VARIABLE-COMPLEXITY RESPONSE SURFACE APPROXIMATIONS FOR WING STRUCTURAL WEIGHT IN HSCT DESIGN*

Matthew Kaufman,† Vladimir Balabanov†, Susan L. Burgee,§ Anthony A. Giunta†
Bernard Grossman,¶ Raphael T. Haftka,‡ William H. Mason,# and Layne T. Watson**

ABSTRACT

A procedure for generating and using a polynomial approximation to wing bending material weight of a High Speed Civil Transport (HSCT) is presented. Response surface methodology is used to fit a quadratic polynomial to data gathered from a series of structural optimizations. Several techniques are employed in order to minimize the number of required structural optimizations and to maintain accuracy. First, another weight function based on statistical data is used to identify a suitable model function for the response surface. In a similar manner, geometric and loading parameters that are likely to appear in the response surface model are also identified. Next, simple analysis techniques are used to find regions of the design space where reasonable HSCT designs could occur. The use of intervening variables along with analysis of variance reduce the number of polynomial terms in the response surface model function. Structural optimization is then performed by the program GENESIS on a 28-node Intel Paragon. Finally, optimizations of the HSCT are completed both with and without the response surface.

1. INTRODUCTION

The use of multidisciplinary optimization techniques in aerospace vehicle design is often limited because of the significant computational expense incurred in the analysis of the vehicle and its many systems. In response to this difficulty, a variable-complexity modeling approach involving the use of refined and computationally expensive models together with simple and computationally inexpensive models has been developed. This variable-complexity technique has been previously applied to the combined aerodynamic-structural optimization of subsonic transport aircraft wings (Unger et al., 1992) and the aerodynamic-structural optimization of the High Speed Civil Transport (HSCT) (Hutchison et al., 1994, Dudley et al., 1995).

In this study, the integration of a response surface approximation to wing bending material weight into the HSCT configuration design process is considered. In previous studies, a series of algebraic functions, formulated using experience-based statistical information, had been used to estimate the bending material weight as a function of the design variables.

* This work was supported by NASA Grants NAG-1-1562 and NAG-1-1160.

† Research Assistant, Dept. of Aerospace & Ocean Eng., Virginia Tech, Blacksburg, VA 24061.

§ Research Assistant, Dept. of Computer Science, Virginia Tech, Blacksburg, VA 24061-0106.

¶ Professor & Dept. Head, Dept. of Aerospace & Ocean Eng., Virginia Tech, Blacksburg, VA 24061.

‡ Professor, Dept. of Aero. Eng., Mech. & Eng. Sci., University of Florida, Gainesville, FL 32611-6250.

Professor, Dept. of Aerospace & Ocean Eng., Virginia Tech, Blacksburg, VA 24061.

** Professor, Depts. of Computer Sci. and Math., Virginia Tech, Blacksburg, VA 24061-0106.

This series of functions is termed the *weight function* (Hutchison et al., 1992). Since the HSCT is a new class of aircraft, the weight function does not account for all features of the design. In a comparison with finite-element-based structural optimization, the weight function was suitable for predicting general trends in structural weight, but unable to accurately model all effects of changing aircraft geometry (Huang et al., 1994).

Following the variable-complexity modeling approach, information from structural optimization was incorporated into the weight function. This was accomplished by generating a scale factor which multiplied the bending material part of the weight function to match structural optimization results (Dudley et al., 1994). This procedure improved the accuracy of the weight function, but derivatives of the scale factor were not available to the configuration design optimizer, making optimal solutions difficult to find. In addition, the scale factor was only updated periodically during the design due to the cost associated with structural optimization.

In view of the weight function deficiencies, a full integration of the structural and configuration optimization was considered. However, this approach was difficult for several reasons. First, results from the structural optimization did not produce smooth functions with respect to the configuration shape parameters. Therefore, a derivative-based optimization would have been difficult to perform. Second, the configuration design process required structural weight information at a large number of design points. This, coupled with the expense of structural optimization, made an integrated optimization infeasible for HSCT design, where multiple optimizations must be performed.

Coupling unrelated optimization processes also introduced code integration problems. Different optimization software packages utilized different design variables, input parameters, and output formats. Additional software, commonly termed *wrappers*, could have been developed to automate code interaction; however, this is rarely a straightforward process. Moreover, wrappers produce inefficiencies on modern high performance computers. Examples of research in this area are described in Woyak et al., 1995 and Malone et al., 1995.

In this paper, a response surface approximation to the wing bending material weight is used to address the concerns of function smoothness, cost, and code integration, while also improving the accuracy of the statistical weight function. Instead of performing structural optimization during the configuration design process, a large number of aircraft geometries are evaluated beforehand. These results are then used to create a response surface to the wing bending material weight. Since the geometries are based on the configuration design variables, code integration problems are eliminated.

While one desires accurate results throughout the entire design space, it is impractical to perform structural optimization for every conceivable HSCT configuration. For this reason, techniques are developed to limit the design space domain and to balance the response surface accuracy with development cost. These methods are not specific to the wing bending material weight and can be applied to many response surface applications.

Three methods specifically improve the accuracy of the response surface. First, a good model function is found by inspecting various statistical weight functions. Second, inexpensive approximate analysis methods are used to find regions of the design space where reasonable HSCT configurations are likely to appear. Development of the response surface is then limited to this portion of design space. Third, analysis of variance is used to remove unnecessary terms from the model function. The absence of these terms reduces the incidence of large errors in the final response surface.

The cost of generating the response surface is addressed in several ways. Once again, a previously developed statistical weight function is used to identify a small set of geometric and loading parameters which characterize the wing bending material weight objective function. This reduced set of variables is then used in place of the HSCT design variables to reduce the complexity of the model function. Finally, using the D-optimality criterion, a small number of reasonable designs, which produced a response surface that accurately reflect the entire reasonable design space, are chosen. Only at these designs are structural optimizations performed.

The large number of structural optimizations required make this problem especially suitable for coarse grain parallelization. A finite-element-based structural optimization code, GENESIS, is used to optimize the numerous design configurations. GENESIS has been modified for the Intel Paragon and reasonable speedups are achieved.

By its nature, the response surface is a simple algebraic expression that provides smooth derivative information. Although initially expensive to establish, the surface uses minimal resources once implemented. This makes it ideal for the HSCT design problem. To demonstrate the response surfaces' suitability to design, optimizations of the HSCT configurations are performed. Results are presented with and without the implementation of the response surface.

2. RESPONSE SURFACE METHODS

Response surface methodology (RSM) is a statistical technique in which smooth functions, typically polynomials, are used to model an objective function. For example, a quadratic response surface model for p variables has the form:

$$y = c_0 + \sum_{1 \le i \le p} c_i x_i + \sum_{1 \le i \le j \le p} c_{ij} x_i x_j \, , \tag{1}$$

where the x_i are the variables, the c_i are the polynomial coefficients, and y is the measured response. For p variables, (1) has $n = (p+1)(p+2)/2$ terms. In such a model the polynomial coefficients may be estimated using the method of least squares (Appendix A).

Most RSM applications are based on quadratic polynomials. While these polynomials are relatively easy to utilize, some issues must be considered before their implementation. For p variables, the number of terms in the model function, n, grows at a rate $\mathcal{O}(p^2)$. Creating a response surface for the n-term polynomial requires a number of points, m, which is larger than, but of the same order of magnitude as n. This seems reasonable; however, to maintain good accuracy, one needs to minimize estimations beyond the domain where analyses were performed. That is, we would like the points where we use the response surface to reside within the convex hull of the data points used to construct the surface. To ensure this condition for a p-dimensional box requires at least 2^p points (at the vertices of the box). For the twenty-eight design variables which describe the HSCT, this corresponds to more than 200 million points, which is certainly impossible with present computational capabilities. This problem is often referred to as the *curse of dimensionality*. Quadratic polynomials also of course cannot model higher order variations.

Estimations outside the design space where the response surface was created may produce accurate results; however, it is well known that if the model is wrong, estimates only slightly outside the known data can be grossly in error. Several measures can be taken to address model error. Although seldom an option, the number of data points used to create the response surface can be increased so that the convex hull defined by the known data points encloses a larger portion of the design space. Another option is to reduce the volume of the design space modeled by RSM. This will lower the distance between points interior to the convex hull and the boundary of the convex hull. Finally, the complexity of the model function can be reduced by eliminating unnecessary terms. A statistical technique, analysis of variance (ANOVA) enables the less significant terms in the polynomial approximation to be identified. ANOVA is described in Appendix A.

Computational resources limit the number of points from an ideal 2^p to the same order of magnitude as n. For a six-term polynomial in two dimensions, results from Giunta et al., 1994, indicate that $1.5n$ function analyses are sufficient to produce a response surface where the error stabilize and the global trends of the underlying function are accurately approximated.

RSM typically employs a structured method such as central composite design (CCD) for selecting analysis points in the design space (Mason et al., 1989). However, CCD is only effective for a low dimensional regularly shaped design space, which is unlikely to appear in this study. For an irregular design space in many dimensions, there is no simple way of creating a finite number of points that span the entire region. Hence, a very large number of points must be produced knowing that many will fall outside of the feasible design space. The infeasible points can then either be perturbed until they fall within the feasible region, or removed. This process will lead to a large number of points inside the feasible region, but whose geometric distribution is irregular. From these points, a small number must be chosen to construct the response surface. In a previous study (Giunta et al., 1994), it was found that the D-optimal criterion (Box and Draper, 1971) provided a rational means for choosing any number of these points.

Once the response surface is generated, its predictive capabilities must be evaluated. This is accomplished by finding the response surface prediction at a series of data points with known responses. Measuring the difference between the known and the predicted response yields the following information: average error, root mean square (RMS) error, and maximum error. Points used to evaluate the response surface should be different from the points used to create the response surface. If not, the error results may be misleading.

