

Simultaneous Optimization of a Multiple-Aircraft Family

Karen Willcox*

Massachusetts Institute of Technology, Cambridge, Massachusetts 02139

and

Sean Wakayama†

The Boeing Company, Huntington Beach, California 92647

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 airplanes 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.

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. When the effects of aerodynamics, structures, propulsion, flight mechanics, and dynamics are considered simultaneously, as well as 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. Whereas 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 nontraditional 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. Whereas 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 because the impact on revenue potential could be so adverse as to outweigh the cost savings. This leads us to 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, for example, longer range, more payload. When commonality with the existing base model is taken advantage of, it is possible to achieve the new mission at a far lower cost than would be incurred if a completely new airplane were designed. Often, the modifications can be substantial, resulting in an almost entirely new airplane. 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 airplanes 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 nonrecurring cost (design effort and tooling cost). The savings from commonality come at a price: The weight of these common airplanes 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 airplanes that satisfies all missions but whose cost saving outweighs the weight penalty. An example of this exists in practice: The Airbus A330 and A340 airplanes share common wings.

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^{5–7} (Fig. 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

Received 25 June 2002; revision received 1 November 2002; accepted for publication 18 December 2002. Copyright © 2003 by Karen Willcox and Sean Wakayama. Published by the American Institute of Aeronautics and Astronautics, Inc., with permission. Copies of this paper may be made for personal or internal use, on condition that the copier pay the \$10.00 per-copy fee to the Copyright Clearance Center, Inc., 222 Rosewood Drive, Danvers, MA 01923; include the code 0021-8669/03 \$10.00 in correspondence with the CCC.

*Assistant Professor, Aerospace Computational Design Laboratory, Aeronautics and Astronautics, Room 37-447, 77 Massachusetts Avenue; kwillcox@mit.edu. Member AIAA.

†Principal Engineer/Scientist, Phantom Works, 5301 Bolsa Avenue, Mail Code H013-A319; sean.r.wakayama@boeing.com. Senior Member AIAA.



Fig. 1 BWB.




Fig. 2 WingMOD BWB model.

large, 475-passenger airplane and a smaller, 272-passenger airplane, 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.

MDO Framework

As described by Wakayama and Kroo⁸ and Wakayama,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 aircraft (Ref. 10) and then went through considerable modification for application to the BWB.^{11,12} The BWB platform is modeled as a series of spanwise elements as shown in Fig. 2. Optimization services for WingMOD are provided by the Genie framework.¹¹

WingMOD uses intermediate fidelity analyses to analyze quickly an aircraft in over 20 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 of these issues in an optimization with over 100 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

$$\begin{aligned} & \text{minimize} && F(\mathbf{x}) \\ & \text{subject to} && c_i(\mathbf{x}) \geq 0, \quad i = 1, 2, \dots, m \end{aligned} \quad (1)$$

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

$$\nabla^2 F(\mathbf{x}) \equiv H(\mathbf{x}) \equiv \begin{bmatrix} \frac{\partial^2 F}{\partial x_1^2} & \cdots & \frac{\partial^2 F}{\partial x_1 \partial x_n} \\ \vdots & & \vdots \\ \frac{\partial^2 F}{\partial x_1 \partial x_n} & \cdots & \frac{\partial^2 F}{\partial x_n^2} \end{bmatrix} \quad (2)$$

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 update (see Ref. 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 \mathbf{x} and the actual optimum \mathbf{x}^* , the error in the computed solution is given approximately by¹⁴

$$\|\mathbf{x} - \mathbf{x}^*\|^2 \approx \frac{2\varepsilon_A}{\mathbf{p}^T H(\mathbf{x}^*) \mathbf{p}} \quad (3)$$

where ε_A is the absolute precision and \mathbf{p} is any perturbation vector of unit length. Equation (3) shows that if $H(\mathbf{x}^*)$ is ill conditioned, then the error in the computed solution can be very large along certain directions. As stated by Gill et al.,¹⁴ because 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 because 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

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

where $\lambda_{\max}(H)$ and $\lambda_{\min}(H)$ are the maximum and minimum eigenvalues of H , respectively. (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.

Therefore, we can ensure that problem (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: When only the diagonal elements of $H(\mathbf{x})$ are considered and each of these are scaled to be $\mathcal{O}(1)$, the problem becomes sufficiently well scaled.




