



# Polynomial Response Surface Approximations for the Multidisciplinary Design Optimization of a High Speed Civil Transport

SERHAT HOSDER, LAYNE T. WATSON, BERNARD GROSSMAN, WILLIAM H. MASON,  
HONGMAN KIM

*Multidisciplinary Analysis and Design Center for Advanced Vehicles, Virginia Polytechnic Institute and State  
University, Blacksburg, VA 24061-0203, USA*

RAPHAEL T. HAFTKA, STEVEN E. COX

*Department of Aerospace Engineering, Mechanics & Engineering Science, University of Florida, Gainesville,  
FL 32611-6250, USA*

*Received March 1, 2001; Revised August 2, 2001*

**Abstract.** Surrogate functions have become an important tool in multidisciplinary design optimization to deal with noisy functions, high computational cost, and the practical difficulty of integrating legacy disciplinary computer codes. A combination of mathematical, statistical, and engineering techniques, well known in other contexts, have made polynomial surrogate functions viable for MDO. Despite the obvious limitations imposed by sparse high fidelity data in high dimensions and the locality of low order polynomial approximations, the success of the panoply of techniques based on polynomial response surface approximations for MDO shows that the implementation details are more important than the underlying approximation method (polynomial, spline, DACE, kernel regression, etc.). This paper selectively surveys some of the ancillary techniques—statistics, global search, parallel computing, variable complexity modeling—that augment the construction and use of polynomial surrogates.

**Keywords:** global optimization, multidisciplinary design, parallel computing, response surface

## 1. Introduction

Multidisciplinary design optimization (MDO) is an approach to formalizing the design process, producing higher quality designs more rapidly than traditional engineering design methods (Sobieszczanski-Sobieski and Haftka, 1997), and Alexandrov and Hussaini (1997). To achieve the MDO goal, high accuracy or high fidelity results from specific disciplines such as aerodynamics and structures are required immediately in the design process. For extremely challenging designs, where the disciplines are highly coupled, use of high fidelity results is critical to establishing the feasibility of the design and making the risk acceptable. The high speed civil transport (HSCT) is an example of this type of design problem. Early experience with including all the analyses in a single computer program (Hutchison et al., 1994) made it clear that it was not possible to include high fidelity computations using the most advanced computational methods in a single program.

Computational issues in MDO systems design have been addressed over the past decade by studying a sequence of example design problems of increasing levels of sophistication.

One result of this work is an approach called *variable complexity* (or *multiple fidelity modeling*), which simultaneously utilizes computational models of different levels of fidelity to reduce the computational effort (Hutchison et al., 1993). These computational models may be history-based or analysis-based, such as an algebraic equation for estimating a vehicle weight contrasted with a finite element structural analysis and corresponding structural optimization. The models may involve different levels of physics such as panel methods and computational fluid dynamic (CFD) Euler or Navier-Stokes solutions. Finally, they may involve different levels of fidelity with the same physics, such as using a refined grid system.

The general approach is to utilize the lower fidelity, computationally less expensive methods to explore the design space and to identify promising regions in which to perform design optimization using the higher fidelity methods. However, computational experience revealed a number of problems associated with numerical noise, discontinuous derivatives, and nonconvex design spaces that may lead to spurious local optima (Burgee et al., 1996). Additional problems arise when considering more than a few design variables ( $\geq 10$ ). In high dimensions, computational costs become a dominant consideration and real local optima also appear. Furthermore, the use of high fidelity models in the MDO process may lead to software problems associated with code integration and disciplinary boundaries.

Giunta et al. (1997a) and many others following this same thread dealt with these problems by providing high fidelity disciplinary results to the MDO process in a simplified form, as an approximation of the true response of a disciplinary analysis to a particular set of design variables. In the work surveyed here, these surrogates for the true responses are in the form of polynomial response surface approximations (RSAs). The development of these response surface approximations requires the use of experimental design methods and statistical analysis. However, once created, the RSAs can be used repeatedly, just as engineers used design charts from handbooks in previous generations. A more detailed description of this process is found in Giunta et al. (1997b).

This body of work shows that response surface approximations can effectively filter out noise, can easily combine multifidelity models, and can alleviate code integration issues. Furthermore, this approach lends itself to coarse grained parallel computing and global optimization. In addition, since large numbers of designs are analyzed to create the response surface approximations, statistical analysis can be used to provide quality and error control.

This paper reviews some MDO work using the HSCT design problem as an example problem. After defining the specific HSCT problem, surrogate models for aircraft weight evaluation using finite element/structural optimization with several levels of fidelity are addressed. This is followed by a description of the variable complexity development of response surface approximations for aircraft drag utilizing analytic approximations, panel methods, and CFD solutions. Next, some work in experimental design and variable complexity modeling is reviewed. The paper is completed with a description of some statistical methods used to characterize the optimization errors, followed by two other key aspects of producing a viable MDO procedure, the use of parallel computation and global search methods.

## 2. Description of the HSCT configuration design problem

The design problem can be described as minimizing the takeoff gross weight (TOGW) of a High Speed Civil Transport (HSCT) aircraft with a 5500 nautical mile range, designed to cruise at Mach 2.4 and carry 251 passengers. The choice of the gross weight as the objective function incorporates structural and aerodynamic considerations. The structural considerations are directly related to the aircraft empty weight and the drag, while the aerodynamic performance dictates the drag and hence the thrust and fuel weight required.

The mathematical formulation of the HSCT design problem can be written as

$$\begin{aligned} \min_x W(x) \quad \text{subject to} \quad & g_i(x) \leq 0, \quad i = 1, \dots, 68, \\ & x_{\min} \leq x \leq x_{\max}, \end{aligned}$$

where  $W(x)$  is the takeoff gross weight,  $x$  is the design variable vector,  $g(x)$  represents the nonlinear inequality constraints, and  $x_{\min}$ ,  $x_{\max}$  are the lower, upper bounds for the design variables.

In the most general case considered here, the HSCT design problem is modeled by using 29 design variables associated with the aircraft geometry and the mission specified (Table 1). The wing planform is created with eight geometric design variables (figure 1): the root chord, location of the leading edge (LE) break point, location of the trailing edge (TE) break point, wing tip LE location, tip chord, and wing semispan. Camber and twist (washout) are derived quantities handled as described in Balabanov (1997) and Balabanov et al. (1998). The airfoil sections are described by using five design variables (figure 2): location of the maximum thickness, the LE radius parameter, and thickness-to-chord ratios at the wing root, LE break, and wing tip. The wing thickness is varied linearly between these three spanwise locations. Fuselage geometry is specified by four radii at four axial stations along the centerline. The shape of the body between these locations is determined by considering it as a minimum wave drag body of a fixed volume. Variables  $x_{22}$  and  $x_{23}$  represent location of the inboard nacelle and the nacelle separation, respectively. The remaining design variables are the mission fuel weight, starting cruise altitude, cruise climb rate, vertical tail area, horizontal tail area, and the thrust per engine.