3. HSCT DESIGN PROBLEM

Successful aircraft configuration optimization requires a simple yet meaningful mathematical characterization of the geometry. The model used here completely defines the HSCT design problem using twenty-eight design variables (Hutchison et al., 1993, Hutchison, 1993). Twenty-five of the design variables describe the geometry of the aircraft and can be divided into five categories: wing planform, airfoil shape, tail areas, nacelle placement, and fuselage shape. The wing planform is described using the root and tip chord lengths, the wing span, and by blending linear line segments at the leading and trailing-edges. The airfoil sections have round leading edges and are defined using an analytic description that incorporates four variables. The horizontal and vertical tail areas are described by two variables. The nacelles move axially with the trailing-edge of the wing, and two variables

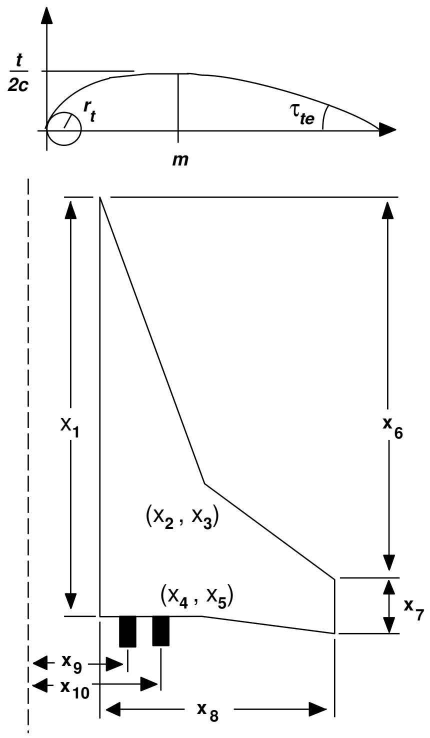

Figure 1. Airfoil thickness distribution parameters (top) and wing planform and nacelle placement parameters (bottom).

define their spanwise locations. The axisymmetric fuselage requires eight variables to specify both the axial positions and radii of the four fuselage restraint locations. Figure 1 and Table 1 define the HSCT geometry and design variables.

In this study, a baseline HSCT is used to provide a point near the interior of the feasible design space. The baseline geometry is from an HSCT configuration previously thought to be optimal by Dudley et al, 1995 (Table 1). Because of modifications and improvements to our analysis methods, this geometry no longer satisfies all of the performance constraints; however, more recent optimal designs exhibit similar characteristics.

The design problem is to minimize the takeoff gross weight of an HSCT configuration with a range of 5,500 nautical miles and a cruise speed of Mach 2.4 while transporting 251 passengers. For this mission, in addition to the geometric parameters, four variables define the idealized cruise mission (Table 1): fuel weight, initial cruise altitude, and the constant climb rate used in the range calculation.

A complete design optimization is composed of a sequence of optimization cycles. Detailed analysis methods are employed at the beginning of each cycle, while simple analyses, scaled to match the initial detailed results, are performed in subsequent calculations during each cycle (Hutchison et al., 1994, Dudley et al., 1995). A typical HSCT design requires approximately twenty-five cycles until an optimal configuration is identified. The optimizer NEWSUMT-A (Grandhi et al., 1985), which employs an extended interior penalty function method, is used for this work. NEWSUMT-A utilizes constraints based on both the simple and detailed analyses along with constraints that limit the movement of the design variables.

Sixty-eight constraints, including geometry, performance, and aerodynamic constraints, prevent the optimizer from creating physically impossible designs (Hutchison, 1993). The

Table 1. Design variables and baseline values.

#	Value	Description
1	181.5	Wing root chord (ft)
2	155.9	LE Break, x (ft)
3	49.2	LE Break, y (ft)
4	181.6	TE Break, x (ft)
5	64.2	TE Break, y (ft)
6	169.5	LE of wing tip, x
7	7.00	Tip chord (ft)
8	75.9	Wing semi-span (ft)
9	0.40	Chordwise location of max. t/c
10	3.69	Airfoil LE radius parameter, r_t
11	2.58	Airfoil t/c at root (%)
12	2.16	Airfoil t/c at LE break (%)
13	1.80	Airfoil t/c at tip (%)
14	2.20	Fuselage restraint 1, x (ft)
15	3.50	Fuselage restraint 1, r (ft)
16	12.2	Fuselage restraint 2, x (ft)
17	3.50	Fuselage restraint 2, r (ft)
18	132.5	Fuselage restraint 3, x (ft)
19	5.34	Fuselage restraint 3, r (ft)
20	248.7	Fuselage restraint 4, x (ft)
21	4.67	Fuselage restraint 4, r (ft)
22	26.2	Nacelle 1, x_9 (ft)
23	33.1	Nacelle 2, x_{10} (ft)
24	322,617	Mission fuel (lbs)
25	64,794	Starting cruise altitude (ft)
26	33.9	Cruise climb rate (ft/min)
27	697.9	Vertical tail area (ft^2)
28	713.0	Horizontal tail area (ft^2)

aerodynamic and performance constraints can only be assessed after a complete analysis of the HSCT design; however, the geometric constraints can be evaluate using algebraic relations based on the twenty-eight design variables. For this reason, they offer a very efficient, albeit rudimentary, means of identifying unreasonable HSCT designs. Reasonable designs are defined to be those that exhibit similar characteristics to feasible designs even though some aerodynamic and performance constraints might be violated. For the most part, however, the geometric constraints only prevent nonsensical configurations such as those where the engine nacelles are not located on the wing or the airfoil chord lengths become negative. All sixty-eight constraints are listed in Table 2.

4. THE WEIGHT FUNCTION

Gross takeoff weight is minimized during the HSCT configuration optimization process. Therefore, weight calculations are numerous and the final HSCT design is highly dependent on the accuracy of these results. To ease the computational expense of the design process, we have implemented the statistical weight function in the weight module of Flight Optimization System (FLOPS) (McCullers, 1984). FLOPS is used to determine takeoff gross

Table 2. Optimization constraints.

#	Description
1	Range $\geq 5,500$
2	Landing angle of attack $\leq 12°$
3	Landing CL ≤ 1.0
4-21	Landing section Cl ≤ 2.0
22	Fuel volume $\leq$ half of wing volume
23-40	Wing chord $\geq$ 7.0ft
41	LE break, $y \leq$ wing semi-span
42	TE break, $y \leq$ wing semi-span
43	Root $t/c \geq 1.5\%$
44	LE break $t/c \geq 1.5\%$
45	TE break $t/c \geq 1.5\%$
46	Fuselage: $x_{rest_1} \geq 5ft$
47	Fuselage: x_{rest_1} + 10ft $\leq x_{rest_2}$
48	Fuselage: x_{rest_2} + 10ft $\leq x_{rest_3}$
49	Fuselage: x_{rest_3} + 10ft $\leq x_{rest_4}$
50	Fuselage: x_{rest_4} + 10ft $\leq$ 300ft
51	Nacelle 1, $y \geq$ side-of-body
52	Nacelle 1, $y \leq$ nacelle 2, y
53	Engine out stability criterion
54	Minimum airfoil section spacing at wing tip
55-56	No engine scrape at landing angle-of-attack
57-58	No engine scrape at landing angle-of-attack, with 5° roll
59	No wing tip scrape at landing angle-of-attack, with 5° roll
60	No wing trailing-edge break point scrape at landing, with 5° roll
61	Crosswind landing capability with aileron deflection only
62	Crosswind landing capability with aileron and rudder deflection
63	Tail deflection $\leq 22.5°$ for landing
64	Takeoff rotation must occur prior to reaching 90% of takeoff velocity
65	No negative sweep of wing inboard trailing-edge
66	Root trailing-edge must not overlap root leading-edge of horizontal tail
67-68	Required engine thrust $\leq$ available thrust

weight and to find the effect of planform geometry changes on structural weight. However, since the HSCT is a new aircraft, the weight function does not account for all features of the design. Furthermore, by its nature, a weight function can only provide an approximate estimate of structural weight. The accuracy of the structural weight predictions from

FLOPS is inadequate for HSCT design. This work focuses on the wing bending material weight since most of the load-dependent wing weight is due to bending.

The general wing weight function in FLOPS is based on an analytic expression to relate wing bending material weight to wing geometry, material properties, and loading. Other terms are added to account for shear material, control surfaces, etc. In addition, constants are included to correlate with a wide range of existing transports and to reflect features such as composite materials, strut braced wings, etc. The wing weight W_w used within FLOPS is given as

$$W_w = \frac{W_g K_e W_b + W_s + W_n}{1 + W_b}, \tag{2}$$

where

$$\begin{aligned} W_s &= 0.68\,(1 - 0.17 f_c)\,(S - S_b)^{0.34}\, W_g^{0.6}, \\ W_n &= 0.35\,(1 - 0.3 f_c)\, S^{1.5}, \\ W_b &= K f_{ul} b\,(1 - 0.4 f_c)\,(1 - .01 f_a), \\ K &= 8.8 B_z \left(1 + (6.25/b)^{0.5}\right) \times 10^{-6}, \\ K_e &= 1.0 - (B_{ze}/B_e)\,(W_{pod}/W_g), \end{aligned}$$

and W_g is the gross takeoff weight (lbs), W_b is the wing bending material weight (lbs), W_s is the wing shear material and flaps weight (lbs), W_n is the wing control surfaces and non-structural weight (lbs), W_b is the wing bending material weight (lbs), W_{pod} is the engine pod weight (lbs), b is the wing span (ft), B_z is the bending material factor, B_{ze} is the engine relief factor, f_{ul} is the ultimate load factor, f_a is the composite material factor, f_c is the aeroelastic tailoring factor, S is the wing area (ft^2), and S_b is the wing box area (ft^2).

The system is closed except for the bending material factor B_z and the engine relief factor B_{ze}. The parameter B_z accounts for the distribution of load on the wing and is calculated by approximately determining the required material volume of the upper and lower skins in a simple wing box description of the wing. The parameter B_{ze} accounts for the reduced amount of structural weight necessary due to the presence of the engines on the wing. Analytic expressions for B_z and B_{ze} are given in Appendix B. The FLOPS wing weight calculation is an iterative process since the wing weight and the gross takeoff weight are dependent on one another.

A close inspection of the FLOPS weight function reveals that, for a design such as the HSCT, the wing weight is entirely based on a set of ten parameters, listed in Table 3. Each of these parameters can be found using the twenty-eight HSCT design variables.

5. RESPONSE SURFACE APPROXIMATION

An effective way to improve the estimate for the wing bending material weight given by FLOPS appears to be finite-element-based structural optimization. For this reason, an integration of the structural and configuration optimization was considered. However, problems with function smoothness, code integration, and limited computational resources prevented a combination of the two optimization processes. Instead, response surface methodologies are implemented to model the wing bending material weight calculated through structural optimization.