Fig. 3 Modular structural breakdown of BWB.

Simultaneous Optimization with Commonality

The MDO framework described in the preceding 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 airplane into components. There were four main structural components: centerbody, inner wing, outer wing, and winglet, as shown in Fig. 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.

Because 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 that 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 airplane 1 is linked to the corresponding chord on airplane 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 because “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 extend range/payload) suggest that the missions of each family member are only incrementally different. When the methodology developed here is used, 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 airplane 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-airplane example, it is likely, although not certain, that the planform arising from optimizing a large airplane will satisfy all requirements on the smaller airplane. Moreover, adding a third member reduces the likelihood of the optimal large airplane planform satisfying requirements of all family members. The second issue is that this sequential design approach results in a suboptimal family. For example, by trading some optimality on the larger airplane, 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.

Family Optimization Examples

Results will be presented for design of a two-member BWB family. The two aircraft satisfy the following mission requirements: airplane 1, 8550-nm range and 475 passengers and airplane 2, 8550-nm range, and 272 passengers. These two airplanes represent what might be the smallest and largest aircraft of a family with more members, that is, 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

To investigate the new multiple-aircraft optimization framework, a test problem was set up. The two airplanes were designed simultaneously, but with no commonality constraints. Obviously, the solution to this problem could be obtained by optimizing each airplane separately. Because 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 earlier. Initially the simultaneous design problem was set up by simply concatenating the design variables and constraints for each airplane and attempting to minimize the sum of the takeoff weights. This combines the 139 design variables and 883 constraints of the small airplane with the 150 design variables and 973 constraints of the large airplane, to generate a system with 289 design variables and 1856 constraints. A detailed description of the design variables and constraints may be found by 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 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 airplane 1 (the larger airplane); however, the solution for airplane 2 was significantly suboptimal 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 airplane 2 were badly scaled because the diagonal entries of the Hessian associated with these variables were $\mathcal{O}(10^{-2})$ and $\mathcal{O}(10^{-3})$. 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 Fig. 4. Shown are the calculated takeoff weights for each plane, normalized by the known point-optimum solution. The results are very similar for airplane 1 in both the scaled and unscaled cases. In the unscaled case, the solution has converged to a suboptimum level for airplane 2.

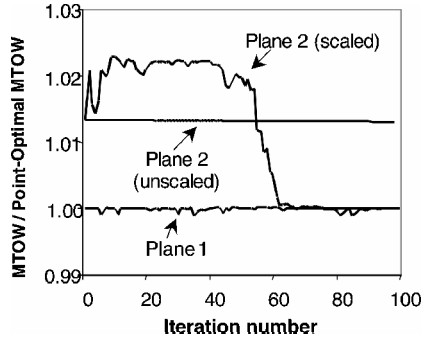


Fig. 4 Objective history for optimization algorithm: calculated weights for each airplane over the last 100 iterations, normalized by the appropriate point-optimum solution.




Fig. 5 Quadratic approximation to a two-dimensional design space.

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 airplane 1 alone is H_1 and for airplane 2 is H_2 , then the Hessian of the combined two-aircraft system is given by

$$H = \begin{bmatrix} H_1 & 0 \\ 0 & H_2 \end{bmatrix} \quad (5)$$

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


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

and for the combined system,


$$\kappa(H) = \frac{\max(\lambda_{\max}^1, \lambda_{\max}^2)}{\min(\lambda_{\min}^1, \lambda_{\min}^2)} \quad (7)$$

From Eq. (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 vs each of the design variables, we would obtain a paraboloid as shown in Fig. 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.



a) Well scaled case



b) Badly scaled case

Fig. 6 Geometric interpretation of quadratic programming line search.

Figure 6 demonstrates how the SQP algorithm might be adversely affected by poor scaling of the design variables. In both Figs. 6a and 6b, the optimum is located at the origin. In the well-scaled case (Fig. 6a), the objective contours are circles, and the line search moves the solution in the correct direction. Now consider the case in Fig. 6b, 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 suboptimal solution represented by the star in Fig. 6b, where the remaining improvement in x_2 is difficult to identify. Figure 6b describes the BWB example very closely. One could consider x_1 as representing all design variables for airplane 1 and x_2 as representing all of those for airplane 2. Because the set of variables in x_2 was poorly scaled, the optimizer achieved the true minimum-weight solution for airplane 1, but converged to a suboptimal solution for airplane 2.

