

SYM-AM-16-054



PROCEEDINGS OF THE THIRTEENTH ANNUAL ACQUISITION RESEARCH SYMPOSIUM

THURSDAY SESSIONS VOLUME II

An Optimization-Based Approach to Determine System Requirements Under Multiple Domain-Specific Uncertainties

Parithi Govindaraju, Graduate Research Assistant, Purdue University
Navindran Davendralingam, Research Scientist, Purdue University
William Crossley, Professor, Purdue University

Published April 30, 2016

Approved for public release; distribution is unlimited.

Prepared for the Naval Postgraduate School, Monterey, CA 93943.



ACQUISITION RESEARCH PROGRAM
GRADUATE SCHOOL OF BUSINESS & PUBLIC POLICY
NAVAL POSTGRADUATE SCHOOL

The research presented in this report was supported by the Acquisition Research Program of the Graduate School of Business & Public Policy at the Naval Postgraduate School.

To request defense acquisition research, to become a research sponsor, or to print additional copies of reports, please contact any of the staff listed on the Acquisition Research Program website (www.acquisitionresearch.net).



ACQUISITION RESEARCH PROGRAM
GRADUATE SCHOOL OF BUSINESS & PUBLIC POLICY
NAVAL POSTGRADUATE SCHOOL

Panel 13. Setting Requirements and Managing Risk in Complex, Networked Projects

Thursday, May 5, 2016	
9:30 a.m. – 11:00 a.m.	<p>Joseph Yakovac, Lieutenant General, USA (Ret.)– Former Principal Military Deputy, Assistant Secretary of the Army for Acquisition, Logistics and Technology</p> <p><i>Acquisition in a World of Joint Capabilities: Methods for Understanding Cross-Organizational Network Performance</i> Mary Brown, Professor, UNCC Zachary Mohr, Assistant Professor, UNCC</p> <p><i>Modeling Uncertainty and Its Implications in Complex Interdependent Networks</i> Anita Raja, Professor, The Cooper Union Mohammad Rashedul Hasan, Assistant Professor of Practice, UNL Robert Flowe, Office of Acquisition Resources & Analysis, OUSD (AT&L) Brendan Fernes, Student, The Cooper Union</p> <p><i>An Optimization-Based Approach to Determine System Requirements Under Multiple Domain-Specific Uncertainties</i> Parithi Govindaraju, Graduate Research Assistant, Purdue University Navindran Davendralingam, Research Scientist, Purdue University William Crossley, Professor, Purdue University</p>



An Optimization-Based Approach to Determine System Requirements Under Multiple Domain-Specific Uncertainties

Parithi Govindaraju—is a PhD student of Dr. William Crossley in the School of Aeronautics and Astronautics at Purdue University specializing in the field of aerospace systems. His research interests center on design optimization and systems engineering. [pgovinda@purdue.edu]

Navindran Davendralingam—is a Research Scientist in the School of Aeronautics and Astronautics at Purdue University. He received his PhD from Purdue University in aerospace engineering in 2011. He is currently conducting research in the Center for Integrated Systems in Aerospace (CISA) led by Dr. Daniel DeLaurentis. [davendra@purdue.edu]

William A. Crossley—is a Professor in the School of Aeronautics and Astronautics, Purdue University, where he has been on the faculty since 1995. His research and teaching interests focus on design and optimization of aerospace vehicles and systems of systems. [crossley@purdue.edu]

Abstract

The task of determining the optimal design requirements of a new system, which will operate along with other existing systems to provide a set of overarching capabilities, is challenging due to the tightly coupled effects that setting requirements on a system's design can have on how the operator uses the system. In this paper, the new system is a strategic military cargo aircraft and the other systems are a fleet of different, existing cargo aircraft; a subset of actual fleet operations from the U.S. Air Force Air Mobility Command defines the example problems in this work. This research builds upon prior efforts to develop a quantitative approach that identifies optimum design requirements of new, yet-to-be-designed systems that, when serving alongside other systems, will optimize fleet-level objectives. The new efforts here address the effect of various uncertainties. The approach incorporates techniques from multidisciplinary design optimization, statistical theory, and robust/reliability-based methods to develop computationally tractable approaches for this kind of problem. The paper also demonstrates the ability to generate tradeoffs between a cost-related metric of fleet-level fuel usage and a performance related metric of fleet-wide productivity. A possible extension for application in commercial air travel also appears in the paper.



Introduction

Nomenclature

AR_x	= aspect ratio of aircraft type X
B_p	= maximum average daily utilization of each aircraft (hours)
BPR_x	= engine bypass ratio of aircraft type X
$BH_{p,k,i,j}$	= number of block hours for k^{th} trip of aircraft p from base i to base j
$\bar{B}H_{p,k,i,j}$	= distribution of block hours for k^{th} trip of aircraft p from base i to base j
$BH^U_{p,k,i,j}$	= upper bound of block hours for k^{th} trip of aircraft p from base i to base j
$Cap_{p,k,i,j}$	= pallet carrying capacity for k^{th} trip of aircraft p from base i to base j
C_{D_0}	= parasite drag coefficient
$dem_{i,j}$	= demand from base i to base j in number of pallets
$dem^U_{i,j}$	= upper bound of demand from base i to base j in number of pallets
$E[]$	= expectation (arithmetic mean) of a distribution []
$FC_{p,k,i,j}$	= fuel consumption coefficient for k^{th} trip of aircraft p from base i to base j
$\bar{F}C_{p,k,i,j}$	= distribution of the fuel consumption coefficient for k^{th} trip of aircraft p from base i to base j
$FC^U_{p,k,i,j}$	= upper bound of fuel consumption coefficient for k^{th} trip of aircraft p from base i to base j
L	= fleet-level productivity limit (knots-lbs)
$O_{p,i}$	= indicates if airport i is the initial location (e.g., home base) of an aircraft p
$Pallet_x$	= number of pallets carried by aircraft type X
$P[]$	= probability of satisfying expression []
$Prod_{p,k,i,j}$	= productivity coefficient for k^{th} trip of aircraft p from base i to base j
$Rangex$	= design range of aircraft type X (nmi)
SFC	= specific fuel consumption (1/hr)
$Speed_x$	= cruise speed of aircraft type X (knots)
$Sweep_x$	= leading edge wing sweep of aircraft type X (deg)
TR_x	= taper ratio of aircraft type X
$(TW)_x$	= thrust-to-weight ratio of aircraft type X
$(W/S)_x$	= wing loading of aircraft type X (lb/ft ²)
$x_{p,k,i,j}$	= binary variable for k^{th} trip flown by aircraft p from base i to base j

Research Issue

The *Better Buying Power 3.0* (Office of the Under Secretary of Defense for Acquisition, Technology, and Logistics [OUSD(AT&L)], 2014) document states, “Defining requirements well is a challenging but essential prerequisite in achieving desired service acquisition outcomes.” Typical acquisition processes focus on development at the *system-level* (e.g., aircraft performance), with little explicit consideration for the impact that the new system will have on the holistic performance of a combined set of existing and new systems. Current acquisition processes (how a decision-maker evaluates and acquires systems) are disjointed from considering operations (the way an end user operates these new systems alongside existing ones), resulting in inefficiencies at the higher aggregate level (Taylor & de Weck, 2007; Mane, Crossley, & Nusawardhana, 2007). As an example, consider the acquisition decision-making process within the Department of Defense (DoD) that traditionally involves identification of alternatives, establishment of requirements, estimation of effectiveness, and cost-benefit analyses (Greer, 2010). These action processes do not involve an exhaustive search of the “requirements space” of the new system, where changes in requirements can affect operations due to how the new, yet-to-be-introduced system will be used in conjunction with other existing systems in the fleet.