In the development of the response surface, a large number of data points spanning the design space are required; however, limits are imposed by the expense of performing structural optimization. For this reason, the bending material weight predicted by FLOPS is used to develop the response surface. Once this is complete, a sequence of structural optimizations is performed to generate the final form of the response surface.

Table 3. Basic parameters used to calculate wing weight in FLOPS.

#	Name	Description
1	S_{ht}	Horizontal tail surface area
2	S_{vt}	Vertical tail surface area
3	$wfuse$	Maximum fuselage diameter
4	b	Wing span
5	$sweep$	Average 1/4 chord sweep angle
7	B_z	Bending material factor
6	B_{ze}	Engine relief factor
8	S_w	Wing surface area
9	W_{fuel}	Weight of fuel at takeoff
10	W_{to}	FLOPS estimated gross weight

5.1 The Model Function

The basic form of statistical weight functions such as those defined within FLOPS involve relationships such as

$$y = C x_1^{c_1} x_2^{c_2} \cdots x_p^{c_p} , \tag{3}$$

where y is an intermediate variable and the x_i are the variables. This suggests a model equation of the form

$$\ln(y) = c_0 + \sum_{i=1}^{p} c_i \ln(x_i) + \sum_{1 \le i \le j \le p} c_{ij} \ln(x_i) \ln(x_j) \ . \tag{4}$$

Essentially, the logarithm of the response, y, is a quadratic function in the logarithms of the variables x_i. The b coefficients are unknown values which must be estimated through the method of least squares (Appendix A). For comparison, model functions of both forms (1) and (4) are used to develop the response surfaces.

5.2 Identification of the Reasonable Design Space

Three of the twenty-eight HSCT design variables — wing leading edge radius, cruise climb rate, and starting cruise altitude — have no effect on the FLOPS estimate for wing bending material weight, and so are ignored. The enormous design space associated with twenty-five variables, as well as the expense of performing structural optimization, forces one to limit the response surface to reasonable regions of the space. Again, a reasonable design refers to one whose characteristics are similar to a feasible design even though some of the aerodynamic and performance constraints might be violated.

The first step in identifying the reasonable design space is to construct a suitably large hypercube, defined by the twenty-five design variables (Table 1), that encompasses this entire region of space. Each of the variables, except the fuel weight, is allowed to assume values between 20% and 180% of its baseline value, given in Table 1. The fuel weight is only allowed to vary between 75% and 125% of its baseline value because of its strong influence on the design's range and therefore feasibility. Using the technique described in Appendix C, 19,651 configurations are found on the boundary of the domain. Of these designs, 83% violate one or more of the HSCT's geometric constraints (Table 2) and a large portion of the remaining designs appear to be unreasonable.

Table 4. Criteria for reasonable designs.

#	Description
1-34	HSCT geometric constraints (Table 2)
35-36	20,000 lbs $< W_{b_F} <$ 120,000 lbs
37-58	Minimum fuselage radius
59	Inboard $\Lambda_{le} >$ Outboard Λ_{le}
60	$\Lambda_{le} > 0$
61-62	5,000 ft$^2 < S_w < 15,000$ ft^2
63-64	$1.0 < AR < 3.2$
65	Inboard $\Lambda_{te} < 40°$
65-83	$c_{y_{i+1}}/c_{y_i} < 1.0$
84	Approximate range > 5,000 n. mi.

Eliminating designs that are unreasonable cannot be accomplished without removing nearly every design in the pool of 19,651 candidate points. For this reason, each unreasonable design, $\mathbf{x}$ is moved so that it resides on the edge of the reasonable design space:

$$\mathbf{x}' = \alpha\left(\mathbf{x} - \mathbf{x}_c\right) + \mathbf{x}_c\,. \tag{5}$$

The parameter α, $0 \leq \alpha \leq 1$, is found using bisection. Computing α in (5) requires a set of criteria to determine whether a design is reasonable or not. These criteria must be selected carefully to avoid a computationally expensive procedure and to ensure that no reasonable designs are inadvertently removed. In order to make use of complex constraints, a series of increasingly expensive evaluations are defined and applied in phases. Initially, the simple criteria are applied to the data and a large percentage of the candidate points are moved toward x_c. However, as the increasingly complex constraints are applied, fewer of the points have to be moved and the expense of the constraint evaluations does not become prohibitive.

Table 4 lists the criteria used to move the data towards the reasonable design space. They are listed in order of application, with the range constraint, which is the most expensive, coming last. During the first phase, the thirty-four geometric constraints listed in Table 2 are applied, resulting in the movement of 16,297 designs. Although these criteria do exclude geometrically impossible designs, they do not preclude unreasonable geometries. Originally, these constraints were intended to work in conjunction with the aerodynamic and performance constraints, which are not considered in this phase. Two unreasonable aircraft planforms that conform to the thirty-four constraints are depicted in Figure 2.

During the next phase of moving points towards the reasonable design space, the bending material weight W_{b_F} is computed by FLOPS. Based on past experience with the weight function, reasonable designs only occur when the FLOPS estimate for wing bending material weight W_{b_F} is between 20,000 lbs and 120,000 lbs, so designs falling outside that range are moved closer to the baseline design. This phase moves a total of 1,210 designs.

To further reduce the design space several new criteria are formulated to address some of the features illustrated in Figure 2. The minimum fuselage radius is set to 4.0 ft along the entire length of the wing root chord. In addition, the fuselage must enclose a cone with its base, having a radius of 4.0 ft, at the leading edge of the root chord and its apex at the nose of the aircraft. A similar cone has to fit within the fuselage aft of the root chord's trailing-edge.

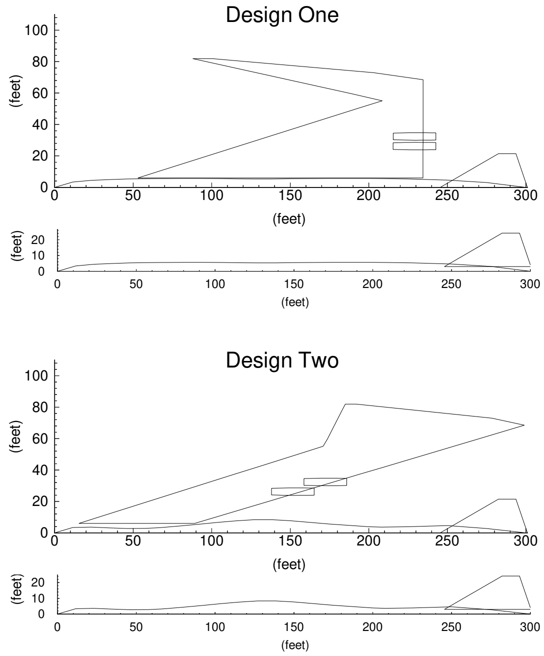

Figure 2. Sample planforms present after the exclusion of infeasible geometries.

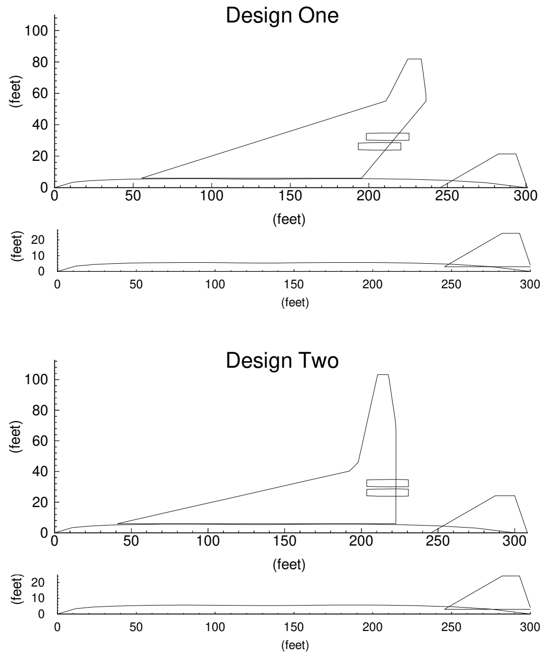

Figure 3. Sample planforms present after the exclusion of unreasonable geometries.

The majority of the new criteria, however, are based on the wing planform. Limits are placed on both the minimum and maximum allowable aspect ratio, AR, and wing planform area, S_w. Leading and trailing-edge sweep angles, Λ_{le} and Λ_{te}, are prevented

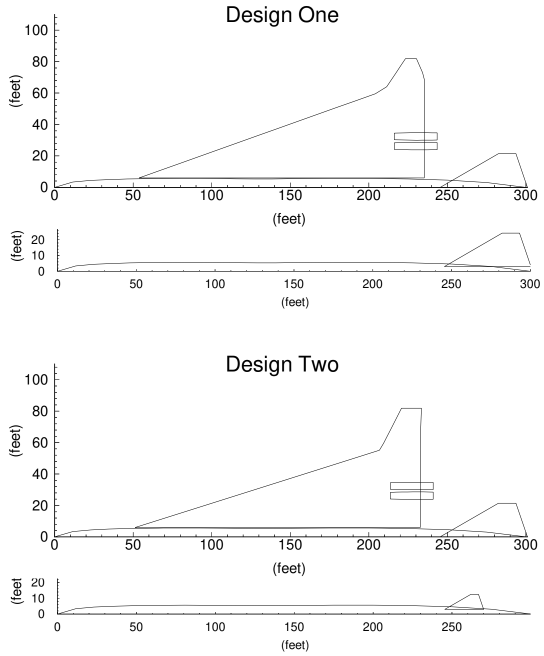

Figure 4. Sample planforms present after the exclusion of designs with insufficient range.

from assuming improbable values. Local taper ratios, $c_{y_{i+1}}/c_{y_i}$, cannot be greater than one. Finally, the wing planform is further constrained from having a forward swept leading edge, or from having an outboard leading edge sweep greater than the inboard leading edge sweep. Applying the new criteria causes 6,473 points to be moved towards the baseline design. Figure 3 shows two designs which appear after the implementation of these criteria.