The preceding discussion 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 because 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 airplanes are constrained to have completely common wings (inner wing, outer wing, and winglet in Fig. 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 and common centerbodies, common wing and distinct centerbodies; and 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 earlier 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 airplane was optimized for minimum takeoff weight as if it were a point design. The smaller plane

Table 1 Optimization results for simultaneous family design

Weight ^a	Total family, %	Airplane 1, %	Airplane 2, %
MTOW	-0.2	+0.1	-0.3
OEW	-0.9	-0.3	-0.6

^aChanges in maximum MTOW and OEW are relative to the sequential design solution and are normalized by the point-design weight of airplane 2.




Fig. 7 Planforms for airplane 1: wings common with airplane 2, centerbody optimized; sequential (black line) and simultaneous (gray shade) designs.

was then optimized to achieve its minimum takeoff weight, but constrained to have an identical wing to its predetermined larger mate. This resulted in what we refer to as the sequential family design. The new optimization framework was used to design the two airplanes 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 operating empty weight (OEW) as appropriate]. The total takeoff weight for the simultaneous family is 0.2% less than that of the sequential family. Although 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 airplane 1, which allows a significant decrease in the takeoff weight of airplane 2, resulting in an overall better solution. 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 airplane 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 of this aircraft decreases while the total takeoff weight increases. Although this means that the takeoff weight for airplane 1 is slightly greater than the point-optimum solution, the reduced wing area has very positive benefits for airplane 2. In the sequential design, the constrained wing area for airplane 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.

Figures 7 and 8 show the two different planforms for airplanes 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 may be found in Wakayama¹³ for a single aircraft optimization. For the family optimization example presented in this paper, similar trends were observed.

Table 2 Optimization results for different objective^a

s_1	s_2	MTOW, %		
		Airplane 1	Airplane 2	Total
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

^aChanges are relative to the sequential design solution and are normalized by the point-design MTOW of airplane 2.




Fig. 8 Planforms for airplane 2: wings common with airplane 1, centerbody optimized; sequential (black line) and simultaneous (gray shade) designs.

Choice of Objective

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

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

where $0 < s_i < 1$ is a weighting factor for the i th 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 ($s_i = 1/N_f$, $i = 1, \dots, N_f$). If other factors are taken into consideration, such as relative market demands or competing aircraft, then an objective should be chosen that 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 multiobjective system.

Investigation into the effect of objective weighting factors was performed for the two-airplane 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.

Whereas the first optimization provides the lowest total takeoff weight, the weight of the smaller airplane can be further reduced by biasing the objective function in its favor. Again, a reduction of the aerodynamic efficiency of the large airplane is traded for reduced wing area. With a large weighting placed on airplane 1, the

simultaneous design begins to approach its point-optimal solution; however, note that even with a smaller weight increase of 0.05% on the larger aircraft, larger savings can still be made on airplane 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 preceding 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 balance rigorously 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 preceding 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. 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 life cycle 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. Although 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. Note that the concept of value as a metric does not preclude the option to focus exclusively on performance, if so desired.

Figure 9 shows 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 life cycle

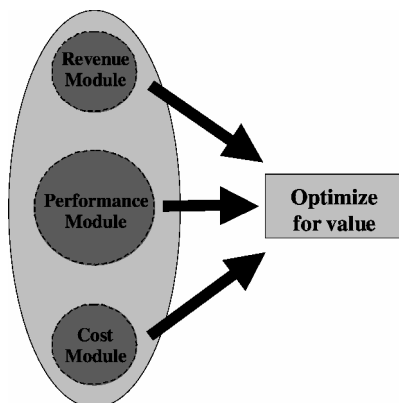


Fig. 9 Expanded MDO framework integrates performance, cost, and revenue considerations.

Table 3 Categories for recurring and nonrecurring cost

Recurring cost	Nonrecurring cost
Labor	Engineering
Support	Tool engineering
Materials and equipment	Tool fabrication
	Other

cost ranging from design and development through manufacture and operation to disposal. Finally, to address quantitatively 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 earlier. The cost model must be developed in such a way that the advantages of commonality can be captured and quantified. To achieve this effectively, 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 subassembly, 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.