This research effort seeks to reduce such “handoffs” between the acquisition phase and the operations phase through leveraging quantitative innovations that reduce such handoffs. However, this coupling of the requirements of new systems, and the resulting system’s impact on operations, brings an added dimension of complexity to the acquisition problem. The complexity of dealing with many variables related to interdependent systems, the impact of changing characteristics of such systems, and the uncertainties related to allocations of such systems becomes cognitively impossible to manage without a decision-support framework. Hence, determining the optimal set of requirements for a new, yet-to-be-designed system presents a need for analytical tools to assist decision-makers with quantitatively supported insights.

This paper presents the methodology and formulations of a quantitative approach that identifies the design requirements for a new aircraft under multiple domain-specific uncertainties through an optimization approach. The paper illustrates the approach via examples derived from reported operations of the U.S. Air Force Air Mobility Command (AMC). The approach treats design requirements of new individual systems as decision variables in an optimization problem formulation under various uncertainties to minimize (or maximize) fleet-level objectives—the solution, based on mathematical techniques, identifies the *new aircraft requirement decision variables* that yield the best *fleet-level objectives*.

Two different types of uncertainty are important to this problem: (1) uncertainty in how a designed system “actually” performs in operations as to opposed the predicted performance in the design phase, and, (2) the variations in how much the operator uses the system, as reflected by, say, changing demand for air transportation of military cargo pallets. The uncertainties in how the designed system performs naturally affects the uncertainties in how much the system is being used. Simultaneously considering the system design problem and resource allocation problem under uncertainty captures most of the coupling and interactions present in these two problems, and capturing these can result in fleet-level improvements. Often, high computational expense accompanies quantifying and addressing uncertainty in multiple dimensions, which can make the design problem intractable. Effectively conducting studies that examine several scenarios using different predictions of demand, cost of operating the fleet, and so forth, requires a computationally efficient approach. The authors’ initial efforts to identify an effective approach explored two different strategies—design of experiments and bounding analysis—to understand the effects of considering both demand and design parameter uncertainties in the coupled aircraft design and fleet assignment problem (Govindaraju, Davendralingam, & Crossley, 2015; Govindaraju & Crossley, 2015).

Furthermore, this paper demonstrates an extension of the approach to consider multiple objectives, thereby enabling the assessment of tradeoffs that choices about design requirements may have on fleet-level metrics of interest (e.g., choice of an aircraft may affect fleet-level *productivity* and *fuel burn*—quantities that are at odds with one another). This can allow decision-makers to view this problem in the context of fleet-level fuel consumption as an independent variable. Two key innovations in this approach are that it

1. Considers the holistic implications that setting design requirements may have on the fleet-level metrics.
2. Relegates the mathematical complexities of considering the design requirements, operations of the fleet and manifestations of uncertainties to sound algorithmic approaches, while retaining exploratory and decision-making elements to the practitioner.



Modeling Military Cargo Air Transportation

The transportation of cargo across the AMC service network requires effective deployment of its fleet of cargo aircraft to meet daily cargo delivery requirements while minimizing fuel consumption and related costs. The choice of which aircraft to operate on individual flight legs to meet the cargo delivery obligations within a scheduled time frame determines the total amount of fuel consumed by the AMC fleet. Fleet-wide fuel consumption is tied to the features of aircraft used and the structure of the routes flown. However, the characteristics of the aircraft (e.g., range of the aircraft) also dictate the kind of network that the fleet can serve, thus making it a closely coupled problem. Because of this, there may be an opportunity to identify design requirements for a new aircraft that can reduce the total fleet fuel consumption and/or improve fleet-level cargo delivery performance. This work extends a deterministic decomposition approach (Mane et al., 2007) to allow for the examination of tradeoffs between objectives of productivity (as a measure of mission effectiveness) and fuel consumption when considering the addition of a new, yet-to-be-acquired aircraft to a fleet of existing aircraft under various domain-specific uncertainties. These two competing objectives of productivity and fuel consumption (maximizing productivity increases fuel consumption and minimizing fuel consumption decreases productivity) play a critical role in determining new system requirements—an analyst can perform acquisition assessments by treating fuel consumption as an independent variable in our approach.

Cargo Demand in the AMC Service Network

The AMC service/demand network differs from commercial airline passenger or cargo networks in that cargo demand fluctuates greatly over time and in that cargo demand is asymmetrical, meaning that the demand for cargo from one base to another is usually very different than the demand in the opposite direction between the same bases. Figure 1a shows the fluctuation in the number of pallets transported daily between a representative base pair in the Global Air Transportation Execution System (GATES) dataset for the year 2006. In this plot, the calendar day appears on the horizontal axis, while the heights of the bars indicate the number of pallets transported each day in one direction. Figure 1b presents a histogram of the number of pallets transported per day for the same representative base pair; this reveals that many days had a demand of 20 or fewer pallets on this route. Twenty pallets might be well below the maximum capacity of a single aircraft used to transport this demand. The AMC fleet must have the flexibility to meet fluctuating demand—the comparatively rare, high-demand scenarios, and the typical, nominal demand scenarios—to address fuel efficiency effectively.



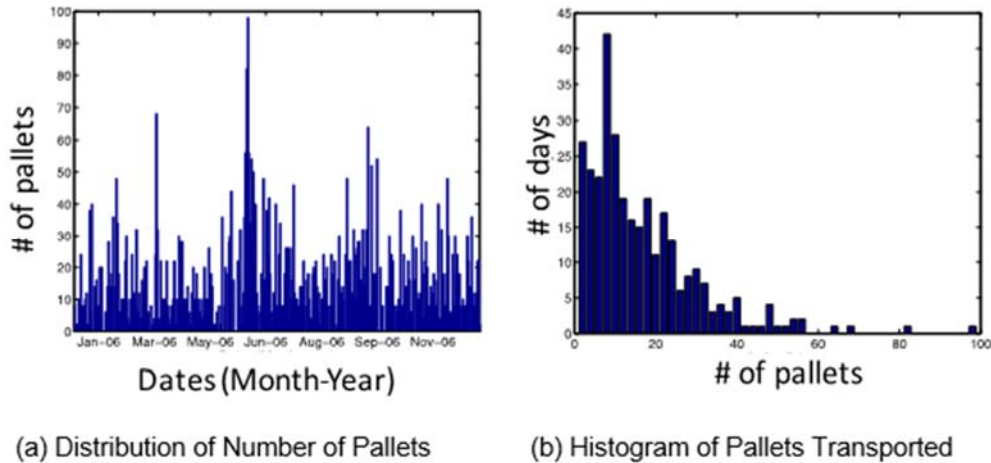


Figure 1. Pallets Transported on a Sample Route From the GATES Dataset

GATES Dataset

The AMC fleet operates on a global network consisting of over 350 bases and in excess of 1750 routes. The GATES dataset provides historical route and cargo demand data, and it contains comprehensive information on palletized cargo and personnel transported by the AMC fleet. From the GATES dataset for 2006, the existing AMC fleet to serve the demand consisted of 92 C-5s, 145 C-17s, and 69 747-Fs. This shows that the AMC transported cargo using C-5 and C-17 aircraft from the strategic fleet and using chartered Boeing 747 Freighter (747-F) aircraft from the Civil Reserve Air Fleet (CRAF) for long range missions. The 2006 GATES data provides a representative cargo flow in the AMC service network and the aircraft used to transport the cargo. For future aircraft design, the demand should be a prediction of future demand; in this work, this historical data takes the place of this future demand prediction.