As a final step, a criterion is added based on aerodynamic analysis to exclude unreasonable designs that cannot be identified with geometric conditions. The approximate range is evaluated for each design using simple methods to estimate total drag on the aircraft (Dudley et al., 1995). Because of the nature of this calculation, designs having an approximate range greater than 5,000 n mi. are considered reasonable. Recall that that final HSCT design must attain a range of at least 5,500 n mi. This criterion moves 4,336 designs towards the baseline configuration. Figure 4 shows two designs which appear after the application of this final criterion.

After each new criteria is applied, a sample response surface is constructed based on the wing bending material weight estimate given by FLOPS. The response surface is created using all 19,651 designs and the associated error is measured using the same designs (Table 5). As the design space shrinks, the accuracy of the response surface improves significantly. In addition, the advantage of using the model with logarithms of the variables versus a polynomial model in the x_i diminishes.

5.3 Selection of Intervening Variables

Using the initial twenty-five variables and the 19,651 design points found after reducing the design space, a response surface model is fit to wing bending material weight estimates from FLOPS. The errors are calculated from the difference between the response surface prediction for the wing bending material weight and the value predicted by FLOPS.

Table 5. Reduced design space response surface errors based on FLOPS weight prediction.

Model Variables	Ave. Err. (%)	RMS Err. (%)	Max. Err. (%)
Exclude Geometrically Infeasible Designs			
x_i	13.22	49.49	5,532
$\ln(x_i)$	2.458	6.762	459.1
Exclude Geometrically Unreasonable Designs			
x_i	0.5249	1.009	16.97
$\ln(x_i)$	0.3624	0.6358	14.73
Exclude Designs with Insufficient Range			
x_i	0.3541	0.6153	11.53
$\ln(x_i)$	0.2727	0.4685	14.73

Table 6. Accuracy of various models based on FLOPS weight prediction.

Model Variables	Ave. Err. (%)	RMS Err. (%)	Max. Err. (%)
Linear in 25 Design Variables			
x_i	3.032	4.868	67.10
$\ln(x_i)$	1.963	3.529	56.59
Quadratic in 25 Design Variables			
x_i	0.3541	0.6153	11.53
$\ln(x_i)$	0.2727	0.4685	14.73
Linear in 10 Intervening Variables			
x_i	0.1895	0.2932	2.232
$\ln(x_i)$	0.0781	0.1391	2.253
Quadratic in 10 Intervening Variables			
x_i	0.0025	0.0056	0.1525
$\ln(x_i)$	0.0062	0.0124	0.5200

Although the results given in Table 6 are quite good, the number of terms in the response surface function is 351. Recall that at least $1.5n$ function evaluations are required to produce an accurate response surface. Therefore, the number of structural optimizations necessary is at least 527 and perhaps several times larger.

A significant reduction in the size of the model function can be obtained by replacing the twenty-five design variables with a smaller set of variables that are more appropriate for a weight analysis. To maintain compatibility with prior work, these variables are also entirely dependent on the design variables. Such variables are called *intervening variables* in the structural optimization literature and the terminology is used here.

Returning to the description of the FLOPS weight function, the wing weight is based on a set of ten basic parameters, which are listed in Table 3. Each parameter can be found using the twenty-five HSCT design variables that contribute to the bending material weight. Using these parameters as the intervening variables reduces the number of terms in the model function from 351 to 66. In addition, the accuracy of the model function actually increases when the intervening variables are introduced. A further improvement in accuracy is realized when (1) is used for the model function in place of (4). Table 6 summarizes these results. Note that a linear approximation in the intervening variables is more accurate than a quadratic approximation in the original variables.

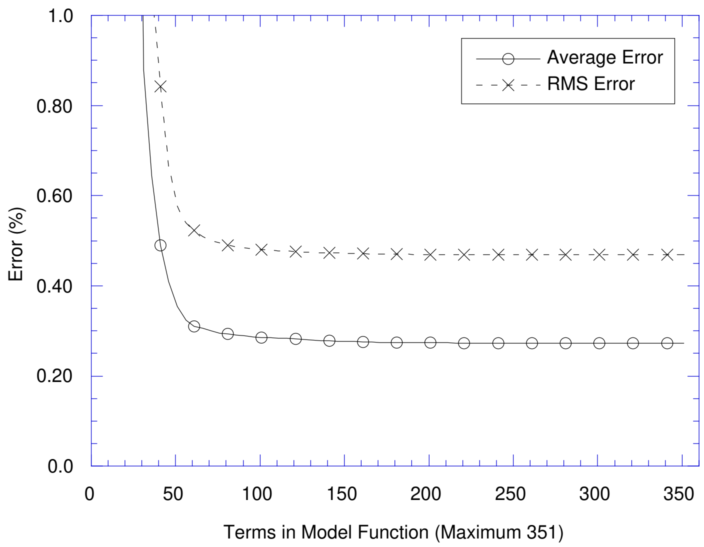

Figure 5. RS error versus terms remaining in design variable model function (FLOPS data).

5.4 Regression Analysis and ANOVA

Continuing with the 19,651 HSCT designs and the FLOPS estimates for the wing bending material weight, regression analysis and analysis of variance (ANOVA) methods are used to identify unnecessary terms in each model function being considered. The intent is not to remove terms from the model functions before performing structural optimization, but simply to study the effect of removing terms. Removing terms based the FLOPS weight estimates may eliminate terms that, while not important for representing the FLOPS estimates, may be influential in modeling the structural optimization data.

Regression analysis and ANOVA are carried out separately for the model functions based on both the twenty-five design variables and the ten intervening variables. The term with the highest coefficient of variation (see Appendix A) is then removed from the model function, and the process of regression analysis and ANOVA is repeated. This sequence of operations is known as one-step-backward elimination (Myers, 1995).

Figures 5 and 6 show the accuracy of the response surface compared to the number of terms remaining in the model function. Note that (4) is used for the twenty-five design variable case, while (1) is used for the ten intervening variable case. 250 terms can be removed from the design variable model function and 26 terms can be removed from the intervening variable model function without significantly affecting accuracy. A similar situation can be expected for the structural optimization data, albeit with different model terms removed.

6. STRUCTURAL OPTIMIZATION

Within the HSCT configuration design process, estimation of wing structural weight using finite-element-based structural optimization has proven to be more accurate than estimates produced by typical weight functions. For this reason, a response surface representation of structural optimization results is developed to replace FLOPS estimates for wing bending material weight.

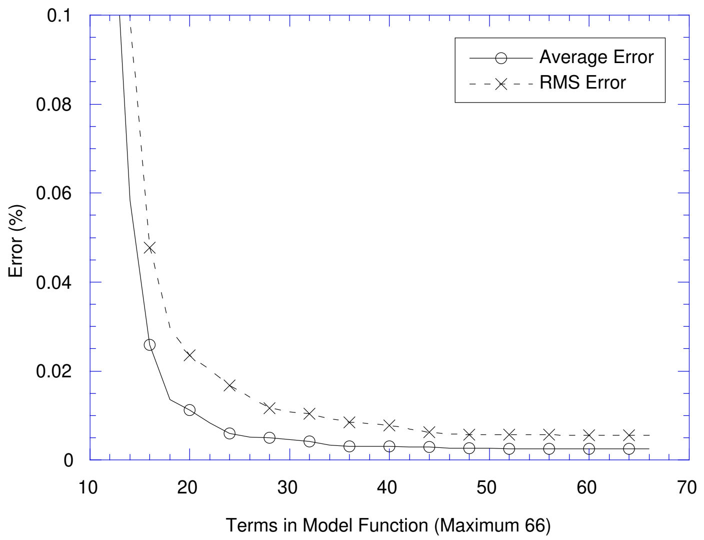

Figure 6. RS error versus terms remaining in intervening variable model function (FLOPS data).

Due to the large number of optimizations that must be performed to create a response surface, a relatively simple structural optimization model is used. For each HSCT design, a fixed arrangement of spars and ribs is generated prior to the optimization while the skin panel thicknesses, spar areas, and rib cap areas appear as design variables. The aircraft is assumed to be built of titanium. Constraints are applied based on Von Mises stress values for each panel, spar, and rib cap element. In addition, local buckling constraints are applied. Even though bending material weight is used for the response surface, the entire wing structural weight is optimized.

6.1 Structural Optimization Model

Because of the large number of designs which are optimized, a special mesh generator is implemented to automatically create a finite element model based on the twenty-eight HSCT design variables (Table 1). In addition to design variables, the number of frames in the fuselage, the number of spars and ribs in the wing, and the chord fractions taken by the leading and trailing-edge control surfaces are specified for each design. The mesh generator creates the finite element nodes and element topology data, estimates the location of nonstructural weights, and predicts the geometry of the wing fuel tanks (Huang, 1994). Fuel is assumed to be stored in thirty-one tanks throughout the aircraft.

Because of symmetry, only half of the aircraft is modeled. A typical finite element model is made up of 963 elements joined at 193 nodes with 1032 total degrees of freedom. The wing and fuselage skin are modeled by membrane elements, the spar and rib cap elements are modeled by rod elements, and the spar and rib webs are modeled by vertical rods. Initial values for the structural optimization design variables are produced using weight estimates provided by FLOPS.

The loads applied to the structural model are composed of the aerodynamic and inertia forces. Inertia loads represent the combined effects of non-structural items, fuel weight, and the distributed weight of the structure. Aerodynamic loads for supersonic flight conditions are determined using a supersonic panel method, and loads for subsonic flight conditions are

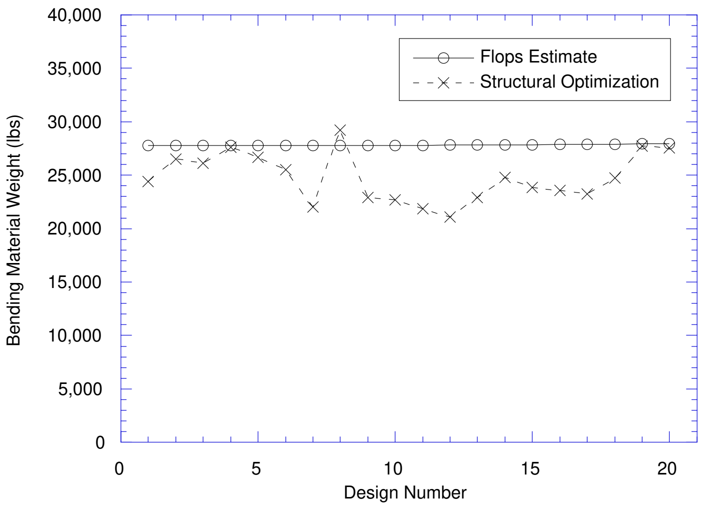

Figure 7. Comparison of FLOPS estimates with structural optimization results.

from a vortex-lattice method. The structure is assumed to be rigid for the determination of aerodynamic forces. Previous studies indicated that structural flexibility did not have a large effect on the loads for this particular configuration (Huang et al., 1994, Barthelemy et al., 1994). For each design, orientation of the aerodynamic loads is governed by camber distributions generated by Carlson's program WINGDES (Carlson and Walkey, 1984). A surface spline interpolation method is used to translate forces between aerodynamic node and structural node locations. More details about the five load cases used can be found in Hutchison et al., 1992.

