standpoint because the branch-and-bound method converges in a small number of iterations. For this 3-route problem, all cases converged in 5-20 node evaluations, and initial studies of 5-route problems show the same results. Furthermore, many of the (NLP-a-m) problems solved during branch-and-bound converge in one iteration or a small number of iterations because the initial point is often the solution.

E. Numerical performance

Figure 9 shows the scaling of an analysis for the (NLP-a-m) formulation and an adjoint solution. It shows the execution time for problems with between 1 and 16 routes. The scaling is between linear and quadratic, but in future studies, the mission analyses can be easily parallelized since there is no coupling between mission analyses for different routes and aircraft.

Figure 10 shows the convergence history for the solution of (NLP-a-m) with 3 routes and 2 new aircraft, resulting

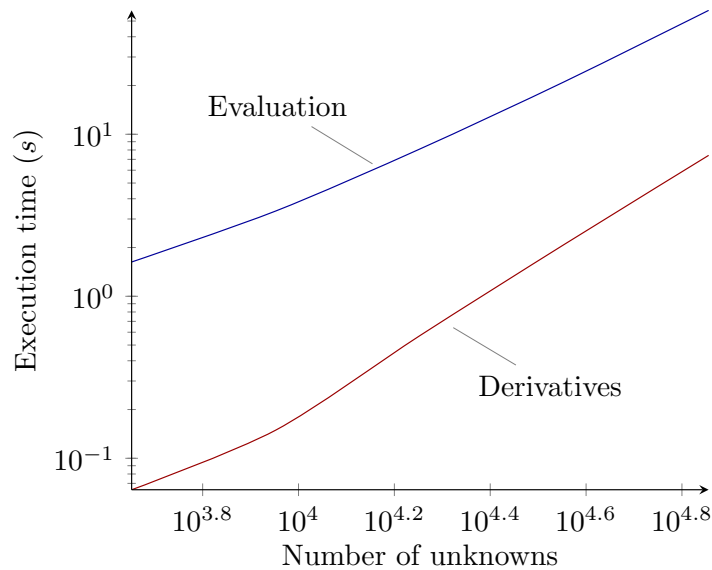


Figure 9: Execution time versus number of total unknowns in the simulation for an analysis and an adjoint solution with problems ranging from 1 route to 16 routes. The scaling is between linear and quadratic.

in 330 design variables in total. The fact that we can achieve such a high level of convergence provides confidence that the derivatives are computed accurately and the design variables and constraints are properly scaled. Table 3 shows the times for solving (MINLP-a-m) for the various cases.

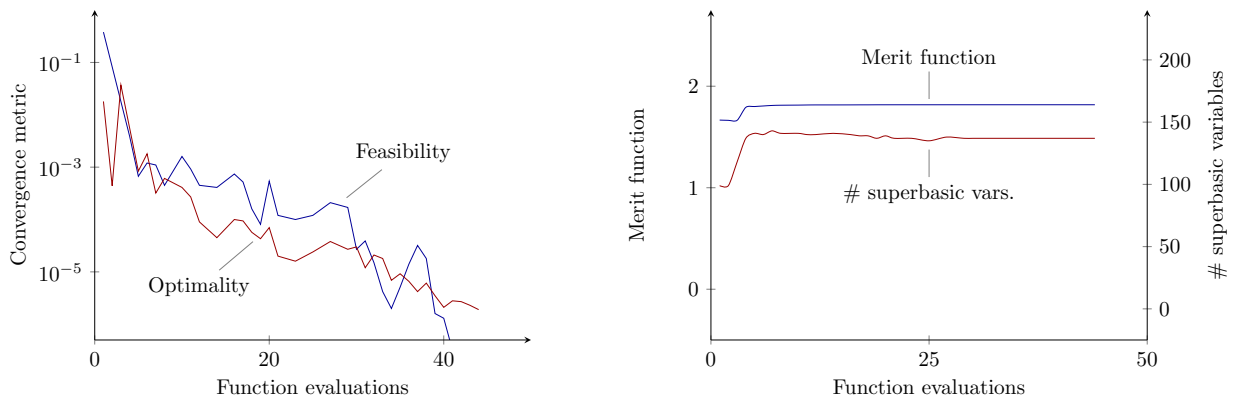


Figure 10: Optimization convergence history for a 3-route, 2-new aircraft problem with 330 design variables. The tight convergence of optimality and feasibility helps verify the accuracy of the derivative computation.