During the optimization process, up to 68 inequality constraints are used (Table 2). These constraints were devised to ensure feasible aircraft geometry and impose realistic performance and control capabilities. Fuel volume and wing chord length limits can be given as examples of geometric constraints. Performance/aerodynamic constraints include, for example, landing angle of attack limits and wing, tail, and engine scrape prevention criteria. Aerodynamic and performance constraints can only be assessed after a complete analysis of a HSCT design, however, the geometric constraints can be evaluated using algebraic relations based on the design variables.

The HSCT configuration design problem was used as a testbed for the evaluation of new design optimization methodologies and techniques developed in the Multidisciplinary Analysis and Design (MAD) Center for Advanced Vehicles. In some of these studies, only a certain subset of the design variables were used while the rest of the variables were kept

Table 1. HSCT configuration design variables.

Index	Description
1	Wing root chord (ft)
2	LE break point, streamwise (ft)
3	LE break point, spanwise (ft)
4	TE break point, streamwise (ft)
5	TE break point, spanwise (ft)
6	LE wing tip, streamwise (ft)
7	Wing tip chord (ft)
8	Wing semispan (ft)
9	Chordwise location of max. thickness (% of chord)
10	LE radius parameter
11	Airfoil $t/c$ at wing root
12	Airfoil $t/c$ at wing break
13	Airfoil $t/c$ at wing tip
14	Fuselage restraint 1, streamwise (ft)
15	Fuselage restraint 1, radius (ft)
16	Fuselage restraint 2, streamwise (ft)
17	Fuselage restraint 2, radius (ft)
18	Fuselage restraint 3, streamwise (ft)
19	Fuselage restraint 3, radius (ft)
20	Fuselage restraint 4, streamwise (ft)
21	Fuselage restraint 4, radius (ft)
22	Inboard nacelle location (ft)
23	Nacelle separation (ft)
24	Fuel weight (lb)
25	Starting cruise altitude (ft)
26	Cruise climb rate (ft/min)
27	Vertical tail area (ft <sup>2</sup> )
28	Horizontal tail area (ft <sup>2</sup> )
29	Thrust per engine (lb)

at fixed values in the optimization process. Also a reduced number of constraints were used in the same studies.

### 3. Response surface approximations for structural analysis

Response surface (RS) approximations are increasingly being used as a tool in MDO. However, as the dimension of the design space increases, the accuracy of the RS approximations deteriorates and their computational cost grows. This difficulty typically limits RS modeling




Figure 1. Wing planform design variables.




Figure 2. Wing airfoil design variables.

to problems with 10–15 design variables. To alleviate this difficulty, Kaufman et al. (1996) and Balabanov et al. (1999) have developed a reasonable design space approach. In this approach, geometric constraints and the constraints based on simple, computationally inexpensive analytical tools were used to eliminate large portions of the design space from consideration. This reduction in the size of the design space improved the accuracy of the RS approximations.

In most of the applications of RS approximations, a box is defined in the design space by lower and upper bounds on the design variables. Without considering the designs at least at

Table 2. HSCT optimization constraints.

Index	Constraint
1	Fuel volume $\leq 50\%$ wing volume
2	Wing root TE $\leq$ Tail LE
3–20	Wing chord $\geq 7.0$ ft
21	LE break within wing semi-span
22	TE break within wing semi-span
23	Root chord $t/c$ ratio $\geq 1.5\%$
24	LE break chord $t/c$ ratio $\geq 1.5\%$
25	Tip chord $t/c$ ratio $\geq 1.5\%$
26–30	Fuselage restraints
31	Wing spike prevention
32	Nacelle 1 inboard of nacelle 2
33	Nacelle 2 inboard of semi-span
34	Range $\geq 5500$ nautical miles
35	$C_L$ at landing speed $\leq 1$
36–53	Section $C_L$ at landing $\leq 2$
54	Landing angle of attack $\leq 12^\circ$
55–58	Engine scrape at landing
59	Wing tip scrape at landing
60	TE break scrape at landing
61	Rudder deflection $\leq 22.5^\circ$
62	Bank angle at landing $\leq 5^\circ$
63	Tail deflection at approach $\leq 22.5^\circ$
64	Takeoff rotation to occur $\leq V_{\min}$
65	Engine-out limit with vertical tail
66	Balanced field length $\leq 11000$ ft
67–68	Mission segments: thrust available $\geq$ thrust required

all the vertices of the design box for the construction of the RS model, sufficient accuracy in the approximation may not be obtained. With an  $n$ -dimensional box having  $2^n$  vertices, it becomes impractical to evaluate the designs at all of the vertices for values of  $n$  on the order of ten or more (the curse of higher dimensionality). However, the reasonable design space approach can reduce the size of the considered design space and render it more similar to a simplex (with only  $n + 1$  vertices) or at least an ellipsoid. This reduction in the size (more precisely, the volume) of the design space enables efficient construction of the RS approximations with desired accuracy over the entire considered design space.

Balabanov et al. (1999) have successfully applied the reasonable design space approach to an HSCT configuration optimization problem. They used quadratic polynomials for the response surface model to approximate the structural bending material weight (SBMW).

The objective function was to minimize the TOGW of the aircraft and 25 out of 29 design variables were used in the optimization (influence of the LE radius parameter, starting cruise altitude, cruise climb rate, and the thrust per engine were found to be negligible in calculating SBMW). The low fidelity approximation of the SBMW and the rest of the weight components were calculated by using the weight functions from the flight optimization system program FLOPS (McCullers, 1984). For the high fidelity approximation, the results of the finite element structural optimization code GENESIS (Vanderplaats et al., 1993) were used. The optimizations were started from a baseline HSCT configuration located in the interior of the feasible design space.

As the first step in identifying the reasonable design space, a box, defined by the bounds of the design variables, that encompasses a suitably large region of the design space was constructed. A partially balanced incomplete block statistical experimental design was used to select 19651 HSCT configurations. Out of these 19651 HSCT configurations, 83% violated one or more of the HSCT geometric constraints and a large portion of the remaining configurations appeared to be unreasonable. Each point in the experimental design set corresponding to an unreasonable HSCT configuration  $x$  was then moved until it resided on the edge of the reasonable design space by using the equation

$$x' = \alpha(x - x_b) + x_b,$$

where  $x_b$  represents the (feasible) baseline HSCT configuration and  $\alpha \geq 0$  is a parameter that is adjusted from 1 toward 0 to make the HSCT configuration reasonable. The shift of  $x$  to the edge of the reasonable design space was performed in three steps in order to minimize the computationally expensive evaluations of the complex constraints and to avoid the removal of any reasonable configurations.