Each data entry in the GATES dataset represents cargo on a pallet or a pallet-train that the AMC actually transported. Each pallet data entry has detailed information about the pallet transported, such as pallet gross weight, departure date and time, arrival date and time, mission distribution system (MDS), aircraft tail number, aerial port of embarkation (APOE), aerial port of disembarkation (APOD), pallet volume, pallet configuration, and so forth. These data enable the reconstruction of the route network, pallet demand characteristics, and existing fleet size for the fleet assignment problem.

Based on the available dataset, this problem investigation uses the following assumptions:

1. The refined route network from the GATES dataset is representative of all AMC cargo operations
 - a. Only routes served by C-5, C-17 and 747-F aircraft are considered. These aircraft types account for a substantial portion ($\approx 75\%$) of the total pallets transported in the year 2006.
 - b. All pallets have fixed dimensions representing the 463L pallet type. Sizing the payload bay and, therefore, the fuselage, of the yet-to-be-acquired aircraft uses these pallet dimensions. In this effort, the problem formulation does not consider any “outsized” cargo capacity requirements.



2. The demand reported in GATES for 2006 is representative of future demand requirements when a new, yet-to-be-designed aircraft would enter into service.

The application problem does not assume any demand growth. The lack of publicly available information coupled with having only one year of operations reported in the GATES dataset prevents the development of a reasonable future pallet demand-forecasting algorithm for the routes operated by the AMC. However, the research methodology is still applicable, and effective, if future demand distributions are available, or if future demand can be estimated using demand forecasting algorithms.

Methodology

We pose the monolithic problem of simultaneously designing an aircraft and its operations as a mathematical programming problem that seeks to minimize (or maximize) a fleet-level objective by searching for the optimal values of a set of decision variables. These decision variables describe the requirements of the new system and the new system design features and determine the assignment of the new and existing systems to meet demand requirements under multi-domain uncertainties. The resulting problem is a stochastic MINLP problem.

- It is stochastic because of the presence of uncertainty in both new system design and pallet demand.
- It is mixed-integer because of the presence of continuous decision variables such as the aircraft design variables of aspect ratio and wing loading, along with integer decision variables such as pallet capacity.
- It is non-linear because of the existence of non-linear objective function and constraints related to the aircraft sizing equations.

Subspace Decomposition Strategy

The monolithic deterministic problem formulation results in an MINLP problem, which is, in general, difficult—if not impossible—to solve; MINLP problems combine the difficulty of nonlinear optimization and the combinatorial nature of mixed integer programs. The decomposition approach, a procedure of solving several domain-specific subproblems linked by a top-level problem, is one procedure that can obtain results for this kind of problem, with some minor modifications. Figure 2 presents the decomposition strategy and shows how information flows between the three smaller subproblems. The subspace problems presented here follow natural boundaries of the domains involved.






Figure 2. Subspace Decomposition of Monolithic Optimization Problem Addressing Uncertainty in Both Aircraft Sizing and Fleet Assignment

The top-level problem explores the requirements space for the new, yet-to-be-introduced aircraft based on fleet-level metrics. The top-level problem chooses candidate values for the top-level decision variables, which then become parameters for the aircraft sizing subproblem. A reliability-based design optimization formulation is used for the aircraft sizing subproblem. After the aircraft sizing subproblem is solved, the outputs of the aircraft sizing problem and the current values of the top-level optimization problem (namely the productivity coefficients and fuel consumption coefficients, pallet capacity and design range) then become inputs to the fleet assignment problem. A hybrid formulation that combines the descriptive sampling approach and interval robust counterpart formulation solves the fleet assignment subproblem. Here, the assignment problem’s objective is to minimize the fleet-level fuel consumption using characteristics of the new, yet-to-be-introduced aircraft (range, pallet capacity, and speed) along with other existing aircraft in the fleet, subject to capacity, demand, fleet-level productivity, and scheduling constraints. The fleet-level values of the performance metrics return to the top top-level problem as the responses of interest.

Top-Level Problem

In this effort, the top-level optimization problem does not include any nonlinear constraints and only has bounds imposed on the top-level decision variables. Equations 1 to 4 describe the deterministic formulation of the top-level problem; the formulation incorporating uncertainty appears later in the paper.

$$\begin{aligned}
\text{Minimize} \quad & \text{Fleet fuel}(Pallet_x, Range_x, Speed_x) && (1) \\
\text{Subject to} \quad & 14 \leq Pallet_x \leq 38 && \text{(Design pallet capacity bounds) (2)} \\
& 2400 \leq Range_x \leq 3800 && \text{(Range at max. payload bounds in nmi) (3)} \\
& 350 \leq Speed_x \leq 550 && \text{(Cruise speed bounds in knots) (4)} \\
& Pallet_x \in Z^+ \quad Range_x, Speed_x \in R^+
\end{aligned}$$

Equation 1 describes the objective function that seeks to minimize the fleet-level fuel consumption using pallet capacity, range and cruise speed of the new, yet-to-be-introduced aircraft type X as decision variables. Equations 2 to 4 describe the bounds for the top-level design variables. The values for the bounds were based on strategic airlift requirements and characteristics exhibited by current cargo transport aircraft (Gertler, 2010; Graham et al., 2003). Here, the design requirement decision variable describing payload capacity uses an integer number of pallets, while the design range and design speed decision variables are continuous.

Aircraft Sizing Subproblem

Uncertainty in Design Parameters

The conceptual phase of the aircraft design process relies upon semi-empirical equations and simplified physics models. The limited knowledge available about the system definition at this phase of the design process combined with the usage of low-fidelity modeling tools results in high uncertainty. Aircraft sizing typically determines the size, weight and performance of an aircraft to meet its design mission based on a set of nominal values on operating conditions (e.g., cruise altitude). However, when evaluating the “operating missions” to determine block time and fuel consumed on the flight, there might be a variation in assigned altitude, routing, speed, and so forth, which would alter the block time and fuel consumed. For instance, there is uncertainty in the prediction of the parasite drag coefficient. In this example, a scaling factor k_{C_D} follows a distribution to represent the uncertainty in the parasite drag prediction, so that the “actual” coefficient relates to the “predicted” coefficient in the following manner:

$$C_{D_0 \text{ actual}} = k_{C_D} \times (C_{D_0 \text{ predicted}})$$

To address the uncertainty related to operations and predictions of the new aircraft performance in the aircraft sizing subspace with reasonable computational expense, the Analysis of Variance (ANOVA) technique, a sensitivity analysis method, determined the subset of the most important parameters that influence the outputs under consideration (Montgomery, 2008). This investigation assumes triangular distributions for the scaling factors of identified parameters listed in Table 1.