V. Conclusion

In this paper, we developed an algorithm for performing simultaneous allocation-mission optimization. This required four components: the use of a computational framework for automatically solving the coupled simulation and computing derivatives using the adjoint method, an efficient mission analysis algorithm that uses surrogates for aerodynamics and propulsion, implementation of the allocation problem within the framework that couples an arbitrary number of mission analyses together, and a method for solving the mixed-integer allocation-mission optimization problem. The key step in this method is the initialization of the branch-and-bound allocation-mission optimization using the solution of the allocation-only optimization. We applied this algorithm to the allocation of 4 existing aircraft

Case	Time [min.]
S-CRM (local)	7.7
S-CRM	20.4
S-BWB (local)	9.0
S-BWB	20.5
S-both (local)	33.3
S-both	40.6

Table 3: The execution times for the solution of (MINLP-a-m).

and 2 next-generation aircraft to a 3-route network and presented three conclusions.

First, the algorithm found significant profit increases if the airline chooses to purchase the next-generation aircraft. Compared to a baseline, purchasing 8 advanced conventional aircraft yielded a 192 % increase in profit, purchasing 8 blended wing body aircraft yielded a 323 % increase in profit, and purchasing 8 of each yielded a 414 % increase in profit with the optimal allocation.

Second, the allocation-mission optimization produced a higher profit than allocation-only optimization. Part of the difference is attributed to the fact that block fuel is computed dynamically during optimization using the allocated number of passengers, rather than pre-computing block fuel assuming full capacity. However, this error is minor, and the true benefit of simultaneously considering mission during the allocation problem is that this enables also incorporating design in the simultaneous optimization in future work.

Third, we found multiple local optima in the mixed-integer allocation-mission problem, and they yielded profit numbers that are 2-3 times apart. However, we showed that our algorithm efficiently finds a good local optimum by using an allocation-only optimization to find a good starting point. This is advantageous because linear programming has a low computational cost and it finds the global optimum. In practice, the solution to the first relaxed allocation-mission optimization during branch-and-bound deviates a small amount from the linear solution and the branch-and-bound method, in turn, finds one of the neighboring integer points as the solution. The significance of this is that many nodes in the branch-and-bound algorithm converge very quickly during the continuous optimization.

The significance of this work is the demonstration of an efficient algorithm for solving simultaneous allocation-mission optimization problems using an approach that naturally extends to design-mission-allocation optimization. Therefore, the next step in this research effort is adding design to the simultaneous optimization problem. As part of this step, an important avenue for future work is testing the current algorithm with the mission analyses parallelized for improved scalability.

VI. Acknowledgments

This work was supported by NASA through grant number NNX14AC73A. The authors would like to thank Justin S. Gray from the NASA Glenn Research Center for his help implementing the linear allocation problem in OpenMDAO and Rhea P. Liem from the University of Toronto for providing the aerodynamic surrogate models for the CRM and BWB.