While acceptable for most HSCT configurations, in some cases the simplified finite element model coupled with a limited number of loading conditions can lead to poor structural weight predictions. For this reason, structural optimization results which differ significantly from FLOPS must be regarded with caution.

6.2 Noise in Structural Optimization

Structural optimization is more accurate than weight function estimates; however, they do not produce smooth functions with respect to the design variables. This is one of several reasons why structural optimization is not integrated directly into the overall configuration optimization of the HSCT. The non-smooth characteristics of structural optimization can be attributed to numerical noise associated with various parts of the optimization process as well as the irregularly shaped design space from which the optimal structure is found.

In this study, output from structural optimization is used as the observed function values for a response surface in wing bending material weight. Thus, understanding the nature of any noise produced during the structural design is critical to evaluating the performance of the response surface. To this end, twenty feasible HSCT designs are found along a line segment centered at the baseline configuration (Table 1). Figure 7 depicts both the FLOPS weight estimate and those produced by structural optimization.

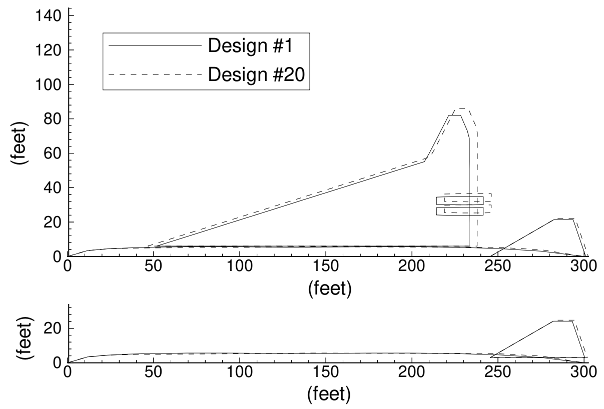

Figure 8. Comparison of extremum designs used for noise analysis.

Figure 8 shows the two extreme designs at the ends of the line segment. All twenty designs are nearly identical in planform shape, and FLOPS predicts that the wing bending material weight varies smoothly from 27,772 lbs to 27,962 lbs, a difference of only 190 lbs. While the structural optimization results are fairly smooth along portions of the line, the weight varies between 21,062 lbs and 27,712 lbs with a spike of almost 30% at the eighth design point.

Several avenues are considered to reduce the noise in the structural optimization. First, parameters which govern the optimizer are altered to promote convergence to an optimum. Initial move limits are reduced from 30% to 10% and the stopping convergence criterion is changed from 1% to 0.1%. Although this has a favorable impact on the noise, it is negligible compared to the large fluctuation at the eighth design point. Attention is next turned towards the objective function of the optimizer. Although wing bending material weight values are used for the response surface, the optimizer uses the weight of the entire wing as its objective function. For this reason, changes are made to the finite element model so that the optimizer manipulates only those portions of the wing which were accounted for in the bending material weight. Together, the changes to the optimizer and and the finite element model reduce the spike in Figure 8 from 30% to 27%.

From the aforementioned results, it appears that the noise is not produced within the structural optimization process. Thus, attention is shifted to the loading data produced by the aerodynamic analysis. Optimizing each design using identical load data completely eliminates the spike at the eighth design point. Further, an analysis of the center of pressure locations based on the loading data for each design reveals a pattern of noise similar to that found in Figure 7. Thus, the most significant noise source is not the structural optimization, but instead the aerodynamic analysis.

Because of time limitations associated with this study, we are unable to investigate the source of the aerodynamic noise. Therefore, we expect the structural optimization results to exhibit properties similar to those depicted in Figure 7. Considering the distribution of data, it is clear that a quadratic polynomial can model the FLOPS data with little or no error. In contrast, errors of as much as 27% can be anticipated if a similar polynomial is used to model the structural optimization data.

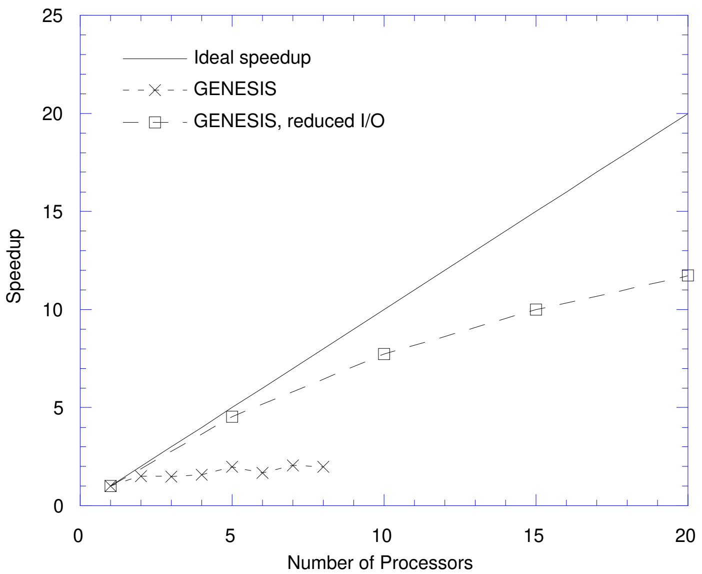

Figure 9. Ideal versus actual speedup for parallel execution of the GENESIS code.

6.3 Parallelized Structural Optimization

Because of the numerous structural optimizations required for the response surface, this problem is well suited for coarse grain parallel computation. Coarse grain parallelization implies that structural optimizations for different HSCT designs are performed simultaneously on separate processors. Each processor maintains its own data, so that interaction between the processors is not necessary.

As the number of processors is increased on a distributed memory architecture machine like the Intel Paragon, disk I/O may become a bottleneck limiting the efficiency of the parallel computations. This fact is the primary reason for choosing the GENESIS structural optimization code for implementation in the parallel environment. GENESIS, a code developed and supported by Vanderplaats, Miura and Associates, Inc (Vanderplaats et al., 1993), is available from the developer in a reduced I/O form making it an effective code to use on the Paragon.

The speedups of a parallel computation is defined as T_s/T_p where T_s is the serial execution time and T_p is the parallel execution time using p processors. In an ideal situation, speedup would be equal to the number p of processors being used. Figure 9 shows the benefits of the reduced I/O version of GENESIS. With the standard version, maximum speedup levels off at 2.3, regardless of the number of processors, while the reduced I/O version achieves a speedup of 11.7 using 20 processors. This is still rather poor, showing how ill-suited packages developed for serial computation, like GENESIS, are for parallel computation.

7. RSM for Structural Optimization

This section concerns the creation of a response surface for the wing bending material weight based on data from GENESIS. As discussed in the last section, there is significant numerical noise within the structural optimization process, perhaps as much as 27%. It is thus difficult to separate noise in the optimized wing bending material weights from fitting error, and thus the *apparent error* in the response surface approximation must be interpreted properly. Indeed, a substantial part of the error may represent a desirable smoothing of the results from the structural optimization.

With the FLOPS weight function, all 19,651 configurations could be used to construct a response surface. However, the expense of structural optimization limits the number of designs which can be considered to a small subset of the 19,651 points. D-optimality (Box and Draper, 1971) provides a rational criterion for choosing this subset. For an n-term model function, the cost of assessing the merit of each candidate design is $\mathcal{O}(n^3)$. Calculating the merit function for the 351-term polynomial in the twenty-five design variables at 19,651 points is computationally prohibitive.

Candidate HSCT configurations are chosen by picking 3,000 points (mostly at random, but also containing previously analyzed designs) from the 19,651 points; the condition number of the least squares matrix $\mathbf{X}$, cond $\mathbf{X}$, (Appendix A) for these 3,000 points is below 10^4.

D-optimal point sets are typically computed by simple exchange algorithms or genetic algorithms. Both classes of algorithms are totally overwhelmed by a 351-term model function and 3,000 points. An ad hoc approach is to decompose the polynomial model into a sum of polynomials with fewer terms, find a D-optimal set for each summand, and then take the union of all these D-optimal sets. Not just any decomposition will yield a well conditioned $\mathbf{X}$ matrix for the union (in particular a decomposition into 7 disjoint sets of summands does not work). Using seven summands, each of which contained all the linear terms and 50 quadratic terms (25 quadratic terms in the seventh summand), and finding 150 D-optimal points for each summand (125 D-optimal points for the last one), yields 1,025 designs, for which cond $\mathbf{X} = 5.7 \times 10^3$.

Structural optimization with GENESIS (Vanderplaats et al., 1993) are performed for each of the 1,025 designs to find the optimum wing bending material weight. 991 of the optimizations are successful and cond $\mathbf{X} = 5.7 \times 10^3$. Using the same data points and the 66-term model function in the intervening variables, cond $\mathbf{X} = 5.9 \times 10^4$. These 991 points, referred to as Set A, are used to create the response surfaces for bending material weight.

A similar procedure is used to identify another set of 971 points (Set B), used for checking the accuracy of the response surface created by Set A. Taking results from both Sets A and B, the ratios of structural optimization values to values predicted by FLOPS are in the range of 0.51 to 2.05, and the average value is 0.89.

Using every design in Set A, two response surfaces are created, one using the 66-term model function in the intervening variables and the other using the 351-term model function in the design variables. To determine what portion of Set A is necessary to accurately model the entire design space, a number of D-optimal sets are found from within that set. For each D-optimal subset of Set A, a response surface is created using the associated model function and then evaluated using Set B. The relationship between the error and the number of points used to create the response surface is depicted in Figure 10. From the figure it appears that the intervening variables require 300–400 points for the average error to stabilize, while for the design variables, 1,000 points may not be enough.