Table 1. Triangular Distributions of the ANOVA Identified Uncertain Parameters in the Aircraft Sizing Subspace

Uncertain Parameters (ξ)	Lower limit	Mode	Upper Limit
C_{D_0} multiplier, k_{c_D}	0.90	1.0	1.10
Specific Fuel Consumption, SFC [hr^{-1}]	0.45	0.5	0.55
Oswald efficiency multiplier, k_{e_0}	0.95	1.0	1.05
Cruise altitude [ft]	32000	35000	38000
Pallet mass [lbs]	7200	7500	7800

The aircraft sizing subproblem seeks to minimize the fuel consumption of the new, yet-to-be-introduced aircraft for the values of design range ($Range_x$), pallet capacity ($Pallet_x$), and cruise speed ($Speed_x$) from the top-level problem. With the top-level objective to minimize fleet-level fuel consumption, and the aircraft sizing objective to minimize the fuel consumed by the new aircraft for its prescribed design range, pallet capacity, and cruise speed, a slight disconnect exists between the objectives of these two levels. The difference in the objectives is that, at each aircraft sizing iteration, the minimization of fuel consumption uses a single combination of fixed values for design range, pallet capacity, and cruise speed—this is the typical case in aircraft design where these quantities are set as requirements for some “representative design mission.” However, the top-level optimization problem drives the question of “what requirements do we need to set in the first place?” by searching through the decision space of the top-level variables to find aircraft requirements that optimize fleet-level operational aspects of how the aircraft is used.

For example, consider the dimension of design range—as the top-level problem searches across values of range, this naturally changes the set of feasible routes that the new aircraft can fly, thereby changing how the fleet comprising existing and new aircraft serves the overall route network. By doing so, the top-level problem seeks additional fleet-wide fuel savings that these operational aspects reflect as a function of the decision variables. Therefore, the aircraft sizing objective can be viewed as a subset of the top-level problem objective. Because the type of aircraft assigned on individual flight segments drives the total amount of fuel consumed by the fleet, an aircraft designed for minimal fuel consumption will lead to improved fleet utilization that reduces fleet-level fuel consumption, when compared to fleet operations using only the fleet of existing aircraft. The approach in this work poses the aircraft design subproblem in the context of Reliability Based Design Optimization problem to account for uncertainty in the design phase.

The Reliability-Based Design Optimization (RBDO) formulation (shown below) represents the aircraft design under uncertainty problem.

$$\underset{x}{\text{Minimize}} \quad E[f(x, \xi)]$$

$$\text{Subject to } P[g(x, \xi) \leq 0] \geq b_i \quad \forall i = 1, 2, \dots, m$$

x : set of decision variables

ξ : set of uncertain parameters

b_i : desired probability of satisfying the i^{th} constraint



Aggregating the outputs for each realization (sample) of the uncertain parameter allows for the estimation of statistical measures such as expectation and probability, which the objective and constraint function evaluations require. The objective of the aircraft sizing subspace is to minimize the fuel consumption of the new aircraft X using the decision variables listed in Table 2. For each function evaluation of the top-level problem, the current values of $Pallet_x$, $Range_x$, and $Speed_x$ become fixed parameters for the aircraft sizing problem. Table 2 summarizes the decision variables, uncertain parameters and constraints in the aircraft sizing optimization problem.

Table 2. Decision Variables and Constraint Limits in the Aircraft Sizing Optimization Problem

Decision variables, (x)	Lower Bound	Upper Bound
Wing Aspect Ratio, AR_x	6.00	9.50
Thrust-to-weight Ratio, $(TW)_x$	0.18	0.35
Wing Loading [lb/ft^2], $(W/S)_x$	65.00	161.00
Engine Bypass Ratio, BPR_x	4.50	14.50
Wing Leading Edge Sweep [deg], $Sweep_x$	10.00	35.00
Wing Taper Ratio, TR_x	0.10	0.40
Constraints	Value	
Takeoff Distance [ft]	≤ 8500	
Landing Distance [ft]	≤ 5500	
Second segment climb gradient	≥ 0.025	
Top-of-climb rate [ft/min]	≥ 500	

The aircraft sizing subproblem includes performance constraints such as limits on takeoff and landing distances, and also upper and lower bounds for the decision variables. The RBDO formulation optimizes the expected performance metric of interest and ensures that the probability of satisfying the performance constraints is greater than or equal to the user-defined reliability level, b_i , considering the uncertainty present in this subproblem.

Fleet Assignment Subproblem

The fleet assignment subproblem identifies the optimal assignment of the fleet's aircraft to meet demand obligations; this includes allocation of the new aircraft—as described by the solution from the preceding aircraft sizing subproblem—along with existing aircraft in the fleet. The following equations describe the deterministic formulation of the fleet assignment problem; a sampling approach, as described later, will address the uncertainty in the fleet assignment subproblem.



Minimize

$$\sum_{p=1}^P \sum_{k=1}^K \sum_{i=1}^N \sum_{j=1}^N x_{p,k,i,j} \times FC_{p,k,i,j} \quad (\text{Fleet-level fuel consumption}) \quad (5)$$

Subject to

$$\sum_{p=1}^P \sum_{k=1}^K \sum_{i=1}^N \sum_{j=1}^N x_{p,k,i,j} \times (\text{Speed}_{p,k,i,j} \times \text{Cap}_{p,k,i,j}) \geq L \quad (\text{Fleet-level productivity limit}) \quad (6)$$

$$\sum_{i=1}^N x_{p,k,i,j} \geq \sum_{i=1}^N x_{p,k+1,i,j} \quad \forall k = 1, 2, 3 \dots K, \quad (\text{Node balance constraints}) \quad (7)$$

$$\forall p = 1, 2, 3 \dots P, \quad \forall j = 1, 2, 3 \dots N$$

$$\sum_{k=1}^K \sum_{i=1}^N \sum_{j=1}^N x_{p,k,i,j} \times BH_{p,k,i,j} \leq B_p \quad \forall p = 1, 2, 3 \dots P \quad (\text{Daily utilization limit}) \quad (8)$$

$$\sum_{p=1}^P \sum_{k=1}^K \text{Cap}_{p,k,i,j} \times x_{p,k,i,j} \geq \text{dem}_{i,j} \quad (\text{Pallet demand constraints}) \quad (9)$$

$$\forall i = 1, 2, 3 \dots N, \forall j = 1, 2, 3 \dots N$$

$$\sum_{i=1}^N x_{p,1,i,j} \leq O_{p,i} \quad \forall p = 1, 2, 3 \dots P, \forall i = 1, 2, 3 \dots N \quad (\text{Starting location constraints}) \quad (10)$$

$$\sum_{i=1}^N \sum_{j=1}^N x_{p,k,i,j} \leq 1 \quad \forall p = 1, 2, 3 \dots P, \forall k = 1, 2, 3 \dots K \quad (\text{Trip limit}) \quad (11)$$

$$x_{p,k,i,j} \in \{0, 1\} \quad (\text{Binary variable})$$

Equation 5 is the objective function that seeks to minimize the fleet-level fuel consumption, where $FC_{p,k,i,j}$ indicates the fuel consumption coefficient of the k^{th} trip for aircraft p from base i to base j . The equation has two parts; the first product inside the square brackets, $x_{p,k,i,j} \times FC_{p,k,i,j}$, represents the fuel consumption of the existing fleet, while the rest of the terms inside the square brackets represents the fuel consumption of assigning the new, yet-to-be-designed aircraft. The fuel consumption characteristics of the new aircraft are a function of aircraft design variables (aspect ratio, thrust-to-weight ratio, etc.) and aircraft design requirements (pallet capacity, design range, and cruise speed). The term $x_{p,k,i,j}$ is a binary decision variable that takes a value of 1 if the k^{th} trip of aircraft p is flown from base i to base j , and it takes a value of 0 otherwise.