References

- [1] Rutherford, D. and Zeinali, M., "Efficiency Trends for New Commercial Jet Aircraft 1960 to 2008," The International Council on Clean Transportation, 2009.
- [2] ICAO Secretariat, "Aircraft technology improvements," *ICAO Environmental Report 2010*, International Civil Aviation Organization, 2010.
- [3] Kennedy, G. J. and Martins, J. R. R. A., "A Comparison of Metallic and Composite Aircraft Wings Using Aerostructural Design Optimization," *14th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*, Indianapolis, IN, Sep. 2012, AIAA-2012-5475.
- [4] Kennedy, G. J., Kenway, G. K., and Martins, J. R. R. A., "A Comparison of Metallic, Composite and Nanocomposite Optimal Transonic Transport Wings," Tech. rep., NASA, March 2014, CR-2014-218185.
- [5] Lyu, Z. and Martins, J. R. R. A., "Aerodynamic Design Optimization Studies of a Blended-Wing-Body Aircraft," *Journal of Aircraft*, Vol. 51, No. 5, September 2014, pp. 1604–1617. doi:[10.2514/1.C032491](https://doi.org/10.2514/1.C032491).
- [6] Kenway, G. K. W. and Martins, J. R. R. A., "Multipoint High-Fidelity Aerostructural Optimization of a Transport Aircraft Configuration," *Journal of Aircraft*, Vol. 51, No. 1, January 2014, pp. 144–160. doi:[10.2514/1.C032150](https://doi.org/10.2514/1.C032150).
- [7] Lyu, Z., Kenway, G. K., and Martins, J. R. R. A., "Aerodynamic Shape Optimization Studies on the Common Research Model Wing Benchmark," *AIAA Journal*, 2014. doi:[10.2514/1.J053318](https://doi.org/10.2514/1.J053318), (In press).
- [8] Cramer, E. J., Dennis, J. E., Frank, P. D., Lewis, R. M., and Shubin, G. R., "Problem Formulation for Multidisciplinary Optimization," *SIAM Journal on Optimization*, Vol. 4, No. 4, 1994, pp. 754–776.
- [9] Martins, J. R. R. A. and Lambe, A. B., "Multidisciplinary Design Optimization: A Survey of Architectures," *AIAA Journal*, Vol. 51, No. 9, September 2013, pp. 2049–2075. doi:[10.2514/1.J051895](https://doi.org/10.2514/1.J051895).
- [10] Liem, R., Kenway, G. K. W., and Martins, J. R. R. A., "Multimission Aircraft Fuel Burn Minimization via Multipoint Aerostructural Optimization," *AIAA Journal*, 2014. doi:[10.2514/1.J052940](https://doi.org/10.2514/1.J052940), (In press).
- [11] Hwang, J. T. and Martins, J. R. R. A., "A parallel hierarchical algorithmic framework for large-scale simulation and optimization," University of Michigan, 2015.
- [12] Hwang, J. T., Lee, D. Y., Cutler, J. W., and Martins, J. R. R. A., "Large-Scale Multidisciplinary Optimization of a Small Satellite's Design and Operation," *Journal of Spacecraft and Rockets*, Vol. 51, No. 5, September 2014, pp. 1648–1663. doi:[10.2514/1.A32751](https://doi.org/10.2514/1.A32751).
