SIMULTANEOUS OPTIMIZATION OF
A MULTIPLE-AIRCRAFT FAMILY

Karen Willcox[*]
Massachusetts Institute of Technology, Cambridge, MA 02139

Sean Wakayama[†]
The Boeing Company, Huntington Beach, CA 92647

Abstract

Multidisciplinary design optimization is considered in the context of designing a family of aircraft. A framework is developed in which multiple aircraft, each with different missions but sharing common parts, can be optimized simultaneously. The new framework is used to gain insight to the effect of design variable scaling on the optimization algorithm. Results are presented for a two-member family whose individual missions differ significantly. Both missions can be satisfied with common designs. Moreover, optimizing both planes simultaneously rather than following the traditional baseline plus derivative approach vastly improves the common solution. A cost modeling framework is outlined that allows the value of commonality to be quantified for design and manufacturing costs. A notional example is presented to show the cost benefit that may be achieved by designing a common family of aircraft.

[*] Assistant Professor of Aeronautics & Astronautics, Aerospace Computational Design Laboratory, Room 37-447, 77 Massachusetts Avenue, Cambridge, MA 02139, kwillcox@mit.edu, Member AIAA.
[†] Principal Engineer/Scientist, Phantom Works, 5301 Bolsa Avenue, M/C H013 - A319, Huntington Beach, CA 92647, sean.r.wakayama@boeing.com, Senior Member AIAA.

American Institute of Aeronautics and Astronautics

Introduction

In today's competitive environment, the aerospace industry is faced with the challenge of designing aircraft not only with superior performance, but also at a lower cost. Multidisciplinary design optimization (MDO) is a tool that has been used successfully throughout the design process to enable improvements in aircraft performance. By simultaneously considering the effects of aerodynamics, structures, propulsion, flight mechanics and dynamics, and the complicated interaction between them, substantially improved performance can be achieved.

Studies show that the aircraft industry has evolved to a "dominant design," and that factors such as cost are becoming increasingly important[1]. In the "Better, Faster, Cheaper" era, the aerospace industry is searching for ways to lower costs without compromising aircraft performance. MDO provides a natural context in which to consider the tradeoffs between cost and performance. In the same way that MDO has been used to find "optimal" trades between aerodynamics and structural dynamics, one can conceive of expanding the optimization framework to consider the trades between, for example, aerodynamics, structural dynamics, cost, and revenue. While in the past, MDO for aircraft design has focused almost exclusively on engineering disciplines and performance-related issues, there is a clear need to expand its scope to include non-traditional disciplines such as cost.

The question then arises regarding the definition of "optimal" in the expanded framework. In the past, "optimal" has been synonymous with "minimum weight"; however, this choice of objective function reflects the performance-driven design mentality. While weight is often used as a surrogate for cost, it is clear that the minimum-weight and minimum-cost aircraft designs are not the same. Designing exclusively for minimum cost is also not a practical approach, since the impact on revenue potential could be so adverse as to outweigh the cost savings. This leads us to

American Institute of Aeronautics and Astronautics

the concept of designing for *maximum value*, where value is defined relative to a particular stakeholder, and captures all relevant considerations (such as performance, cost, revenue, environmental impact, *etc.*). The definition of a value metric for aircraft design has been explored recently by several authors[2,3].

One way to reduce costs is to conceive of a family of aircraft that share common parts and characteristics, such as planform and systems, but each aircraft satisfies a different mission requirement. Traditionally this has been achieved through the use of derivatives. A baseline aircraft is designed and subsequently modified to produce a number of derivatives to satisfy different missions (*e.g.* longer range, more payload). By taking advantage of commonality with the existing base model, it is possible to achieve the new mission at a far lower cost than would be incurred if a completely new plane were designed. Often, the modifications can be substantial, resulting in an almost entirely new plane. For example, the 737 Next Generation has a completely new wing.

In this work, the concept of commonality is taken a step further. We consider not just commonality between derivative aircraft whose missions are similar, but between two planes whose missions differ significantly. For example, if one could design a small capacity and a large capacity aircraft with common characteristics (for example a common wing), substantial savings could be realized in both recurring cost (manufacturing) and non-recurring cost (design effort and tooling cost). The savings from commonality come at a price: the weight of these common planes will be higher than if each were optimized separately for its own mission. The question then, is whether it is possible to design a family of common planes that satisfies all missions but whose cost saving outweighs the weight penalty. An example of this exists in practice: the Airbus A330 and A340 planes share common wings.

American Institute of Aeronautics and Astronautics

Fujita et al.[4] discuss the simultaneous optimization of a family of products; however, they assume that a baseline has been designed, and then consider the design of derivatives from this baseline. The mission of these derivatives is fairly close to the original: in the example they present, the only change to the derivative mission is to extend the range. In all cases, the derivative carries the same payload and retains the original fuselage.

In this paper, we will design a common family from a more fundamental approach. We consider the Blended-Wing-Body (BWB), a revolutionary concept for transport that integrates wing, fuselage, engines, and tail to achieve a substantial improvement in performance over a conventional transport[5,6,7] (see Figure 1). The specifics of the MDO approach are described, and then the simultaneous optimization and commonality structure are outlined. An example is presented in which a family of two BWB aircraft is designed: a large, 475-passenger plane and a smaller, 272-passenger plane, both with 8550 nm range. The framework for value optimization is described, and an example of program cost analysis is presented. Finally, we present some conclusions and directions of ongoing and future work.

Multidisciplinary Design Optimization Framework

As described by Wakayama et al.[8,9], WingMOD is an MDO code that optimizes aircraft wings and horizontal tails subject to a wide array of practical constraints. WingMOD was initially applied to the design of a composite wing for a stretched MD-90[10] and then went through considerable modification for application to the BWB[11,12]. The BWB planform is modeled as a series of spanwise elements as shown in Figure 2. Optimization services for WingMOD are provided by the Genie framework[11].