In the first step, geometrically infeasible designs were considered and a large percentage of the candidate points were moved towards a nominal feasible point (HSCT baseline configuration)  $x_b$ . In the second step, constraints that were more complex but did not require complete performance and aerodynamic analysis of the HSCT configurations were applied. In the last step, the most complex constraints were evaluated to relocate the remaining unreasonable configurations. The number of points moved in the last step was small enough so that the constraint evaluations did not become prohibitive. After all 19651 points in the design space were moved to the edge of the reasonable design space, the D-optimality criterion was used to select two sets of HSCT configurations, each having 1000 points. Structural optimization was performed for each HSCT configuration in the two sets and SBMW was estimated from the results of the structural optimization. One set of HSCT configurations was used to create a RS approximation model and the other was used to check the accuracy of the RS approximation. In the accuracy check, both the RS approximations and the results obtained from FLOPS weight functions were compared with the structural optimization results from GENESIS. The comparison showed that the RS model was more accurate than the FLOPS weight function in approximating the SBMW. The average error in the RS approximation was 4.3% while FLOPS gave an average error of 29.8%.

In the overall multidisciplinary optimization of the HSCT configuration, the SBMW in the reasonable design space was determined from the RS model. FLOPS weight functions were used to obtain the SBMW outside the reasonable design space and the other weight

components in the whole design space. Optimization results showed that the optimum design point obtained by using the response surface approximation was slightly more accurate than the one obtained by using the FLOPS weight function. Also a 4400 lbs. saving of TOGW in the RS approximation case was obtained compared to the FLOPS case.

The use of a reasonable design space approach for constructing the RS approximation to the SBMW in HSCT configuration optimization demonstrated the following advantages: (1) the reasonable design space approach reduces the volume of the design space that must be considered significantly, thus enabling the computationally efficient construction of the RS model; (2) significant improvements in the accuracy of the response surface approximations can be achieved.

#### **4. Response surface approximations for aerodynamics**

The computational expense of using more accurate CFD predictions in the aircraft design optimization process motivates the need to develop new techniques that enable the efficient implementation of such CFD methods into high dimensional, highly constrained multidisciplinary design optimization (MDO) procedures. Knill et al. (1999) have developed a method to efficiently implement supersonic aerodynamic predictions from Euler solutions into a highly constrained MDO of an HSCT aircraft. Efficient application of the accurate CFD analysis to the optimization process has been achieved through the use of variable complexity modeling (VCM) techniques, response surface (RS) methodologies, and coarse grained parallel computing.

The design problem used as a testbed for the method was formulated by using subsets of the design variables and the constraints described for the general HSCT configuration. To evaluate the method developed, a series of optimization problems with 5, 10, 15, and 20 variables were solved. The objective function in these optimization problems was to minimize the takeoff gross weight (TOGW) of the aircraft as described in the general HSCT design problem.

The main component of the method involved utilizing information gained from inexpensive lower-fidelity aerodynamic models to more efficiently create quadratic RS models for quantities used to evaluate performance/aerodynamic related constraints. The lower-fidelity aerodynamic predictions were calculated by using supersonic linear theory codes. These gave the volumetric wave drag and the drag due to lift. The viscous drag estimates were obtained from standard algebraic estimates of the skin friction assuming turbulent flow. The viscous drag predictions were also added to the Euler solutions that were obtained by using the General Aerodynamic Simulation Program (GASP) (McGrory et al., 1993) The numerical scheme used in this program is a third order upwind-biased interpolation of the Roe fluxes. The constrained optimization was performed using sequential quadratic programming and central difference gradient approximations in the software package Design Optimization Tools (DOT) (Vanderplaats Research & Development, 1995).

Knill et al. (1999) created an approximate model of the supersonic aerodynamics in order to avoid using relatively expensive Euler solutions for the large number of constraint evaluations required by optimization. The advantages of using response surface approximations for constraint evaluation were shown in the same work. The approximate models



filter out numerical noise present in the analysis, which distorts gradient information and can induce artificial minima in the design space. By separating analysis codes from the optimization routine, the need for the integration of large production level grid generators, analysis codes, and post processing utilities with the optimizer can be eliminated. Also the design trade-offs, sensitivities to certain variables, and insight into a highly constrained, nonconvex feasible set can be easily obtained by using RS approximation models.

In the work of Knill et al. (1999), additional accuracy was gained by creating response surface approximations for separate portions of the drag coefficient. Instead of modelling the drag coefficient as a quadratic function of design variables, Knill et al. (1999) approximated the drag coefficient by the well-known approximation

$$C_D(x) = C_{D_0}(x) + K(x)C_L^2.$$

$C_{D_0}(x)$  and  $K(x)$  were considered as intervening functions and approximated by quadratic models. A quadratic model in  $m$  variables has the form

$$y = c_0 + \sum_{1 \leq j \leq m} c_j x_j + \sum_{1 \leq j < k \leq m} c_{jk} x_j x_k.$$

This quadratic model has  $(m + 1)(m + 2)/2$  coefficients in  $m$  variables. Computing the coefficients for quadratic models for both  $C_{D_0}(x)$  and  $K(x)$  for a 30 variable problem would require 4184 CFD analyses. On a single processor of an SGI Power Challenge, this would require over 46 days of CPU time. This clearly shows the necessity of developing a method that would create accurate RS approximations with a moderate number of CFD analyses.

To select the points used in creating RS approximation models in a systematic way, the D-optimality criterion for experimental design was used. Details concerning the application of design of experiments theory to the point selection procedure are discussed in Knill et al. (1999).

As the first step for obtaining reduced-term RS approximation models of Euler solutions, linear theory RS models were obtained by using all the coefficients in the quadratic model. This procedure was computationally inexpensive. Then stepwise regression analysis was used in a systematic way to remove the coefficients that had little or no impact on the response from the RS models. Terms were eliminated by using a measure of the significance level of the term called the  $p$  value, which represents the probability that the coefficient of a particular term is actually zero. After obtaining the reduced-term RS models for the linear theory, the same reduced quadratic model form was used to calculate correction RS models for  $\Delta C_{D_0}(x)$  and  $\Delta K(x)$  representing the difference between linear theory values and Euler values of the intervening functions. Incremental RS models for the Euler solutions were then created by adding the correction RS models to the full quadratic linear theory RS models. Optimization results using the linear theory RS models were also used to predict a design bounding box within which the optimum from Euler analysis should lie. This improved the accuracy of the RS models by allowing smaller ranges on the design variables than would have been possible if no information was available on the general location of the optimal design.

The optimization results showed that the stepwise regression technique can eliminate unimportant terms from the RS models with little or no effect on the error. An important

conclusion is that the linear theory analysis does reveal the terms that are important to the Euler analysis and no important non linear effects have been masked. Compared with the cost of creating full-term RS models, creating the reduced term RS models resulted in savings of 11 out of 16 h, 47 out of 74 h, 115 out of 192 h, and 255 out of 392 h of CPU time on a single 75-MHz (P2) processor of an SGI Power Challenge for the 5-, 10-, 15-, and 20-variable design problems, respectively. Errors in the reduced-term incremental RS model cruise drag predictions for the optimal designs based on actual Euler calculations ranged from 0.1 to 0.8 counts.