Cost Analysis of Two-Aircraft Family

Representative development and manufacturing costs for the two-aircraft family designed in the preceding 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 airplane 1 and 2% for airplane 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 airplane 1 and 100 units per year for airplane 2. In the beginning of the program, cost savings are realized for the common family because design work and tooling requirements are reduced. During production, further benefit is obtained

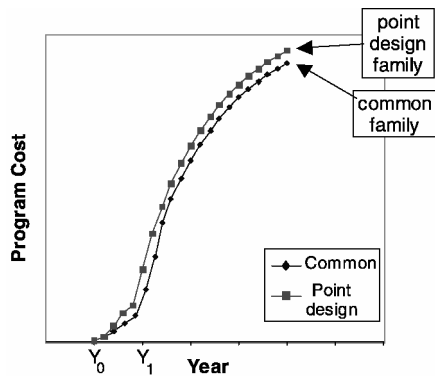


Fig. 10 Notional recurring and nonrecurring costs for two-airplane family.

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 Fig. 10 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 close the loop between performance and finance completely, 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 shown in Fig. 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

Simultaneous optimization of multiple aircraft offers considerable benefit to the design process. By the design of 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 that 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.

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 nonengi-

neering 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 that 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 feedback 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.

Acknowledgments

The authors gratefully acknowledge the contributions of the following people who helped originate and develop many of the concepts described in this paper: John Allen, George Busby, and Jacob Markish.

References

- Murman, E. M., Walton, M., and Rebentisch, E., "Challenges in the Better, Faster, Cheaper Era of Aeronautical Design, Engineering and Manufacturing," *Aeronautical Journal*, Vol. 104, No. 1040, Oct. 2000, pp. 481–89.
- Browning, T. R., "Modeling and Analyzing Cost, Schedule, and Performance in Complex System Product Development," Ph.D. Dissertation, Dept. of Aeronautics and Astronautics, Massachusetts Inst. of Technology, Cambridge, MA, Feb. 1999.
- Markish, J., "Valuation Techniques for Commercial Aircraft Program Design," M.S. Thesis, Dept. of Aeronautics and Astronautics, Massachusetts Inst. of Technology, Cambridge, MA, June 2002.
- Fujita, K., Akagi, S., Yoneda, T., and Ishikawa, M., "Simultaneous Optimization of Product Family Sharing System Structure and Configuration," ASME Design Engineering Tech. Conf., American Society of Mechanical Engineers, Paper DETC98/DFM-5722, Sept. 1998.
- Liebeck, R. H., "Design of the Blended-Wing-Body Subsonic Transport," AIAA Paper 2002-0002 Jan. 2002.
- Liebeck, R. H., Page, M. A., and Rawdon, B. K., "Blended-Wing-Body Subsonic Commercial Transport," AIAA Paper 98-0438, Jan. 1998.
- Liebeck, R. H., and Page, M. A., et al., "Blended-Wing-Body Technology Study," Final Report, NASA Contract NAS1-20275, Boeing Rept. CRAD-9405-TR-3780, The Boeing Co., Long Beach, CA, Oct. 1997.
- Wakayama, S., and Kroo, I., "Subsonic Wing Planform Design Using Multidisciplinary Optimization," *Journal of Aircraft*, Vol. 32, No. 4, 1995, pp. 746–753.
- Wakayama, S., "Lifting Surface Design Using Multidisciplinary Optimization," Ph.D. Dissertation, Dept. of Aeronautics and Astronautics, Stanford Univ., Stanford, CA, Dec. 1994.
- Wakayama, S., Page, M., and Liebeck, R., "Multidisciplinary Optimization on an Advanced Composite Wing," AIAA Paper 96-4003, Sept. 1996.
- Wakayama, S., and Kroo, I., "The Challenge and Promise of Blended-Wing-Body Optimization," *Proceedings of 7th AIAA Symposium on Multidisciplinary Analysis and Optimization*, AIAA, Reston, VA, 1998, pp. 239–250.
- Wakayama, S., "Multidisciplinary Optimization of the Blended-Wing-Body," *Proceedings of 7th AIAA Symposium on Multidisciplinary Analysis and Optimization*, AIAA, Reston, VA, 1998, pp. 1771–1779.
- Wakayama, S., "Blended-Wing-Body Optimization Problem Setup," AIAA Paper 2000-4740, Sept. 2000.
- Gill, P. E., Murray, W., and Wright, M. H., *Practical Optimization*, Academic Press, London, 1981, pp. 346–354.