American Institute of Aeronautics and Astronautics

WingMOD uses intermediate fidelity analyses to quickly analyze an aircraft in over twenty design conditions that are needed to address issues from performance, aerodynamics, loads, weights, balance, stability and control. The low computational cost of the intermediate fidelity analyses allows the examination of all these issues in an optimization with over a hundred design variables while achieving reasonable computation time.

The basic WingMOD method models an aircraft wing and tail with a simple vortex-lattice code and monocoque beam analysis, coupled to give static aeroelastic loads. The model is trimmed at several flight conditions to obtain load and induced drag data. Profile and compressibility drag are evaluated at stations across the span of the wing with empirical relations using the lift coefficients obtained from the vortex lattice code. Structural weight is calculated from the maximum elastic loads encountered through a range of flight conditions, including maneuver, vertical gust, and lateral gust. The structure is sized based on bending strength and buckling stability considerations. Maximum lift is evaluated using a critical section method that declares the wing to be at its maximum useable lift when any section reaches its maximum lift coefficient, which is calculated from empirical data. Balance is evaluated by distributing weight over the planform[13].

The optimization algorithm used is sequential quadratic programming (SQP). This algorithm has been shown to work effectively on the large, nonlinear problems encountered in engineering applications. In particular, several other unconstrained algorithms were found to have convergence difficulties with the large number of constraints in problems of interest. The nonlinear problem can be stated as

$$\text{minimize} \qquad F(x)$$

$$\text{subject to} \qquad c_i(x) \geq 0, \quad i = 1, 2, \ldots, m \qquad (1)$$

American Institute of Aeronautics and Astronautics

where the vector x contains the n design variables, $F(x)$ is the objective function and c_i are the m constraint functions. The Hessian matrix of $F(x)$ contains the second order variations of the objective function with respect to each design variable, and is given by the n x n matrix $H(x)$:

$$\nabla^2 F(x) \equiv H(x) \equiv \begin{bmatrix} \dfrac{\partial^2 F}{\partial x_1^2} & \cdots & \dfrac{\partial^2 F}{\partial x_1 \partial x_n} \\ \vdots & & \vdots \\ \dfrac{\partial^2 F}{\partial x_1 \partial x_n} & \cdots & \dfrac{\partial^2 F}{\partial x_n^2} \end{bmatrix} . \tag{2}$$

The SQP algorithm uses a sequence of line searches to determine the optimum solution to the nonlinear problem (1). The design space is modeled as a quadratic objective with linear constraints, using finite difference gradient calculations. An approximate Hessian matrix is constructed from information gathered over the sequence of iterations using the Broyden-Fletcher-Goldfarb-Shanno (BFGS) update[14]. A quadratic programming problem is solved in the approximate design space to determine an estimated best direction for improvement. A line search is then executed in the actual design space, which seeks improvement in the solution.

Amongst other things, convergence of the algorithm is critically dependent on the conditioning of the Hessian matrix. Given the computed optimal solution $\bar{x}$ and the actual optimum $x*$, the error in the computed solution is given approximately by[14]

$$\left\| \bar{x} - x* \right\|^2 \approx \frac{2\varepsilon_A}{p^T H(x*)p}, \tag{3}$$

where ε_A is the absolute precision and p is any perturbation vector of unit length. Equation (3) shows that if $H(x*)$ is ill-conditioned, then the error in the computed solution can be very large

along certain directions. As stated by Gill et al.[14], since the objective function will vary much more rapidly in some directions than in others, an ill-conditioned Hessian is a form of bad scaling. This scaling problem may also have an adverse effect on the optimization algorithm itself, since the objective may vary extremely slowly along directions associated with a small eigenvalue. In this case, changes in the objective that are significant may be lost, and the algorithm will have trouble converging to the exact solution.

The ill-conditioning of the Hessian matrix can be quantified by its condition number given by

$$\kappa(H) = \frac{\lambda_{\max}(H)}{\lambda_{\min}(H)} \geq 1, \tag{4}$$

where $\lambda_{\max}(H)$ and $\lambda_{\min}(H)$ are respectively the maximum and minimum eigenvalues of H. (Note that because H is a symmetric matrix its eigenvalues and singular values are the same.) A matrix is said to be well-conditioned if its condition number is small (~1), and ill-conditioned if κ is large.

We can therefore ensure that (1) is a well-scaled problem by choosing a linear transformation of the design variables that minimizes the condition number of the Hessian matrix at the solution. In practice this is done a posteriori: once the algorithm has converged to the calculated optimum, the Hessian matrix is inspected and the design variables are scaled if necessary. Experience has shown that the approach can be simplified: by considering only the diagonal elements of $H(\bar{x})$ and scaling each of these to be $O(1)$, the problem becomes sufficiently well scaled.

Simultaneous Optimization with Commonality

The MDO framework described in the previous section was altered to allow the simultaneous optimization of multiple aircraft with varying levels of commonality. Constraints arising from

each of the disciplines were generated for each aircraft in the family. Commonality between family members was defined by dividing each plane into components. There were four main structural components: centerbody, inner wing, outer wing and winglet as shown in Figure 3. Each of these components can be specified to be either common or uncommon between family members. If just a particular component is allowed to vary, then the interface with a neighboring common component is kept common. An example might be allowing the centerbody to vary between family members, but keeping a common wing. In this case, the interface between the centerbody and the wing will be kept common.