The definite success of polynomial RS approximations for CFD can be attributed to several factors: (1) using low fidelity models to identify the important polynomial terms and then reduced-term polynomial models for the expensive response surfaces; (2) modeling well chosen intervening functions rather than the ultimate quantities of interest; (3) using optimization based on low order models to bound the domain over which the high fidelity RS models are employed; (4) approximating the difference between low and high fidelity values rather than the high fidelity data directly (Balabanov et al., 1998) approximated the ratio of high to low fidelity values for structures); (5) numerical optimization using polynomial functions is much more efficient than when using expensive, noisy high fidelity analysis values directly; (6) construction of response surface approximations efficiently utilizes massively parallel computation.

## 5. Experimental design and variable complexity modeling

This section elaborates on two core subjects, variable complexity modeling and the statistical theory of experimental designs, that are used, to some degree, by all the approaches in the previous sections. Computationally efficient application of RS approximations to structural analysis and fluid dynamics problems in MDO with desired accuracy has been achieved through the use of variable complexity modeling (VCM). The term variable complexity refers to a design procedure in which refined computationally expensive analysis techniques are combined with simple, computationally inexpensive techniques.

The construction of RS approximations by using the VCM technique may be viewed as a series of steps to be completed before the HSCT optimization is performed. These steps were outlined in Giunta et al. (1997a) as follows: (1) start with an initial HSCT design in the feasible design space; (2) define design space boundaries around the initial design; (3) use appropriate experimental design methods to create a certain combination of numerical experiments (analyses) in which the design variables are prescribed at specific values or levels; (4) perform low fidelity analyses; (5) determine the reasonable design space; (6) create a D-optimal experimental design; (7) perform medium/high fidelity analyses; (8) create response surface models; (9) perform HSCT optimization using the response surface models.

Before the determination of the reasonable design space by the procedure described in Section 3, low fidelity analyses are performed at the design points specified by different experimental methods. Prior to creating an experimental design, the allowable range of each of the  $m$  variables is defined by lower and upper bounds. The allowable range is then discretized at equally spaced levels. For numerical stability and for ease of notation,

each variable is scaled to the interval  $[-1, 1]$ . The region enclosed by the lower and upper bounds on the variables is termed the design space, the vertices of which determine an  $m$ -dimensional cube. If each of the variables is specified at only the lower and upper bounds (two levels), the experimental design is called a  $2^m$  full factorial. Similarly, a  $3^m$  full factorial design is created by specifying lower bound, midpoint, and upper bound (three levels) for each of the  $m$  variables. A  $3^m$  full factorial design provides ample response evaluations to permit the estimation of the RS model coefficients. For example, fitting a quadratic response surface model in three variables ( $m = 3$ ) requires at least ten evaluations, and a  $3^3$  full factorial design provides 27 evaluations. However, as  $m$  becomes large, the evaluation of both  $2^m$  and  $3^m$  full factorial designs becomes impractical (e.g.,  $2^{30} = 1.1 \times 10^9$ ). A full factorial design is typically used for ten or fewer variables.

Giunta et al. (1997a) used a full factorial experimental design with three levels in each variable for a 10 variable HSCT wing design optimization problem. Kaufman et al. (1996) and Balabanov et al. (1999) applied a partially balanced incomplete block (PBIB) experimental design method to a 25 variable HSCT wing bending material weight optimization in order to select the design points at which the low fidelity analyses were performed. They created a pattern of blocks, each of which contains a fraction of the total number of variables. The variables within a block were evaluated at the two levels  $\pm 1$ , while the variables outside the block were held fixed at the third level 0. Three different blocking systems were incorporated to produce a satisfactory number of points. Every block pattern containing one, two, and three variables was considered, as well as the center point  $(0, \dots, 0)$ . For twenty-five variables, 19651 points were produced using the three blocking systems and one center point.

After a significant reduction of the size of the design space with the reasonable design space approach, the D-optimality criterion for experimental design is used for selecting the points at which the medium/high fidelity analyses are performed. The reasonable design space approach usually creates an irregularly shaped boundary point set, which complicates the application of full factorial or PBIB experimental design methods. However, the D-optimality criterion provides a rational means for creating experimental designs inside an irregularly shaped design space. The objective of the D-optimality criterion is to select the set of  $p$  locations from a pool of  $q$  candidate locations ( $q \geq p$ ) such that the quantity  $\det(X^T X)$  is maximized. Here  $X$  is the  $p \times n$  coefficient matrix assumed to have rank  $n$  in the least squares approximation problem  $Xc \approx Y$ , where  $c$  is the  $n$ -vector of the  $n$  unknown polynomial coefficients in the quadratic RS polynomial approximation, and  $Y$  is a  $p$ -vector of observed response values. A set of  $p$  locations for which  $\det(X^T X)$  is maximum is called a D-optimal experimental design. The statistical reasoning behind the creation of a D-optimal design is that it leads to response surface models for which the uncertainty in the estimated coefficients is minimized. Kaufman et al. (1996) used genetic algorithms to create D-optimal experimental designs, while Giunta et al. (1997a) applied the  $k$ -exchange method of Mitchell to select a set of  $p$  D-optimal locations from a user supplied list of  $q$  candidate locations.

After the D-optimal locations are determined, the quadratic polynomial RS approximations are created by performing medium/high fidelity analyses at these selected design points. Further simplification without loss of accuracy can be achieved by eliminating the

terms of the quadratic polynomial that have negligible effect on the approximation. Kaufman et al. (1996) and Giunta et al. (1997a) used statistical regression analysis and analysis of variance (ANOVA) to determine the unnecessary terms in the quadratic response surface polynomials. As the statistical measure for removing terms from the polynomial model, they used the adjusted  $R^2$  value ( $R^2_{adj}$ ), which can be calculated from the ANOVA results. As described in the previous section, Knill et al. (1999) used stepwise regression analysis and the significance level of the  $p$  value as the statistical measure for removing terms from the polynomial model obtained by using low fidelity calculations.

The application of different experimental design methods to MDO enables the selection of the design points used in the RS construction in a systematic and efficient way. The VCM technique incorporates the computationally inexpensive low fidelity analyses with more accurate but computationally expensive high fidelity methods in a hierarchical way. Via the VCM method in RS modeling, a computationally efficient construction of the polynomial RS approximations can be achieved with acceptable accuracy.

Another common theme in variable complexity modeling is to approximate the ratio or difference between models of different fidelities. For instance, Balabanov et al. (1998) constructed a polynomial surrogate for the ratio of high to low fidelity wing bending material weights (WBMW). Thousands of structural optimizations were performed with a coarse finite element model of the HSCT. These optima were fit with a quadratic function of the configuration design variables. In addition, about 100 higher fidelity structural optimizations were performed with a refined finite element model. The ratio and the difference of the WBMW from the two optimizations were approximated by linear least squares fits. The lower fidelity quadratic surrogate together with a linear correction surrogate allowed the use of a relatively small number of high fidelity structural optimizations without loss of accuracy.