Equation 6 accounts for the multi-objective nature of this problem. This forces the fleet-level productivity to be greater than a pre-defined limit, L ; the limit is varied and the problem is re-solved for each varied value of the limit to generate a set of Pareto optimal solutions. The term $x_{p,k,i,j} \times FC_{p,k,i,j}$ in Equation 6 refers to the productivity (speed of payload delivered) of utilizing aircraft type p for the k^{th} trip from base i to base j .



Equation 7 is the balance and sequencing constraint that enables the $(k+1)^{th}$ trip of an aircraft out of a base, i , to occur only after the k^{th} trip of that aircraft into base i . This constraint ensures that an aircraft needs is present at a base prior to completing a subsequent segment trip out of the same base.

Equation 9 limits flights to only occur within the daily utilization limit, B_p (here, this uses an assumption of 16 hours per day to account for loading, unloading, servicing, maintenance, etc.), of the aircraft, where $BH_{p,k,i,j}$ indicates the block hour of the k^{th} trip for aircraft p from base i to base j .

Equation 9 ensures that the carrying capacity of the combined trips meets or exceeds the pallet demand on each route, where $Cap_{p,k,i,j}$ indicates the pallet carrying capacity of the k^{th} trip for aircraft p from base i to base j .

Equation 10 ensures that the first trip of each aircraft p originates at its initial location (this is considered the aircraft's home or starting base for the day of operations); this initial location is randomly generated. Because the GATES dataset does not clearly indicate the starting location of aircraft each day, the problem formulation here uses a random distribution for each aircraft's starting location. The term $O_{p,i}$ is a binary variable that indicates if base i is the initial location for aircraft p .

Equation 11 ensures that each aircraft p flies at most one trip for its k^{th} segment.

The motivation for the "scheduling-like" formulation is to represent the scheduling and operations decisions made by the Air Mobility Command; it does not explicitly consider pilot scheduling (this 16 hours per day of available aircraft time could represent this, in part), nor does it account for the prioritization of cargo (this is not addressed in this formulation). This formulation, using node balance constraints, allows individual aircraft to make multiple flight segments in one day (as long as these fit within a prescribed time limit), allows for pallets to be carried from their origin to destination on possibly multiple aircraft, and tracks each individual aircraft by "tail number." These features more directly model AMC operations than some of the previous models of the authors and their colleagues when considering passenger airline transportation (Mane et al., 2007; Govindaraju et al., 2015).

Uncertainty in Fleet Operations

The uncertainty associated with the performance of the newly designed aircraft (type X) propagates to the fleet assignment subspace through the distributions of the new aircraft's predicted fuel consumption, $\bar{F}C_{p,k,i,j}$, and flight block hours, $\bar{B}H_{p,k,i,j}$, on given routes in the network; only aircraft "tail numbers" p that are associated with type X aircraft have these distributions. Additionally, the AMC service network has inherent pallet demand uncertainty, as described above. Hence, the fleet assignment problem now includes uncertainty in both the performance of the new aircraft and the pallet demand in the service network. In this paper, a hybrid formulation that combines the interval robust counterpart formulation (Lin, Janak, & Floudas, 2004) for user-defined tolerance parameters (δ) and the descriptive sampling technique (Saliby, 1990) solves the fleet assignment problem under uncertainty.

Lin et al. (2004) proposed a robust optimization approach for bounded uncertainty to overcome the large computational expense incurred by scenario/sampling-based frameworks. Their approach produces "robust" solutions that are immune against uncertainties in both the coefficients and right-hand-side parameters of the inequality constraints of the Mixed Integer Linear Programming (MILP) problems. Lin et al. (2004) term a solution to be robust if it satisfies the following conditions:



- The solution is feasible for the nominal values of the uncertain parameters.
- For any value of the uncertain coefficients in the objective function and the uncertain parameters in the right-hand side of the constraints, the solution must satisfy the i^{th} inequality constraint or, at worst, violate the constraint with an error of at most $\delta \times \max[1, |p_i|]$. In this expression, δ is a user-selected infeasibility tolerance coefficient, and p_i is the right-hand-side limit of the linear inequality constraint.

Applying the interval robust counterpart model to the deterministic formulation of the fleet assignment subproblem described above results in two additional sets of constraints and a modified objective function where an auxiliary variable (*Fleet fuel*) is introduced to enable introduction of the original objective function represented by Equation 5 as a constraint—thereby making it amenable to robust optimization strategies. The reformulation of the original objective function (Equation 5) is now as follows:

$$\text{Minimize Fleet fuel} \quad (12)$$

$$\text{Subject to } \sum_{p=1}^P \sum_{k=1}^K \sum_{i=1}^N \sum_{j=1}^N x_{p,k,i,j} \times FC_{p,k,i,j}^U \leq \text{Fleet fuel} (1 + \delta) \quad (13)$$

where $FC_{p,k,i,j}^U$ is the upper bound of the fuel consumed by aircraft p on the k^{th} trip from base i to base j . Evaluating the performance of the new aircraft for different samples of the aircraft sizing uncertain parameters (ξ) generates distributions of the performance metrics such as the fuel consumption coefficient. The upper bound, $FC_{p,k,i,j}^U$, is then determined from the distribution of the fuel consumption coefficient $\widetilde{FC}_{p,k,i,j}$ applied to only aircraft p that are of the newly-designed type X . δ is the user-defined, infeasibility tolerance parameter that can take values between 0 and 1. For example, setting δ to 0.1 for a particular constraint indicates that 10% violation of the worst-case scenario of that constraint is acceptable. Using Equation 13, if all of the uncertain fuel consumption coefficients for the new aircraft are at their upper bound (i.e., the aircraft burns the most possible fuel from the distribution, $(\widetilde{FC}_{p,k,i,j})$), then the total fuel consumed by the fleet is no more than 10% above the user-defined limit for fleet fuel consumption. The daily utilization limit constraint (Equation 8) is modified as follows:

$$\sum_{k=1}^K \sum_{i=1}^N \sum_{j=1}^N x_{p,k,i,j} \times BH_{p,k,i,j}^U \leq B_p (1 + \delta) \quad \forall p = 1, 2, 3 \dots P \quad (14)$$

where $BH_{p,k,i,j}^U$ is the upper bound of the distribution of block hours of aircraft p (restricted to only aircraft of type X) on the k^{th} trip from base i to base j . The deterministic robust counterpart fleet assignment problem now includes Equations 12, 13, and 14 in addition to Equations 5 to 11 from the original deterministic formulation of the fleet assignment problem.