Since making parts common is a discrete decision, the effect of varying levels of commonality is assessed via trade studies rather than being determined through the optimizer. In the problem setup, parts can be made common by either enforcing explicit constraints on their dimensions or linking the dimensions so they always have the same value. The ability to link variables was added to the Genie framework that supports WingMOD. When a chord at a spanwise location on Plane 1 is linked to the corresponding chord on Plane 2, Genie causes changes on either chord to be immediately reflected on the other. The optimizer then views each set of linked variables as a single quantity. This approach is more efficient than having a design variable for each chord and an explicit commonality constraint to force the variables to be equal: it eliminates a design variable and constraint for each linked variable, thus reducing the overall size of the combined optimization problem.

Fujita et al.[4] state that when the difference between product characteristics is large, two independent products must be designed, since "commonalization of parts cannot meet with performance requirement, even though it is effective for cutting cost." They do not attempt to quantify a "large difference," although the examples they discuss (stretching a baseline design to

American Institute of Aeronautics and Astronautics

extend range/payload) suggest that the missions of each family member are only incrementally different. Using the methodology developed here, a solution can be determined that does satisfy all performance requirements, even when the product characteristics are significantly diverse. Subsequently, a cost analysis can be applied to determine whether such a design is a viable option.

One approach to designing a common family might be to first optimize one family member and then force subsequent family members to share appropriate common features with the established planform. For example, in designing a small and a large plane with common planforms, one could first optimize the large aircraft, and then construct a smaller aircraft from the resulting planform. There are two problems with this approach. First, as pointed out by Fujita et al.[4], there is no guarantee that the solution obtained from the first optimization will satisfy all constraints on other family members. In the two-plane example, it is likely, although not certain, that the planform arising from optimizing a large plane will satisfy all requirements on the smaller plane. Moreover, adding a third member reduces the likelihood of the optimal large plane planform satisfying requirements of all family members. The second issue is that this sequential design approach results in a sub-optimal family. For example, by trading some optimality on the larger plane, significant improvement could be obtained in the smaller plane, resulting in an overall "better" family solution. This raises the issue of what the objective should be in the family optimization: should the combined weight of the family be minimized, or should the objective be more heavily biased towards a certain family member? The results presented in the following section will demonstrate that simultaneous optimization of an aircraft family overcomes the limitations of sequential design. A discussion of the choice of appropriate objectives will also follow.

American Institute of Aeronautics and Astronautics

Family Optimization Examples

Results will be presented for design of a two-member BWB family. The two aircraft satisfy the following mission requirements:

Plane 1: 8550 nm range, 475 passengers

Plane 2: 8550 nm range, 272 passengers

These two planes represent what might be the smallest and largest aircraft of a family with more members, *i.e.* this would be the greatest mission difference of interest. In all cases, the objective is to minimize the maximum takeoff weight (MTOW).

Example 1: Simultaneous Point Optimization of Two Aircraft

In order to investigate the new multiple-aircraft optimization framework, a test problem was set up. The two planes were designed simultaneously, but with no commonality constraints. Obviously, the solution to this problem could be obtained by optimizing each plane separately. Since the actual solution of the problem can be determined using the conventional approach, validation of and insight to the new framework can be gained.

This example highlights the importance of design variable scaling discussed above. Initially the simultaneous design problem was set up by simply concatenating the design variables and constraints for each plane, and attempting to minimize the sum of the takeoff weights. This combines the 139 design variables and 883 constraints of the small plane with the 150 design variables and 973 constraints of the large plane, to generate a system with 289 design variables and 1856 constraints. A detailed description of the design variables and constraints can be found in Wakayama[13]. Both single-aircraft problems were also optimized individually to determine the minimum-weight solution for each plane. In each individual case, the optimizer converged

American Institute of Aeronautics and Astronautics

without difficulty. Despite identical systems being used in the simultaneous setting with no coupling between them, the optimizer could not converge to the correct optimal solution. The minimum-weight solution was obtained for Plane 1 (the larger plane); however, the solution for Plane 2 was significantly sub-optimal when the algorithm claimed to have converged. At the conclusion of optimization, the Genie optimization framework reports finite difference estimates for the diagonal entries of the Hessian matrix. This information showed that geometric variables associated to Plane 2 were badly scaled because the diagonal entries of the Hessian associated with these variables were $O(10^{-2})$ and $O(10^{-3})$. In particular, the thickness and incidence angle variables for the spanwise elements on the smaller aircraft showed poor scaling.

These troublesome variables were rescaled, and the simultaneous optimization was performed again. Now the optimizer had no trouble converging to a solution that agreed with the individually obtained optimal designs. The calculated weights for the final 100 iterations are plotted in Figure 4. Shown are the calculated takeoff weights for each plane, normalized by the known point-optimum solution. The results are very similar for Plane 1 in both the scaled and unscaled cases. In the unscaled case, the solution has converged to a sub-optimum level for Plane 2.

This result suggests that in the simultaneous framework, the optimization algorithm is much more sensitive to poor scaling of the design variables. This can be explained mathematically by comparing the Hessian matrices of the individual and combined systems. If the Hessian of the system for Plane 1 alone is H_1 and for Plane 2 is H_2, then the Hessian of the combined two-aircraft system is given by

$$H = \begin{bmatrix} H_1 & 0 \\ 0 & H_2 \end{bmatrix} \tag{5}$$

since there is no coupling between the planes. The condition numbers of the single-plane Hessians are

$$\kappa(H_1) = \frac{\lambda_{max}^1}{\lambda_{min}^1} \ , \ \kappa(H_2) = \frac{\lambda_{max}^2}{\lambda_{min}^2} \tag{6}$$

and for the combined system:

$$\kappa(H) = \frac{\max\left(\lambda_{max}^1, \lambda_{max}^2\right)}{\min\left(\lambda_{min}^1, \lambda_{min}^2\right)}. \tag{7}$$

From equation (7) we can see that in the best case the condition number of the combined system will be equal to the worst of $\kappa(H_1)$ or $\kappa(H_2)$, and could in fact exceed both. Note that in practice, the full Hessian matrix is not computed, but rather the scaling is determined by considering only the diagonal elements of the Hessian; however, considering the eigenvalue structure and condition number of the matrix can lend insight to the numerical issues.