Returning to the ratio of the number of analysis points to the number of terms in the model function discussed earlier, for the full model function in the intervening variables this ratio appears to be $300/66 \approx 4.5$. The ratio for the full design variable model function appears to be at least $991/351 \approx 2.8$ and most likely much higher. Recall that a smaller value, 1.5, was predicted using the six-term polynomial in two dimensions.

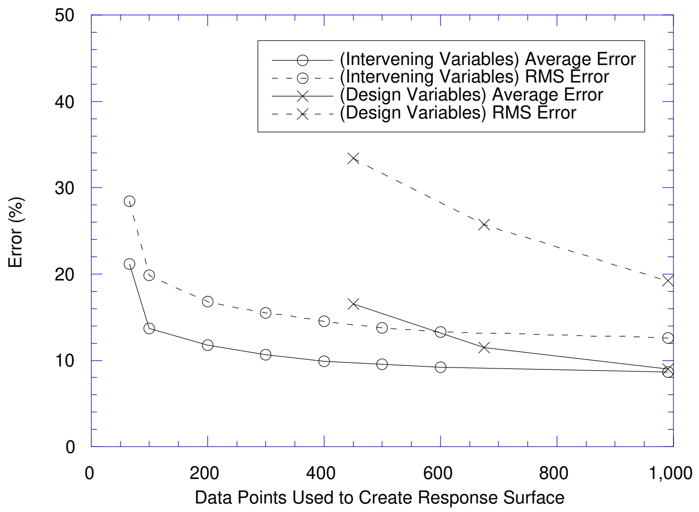

Figure 10. Apparent accuracy of full RS in both design and intervening variables.

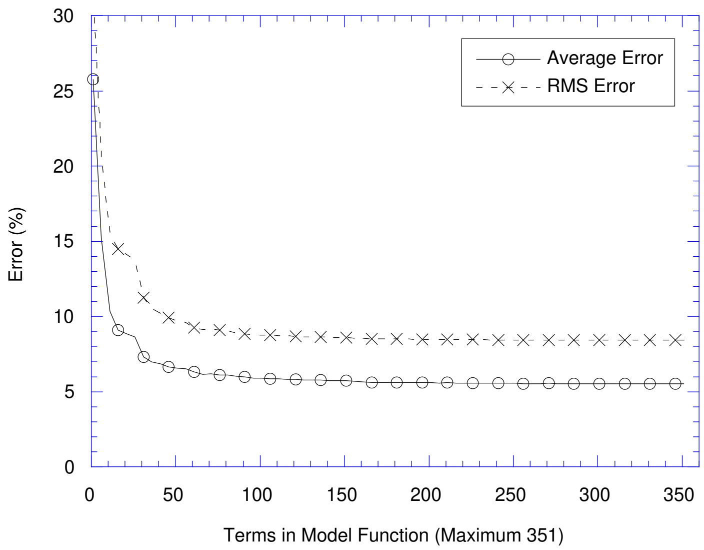

Figure 11. RS apparent error versus terms in design variable model function. (GENESIS Data)

ANOVA and regression analysis with data from Set A are used to remove unnecessary terms from the 66-term model in the ten intervening variables and the 351-term model in the twenty-five design variables. Results are plotted in Figures 11 and 12, which are similar to those found using the FLOPS estimates for wing bending material weight (Figures 5 and 6). The 66-term intervening variable model function can be reduced to 15 terms and the 351-term design variable model function to 61 terms without adversely affecting the approximation.

These reduced term polynomials are now used to determine the number of points required to produce a response surface that accurately reflects wing bending material weight over the entire design space. Referring to Figure 13, for each number k of data points and both reduced term models, a k point D-optimal subset of Set A is found and used to compute a response surface, whose quality is measured over Set B. For both the intervening

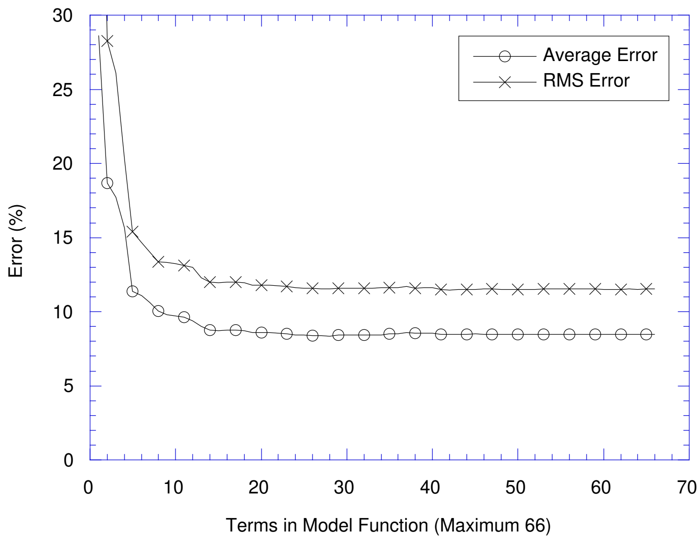

Figure 12. RS apparent error versus terms in intervening variable model function. (GENESIS Data)

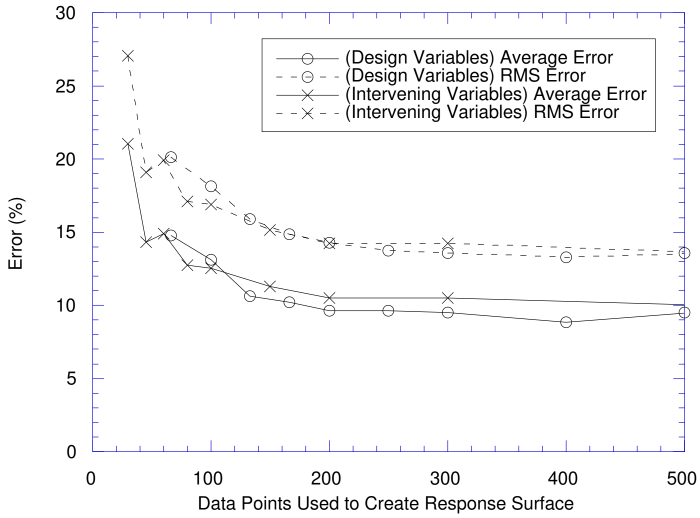

Figure 13. Apparent accuracy of reduced RS in both design and intervening variables.

variables and the design variables, Figure 13 indicates that 200–250 points are needed for the RMS error to stabilize. For the design variables, this translates into a ratio of the number of analysis points to the number of terms in the model function of $200/61 \approx 3.3$, and $200/15 \approx 13.3$ for the intervening variables.

Numerical results are given in Table 7 for response surfaces created using all of Set A and evaluated using all of Set B. The response surface created using the 61-term polynomial in the design variables will henceforth be referred to as RS1, and RS2 refers to the 15-term polynomial in the intervening variables. RS3 is the response surface using

$$y = c_0 + c_1 W_{b_F} \tag{6}$$

as the model, essentially, a linear function in the FLOPS prediction of the bending material weight. Comparing results from Tables 6 and 7 highlights the decrease in response surface accuracy accompanying a move from the FLOPS data to the GENESIS data. Focusing

Table 7. Apparent accuracy of RS based on GENESIS weight predictions.

Name	Ave. Err. (%)	RMS Err. (%)	Max. Err. (%)
RS1	7.72	12.78	114.58
RS2	8.84	12.11	63.67
RS3	13.84	20.49	72.13

on the intervening variables, one could postulate that the variables are only suited to model the FLOPS estimates. Indeed, the ten intervening variables are straight out of the equations used within FLOPS. It is possible that terms such as B_{ze} and B_z (Appendix B) are inappropriate for modeling the structural optimization results. Still, it is reasonable to assume that the twenty-five design variables can model either the FLOPS estimates or the structural optimization results with similar accuracy. This is not the case. In fact, a strong case can be made for using the intervening variables over the design variables. Both the RMS and maximum errors associated with the intervening variables are less than those associated with the other two response surfaces.

There are two plausible explanations for the results presented in Tables 6 and 7: (1) a quadratic model is inadequate for the wing bending material weight over the entire domain being considered here, or (2) the GENESIS optima are incorrect, due to noisy aerodynamic loads or some other source of noise. Section 6 offers evidence in support of (2), but does not rule out (1). Possibly both (1) and (2) occur.

8. HSCT DESIGN OPTIMIZATION

Complete HSCT design optimizations are performed to evaluate the response surfaces to wing bending material weight, RS1 and RS2. At the completion of the optimizations, results are compared with with structural optimization results W_{b_G}. Implementation of each response surface is accomplished by modifying the gross takeoff weight calculations within the weight module of FLOPS. In place of FLOPS estimates for wing bending material weight W_{b_F}, response surface predictions $W_{b_{RS1}}$ and $W_{b_{RS2}}$ are used. RS1 and RS2 are intended for use in the reasonable design space and their predictions cannot be relied upon outside this region. Therefore, all calculations outside the reasonable design space are done with W_{b_F}.

With two exceptions, the criteria listed in Table 4 are used to identify the reasonable design space. First, the range criterion is not evaluated because of its relatively high computational expense. Second, an additional criterion is added to exclude unreasonable predictions by the response surface. Similar to the criterion on W_{b_F}, if the response surface value falls outside the range 20,000 lbs $< W_b <$ 120,000 lbs, the design is considered unreasonable.

At the edge of the reasonable design space, a smoothing function is used to prevent noise associated with the jump between W_{b_F} and the response surface predictions. The smoothing function

$$W_b = \begin{cases} W_{b_{RS}}, & r_{max} \leq 0, \\ f(W_{b_{RS}}, W_{b_F}), & 0.10 > r_{max} > 0, \\ W_{b_G}, & r_{max} \geq 0.10, \end{cases} \tag{7}$$

where

$$f(a,b) = a\left[1 - \sin^2\left(5 r_{max}\pi\right)\right] + b \sin^2\left(5 r_{max}\pi\right) . \tag{8}$$

Table 8. Comparison of HSCT optimal designs.