The interval robust counterpart model is also applicable for the demand constraint (Equation 9) in the deterministic formulation, but this leads to a very conservative (protected against the maximum demand scenario) solution because the right hand side constraint limit, $dem_{i,j}$, is set to its upper bound or maximum value, $dem_{i,j}^U$, for each route as shown in Equation 15 below. For this constraint, the GATES dataset provides the values for the upper bound of the pallet demand on each route.



$$\sum_{p=1}^P \sum_{k=1}^K Cap_{p,k,i,j} \times x_{p,k,i,j} \geq dem_{i,j}^U \quad (15)$$

$$\forall i = 1, 2, 3 \dots N, \forall j = 1, 2, 3 \dots N$$

Instead, because of the AMC service network's high fluctuations in pallet demand and the on-demand nature of military cargo transport, the approach here employs a descriptive sampling approach (Saliby, 1990) to incorporate the stochastic nature of the demand. The method of descriptive sampling involves a deliberate collection of sample values that closely describes the represented distribution. The descriptive sampling approach samples more values from regions of higher density and fewer values from regions of lower density. The purposeful collection of sample values at specific quantile levels helps to match closely the actual or reported discrete demand distributions using a reduced number of samples, thus reducing the computational expense.

The deterministic robust counterpart formulation is solved multiple times for each demand sample vector generated through the descriptive sampling approach. From these multiple solutions, the expected value of the fleet-level performance metrics (fleet-level fuel consumption and/or fleet-level productivity) now return to the top-level optimization problem as the responses of interest. The robust counterpart formulation accounts for the propagation of uncertainty from the aircraft sizing to the fleet assignment subspace, while the descriptive sampling approach addresses the stochastic nature of pallet demand in the service network.

25-Base Network Problem

This section demonstrates how the subspace decomposition approach can identify the best new aircraft requirements and subsequent aircraft design to address fleet-level metrics under uncertainty. By treating this problem as a multiobjective problem, the approach can also generate tradeoffs between fleet-level metrics of interest; from these best tradeoff solutions, a decision-maker can also observe how the optimum design requirements for the new aircraft change for these different tradeoff opportunities.

Network Description

This study uses a subset of the AMC route network and fleet, comprising 25 bases and 219 directional routes, to demonstrate the approach. Figure 3 depicts the geographical locations and routes of the 25-base network. For the 25-base network, the existing fleet of AMC comprises 28 C-5s, 44 C-17s, and 21 chartered 747-Fs. The existing fleet serves as a "baseline" to measure the improvements due to the introduction of the new aircraft. This study assumes that five new, yet-to-be-designed-aircraft (all of type X) are introduced into the fleet. This assumption reflects an external decision made by the user or the decision-maker that specifies the number of new aircraft that are added to the fleet.





Figure 3. 25-Base Network

Note. Illustration was generated using <http://www.gcmap.com/>.

The 25 bases in the network are either the origin or the destination locations that transported the largest number of pallets in the AMC service network for the year 2006. The routes span the continents of North America, Asia, and Europe. Figure 4a shows the average and the minimum/maximum of the directional daily pallet demand for 50 routes in the network. Figure 4b shows the distribution of the number of routes based on the average daily pallet demand. The histogram indicates that the demand distribution is right-skewed and that several of the routes have an average daily demand of less than 20 pallets.

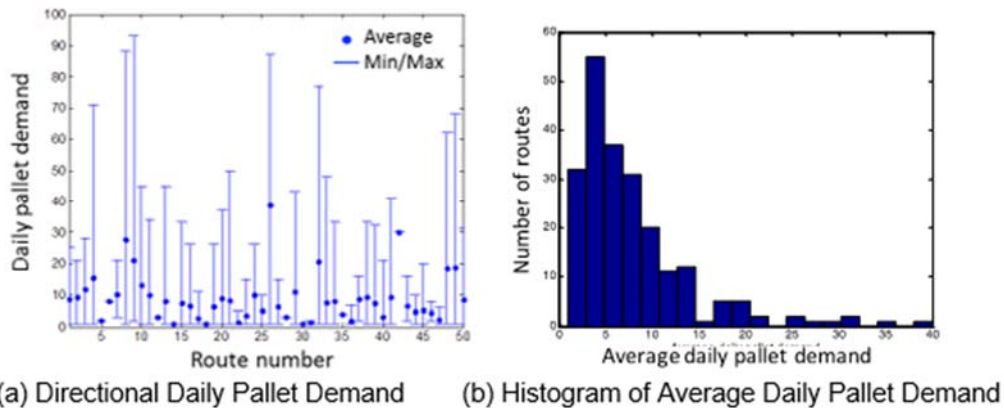


Figure 4. Pallet Demand Characteristics of the 25-Base Network

Results

In this study, the top-level optimization problem (refer to Figure 2) chooses candidate values for the decision variables of pallet capacity, design range, and cruise speed. These candidate values then become inputs to the aircraft sizing subproblem. The RBDO formulation of the aircraft sizing subproblem uses 50 samples for the uncertain parameters; this is a small number, but it allows for a tractable computational time. The reliability level, b_i , is set to 0.90 for the performance constraints in the aircraft.

After the aircraft sizing problem is solved, the outputs of the top-level subspace such as pallet capacity, and the outputs of the aircraft sizing subspace, the uncertain performance coefficients such as $\widetilde{F}C_{p,k,i,j}$, then become inputs to the fleet assignment subspace. The interval robust counterpart formulation of the fleet assignment problem is solved for 20 samples of demand across the network generated through the descriptive sampling approach. In this study, the infeasibility tolerance parameter, δ_i , is set to 0.10 for all the appropriate constraints. The expected values of the fleet-level performance metrics, calculated from the different solutions of the robust counterpart formulation, now return to the top-level subspace, and this process continues until convergence at the top-level.

Then, to identify tradeoffs between fleet-level fuel consumption and productivity, the entire process repeats with a different limit value on the productivity constraint. Minimizing the fleet-level fuel consumption with several different limits on fleet productivity leads to a number of tradeoff solutions. Figure 6 shows the results from the multi-objective analyses of the 25-base network problem. The plot shows the normalized expected values of the fleet-level metrics. Using normalized fleet-level responses helps to identify the trends and helps to show the relative variations in fleet-level responses for different solutions to the multi-objective optimization problem. The fleet-level responses have been normalized with respect to the lowest expected values from the results of the scenario labeled “Fleet with five new A/C.” Each point in the “Fleet with five new A/C” scenario describes the optimal design of the new aircraft required to meet the specific fleet-level objectives. These results show the collection of optimal aircraft designs that would meet the fleet’s operational needs at each level of permitted fuel consumption or at each level of required fleet-wide productivity.