In particular, one can gain an understanding of how poor scaling might affect the convergence of the optimization algorithm by considering a simple geometric representation of the problem. Consider a system with two design variables. In the SQP algorithm, at each step the design space is modeled as a quadratic objective with linear constraints. If we were to plot the objective versus each of the design variables, we would obtain a paraboloid as shown in Figure 5. If the design variables are perfectly scaled so that the Hessian matrix has a condition number of unity, then the cross-sectional slices of the paraboloid are circles. If one (or both) of the variables is badly scaled, then the cross sections of the paraboloid are elliptical. In fact, the relative lengths of the major and minor axes of the ellipse are described by the minimum and maximum eigenvalues of the Hessian matrix.

Figure 6 demonstrates how the SQP algorithm might be adversely affected by poor scaling of the design variables. In both diagrams, the optimum is located at the origin. In the well scaled case (a), the objective contours are circles and the line search moves the solution in the correct direction. Now consider the case in Figure 6(b) where design variable x_1 is well scaled, but x_2 is badly scaled. The line search will choose a direction that captures the correct behavior for x_1, but achieves very little improvement in x_2. If the scaling of the problem is sufficiently poor, the optimizer will converge to a sub-optimal solution represented by the star in Figure 6(b) where the remaining improvement in x_2 is difficult to identify. Figure 6(b) describes the BWB example very closely. One could consider x_1 as representing all design variables for Plane 1 and x_2 as representing all those for Plane 2. Because the set of variables in x_2 was poorly scaled, the optimizer achieved the true minimum-weight solution for Plane 1, but converged to a sub-optimal solution for Plane 2.

The discussion above describes the effect for the SQP algorithm, which uses a quadratic approximation to the design space at each iteration. Other gradient-based optimization methods use a different approach to search the design space; however, in general they will also be affected by design variable scaling, since computational arithmetic is not exact. Achieving good scaling within a problem is an important aspect of applying optimization methods successfully, but is often overlooked.

Example 2: Optimizing a Two-Aircraft Common Family

We now present results for the two-aircraft family with commonality. The planes are constrained to have completely common wings (inner wing, outer wing and winglet in Figure 3), but different centerbodies. The interface between the inner wing and the centerbody is also common. This example is one family combination that was investigated in a larger study. Varying levels of

commonality were considered for the two-aircraft family, including the following: common wing, common centerbodies; common wing, distinct centerbodies; common inner wing, common centerbodies, distinct outer wing. The question of how to choose the level of commonality is very interesting in itself, but will not be addressed here. Instead, the results from the family optimization described above will be used to assess the value of the new, simultaneous design approach.

The two-aircraft family was first designed using a conventional, sequential technique. The larger plane was optimized for minimum takeoff weight as if it were a point design. The smaller plane was then optimized to achieve its minimum takeoff weight, but constrained to have an identical wing to its pre-determined larger mate. This resulted in what we refer to as the *sequential family* design. The new optimization framework was used to design the two planes simultaneously by minimizing the sum of the takeoff weights, resulting in the *simultaneous family* design. Table 1 summarizes the results for the simultaneous optimization. Weights are normalized by the weight of the point-design small aircraft (MTOW or OEW as appropriate). The total takeoff weight for the simultaneous family is 0.2% less than that of the sequential family. While the percentage is small, in actual pounds the reduction is significant. This weight reduction is achieved by a small increase in the takeoff weight of Plane 1, which allows a significant decrease in the takeoff weight of Plane 2, resulting in an overall better solution. It is interesting to note in Table 1 that the empty weight of both aircraft is decreased.

Further interrogation of the solutions shows how the optimizer has made tradeoffs to achieve this weight reduction. The wing area (and hence structural weight) of Plane 1 has been slightly reduced at the expense of aerodynamic efficiency. A decrease in structural weight of 0.45% is traded for a reduction in average cruise lift to drag ratio of 0.35%. Therefore, the empty weight

American Institute of Aeronautics and Astronautics

Table 1: Optimization results for simultaneous family design. Changes in maximum takeoff weight (MTOW) and operating empty weight (OEW) are relative to the sequential design solution and are normalized by the point-design weight of Plane 2.

	Total Family	Plane 1	Plane 2
MTOW	-0.2%	+0.1%	-0.3%
OEW	-0.9%	-0.3%	-0.6%

of this aircraft decreases while the total takeoff weight increases. Although this means that the takeoff weight for Plane 1 is slightly greater than the point-optimum solution, the reduced wing area has very positive benefits for Plane 2. In the sequential design, the constrained wing area for Plane 2 is significantly larger than is actually required. Reduction of the area lowers the structural weight without compromising aerodynamic efficiency, which means that a further cut in weight is achieved via a lowered fuel requirement.

Figure 7 and Figure 8 depict the two different planforms for Planes 1 and 2 respectively. Although very little difference can actually be observed, the comparison shows how the diminished wing area has been achieved: reduced chords and slightly increased sweep. There are a total of 195 constraints active on the two aircraft. A detailed description of which constraints are critical can be found in Wakayama[13] for a single aircraft optimization. For the family optimization example presented in this paper, similar trends were observed.