## 6. Statistical techniques for surrogate construction

The availability of a large amount of data used for surrogate construction provides opportunities for detecting bad data and for assessing the average error in the data. For example, in MDO the response surface data may come from disciplinary optimization, and optimization is typically an iterative process that is rarely allowed to converge to high precision due to computational cost considerations. Consequently, optimization results are usually a noisy function of the parameters of the design problem. For example, the structural optimization adopted to estimate optimal wing structural weight ( $W_s$ ) of the HSCT resulted in noisy  $W_s$  responses (Kaufman et al., 1996). Although it might be difficult to find the error for a single optimization, when many optimization results are available such as in constructing response surface approximations, statistical techniques can be used to estimate the mean error in the optimization results. Convergence difficulties during optimization runs may result in data points with large error, called *outliers*. Robust statistical techniques can be used to identify those outliers, permitting their possible repair and reincorporation into the response surface approximation.

To study the statistical distribution of optimization error incurred by using GENESIS to find the optimal structural weight  $W_s$  of the HSCT, Kim et al. (2000a) selected the

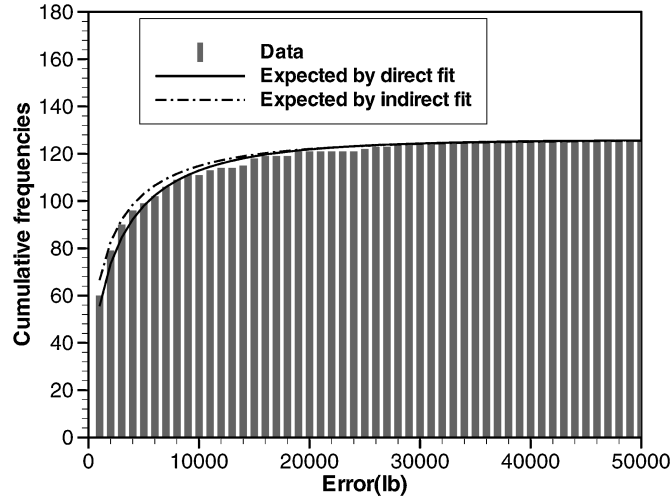


Figure 3. Comparison of cumulative frequencies between direct fit and indirect fit of the Weibull model.

Weibull distribution as a candidate model. The straightforward approach is to fit the model distribution to the optimization error

$$W_s - W_s^t,$$

where  $W_s$  is the calculated optimal wing structural weight and  $W_s^t$  is the (unknown) true optimum.  $W_s^t$  was estimated as the best of six runs using different convergence criteria settings. This approach is denoted as a direct fit. In figure 3, the direct fit using the Weibull model is compared with data in terms of cumulative frequencies; the predicted cumulative frequencies by the Weibull fit is in good agreement with the data. The  $p$ -value for the Chi-square goodness of fit test was 0.0925.

The problem associated with the direct fit is that it can be expensive because  $W_s^t$  has to be estimated by performing high fidelity runs. Moreover, such high fidelity data is not always achievable. Readily available are optimization runs with two different convergence settings, where one is not particularly better than the other. Then, the distribution of the difference of the two optima from the different convergence settings, instead of the optimization error itself, can be estimated.

For two optimization results,  $W_s^1$  with convergence setting 1 and  $W_s^2$  with convergence setting 2, model the optimization errors as random variables  $s$  and  $t$ ,

$$s = W_s^1 - W_s^t, \quad t = W_s^2 - W_s^t. \quad (1)$$

The difference of  $s$  and  $t$  is defined as the *optimization difference*

$$x = s - t = (W_s^1 - W_s^t) - (W_s^2 - W_s^t) = W_s^1 - W_s^2. \quad (2)$$

If  $s$  and  $t$  are independent, the probability density function (PDF) of  $x$  can be obtained by a convolution of the PDF functions  $g(s; \beta_1)$  and  $h(t; \beta_2)$ ,

$$f(x; \beta_1, \beta_2) = \int_{-\infty}^{\infty} g(\tau - x; \beta_1)h(\tau; \beta_2) d\tau, \quad (3)$$

where  $\beta_1$  and  $\beta_2$  are the parameter vectors for the two Weibull PDFs. The import of (3) is that the statistical properties of the errors  $W_s^1 - W_s^t$  and  $W_s^2 - W_s^t$  can be estimated even without knowledge of the true optimum  $W_s^t$ . Note that the optimization difference  $x$  is easily calculated from the known  $W_s^1$  and  $W_s^2$ . Denote this approach as an indirect fit.

The results from the indirect fit are shown together with the direct fit results in figure 3. Kim et al. (2000a) showed that the indirect approach gave comparable fits to the direct approach. Even when only low fidelity simulations are available, the indirect fit can give reasonable estimates of the mean error of the optimization, validating the use of a random variable model for optimization error.

The standard least squares fit can be greatly affected by a few bad data points. Robust regression techniques (Rousseuw and Leroy, 1987) give reasonable fits even if the data is contaminated with outliers. In addition, robust regression can identify the outliers, allowing repair whenever possible.  $M$ -estimation is a robust regression method, where the least squares equation  $X\hat{\beta} \approx y$  is generalized to a weighted form

$$X^t W(r) X \hat{\beta} = X^t W(r) y, \quad (4)$$

where  $W(r) = \text{diag}(w(r_1), w(r_2), \dots, w(r_n))$ . The weighting  $w(r_i)$  is a function of the scaled residual error

$$r = \frac{y - X\hat{\beta}}{s}, \quad (5)$$

where  $s$  is a normalizing scale factor. (4) is a nonlinear system of equations typically solved using iteratively reweighted least squares (IRLS) (Holland and Welsch, 1977).

The weighting functions in the IRLS procedure are usually even functions, i.e.,  $w(r_i) = w(-r_i)$ , because the random error is assumed to be symmetric. However, the error of optimization tends to be skewed due to premature convergence, and the mean error cannot be zero. Therefore, it might be desirable to use a nonsymmetric weighting function rather than a symmetric one (Kim et al., 2000b). Figure 4 shows various weighting functions for  $M$ -estimation. The two symmetric weighting functions, Huber's and biweight, may be combined to devise a nonsymmetric one. Note that biweight penalizes the outliers more severely than Huber's; biweight gives zero weighting to outliers when the absolute residual error is greater than a certain value and removes them completely from the fit, while Huber's function gives small but finite weightings. The nonsymmetric weighting function for IRLS is denoted as NIRLS in figure 4. NIRLS downweights the negative residual points according to Huber's, but downweights positive residual points according to biweight.

Table 3 compares outlier identification of the biweight and NIRLS weighting functions for the HSCT  $W_s$  optimization. Without  $M$ -estimation, the response surface fit using a

Table 3. Results of outlier detection and repair.