Parameter	FLOPS	RS1	RS2
Planform Geometry			
Root chord (ft)	153.6	148.5	157.9
Tip chord (ft)	7.66	8.08	8.01
LE sweep (°)	73.23	69.78	73.01
Aspect Ratio	1.71	1.85	1.46
Wing Area (ft^2)	10,178	11,046	10,275
Performance Data			
Range (n. mi.)	5,534	5,500	5,520
Landing AOA (°)	11.93	11.24	12.00
Weight Data			
W_g (lbs)	672,939	662,865	634,758
W_w (lbs)	91,102	85,770	81,298
W_{b_F} (lbs)	30,502	34,425	22,901
$W_{b_{RS1}}$ (lbs)	81,614	20,010	62,391
$W_{b_{RS2}}$ (lbs)	31,812	35,403	20,955
W_{b_G} (lbs)	29,544	24,097	17,633

is based on the maximum violation r_{max} among the aforementioned criteria.

Three HSCT configuration optimizations are performed, each starting from the baseline design detailed in Table 1. During the first optimization, the FLOPS wing bending material weight is used to find the takeoff gross weight. The next two optimizations are performed using RS1 and RS2 respectively, in conjunction with (7). Results from these optimizations are given in Table 8 and the planforms are plotted in Figures 14, 15, and 16.

From Table 8 it is seen that the best design is obtained using RS2 (the takeoff gross weight W_g is substantially lower than those obtained using FLOPS or RS1). As expected, for all three designs, the bending material weight predicted by the response surface used for the optimization is lower than the values predicted by the other two response surfaces (here the word 'response surface' is also used for the FLOPS function). That is, the optimizer drifts into a region where the particular response surface is more optimistic. However, for both FLOPS and RS2 the structural optimization results are lighter than the response surfaces, while for RS1 the structural optimization weight is 20% heavier. This indicates that the optimizer took more advantage of RS1's weakness than of any weaknesses in the other two.

Additionally, while FLOPS and RS2 are in close agreement on all three designs, the results of RS1 are wildly different, providing another indication of its weakness. It should be noted, however, that the optimization based on RS1 managed to find a planform with the smallest wing bending material weight.

If we look at the optimum design found by each method and compare the results there with GENESIS, then the agreement is quite good as compared to the agreement in other regions of the design space. For FLOPS this may reflect the fact that its weight function was created based on the weight of actual wings and the structural optimization weight of good aerodynamic designs. It is not surprising then that it agrees better with GENESIS for good designs than for poor ones. For RS1 and RS2 a similar conclusion may be drawn in that the bulk of the points used to create these response surfaces represent reasonable designs. However, the poorer results obtained with RS1 may be attributed to one or a combination

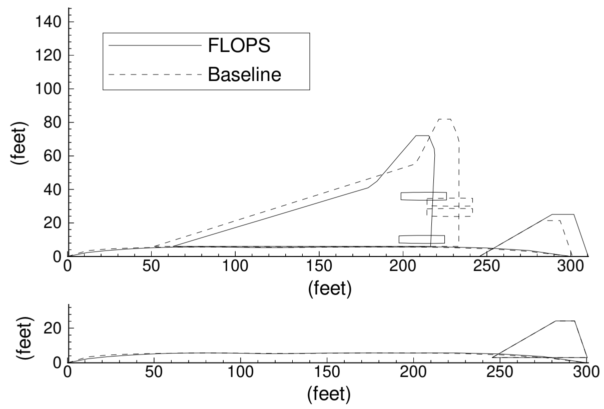

Figure 14. Optimal HSCT planform using FLOPS compared to baseline configuration.

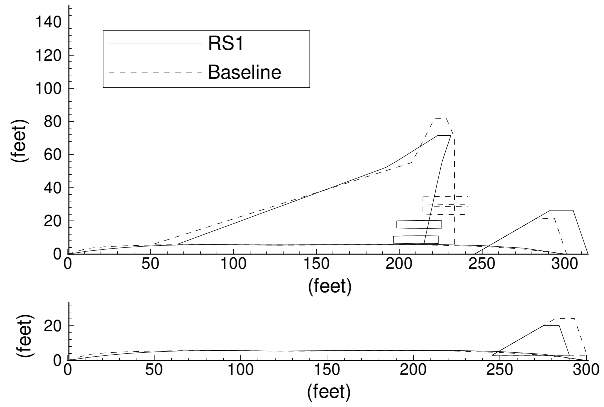

Figure 15. Optimal HSCT planforms using RS1 compared to baseline configuration.

of the following two problems. First, a quadratic model in the original variables may be a poor model for the weight, a conclusion supported by the maximum errors in Tables 6 and 7. Second, the larger number of terms used by RS1 may have resulted in more opportunity to fit some of the noise in the structural optimization and to produce poor predictions even for good designs.

9. CONTINUING WORK

The sources of inaccuracy in the response surface approximations for wing bending material weight need to be positively identified and dealt with. Alternatives to the D-optimal point selection process provide a means for minimizing errors associated with deficiencies in the model function, whereas the D-optimality approach assumes that noise in the data rather than the inadequacy of the response surface model is of paramount concern. The process of response surface generation must be extended to other disciplines, including estimates for drag, stability derivatives, and various performance parameters.

10. CONCLUDING REMARKS

Variable-complexity response surface methods were developed for the calculation of wing bending material weight used in a HSCT configuration design. The response surface

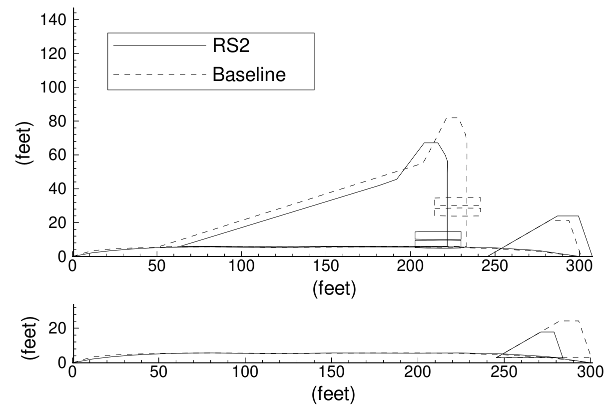

Figure 16. Optimal HSCT planforms using RS2 compared to baseline configuration.

was based on the results of structural optimization. However, simple conceptual-design level models were used to identify good variables for the response surface and to limit the extent of the design space. In particular, the variables used in algebraic weight equations in the FLOPS program were used to reduce the number of variables in the response surface from 25 to 10. The use of the simple models to reduce the extent of the design space was shown to improve the accuracy of the response surface by several orders of magnitude.

Parallel computation permitted the execution of thousands of structural optimizations which were used for the creation of the response surface. Substantial noise in the results of the structural optimization, of the order of 20–30% was found and traced to small variations in the aerodynamic loads. The average difference between the response surface and the structural optimization results was about 10%, with part of this difference attributed to the noise rather than to inaccuracies of the response surface. Several HSCT configuration optimizations were performed, and the optimization with the response surface produced superior results to those obtained with the FLOPS weight equation.

Variable-complexity response surface approximations to the wing bending material weight for HSCT configuration optimization have been shown to be an effective way to deal with the expense of structural optimizations, the nonsmoothness of structural optima, and the practical difficulties of code integration. Practical considerations that balance response surface accuracy with its development cost are nontrivial. Achieving both acceptable accuracy and acceptable cost for response surface approximations to wing bending material weight remains an open question. The techniques described here are not limited to the wing bending material weight response and can be applied to other functions involved in HSCT design.

10. REFERENCES

1 Barthelemy, J. F. M., Wrenn, G. A., Dovi A. R., Coen, P. G., and Hall, L. E., "Supersonic Transport Wing Minimum Weight Design Integrating Aerodynamics and Structures," *J. of Aircraft*, Vol. 31, No. 2, March–April 1994.

2 Box, G. E. P. and Behnken, D. W., "Some New Three Level Designs for the Study of Quantitative Variables," *Technometrics*, Vol. 2, No. 4, November, 1960, pp. 455–475.

3 Box, M. J. and Draper, N. R., "Factorial Designs, the $|\mathbf{X}^T\mathbf{X}|$ Criterion, and Some Related Matters," *Technometrics,* Vol. 13, No. 4, 1971, pp. 731–742.

19 Carlson, Harry W. and Walkley, Kenneth B., "Numerical Methods and a Computer Program for Subsonic and Supersonic Aerodynamic Design and Analysis of Wings With Attainable Thrust Corrections," NASA CR 3808, 1984.

5 Dudley, J., Huang, X., Haftka, R. T., Grossman, B., and Mason, W. H., "Variable-Complexity Interlacing of Weight Equation and Structural Optimization for the Design of the High Speed Civil Transport," AIAA Paper 94–4377, 1994.

6 Dudley, J., Huang, X., MacMillin, P. E., Grossman, B., Haftka, R. T., and Mason, W. H., "Multidisciplinary Optimization of the High-Speed Civil Transport," AIAA Paper 95–0124, 1995.

7 Giunta, A. A., Dudley, J. M., Narducci, R., Grossman, B., Haftka, R. T., Mason, W. H., and Watson, L. T., "Noisy Aerodynamic Response and Smooth Approximations in HSCT Design," AIAA Paper 94–4376, 1994.

8 Grandhi, R. V., Thareja, R., and Haftka, R. T., "NEWSUMT-A: A General Purpose Program for Constrained Optimization Using Approximation," *ASME J. Mechanisms, Transmissions and Automation in Design*, Vol. 107, 1985, pp. 94-99.

9 Hinkelman, Klaus, *Design and analysis of experiments*, John Wiley & Sons, Inc., 1994.

10 Huang, X., "Structural Optimization and its Interaction with Aerodynamic Optimization for a High Speed Civil Transport," Ph.D. Dissertation, VPI&SU, November 1994.

11 Huang, X., Haftka, R. T., Grossman, B., and Mason, W., "Comparison of Statistical-based Weight Equations with Structural Optimization for Supersonic Transport Wings," AIAA Paper 94–4379, 1994.

12 Hutchison, M. G., "Multidisciplinary Optimization of High-Speed Civil Transport Configurations Using Variable Complexity Modeling," Ph.D. Dissertation, VPI&SU, March 1993.

13 Hutchison, M. G., Huang, X., Mason W. H., Haftka, R. T., and Grossman, B., "Variable-Complexity Aerodynamic-Structural Design of a High-Speed Civil Transport Wing," *4th AIAA/USAF/NASA/ OAI Symposium on Multidisciplinary Analysis and Optimization*, Cleveland, OH, Sept. 21-23, 1992.