Choice of Objective

When designing more than one aircraft, the objective should be chosen by careful consideration of the problem at hand. As mentioned previously, commonality between family members results

American Institute of Aeronautics and Astronautics

in a tradeoff between increased operating cost (weight) and reduced acquisition cost (manufacturing/development). Ultimately, cost and revenue models will be incorporated to the MDO framework and the family will be optimized by maximizing profit rather than by minimizing weight. In this expanded framework, the appropriate objective will be clear: maximize the overall profit associated to the family. In the current model where we consider takeoff weight, it is necessary to choose an appropriate objective that includes a weighted contribution from each family member. For a family with N_f members, the objective takes the form

$$\min\left\{\sum_{i=1}^{N_f} s_i W_i\right\} \tag{8}$$

where $0 < s_i < 1$ is a weighting factor for the ith plane and W_i is its takeoff weight. In many cases, simply minimizing the sum of the family takeoff weights will be a suitable approach ($s_i = 1/N_f$, $i = 1...N_f$). If other factors are taken into consideration, such as relative market demands or competing aircraft, then an objective should be chosen which favors the most critical family members. The usual caveats apply to using a weighted sum objective function, and a better approach, which will be investigated in the future, is to formulate the problem as a multi-objective system.

Investigation into the effect of objective weighting factors was performed for the two-plane common-wing family. Three optimizations were carried out with varying weighting factors, s_i, on the takeoff weights of each plane. The results are summarized in Table 2.

Table 2: Optimization results for different objectives. Changes are relative to the sequential design solution and are normalized by the point-design MTOW of Plane 2.

American Institute of Aeronautics and Astronautics

s_1	s_2	MTOW Plane 1	MTOW Plane 2	Total MTOW
0.5	0.5	+0.1%	-0.3%	-0.2%
0.1	0.9	+0.3%	-0.4%	-0.1%
0.9	0.1	+0.05%	-0.1%	-0.05%

While the first optimization provides the lowest total takeoff weight, the weight of the smaller plane can be further reduced by biasing the objective function in its favor. Again, a reduction of the aerodynamic efficiency of the large plane is traded for reduced wing area. With a large weighting placed on Plane 1, the simultaneous design begins to approach its point-optimal solution; however, we note that even with a smaller weight increase of 0.05% on the larger aircraft, larger savings can still be made on Plane 2. Note that it is not advised to try a case with any s_i set to zero. The sequential design solution can be approached in the above example by making s_1 close to unity and s_2 close to zero; however, if s_2 is chosen to be identically zero, there will be a number of design variables with no influence on the objective function, which could lead to convergence problems. Within the simultaneous optimization framework, the objective choice can be used to rigorously balance the compromise between family members. It would be extremely difficult to identify these tradeoffs when working with a sequential design method.

Value Optimization and Analysis

The above methodology describes an effective way to design multiple aircraft sharing common characteristics. However, finding the best possible family design using simultaneous optimization does not answer the following question: is commonality desirable, and if so, at what

level? This question cannot be answered by considering performance alone. Instead, one must determine whether the cost benefits of commonality outweigh the performance penalty incurred. In order to achieve this, financial elements must be incorporated to the design framework.

To date, MDO has been applied primarily in an engineering context, and objectives are typically a technical measure of performance such as takeoff weight. With the increasing focus on cost and affordability, it is a logical extension to consider lifecycle cost as yet another discipline to be incorporated to the multidisciplinary framework. However, this raises the question of what to choose as an objective function. Focusing on performance alone does not capture cost considerations; however, designing to minimum cost could easily result in an aircraft that is unappealing to airlines and thus has a low potential for revenue generation. It is important to balance performance and cost in a quantitative way. This can be achieved by introducing the concept of value.

Several definitions of a value metric have been explored for aircraft design[2,3]. Many candidate metrics draw on financial concepts such as net present value. For the purposes of this paper, we emphasize that the specific definition of value is flexible. It will vary depending on the stakeholder considered and on the particular project at hand. While many value frameworks might include performance, cost, and revenue considerations, more general frameworks that include other factors, such as noise and emissions, could also be constructed. It should also be noted that the concept of value as a metric does not preclude the option to focus exclusively on performance, if so desired.

Figure 9 depicts one possible value framework for design in an MDO context. The performance module contains the traditional engineering disciplines that are currently considered in MDO. The complete cost module contains all relevant pieces of the lifecycle cost ranging from design

American Institute of Aeronautics and Astronautics

and development through manufacture and operation to disposal. Finally, in order to quantitatively address the tradeoff between cost and performance, it is necessary to include a revenue module. This module relates the performance and cost characteristics of an aircraft family to the revenue that can be generated, and may also include factors such as market variation and competition.

In this paper, we begin to assess the value of commonality by considering the development and manufacturing cost savings for the family design presented above. The cost model must be developed in such a way that the advantages of commonality can be captured and quantified. In order to effectively achieve this, a shift in thinking is required: rather than considering a fleet to be composed of various aircraft, it must be thought of as composed of various parts. (Here, a part is taken to be a sub-assembly, which can range in size and complexity.) For example, a fleet might be thought of as comprising 25 of Aircraft 1 and 50 of Aircraft 2, with details on the part breakdown and content captured at a secondary level for each of these aircraft. Here we take a different approach--the fundamental constituents of the fleet are not aircraft, but rather parts. For example, the fleet might be thought of as comprising 25 of Part A, 50 of Part B, 75 of Part C, and so on. Part A would be unique to Aircraft 1, Part B unique to Aircraft 2, and Part C common to both. In this way, one considers not the design and manufacture of aircraft, but rather the design and manufacture of the parts. A part can also range in size and level of detail. For example, a part could be an entire wing, or it could be a particular rivet on the wing. Moreover, it is not necessary that the level of detail be constant throughout the model.