Weighting function	(Detected outliers)/(total outliers)	RMSE (%)	$R^2$
$w(r) \equiv 1$	NA	8.7	0.9297
Repaired biweight	16/22	4.1	0.9828
Repaired NIRLS	22/22	3.2	0.9902

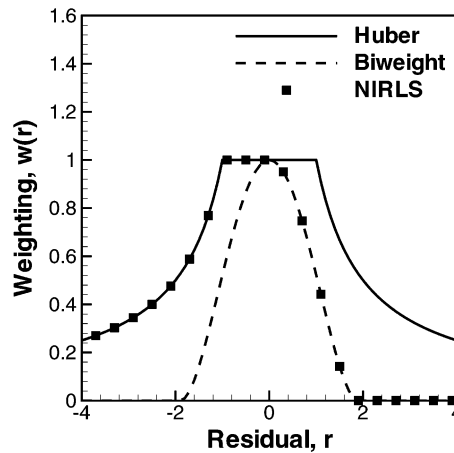


Figure 4. Weighting functions for robust  $M$ -estimation.

quadratic model has 8.7% root mean square error (RMSE) and  $R^2$  was 0.9297. NIRLS found all 22 big (error greater than 10%) outliers. The IRLS/NIRLS procedure removes the identified outliers by downweighting them, so it may result in a poor response surface approximation in the region where the outliers were removed. Therefore, consider repairing the identified outliers using  $W'_s$ , and a least squares fit to the repaired data set. Table 3 shows RMSE and  $R^2$  for the fits from such repaired data. Kim et al. (2000b) concluded that a more accurate response surface fit was obtained by repairing instead of deleting the outliers, and NIRLS performed better than IRLS in identifying outliers.

### 7. Parallel computation

The VCM and RS approximation techniques enable the efficient use of massively parallel computing to reduce the wall clock time for solving large MDO problems. Krasteva et al. (1999) successfully used a highly distributed, massively parallel computing architecture for the reasonable design space determination in the RS model construction for an HSCT design optimization problem.

The way the reasonable design space approach was applied to HSCT design was to generate a PBIB design of a specific size, evaluate all the configurations in the block design

around a nominal configuration design point using low fidelity analyses and constraints, and then move infeasible designs towards the center until they become feasible. Therefore, the number of analyses and computation time needed vary for each configuration. Since the configuration analyses are independent, they can be done in parallel. If a static distribution of the analyses is used, even though each processor gets roughly the same number of configurations initially, a load imbalance is likely to occur later in the process due to the variability of the times required to move points. Thus dynamic load balancing strategies that allow a processor to start searching for more work when it has under a certain threshold of configurations left to evaluate have to be developed. A termination detection algorithm is also needed to assert global termination of the system when all work has been performed. Krasteva et al. (1999) tested two dynamic load balancing algorithms: random polling (RP) and global round robin with message combining (GRR-MC). In RP, when a processor runs out of work, it sends a request to a randomly selected processor. This continues until the processor finds work or there is no more work in the system and termination is established. Each processor is equally likely to be selected. This is a totally distributed algorithm, and has no bottlenecks due to centralized control. A detailed description of the GRR-MC method can be found in Krasteva et al. (1999) or Tel (1994).

For termination detection, two algorithms were implemented: global task count (GTC) and token passing (TP). Following Krasteva et al. (1999), the TP algorithm is described as follows. In TP, defined as a wave algorithm for a ring topology, a token is passed around the ring where all the processors have asynchronously testified to being idle. This is not enough to claim termination, since all nodes are polled at different times, and with dynamic load balancing it is uncertain if they remained idle or later became busy. A second wave is needed to ascertain that there has been no change in the status of any processor. Termination is detected in at most two waves or  $2N$  messages ( $N$  is the number of processors available for computation), after it occurs. The total number of messages used depends on the total number of times the token is passed around the ring, but is bounded below by  $2N$ . Each processor  $P_i$  keeps track of its state in a local flag  $idle_i$ . Initially,  $idle_i$  is set to *false* if a processor starts off with some load, otherwise it is set to *true*. Consequently,  $idle_i$  is set to *false* every time a processor receives more work as a result of dynamic load balancing. A token containing a counter  $T_c$  is being passed among all processors (organized in a ring topology) in a circular fashion. Upon receiving the token, a processor holds it until it has finished all its pending work, is not expecting replies to work requests, and has made more than a certain number of unsuccessful attempts to find work. At that point,  $P_i$  checks the value of its flag  $idle_i$ . If  $idle_i$  is *true*, the processor increments the token counter  $T_c$  by 1 and passes the token along to its neighbor in the ring. If  $idle_i$  is *false*, then  $T_c$  is reset to 0 and passed along in the ring, and the value of  $idle_i$  is set to *true*. After this, if the token counter happens to be equal to the number of processors  $N$ , termination is established and all processors are notified. A rigorous correctness proof of this termination detection scheme is in Krasteva et al. (1999).

Krasteva et al. (1999) evaluated the parallel computing performance on up to 1024 processors (on an Intel Paragon XP/S 150) for all (four) combinations of dynamic load balancing and termination detection schemes. The Intel Paragon parallel computing times for low fidelity analysis of 2,026,231 HSCT designs for different numbers of processors showed a






Figure 5. Snapshot of processor utilization for RP dynamic load balancing.

35–50% improvement over a static distribution, indicating the effectiveness of the dynamic load balancing strategy. The effectiveness of fully distributed control, implemented via random polling and token passing, is illustrated in figure 5, which shows the idle and compute times during a parallel run for a portion of the processors. The corresponding figure for a static distribution has some of the processors idle for over 50% of the time. Most importantly, a dramatic decrease in the time spent identifying the reasonable design space was achieved: with 32 processors and a static distribution it took 12.8 hours to perform the low fidelity analysis, while 16 minutes was required to finish the same job with RP dynamic load balancing and TP termination detection on a 1024 processor Intel Paragon. The results convincingly demonstrate that (1) massively parallel computation can be effectively used with the RS approach to reduce the computation time in the MDO process; (2) efficient use of massively parallel hardware requires sophisticated computer science techniques such as fully distributed control, dynamic load balancing, and termination detection.

## 8. Global search methods

The use of response surface techniques permits the large number of function evaluations that are needed in order to find global optima. In addition, the response surface approximations filter out some of the numerical noise that can create spurious local optima. However, because not every analysis is replaced with a response surface approximation, some of the numerical noise remains, so that the global optimization has to contend with optima generated due to the nonconvexity of the underlying design problem as well as the spurious optima created by noise. Cox et al. (to appear) therefore examined how well different global optimizers deal with two types of local optima. Three global optimization methods were compared: the first method employs sequential quadratic programming (SQP) as implemented in DOT (Vanderplaats Research & Development, 1995) from multiple random starting points. The second method, Snyman's dynamic search algorithm (Snyman, 1983), is capable of passing through shallow local minima to locate a better optimum but still requires multiple starting points. The third method is DIRECT, a Lipschitzian optimization method (Jones et al., 1993; Watson and Baker, 2001), that needs to be run only once.