14 Hutchison, M. G., Unger, E. R., Mason, W. H., Grossman, B., and Haftka, R. T., "Aerodynamic Optimization of an HSCT Wing Using Variable-Complexity Modeling," Paper AIAA 93–0101, Jan. 1993.

15 Hutchison, M. G., Unger, E. R., Mason, W. H., Grossman, B., and Haftka, R. T., "Variable Complexity Aerodynamic-Structural Design of a High Speed Civil Transport Wing," AIAA Paper 92–4695, Sept. 1992.

16 Hutchison, M. G., Unger, E. R., Mason, W. H., Grossman, B., and Haftka, R. T., "Variable-Complexity Aerodynamic Optimization of an HSCT Wing Using Structural Wing-Weight Equations," *J. Aircraft,* Vol. 31, No. 1, 1994, pp. 110–116.

17 Khuri, A. I., and Cornell, J. A., *Response Surfaces: Designs and Analyses*, Marcel Dekker, New York, NY, 1987.

18 Malone, B., and Woyak, S. A., "An Object-Oriented Analysis and Optimization Control Environment for the Conceptual Design of Aircraft," AIAA Paper 95-3862, 1st AIAA Aircraft Engineering, Technology, and Operations Congress, Los Angeles, CA, Sept. 19-21, 1995

19 Mason, R. L., Gunst, R. F., and Hess, J. L., *Statistical Design and Analysis of Experiments*, John Wiley & Sons, New York, N. Y., 1989, pp. 215–221.

20 McCullers, L. A., "Aircraft Configuration Optimization Including Optimized Flight Profiles", *Proceedings of a Symposium on Recent Experiences in Multidisciplinary Analysis and Optimization*, J. Sobieski, compiler, NASA Cp–2327, April 1984, pp. 395–412.

21 Myers, R. H., Montgomery, D. C., *Response Surface Methodology: Process and Product Optimization Using Designed Experiments*, John Wiley & Sons, New York, N. Y., 1995, pp. 650-651.

22 Unger, E. R., Hutchison, M. G., Rais-Rohani, M., Haftka, R. T., and Grossman, B., "Variable-Complexity Design of a Transport Wing," *Intl. J. Systems Automation: Res. and Appl. (SARA),* No. 2, 1992, pp. 87–113.

23 Vanderplaats, Miura and Associates, Inc.,*GENESIS* User Manual, Version 1.3, Dec. 1993

24 Woyak, Scott A., Malone, B., and Myklebust, A., "An Architecture for Creating Engineering Applications: The Dynamic Integration System," Proceedings of the Computers in Engineering Conference and the Engineering Database Symposium, ASME, Sept. 17-20, 1995, Boston, MA, pp. 1-8.

APPENDIX A. Least Squares Method and Analysis of Variance

Statistical techniques known as regression analysis and analysis of variance (ANOVA) provide the means to identify and remove the less important terms in the response surface polynomial models based on simple analyses, thereby reducing the number of detailed analyses needed for later construction of accurate response surfaces. Regression analysis is the procedure by which the c_i coefficients for the response surface model are obtained and typically involves the method of least squares, as follows. Let there be m measurements Y_i at m distinct design points $(x_1^{(i)}, \ldots, x_p^{(i)})$, $i = 1, \ldots, m$, and let the basis functions for the regression model be $\varphi_1(x), \ldots, \varphi_n(x)$, where $x = (x_1, \ldots, x_p)$. A choice for the φ_i might be

$$1, x_1, x_2, \ldots, x_p, x_1^2, x_1 x_2, \ldots, x_p^2,$$

corresponding to a quadratic model. The regression model is

$$y \approx c_1\varphi_1(x) + \cdots + c_n\varphi_n(x),$$

which leads to the overdetermined matrix problem

$$Y \approx \mathbf{X}c, \tag{A.1}$$

where $Y = (Y_1, ..., Y_m)$, $c = (c_1, ..., c_n)$, and $\mathbf{X}$ is the $m \times n \quad (m \geq n)$ matrix

$$\begin{pmatrix} \varphi_1(x^{(1)}) & \ldots & \varphi_n(x^{(1)}) \\ \vdots & \ddots & \vdots \\ \varphi_1(x^{(m)}) & \ldots & \varphi_n(x^{(m)}) \end{pmatrix}.$$

The least squares solution is that (unique, assuming rank $\mathbf{X} = n$) vector $\tilde{c}$ which minimizes the 2-norm $||Y - \mathbf{X}c||_2$, or the sum of squares of the components of $Y - \mathbf{X}c$, the *errors.* The regression or response surface model of the data is then taken as

$$\tilde{c}_1\varphi_1(x) + \cdots + \tilde{c}_n\varphi_n(x).$$

The matrix $(\mathbf{X}^T\mathbf{X})^{-1}$ is called the variance-covariance matrix. The n diagonal elements in this matrix are the variance values associated with the n respective coefficients (Khuri and Cornell, 1987). The standard deviation σ_k of each coefficient is the square root of its variance.

Analysis of variance (ANOVA) involves estimating the variance of the predicted polynomial coefficients and uses the variance-covariance matrix $(\mathbf{X}^T\mathbf{X})^{-1}$. The diagonal terms ${\sigma_i}^2$ in this square matrix multiplied by the variance $\mathrm{Var}(\epsilon)$ in the measured function values Y_i are the variances of the respective coefficients c_i in the response surface polynomial model. The coefficient of variation Υ for each term in the polynomial is calculated as

$$\Upsilon = \begin{cases} \dfrac{100|\sigma_i|\sqrt{\mathrm{Var}(\epsilon)}}{c_i}, & c_i \neq 0, \\ 100, & c_i = 0, \end{cases} \tag{A.2}$$

where the factor of 100 expresses the coefficient of variation as a percentage. The term $\sqrt{\mathrm{Var}(\epsilon)}$ is usually estimated by the RMS error of the least squares approximation at the m data points. For the coefficient of variation calculations given below, $\mathrm{Var}(\epsilon)$ is taken to be unity. Terms in the polynomial model having large coefficients of variation, typically over 0.10, may be dropped from the polynomial without significantly affecting the fidelity of the response surface fit.

APPENDIX B. FLOPS Weight Function Details

The wing weight function in FLOPS is based on an analytic expression to relate wing bending material weight to wing geometry, material properties and loading. Two terms in the function account for the load distribution on the wing, the bending material factor, B_z, and the engine relief factor, B_{ze}.

B_z accounts for the distribution of load on the wing, neglecting the engines, and is calculated by approximately determining the required material volume of the upper and lower skins in a simple wing box description of the wing. We begin by finding the weighted average of the load sweep angle at 75% chord position

$$\Lambda_L = \int_0^1 (1 + 2y)\Lambda(y)\, dy.$$

Next it is necessary to determine the bending moment assuming a simple elliptic pressure distribution.

$$M(y) = \int_0^y p(\xi)\xi\, d\xi.$$

With this bending moment distribution it is possible to calculate the necessary flange are as

$$A(y) = \frac{M(y)}{t(y)c(y)\cos\Lambda(y)},$$

and the required volume as

$$V = \int_0^1 A(y)\, dy.$$

The total load is simply

$$L = \int_0^1 p(y)\, dy,$$

and the bending material factor is finally given as

$$B_z = \frac{2V}{Ld},$$

where

$$d = AR^{0.25}[1 + (0.5f_a - 0.16)\sin^2\Lambda_L + 0.3C_y$$

$$(1 - 0.5 f_a) \sin \Lambda_L].$$

B_{ze} accounts for the reduced amount of structural weight necessary due to the presence of the engines on the wing.

$$B_{ze} = 16 \int_{y_{root}}^{y_{tip}} NE \left(\frac{y_{tip} - y}{\cos^2 \Lambda(y)} \right) \left(\frac{y_{tip} - y_{root}}{c(y) Max(t/c)} \right) dy,$$

where NE is given by

$$NE = \begin{cases} 2, & y < y_{nacelle_1}, \\ 1, & y_{nacelle_1} < y < y_{nacelle_2}, \\ 0, & y > y_{nacelle_2}. \end{cases}$$

APPENDIX C. Generation of Candidate Designs

In the development of a response surface, a large number of data points that *cover* the design space are required. A systematic approach for producing the locations of these points has been developed assuming that the initial design space is a p-cube centered at the origin with vertices $(\pm 1, \ldots, \pm 1)$ and a quadratic (1) is used for the model function. The level of a variable in a point selection system refers to the number of values that the variable is allowed to take on within a set of data points; for example, if a variable assumes the values $\{-1, 0, +1\}$ within a set of points, then it is said to have a level of three (Box and Behnken, 1960).

The most general three-level point selection system is the full factorial method, where every variable assumes any of three possible values. For twenty-five variables, this requires the evaluation of $3^{25} \approx 8.5 \times 10^{11}$ designs, which is unrealistic. Consider instead a combination of two level-point selection systems based on the Partially Balanced Incomplete Box design (PBIB) (Hinkelman, 1994). For a two-level system, each variable can take on the values $\{-1, +1\}$.

Similar to the PBIB design, a pattern of blocks is created, each of which contains a fraction of the total number of variables. The variables within a block are evaluated at the two levels ± 1, while the variables outside of the block are held fixed at the third level 0. Three different blocking systems are incorporated to produce a satisfactory number of points. Every block pattern containing one, two, and three variables is considered, as well as the center point $(0, \ldots, 0)$.

All possible two-variable block patterns for a problem in three variables, and the corresponding design points are:

	x_1	x_2	x_3
Block #1	*	*	
Block #2		*	*
Block #3	*		*

	x_1	x_2	x_3
Block #1	1	1	0
	-1	1	0
	1	-1	0
	-1	-1	0
Block #2	0	1	1
	0	-1	1
	0	1	-1
	0	-1	-1
Block #3	1	0	1
	-1	0	1
	1	0	-1
	-1	0	-1

For twenty-five variables, 19,651 points are produced using the three blocking systems and one center point. For an arbitrary number p of variables the total number of points created is

$$1 + \sum_{i=1}^{3} 2^i \frac{p!}{i!\,(p-i)!}.$$