For the results presented here, a simple weights-based costing method was used. In the cost model, each part has certain properties associated with it, including its weight, geometric data, and cost information. For each part, costs per pound were defined in the categories shown in

Table 3. Within each category of recurring cost, the costs may be further broken down. For example, labor may be further divided into fabrication, major assembly, minor assembly, and assembly and integration. The level of detail can vary according to the information available.

Calculating aircraft component cost as a linear function of weight is a simple approximation and more sophisticated cost estimating relationships will be investigated in the future. Although the baseline component costs vary linearly with weight, the total cost is built up as a summation of these baseline values modified by learning curve and commonality effects. The total development and manufacturing costs of the aircraft are thus not simple linear functions of total weight.

Table 3: Categories for recurring and non-recurring cost.

Recurring Cost	Non-Recurring Cost
Labor	Engineering
Support	Tool Engineering
Materials & Equipment	Tool Fabrication
	Other

Cost Analysis of Two-Aircraft Family

Representative development and manufacturing costs for the two-aircraft family designed in the previous section will be presented here. Two options are considered: the simultaneously optimized family with common wings and minimum combined takeoff weight, and the point-design family in which each aircraft is designed for minimum takeoff weight. The takeoff weight penalties of the common family compared to the point designs are approximately 0.05% for

Plane 1 and 2% for Plane 2, where the percentages are relative to the combined takeoff weight of the point-design family.

Figure 10 shows notional design and manufacturing program costs for the aircraft family. Design of the family begins in year Y_0 and production begins in year Y_1. Production rates are assumed constant at 50 units per year for Plane 1 and 100 units per year for Plane 2. In the beginning of the program, cost savings are realized for the common family, since design work and tooling requirements are reduced. During production, further benefit is obtained from commonality due to increased learning curve effects; however, the baseline unit cost of the common family is higher due to increased wing weight. For the notional example shown, the benefits outweigh the cost penalty. Note that this plot is a strictly hypothetical case to describe the kind of output the cost analysis might generate, and is not intended to represent actual BWB costs.

As demonstrated, the cost tool enables the development and manufacturing cost of various design options to be assessed and compared; however, the question of what level of commonality is desirable has still not been answered. To completely close the loop between performance and finance, a revenue model is required to estimate the penalty associated with the lower performance common family. The value of commonality may then be assessed by comparing this penalty with the cost savings such as those depicted in Figure 10. This is the subject of ongoing research. Furthermore, the cost analysis shown here is performed a posteriori once the simultaneous optimization has been completed (based only on performance considerations), and thus the cost analysis has no direct impact on the simultaneous optimization framework. In an expanded framework, financial models will be a part of the optimization so that they can impact the optimal family solution directly.

Conclusions

American Institute of Aeronautics and Astronautics

Simultaneous optimization of multiple aircraft offers considerable benefit to the design process. By designing family members with common parts, substantial savings can be realized in manufacturing and development costs. The simultaneous optimization methodology presented here not only ensures that a common design can be found which satisfies all constraints on each family member, but also determines the best overall family solution by applying appropriate tradeoffs between aircraft. The results presented demonstrate that a common solution can be found even when mission requirements differ significantly between family members. Substantial weight savings can be achieved by designing family members simultaneously and allowing the optimizer to make apt tradeoffs.

In order to determine the most appropriate level of commonality between aircraft family members, financial models must be developed and applied to the problem. Much benefit can be gained through expanding the role of MDO in aircraft design and incorporating non-engineering disciplines, such as cost and revenue. A cost model has been developed that can be applied to a multiple aircraft family to determine the cost benefits of commonality. Preliminary cost analysis indicates that the weight penalties incurred by the common design are more than offset by savings in manufacturing and development cost. A revenue model is currently under development, which will enable the tradeoff between cost and performance to be quantified more accurately. It is also desirable to close the design loop and allow financial considerations to feed back and impact engineering design decisions. In future work, the cost and revenue models will be incorporated to the MDO framework in such a way that the optimizer can make decisions based on performance, cost, and revenue.
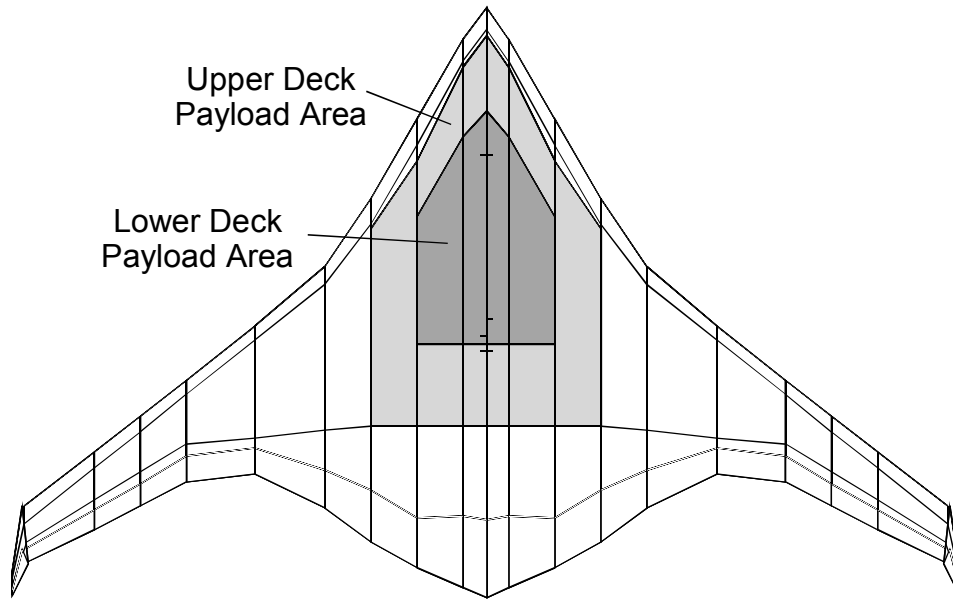
Acknowledgment

Figure 1: Blended-Wing-Body.