For three different solutions from the “Fleet with five new A/C” results, Figure 5 contains callout boxes that describe the values of the new aircraft requirement decision variables along with the values of the aircraft design variables. The trends in the fleet-level responses are as expected, with fuel consumption increasing as productivity increases. There appears to be a trend in the “size” of the optimal aircraft along the Pareto frontier for increasing productivity/fuel consumption values. For a normalized expected productivity and normalized expected fuel consumption value of 1.0, the optimal requirement decision variables of the new aircraft X are at the lower bounds for pallet capacity (16) and design range (3800 nmi). Moving from this point on the tradeoff plot towards solutions with increasing fleet-level productivity, the results suggest that larger pallet capacities for the new aircraft X can best meet the fleet-level objectives. There is not substantial evidence to determine whether these trends would generalize to other route networks or other similar design problems; however, the behavior is not unexpected because the aircraft pallet capacity strongly drives the fleet-level productivity metric. Though it is intuitive that a larger aircraft would increase productivity, the optimal design features of the new aircraft X, such as the aspect ratio (AR_x), the wing loading ($(W/S)_x$), the thrust-to-weight ratio ($(T/W)_x$), etc., are reflective of the specific existing fleet and demand characteristics of the service network. For each solution in the plot, the assignments of the fleet of aircraft to routes are different to meet the actual demands better. The introduction of the five new aircraft (of type X) results



in fleet-level fuel savings between 2.79% and 6.48% for the same normalized expected fleet productivity values when compared to the case where only the existing fleet operates in the network.

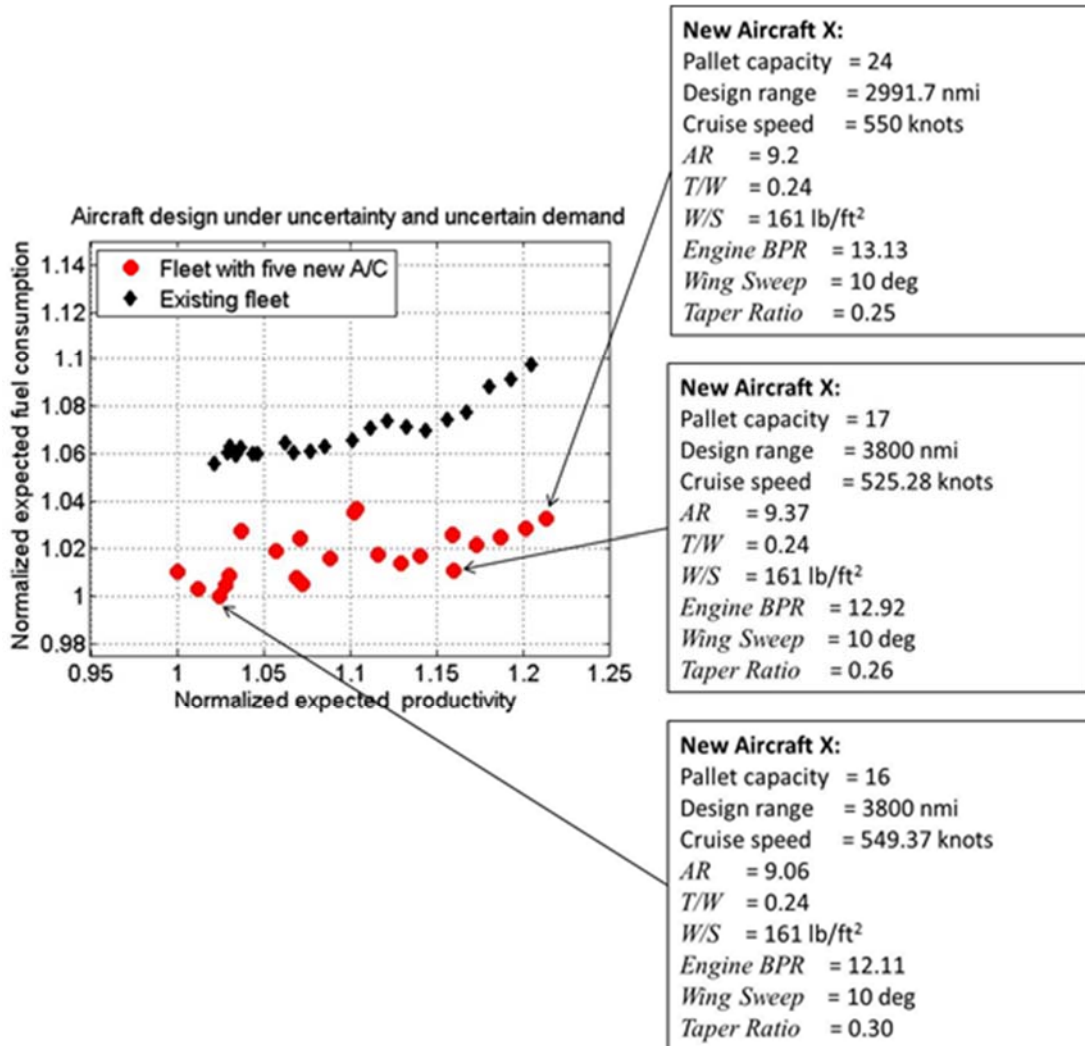


Figure 5. Results From Multi-Objective Analyses of 25-Base Network Problem

The solutions to multi-objective analyses present a way to perform “fuel/cost as an independent variable” type of trade-space analysis; this might be more obvious by switching the axes in the plot from Figure 5. These types of plots can help decision-makers/acquisition planners to analyze the trade-space and select the optimal requirements and design of the new aircraft that would achieve the desired level of fleet fuel consumption and productivity. For instance, a decision-maker can determine the level of fleet productivity available for a specific level of fleet fuel consumption; this fleet-level productivity value can then be translated to a specific (or bounded) level for the mobility airlift requirements that are set by the DoD in terms of tonnage of cargo transported per day. Having established the goals for the fleet-level productivity and fuel consumption, the collection of optimal aircraft designs required to achieve these fleet-level goals can be determined from plots such as those shown in Figure 5.



Decision-makers/acquisition planners can use such results to perform comprehensive exploratory analysis of the design space and identify regions in this design space that present significant viable opportunities to reduce the fleet fuel consumption. For instance, the AMC may need to incur “switching costs” (additional cost for training, maintenance and infrastructure due to the addition of a new aircraft type into the fleet) of integrating a new aircraft type into the fleet for a relatively small decrease in fuel burn; however, the trade-space analysis (Figure 5) can help identify promising designs and “inflection points,” if they exist, where the decision to acquire a new aircraft type could provide significant benefits.

Modeling and Solution Procedure for Commercial Air Travel Applications

In an effort to explore broader applications of the approach for similar acquisition-related issues, the authors conjecture how the approach could help decision-makers consider the best requirements for a new passenger transport aircraft. The nature and structure of uncertainty in commercial passenger air travel differs from the characteristics of data for the AMC service network. Adapting the subspace decomposition framework to commercial applications requires proper attention to the differences in the nature of the uncertainty that manifests in the data.

Similar to the approach used to extract an example problem from the GATES data, data from the Bureau of Transportation Statistics (BTS) T100 Segment database for non-stop monthly passenger demand can provide the basis for a commercial passenger airline problem.