- [13] Lambe, A. B. and Martins, J. R. R. A., "Extensions to the Design Structure Matrix for the Description of Multidisciplinary Design, Analysis, and Optimization Processes," *Structural and Multidisciplinary Optimization*, Vol. 46, August 2012, pp. 273–284. doi:[10.1007/s00158-012-0763-y](https://doi.org/10.1007/s00158-012-0763-y).
- [14] Ferguson, A. R. and Dantzig, G. B., "The allocation of aircraft to routes—an example of linear programming under uncertain demand," *Management science*, Vol. 3, No. 1, 1956, pp. 45–73.
- [15] Gass, S., "George B. Dantzig," *Profiles in Operations Research*, edited by A. A. Assad and S. I. Gass, Vol. 147 of *International Series in Operations Research & Management Science*, Springer US, 2011, pp. 217–240. doi:[10.1007/978-1-4419-6281-2_13](https://doi.org/10.1007/978-1-4419-6281-2_13).
- [16] Simpson, R. W., "Computerized schedule construction for an airline transportation system," *Department of Aeronautics and Astronautics, Massachusetts Institute of Technology*, Vol. 232, 1966.
- [17] Abara, J., "Applying integer linear programming to the fleet assignment problem," *Interfaces*, Vol. 19, No. 4, 1989, pp. 20–28.
- [18] Barnhart, C., Farahat, A., and Lohatepanont, M., "Airline fleet assignment with enhanced revenue modeling," *Operations research*, Vol. 57, No. 1, 2009, pp. 231–244.
- [19] Dumas, J., Aithnard, F., and Soumis, F., "Improving the objective function of the fleet assignment problem," *Transportation Research Part B: Methodological*, Vol. 43, No. 4, 2009, pp. 466–475.
- [20] Johnson, E. L., "Robust airline fleet assignment: imposing station purity using station decomposition," *Algorithmic Applications in Management*, Springer, 2005, pp. 1–2.
- [21] Bélanger, N., Desaulniers, G., Soumis, F., Desrosiers, J., and Lavigne, J., "Weekly airline fleet assignment with homogeneity," *Transportation Research Part B: Methodological*, Vol. 40, No. 4, 2006, pp. 306–318.
- [22] Lohatepanont, M. and Barnhart, C., "Airline schedule planning: Integrated models and algorithms for schedule design and fleet assignment," *Transportation Science*, Vol. 38, No. 1, 2004, pp. 19–32.
- [23] Clarke, M. and Smith, B., "Impact of operations research on the evolution of the airline industry," *Journal of Aircraft*, Vol. 41, No. 1, 2004, pp. 62–72.
- [24] Barnhart, C., Belobaba, P., and Odoni, A. R., "Applications of operations research in the air transport industry," *Transportation science*, Vol. 37, No. 4, 2003, pp. 368–391.
- [25] Subramanian, R., Scheff, R. P., Quillinan, J. D., Wiper, D. S., and Marsten, R. E., "Coldstart: fleet assignment at delta air lines," *Interfaces*, Vol. 24, No. 1, 1994, pp. 104–120.