Upper Deck
Payload Area

Lower Deck
Payload Area

Figure 2: WingMOD BWB Model.

American Institute of Aeronautics and Astronautics

Figure 3: Modular structural breakdown of BWB.

American Institute of Aeronautics and Astronautics

Figure 4: Objective history for optimization algorithm. Shown are the calculated weights for each plane over the last 100 iterations, normalized by the appropriate point-optimum solution.

American Institute of Aeronautics and Astronautics

Figure 5 : Quadratic approximation to a two-dimensional design space.

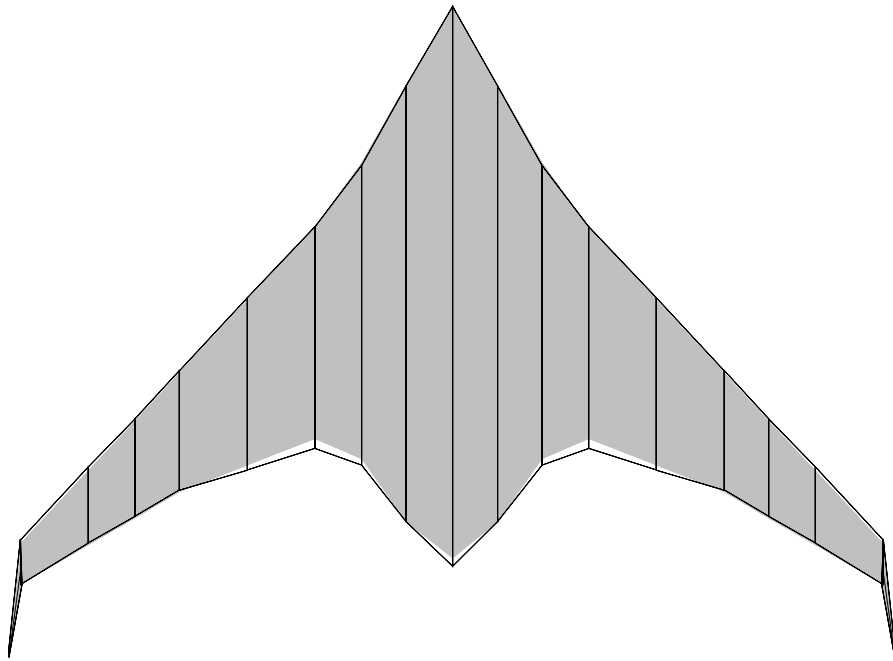American Institute of Aeronautics and Astronautics

Figure 6: Geometric interpretation of QP line search. (a) Well scaled and (b) badly scaled cases.
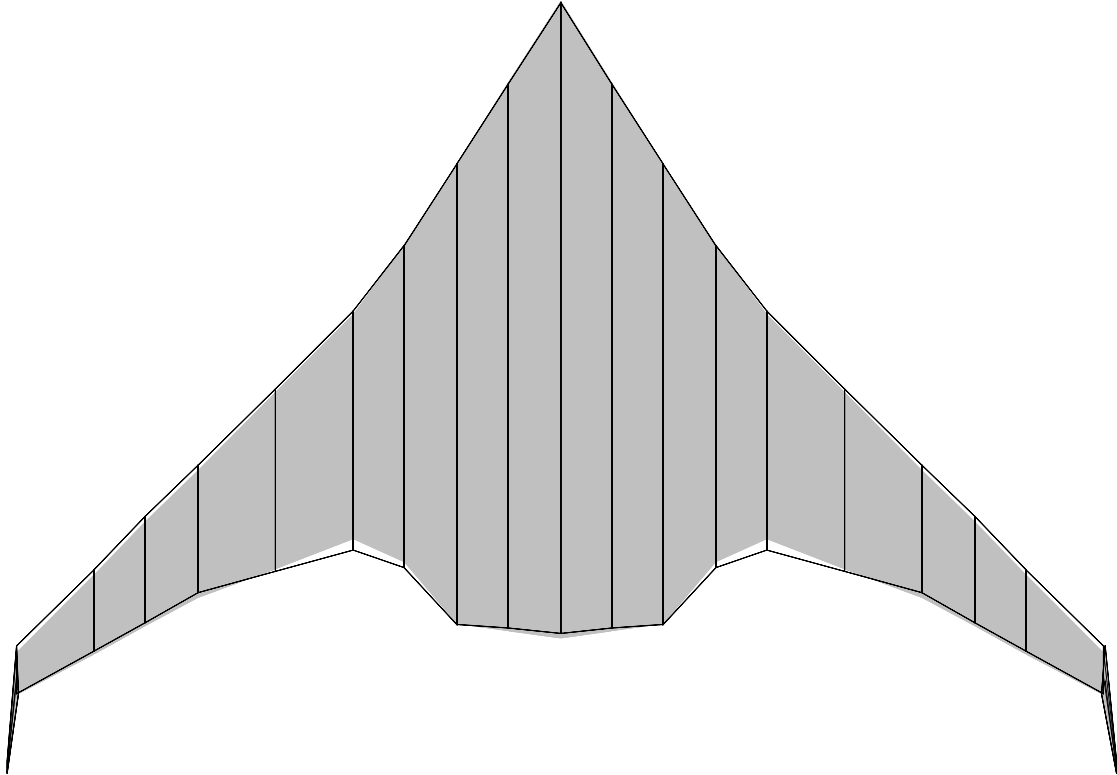
Figure 7: Planforms for Plane 1: wings common with Plane 2, centerbody optimized.