The commercial program Design Optimization Tools (DOT) (Vanderplaats Research & Development, 1995) was used to optimize the HSCT design using SQP. SQP forms a

quadratic approximation of the objective function and a linear approximation for the constraints and moves towards the optimum within given move limits. It then forms a new approximation and repeats the process until it reaches a local optimum. Due to the use of approximations, DOT is relatively quick to perform a single optimization and will handle a limited amount of noise without becoming stuck in spurious local minima. DOT has been successfully used in the past with the HSCT problem, and compared favorably to other local optimizers (Haim et al., 1999).

Snyman's dynamic search method, Leap Frog Optimization Procedure with Constraints, Version 3 (Snyman, 1983), is a semiglobal optimization method that is capable of moving through shallow local minima. The method simulates a particle rolling down a hill. As the particle moves down the gradient of the objective function, it builds momentum, which carries it out of small dips in its path. When the optimizer is moving up the gradient, a damping strategy is used to extract energy from the particle to prevent endless oscillation about a minimum. LFOPCV3 handles constraints with a standard quadratic penalty function approach.

The DIRECT algorithm (Jones et al., 1993) is a variant of Lipschitzian optimization that uses all values for the Lipschitz constant. DIRECT uses the function value at the center of each box and the box size to find the boxes which potentially could contain the optimum. A box is selected if using some Lipschitz constant  $K$  that box could contain the lowest function value. It can be shown that this requires the box to lie on the bottom part of the convex hull of the scatter plot of the points (box diameter, objective function value). DIRECT divides the potentially optimal boxes and calculates the next set of potentially optimal boxes until convergence is satisfied (see figure 6). DIRECT was found to be quick to locate regions of local optima but slow to converge. To speed up the convergence, DIRECT is stopped once the smallest box reaches a specified percentage of the original box size and DOT is used for the final optimization.

Two test functions were first used to explore the performance of the optimizers. The Griewank function was used as an example of a problem with multiple local optima caused by noise (figure 7), while a quartic function was used as a problem where there are widely separated local optima caused by a nonconvex objective function (figure 7). For these problems it was found that DOT was able to smooth out noise in the objective function through its use of approximations but it is not good at moving between widely separated local optima. LFOPCV3 was able to handle smaller amplitude noise well but was the most sensitive to increases in the amplitude of the noise and was much worse than DOT for the quartic function. DIRECT was less efficient than DOT for the Griewank function, while it was clearly the best at dealing with the quartic function and found the global optimum in every optimization.

The optimizers were then compared on the HSCT problem for a series of problems with the number of design variables ranging from five to 20. DOT and DIRECT performed about the same. They both located the same minimum weights, but for 90% confidence in locating the global optimum, DOT was cheaper for the lower dimensional cases while DIRECT became cheaper when the problem reached 20 design variables.

Baker et al. (1998, 2000) studied the use of response surface approximations in HSCT design to reduce the effect of noise and the computational effort for multiple local optimizations. This was compared with a simple version of variable complexity modeling (VCM)

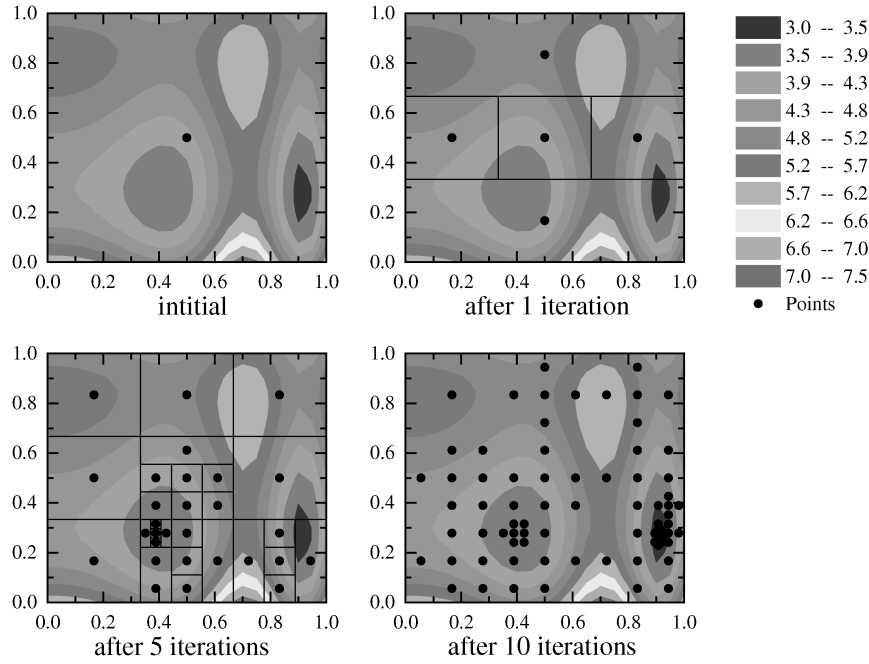


Figure 6. Example of progression of box subdivision by DIRECT.




Figure 7. Griewank (left) and quartic (right) functions in one dimension.

based on scaling of the response from a single high fidelity analysis. The response surface approximation (RSA) method was successful in reducing the impact of noise on the optimizations, and although it required a large number of analyses to create the polynomial surrogate, when multiple local optimizations were used, the RSA method was less expensive than the VCM method. However, the VCM method was better able to jump from one feasible region to another, and resulted in a better optimum than the RSA method for most of the starting points.

## 9. Conclusions

Polynomial response surface approximations have been successfully used for multidisciplinary optimization of aircraft. The practical use of polynomial surrogates with up to about 30 variables requires several techniques to combat the steep rise in computational cost associated with high dimensional polynomial approximations. (1) A reasonable design space approach using inexpensive constraints reduces substantially the volume of the approximation domain. (2) Lower fidelity analyses are used to identify good intervening functions and polynomial terms that can be discarded, thus enabling a more efficient functional representation of the response by a polynomial. (3) The approximations are computed for corrections from low fidelity to high fidelity analyses instead of for the high fidelity data directly. (4) Substantial and efficient use of parallel computation is naturally achieved for surrogate construction.

The use of response surface approximations also provides several advantages over traditional MDO techniques: (1) the integration of multiple disciplinary codes with an optimization code becomes more manageable; (2) statistical methods can be used to detect and repair bad data and to estimate the average error in the data; (3) global optimization become affordable, enabling the comparison of several global optimization algorithms.

## Acknowledgments

This work was supported in part by National Aeronautics and Space Administration Grant NAG-2-1180, Air Force Office of Scientific Research Grant F496320-99-1-0128, and National Science Foundation Grants EIA-9974956 and DMI-9979711.