The operations of the commercial air travel industry differ from military airlift operations, such as those managed by the Air Mobility Command (AMC) of the U.S. Air Force. The primary difference lies in the fact that commercial aviation operators such as airlines publish their schedules several weeks in advance of operating the flights, limiting the opportunities for modifications in the face of uncertain passenger demand. However, the AMC has a higher flexibility to modify their flight schedules due to the on-demand nature of palletized cargo transportation. The airline planning process typically involves a chronological sequence of decision-making phases. The planning process starts with schedule planning and development followed by four concurrent routines, namely, crew scheduling, revenue management, airport resource management, and aircraft maintenance routing. The schedule planning and development phase comprises market forecasting, schedule construction, capacity planning, fleet assignment, and schedule evaluation procedures.

For the purposes of strategic fleet planning and acquisition decision-making, the fleet allocation formulation for the commercial air travel case study integrates the schedule creation and fleet assignment procedures into a single mathematical programming problem. Figure 6 shows the modified subspace decomposition framework addressing uncertainty in both the design of the new aircraft and passenger demand for the commercial air travel case study.






Figure 6. Subspace Decomposition Framework Addressing Uncertainty in Both the Aircraft Sizing and Airline Allocation Subspace for Commercial Air Travel Application

The top-level problem explores the “requirements space” for the new, yet-to-be-designed aircraft using passenger capacity, range and cruise speed as the top-level decision variables. The aircraft sizing subspace accounts for the inherent uncertainty present in the conceptual phase of the design process through an RBDO formulation.

The airline allocation subspace is solved in two steps. The first step involves determining the minimum cost schedule based on the maximum number of passengers transported on each route. To use the framework in practice, the maximum anticipated demand would rely upon internal analysis performed by the airline to predict this; therefore, the demand is a point of input from the analyst and reflects a priori beliefs on the future state of demand. Here, our implementation will use the reported historical demand served on each route as available from the Bureau of Transportation Statistics, and the highest demand served will take the place of what we would expect the airline to predict. Solving the airline allocation problem generates an optimum schedule for the aircraft in the fleet to service the passenger demand in the route network. The second step involves a Monte Carlo simulation to account for the uncertainty in actual/realized passenger demand. The two-step procedure mimics the decision-making process of an airline, where a specific number of seats are allocated first for each route, followed by passengers buying tickets from the airline for traveling on those routes. Using the data from the demand distribution plots, the Monte Carlo simulation calculates the profits for the various simulated instances of passenger demand. From these numerous samples, the average (expected) fleet profit

values are then estimated. If required, the airline allocation subspace can be solved for four different quarters of demand data to reflect the seasonal variations in demand and the tactical planning timeframe usually employed by airlines. The expected fleet profit values return to the top-level subspace as the metrics of interest. The process continues until the top-level converges. At convergence, the solution describes the optimal aircraft requirements, the optimal description of the new aircraft, and the optimal allocation of the new and existing fleet of aircraft.

Concluding Statements and Future Work

The approach presented in this paper allows investigation of tradeoffs between fleet-level fuel usage, performance metrics and acquisition alternatives for a conceptual problem based on operations of the U.S. Air Force Air Mobility Command (AMC) under domain-specific uncertainties. The approach, while applied to the AMC case study, appears to be domain agnostic. Results from the AMC case study describe a collection of optimal aircraft design requirements and subsequent aircraft design descriptions that reduce fleet-level fuel consumption while satisfying the operational requirements under uncertainty in the new system design and uncertainty in the service network demand. A reliability-based design optimization formulation addresses uncertainty in the design of the new aircraft. A hybrid fleet assignment formulation that combines the interval robust counterpart model and the descriptive sampling approach addresses both the propagation of uncertainty from the aircraft sizing subspace to the fleet assignment subspace and the demand uncertainty in the service network. The immediately preceding section describes modification of the approach to address a commercial passenger airline application.

The methodology described in this paper can help guide decision-makers and acquisition planners to determine optimal design requirements for new, yet-to-be-introduced aircraft to reduce fleet-level fuel consumption. Solutions from these “design under uncertainty” problems provide insight (expected performance gain and costs incurred) about new systems, and these insights can inform acquisition decisions related to setting the right design requirements for the new system. Addressing uncertainty explicitly in this quantitative approach allows for a more “robust” selection of these new system requirements.

Using the approach to address this as a multi-objective problem enables tradeoffs in the context of “fuel/cost as an independent variable.” Generating the new design requirement and new aircraft design solutions should facilitate discussion and understanding about what features this kind of process should entail under various operational scenarios. The results from the 25-base network problem demonstrate the quantitative framework’s applicability in guiding potential acquisition decisions under uncertainty for the AMC case study, and the computational tractability of the approach to solve large-scale real-world problems. A preliminary framework for adapting the approach to commercial aviation application is presented as well. Future work will focus on extending the decomposition approach to solve the combined aircraft design and fleet assignment problem under commercial aviation specific uncertainties for the commercial travel case study.



References

- Gertler, J. (2010). *Air Force C-17 aircraft procurement: Background and issues for Congress*. DIANE.
- Govindaraju, P., & Crossley, W. A. (2015). Concurrent aircraft design and trip assignment under both design parameter and demand uncertainty. In *15th AIAA Aviation Technology, Integration, and Operations Conference (ATIO)*.
- Govindaraju, P., Davendralingam, N., & Crossley, W. (2015). A framework to determine new system requirements under design parameter and demand uncertainties. In *Proceedings of the 11th Annual Acquisition Research Symposium*. Monterey, CA: Naval Postgraduate School.
- Graham, D., Bracken, J., Dalfonzo, J., Fedorochko, W., & Hilton, R. (2003). *Sustaining the Civil Reserve Air Fleet (CRAF) program* (No. IDA/HQ-P-3819). Alexandria, VA: Institute for Defense Analyses.
- Greer, W. L. (2010). *An application of cost-effectiveness analysis in a major defense acquisition program: The decision by the U.S. Department of Defense to Retain the C-17 transport aircraft* (No. IDA-D-4218). Alexandria, VA: Institute for Defense Analyses.
- Lin, X., Janak, S. L., & Floudas, C. A. (2004). A new robust optimization approach for scheduling under uncertainty: I. Bounded uncertainty. *Computers & Chemical Engineering*, 28(6), 1069–1085.
- Mane, M., Crossley, W. A., & Nusawardhana, A. (2007). System-of-systems inspired aircraft sizing and airline resource assignment via decomposition. *Journal of Aircraft*, 44(4), 1222–1235.
- Montgomery, D. C. (2008). *Design and analysis of experiments*. John Wiley & Sons.
- Office of the Under Secretary of Defense for Acquisition, Technology, and Logistics (OUSD[AT&L]). (2014). *Better buying power 3.0* (White paper). Washington, DC: Author.
- Saliby, E. (1990). Descriptive sampling: A better approach to Monte Carlo simulation. *Journal of the Operational Research Society*, 41(12), 1133–1142.
- Taylor, C., & de Weck, O. D. (2007). Coupled vehicle design and network flow optimization for air transportation systems. *Journal of Aircraft*, 44(5), 1479–1486.





ACQUISITION RESEARCH PROGRAM
GRADUATE SCHOOL OF BUSINESS & PUBLIC POLICY
NAVAL POSTGRADUATE SCHOOL
555 DYER ROAD, INGERSOLL HALL
MONTEREY, CA 93943

www.acquisitionresearch.net