Sequential (black line) and simultaneous (gray shade) designs.

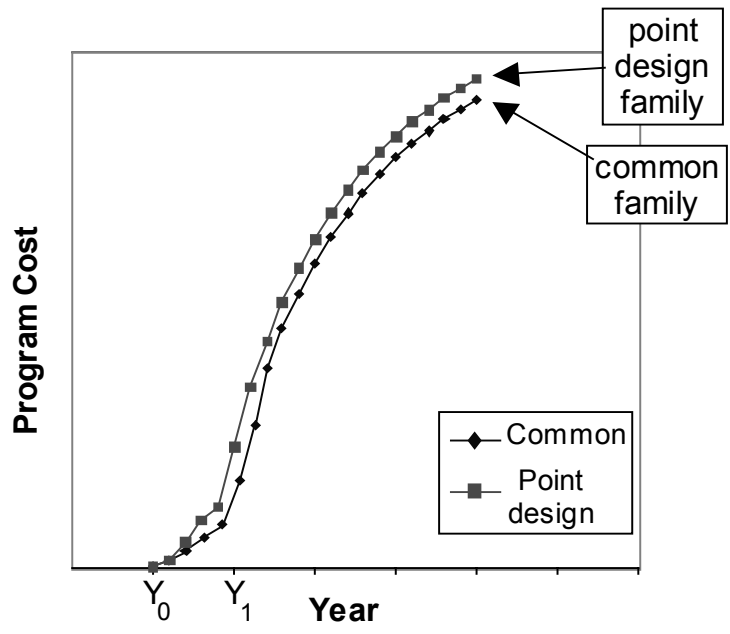American Institute of Aeronautics and Astronautics

Figure 8: Planforms for Plane 2: wings common with Plane 1, centerbody optimized.

Sequential (black line) and simultaneous (gray shade) designs.

American Institute of Aeronautics and Astronautics

Figure 9: Expanded MDO framework integrates performance, cost and revenue considerations.

American Institute of Aeronautics and Astronautics

Figure 10: Notional recurring and non-recurring costs for two-plane family.

American Institute of Aeronautics and Astronautics

References

1 Murman, E.M., Walton, M. and Rebentisch, E., "Challenges in the Better, Faster, Cheaper Era of Aeronautical Design, Engineering and Manufacturing," *The Aeronautical Journal*, Vol. 105, October 2000, pp 481-89.

2 Browning, T.R. *Modeling and Analyzing Cost, Schedule, and Performance in Complex System Product Development*, Ph.D Thesis, Dept. of Aeronautics & Astronautics, Massachusetts Institute of Technology, 1998.

3 Markish, J., *Valuation Techniques for Commercial Aircraft Program Design*, S.M. Thesis, Dept. of Aeronautics & Astronautics, Massachusetts Institute of Technology, June 2002.

4 Fujita, K., Akagi, S., Yoneda, T. and Ishikawa, M., "Simultaneous Optimization of Product Family Sharing System Structure and Configuration," Proceedings of ASME Design Engineering Technical Conference, September 1998, Atlanta, GA.

5 Liebeck, R.H., "Design of the Blended-Wing-Body Subsonic Transport," 2002 Wright Brothers Lecture, AIAA Paper 2002-0002 presented at 40th Aerospace Sciences Meeting and Exhibit, Reno, NV, January 14-17, 2002.

6 Liebeck, R. H., Page, M. A., Rawdon, B. K., "Blended-Wing-Body Subsonic Commercial Transport," AIAA Paper 98-0438 presented at 36th Aerospace Sciences Meeting and Exhibit, Reno, NV, January 12-15, 1998.

7 "Blended-Wing-Body Technology Study," Final Report, NASA Contract NAS1-20275, Boeing Report CRAD-9405-TR-3780, Oct. 1997.

8 Wakayama, S., Kroo, I., "Subsonic Wing Planform Design Using Multidisciplinary Optimization," *Journal of Aircraft*, Vol. 32, No. 4, Jul.-Aug. 1995, pp.746-753.

9 Wakayama, S., *Lifting Surface Design Using Multidisciplinary Optimization*, Ph.D. Thesis, Dept. of Aeronautics & Astronautics, Stanford University, Dec. 1994.

10 Wakayama, S., Page, M., Liebeck, R., "Multidisciplinary Optimization on an Advanced Composite Wing," AIAA Paper 96-4003 presented at 6[th] AIAA Symposium on Multidisciplinary Analysis and Optimization, Bellevue, WA, September 4-6, 1996.

11 Wakayama, S., Kroo, I., "The Challenge and Promise of Blended-Wing-Body Optimization," AIAA Paper 98-4736, in Proceedings of 7[th] AIAA Symposium on Multidisciplinary Analysis and Optimization, St Louis, MO, September 1998, pp. 239-250.

12 Wakayama, S., "Multidisciplinary Optimization of the Blended-Wing-Body," AIAA Paper 98-4938, in Proceedings of 7[th] AIAA Symposium on Multidisciplinary Analysis and Optimization, St Louis, MO, September 1998, pp. 1771-1779.

13 Wakayama, S., "Blended-Wing-Body Optimization Problem Setup," AIAA Paper 2000-4740 4003 presented at 8[th] AIAA Symposium on Multidisciplinary Analysis and Optimization, Long Beach, CA, September 6-8, 2000.

14 Gill, P.E, Murray, W. and Wright, M. H., *Practical Optimization*, Academic Press, London, 1981, pp. 346-354.