- [26] Kontogiorgis, S. and Acharya, S., "US Airways automates its weekend fleet assignment," *Interfaces*, Vol. 29, No. 3, 1999, pp. 52–62.
- [27] Zhao, J., Tetzloff, I. J., Tyagi, A., Dikshit, P., Mane, M., Agusdinata, D., Crossley, W. A., and DeLaurentis, D., "Assessing New Aircraft and Technology Impacts on Fleet-Wide Environmental Metrics including Future Scenarios," *48th AIAA Aerospace Sciences Meeting Including the New Horizons Forum and Aerospace Exposition, Orlando, Florida*, 2010.
- [28] Mane, M., Tetzloff, I., Zhao, J., Noth, K., Agusdinata, D., DeLaurentis, D., and Crossley, W., "Fleet-wide environmental impact of future aircraft technologies," *10th AIAA Aviation Technology, Integration, and Operations Conference (ATIO)*, 2010.
- [29] Crossley, W. and DeLaurentis, D., "System-of-systems approach for assessing new technologies in NGATS," *Purdue University*, 2007.
- [30] Govindaraju, P. and Crossley, W. A., "Profit Motivated Airline Fleet Allocation and Concurrent Aircraft Design for Multiple Airlines," *Aviation Technology, Integration, and Operations Conference*, Aug 2013, AIAA 2013-4391.
- [31] Betts, J. T., "Survey of numerical methods for trajectory optimization," *Journal of guidance, control, and dynamics*, Vol. 21, No. 2, 1998, pp. 193–207.
- [32] Betts, J. T. and Huffman, W. P., "Application of sparse nonlinear programming to trajectory optimization," *Journal of Guidance, Control, and Dynamics*, Vol. 15, No. 1, 1992, pp. 198–206.
- [33] Betts, J. T. and Cramer, E. J., "Application of direct transcription to commercial aircraft trajectory optimization," *Journal of Guidance, Control, and Dynamics*, Vol. 18, No. 1, 1995, pp. 151–159.
- [34] Park, S. G. and Clarke, J.-P., "Vertical Trajectory Optimization for Continuous Descent Arrival Procedure," *AIAA Guidance, Navigation, and Control Conference, Minneapolis (Minnesota-USA), Paper*, Vol. 4757, 2012, p. 2012.
- [35] Clarke, J.-P. B., Ho, N. T., Ren, L., Brown, J. A., Elmer, K. R., Zou, K., Hunting, C., McGregor, D. L., Shivashankara, B. N., Tong, K.-O., et al., "Continuous descent approach: Design and flight test for Louisville International Airport," *Journal of Aircraft*, Vol. 41, No. 5, 2004, pp. 1054–1066.
- [36] Shresta, S., Neskovic, D., and Williams, S., "Analysis of continuous descent benefits and impacts during daytime operations," *8th USA/Europe Air Traffic Management Research and Development Seminar (ATM2009), Napa, CA*, 2009.
- [37] Clarke, J.-P., Brooks, J., Nagle, G., Scacchioli, A., White, W., and Liu, S., "Optimized Profile Descent Arrivals at Los Angeles International Airport," *Journal of Aircraft*, Vol. 50, No. 2, 2013, pp. 360–369.
- [38] Micallef, M., Chircop, K., Zammit-Mangion, D., and Sammut, A., "Revised approach procedures to support optimal descents into Malta International Airport," *CEAS Aeronautical Journal*, 2014, pp. 1–15.
- [39] Martins, J. R. R. A. and Hwang, J. T., "Review and Unification of Methods for Computing Derivatives of Multidisciplinary Computational Models," *AIAA Journal*, Vol. 51, No. 11, November 2013, pp. 2582–2599. doi:[10.2514/1.J052184](https://doi.org/10.2514/1.J052184).
- [40] Balay, S., Gropp, W. D., McInnes, L. C., and Smith, B. F., "Efficient Management of Parallelism in Object Oriented Numerical Software Libraries," *Modern Software Tools in Scientific Computing*, edited by E. Arge, A. M. Bruaset, and H. P. Langtangen, Birkhäuser Press, 1997, pp. 163–202.
- [41] Kreisselmeier, G. and Steinhauser, R., "Systematic Control Design by Optimizing a Vector Performance Index," *International Federation of Active Controls Symposium on Computer-Aided Design of Control Systems, Zurich, Switzerland*, 1979.
- [42] Bertsimas, D. and Tsitsiklis, J. N., "Introduction to linear optimization," 1997.
- [43] Gill, P. E., Murray, W., and Saunders, M. A., "SNOPT: An SQP algorithm for large-scale constrained optimization," *SIAM Journal of Optimization*, Vol. 12, No. 4, 2002, pp. 979–1006. doi:[10.1137/S1052623499350013](https://doi.org/10.1137/S1052623499350013).
- [44] Perez, R. E., Jansen, P. W., and Martins, J. R. R. A., "pyOpt: a Python-Based Object-Oriented Framework for Nonlinear Constrained Optimization," *Structural and Multidisciplinary Optimization*, Vol. 45, No. 1, January 2012, pp. 101–118. doi:[10.1007/s00158-011-0666-3](https://doi.org/10.1007/s00158-011-0666-3).
- [45] Gray, J., Moore, K. T., Hearn, T. A., and Naylor, B. A., "Standard Platform for Benchmarking Multidisciplinary Design Analysis and Optimization Architectures," *AIAA Journal*, Vol. 51, No. 10, 2013, pp. 2380–2394. doi:[10.2514/1.J052160](https://doi.org/10.2514/1.J052160).
- [46] Vassberg, J. C., DeHaan, M. A., Rivers, S. M., and Wahls, R. A., "Development of a Common Research Model for Applied CFD Validation Studies," 2008, AIAA 2008-6919.
- [47] Liebeck, R. H., "Design of the Blended Wing Body Subsonic Transport," *Journal of Aircraft*, Vol. 41, No. 1, 2004.