## References

- N. M. Alexandrov and M. Y. Hussaini (eds.), *Multidisciplinary Design Optimization: State of the Art*, SIAM: Philadelphia, PA, 1997.
- C. A. Baker, "Parallel global aircraft configuration design space exploration," Technical Report MAD 2000-06-28, Virginia Polytechnic Institute and State University, Blacksburg, VA, 2000.
- C. A. Baker, B. Grossman, R. T. Haftka, W. H. Mason, and L. T. Watson, "HSCT configuration design space exploration using aerodynamic response surface approximations," in *Proc. of 7th AIAA/USAF/NASA/ISSMO Symp. on Multidisciplinary Anal. and Optim.*, Saint Louis, MO, 1998, pp. 769–777.
- C. A. Baker, L. T. Watson, B. Grossman, R. T. Haftka, and W. H. Mason, "Study of a global design space exploration method for aerospace vehicles," in *Proc. AIAA/USAF/NASA/ISSMO Symp. on Multidisciplinary Anal. and Optim.*, AIAA Paper 2000-4763, Long Beach, CA, September 6–8, 2000.
- C. A. Baker, L. T. Watson, B. Grossman, W. H. Mason, and R. T. Haftka, "Parallel global aircraft configuration design space exploration," *Internat. J. Comput. Res.*, to appear.
- V. O. Balabanov, "Development of approximations for HSCT wing bending material weight using response surface methodology," Ph.D. Dissertation, Virginia Polytechnic Institute and State University, Blacksburg, VA, 1997.
- V. Balabanov, A. A. Giunta, O. Golovidov, B. Grossman, W. H. Mason, L. T. Watson, and R. T. Haftka, "Reasonable design space approach to response surface approximation," *J. Aircraft* vol. 36, pp. 308–315, 1999.
- V. O. Balabanov, R. T. Haftka, B. Grossman, W. H. Mason, and L. T. Watson, "Multidisciplinary response surface model for HSCT wing bending material weight," in *Proc. 7th AIAA/USAF/NASA/ISSMO Symp. on Multidisciplinary Anal. and Optim.*, AIAA Paper 98-4804, St. Louis, MO, 1998, pp. 778–788.

- S. Burgee, A. A. Giunta, V. Balabanov, B. Grossman, W. H. Mason, R. Narducci, R. T. Haftka, and L. T. Watson, "A coarse grained parallel variable-complexity multidisciplinary optimization paradigm," *Internat. J. Super-computer Appl. High Performance Comput.* vol. 10, pp. 269–299, 1996.
- S. E. Cox, R. T. Haftka, C. A. Baker, B. Grossman, W. H. Mason, and L. T. Watson, "Global multidisciplinary optimization of a high speed civil transport," *J. Global Optim.*, to appear.
- A. A. Giunta, "Aircraft multidisciplinary design optimization using design of experiments theory and response surface modeling methods," Ph.D. Dissertation, Virginia Polytechnic Institute and State University, Blacksburg, VA, 1997.
- A. A. Giunta, V. Balabanov, S. Burgee, B. Grossman, R. T. Haftka, W. H. Mason, and L. T. Watson, "Multidisciplinary optimisation of a supersonic transport using design of experiments theory and response surface modelling," *Aeronautical J.* vol. 101, no. 1008, pp. 347–356, 1997a.
- A. A. Giunta, J. M. Dudley, R. Narducci, B. Grossman, R. T. Haftka, W. H. Mason, and L. T. Watson, "Noisy aerodynamics response and smooth approximations in HSCT design," in *5th AIAA/USAF/NASA/ISSMO Symp. on Multidisciplinary Anal. and Optim.*, AIAA Paper 94-4376, Panama City, FL, Sept. 1994.
- A. A. Giunta, O. Golovidov, D. L. Knill, B. Grossman, W. H. Mason, L. T. Watson, and R. T. Haftka, "Multidisciplinary design optimization of advanced aircraft configurations," in *Lecture Notes in Physics*, 490, P. Kutler, J. Flores, and J.-J. Chattot, eds., Springer-Verlag: Berlin, pp. 14–34, 1997b.
- O. Golovidov, "Variable-complexity response surface approximations for aerodynamic parameters in HSCT optimization," MS Thesis, Virginia Polytechnic Institute and State University, Blacksburg, VA, 1997.
- D. Haim, A. A. Giunta, M. M. Holzwarth, W. H. Mason, L. T. Watson, and R. T. Haftka, "Comparison of optimization software packages for an aircraft multidisciplinary design optimization problem," *Design Optim.* vol. 1, pp. 9–23, 1999.
- P. H. Holland and R. E. Welsch, "Robust regression using iteratively reweighted least squares," *Comm. Stat.: Theory and Methods* vol. 6, pp. 813–827, 1977.
- M. G. Hutchison, W. H. Mason, R. T. Haftka, and B. Grossman, "Aerodynamic optimization of an HSCT configuration using variable-complexity modeling," in *AIAA 31st Aerospace Sciences Meeting and Exhibit*, AIAA Paper 93-0101, Reno, NV, 1993.
- M. G. Hutchison, E. R. Unger, W. H. Mason, B. Grossman, and R. T. Haftka, "Variable-complexity aerodynamic optimization of a high-speed civil transport wing," *J. Aircraft* vol. 31, no. 1, pp. 110–116, 1994.
- D. R. Jones, C. D. Perntunen, and B. E. Stuckman, "Lipschitzian optimization without the Lipschitz constant," *J. Optim. Theory Appl.* vol. 79, no. 1, pp. 157–181, 1993.
- M. Kaufman, V. Balabanov, S. L. Burgee, A. A. Giunta, B. Grossman, R. T. Haftka, W. H. Mason, and L. T. Watson, "Variable-complexity response surface approximations for wing structural weight in HSCT design," *Comput. Mech.* vol. 18, pp. 112–126, 1996.
- H. Kim, R. T. Haftka, W. H. Mason, L. T. Watson, and B. Grossman, "A study of the statistical description of errors from structural optimization," in *Proc. 8th AIAA/USAF/NASA/ISSMO Symp. on Multidisciplinary Anal. and Optim.*, AIAA paper 2000-4840, Long Beach, CA, 2000a, 18 pages.
- H. Kim, M. Papila, W. H. Mason, R. T. Haftka, L. T. Watson, and B. Grossman, "Detection and correction of poorly converged optimizations by iteratively reweighted least squares," in *Proc. 41st AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conf.*, AIAA paper 2000-1525, Atlanta, GA, 2000b, 18 pages.
- D. L. Knill, A. A. Giunta, C. A. Baker, B. Grossman, W. H. Mason, R. T. Haftka, and L. T. Watson, "Response surface models combining linear and Euler aerodynamics for supersonic transport design," *J. Aircraft* vol. 36, no. 1, pp. 75–86, 1999.
- D. T. Krasteva, C. Baker, L. T. Watson, B. Grossman, W. H. Mason, and R. T. Haftka, "Distributed control parallelism in multidisciplinary aircraft design," *Concurrency: Practice and Experience* vol. 11, pp. 435–459, 1999.
- R. M. Lewis, V. Torczon, and M. W. Trosset, "Direct search methods: Then and now," *J. Comput. Appl. Math.* vol. 124, pp. 191–208, 2000.
- P. MacMillan, O. Golovidov, W. Mason, B. Grossman, and R. Haftka, "Trim, control, and performance effects in variable-complexity high-speed civil transport design," Technical Report MAD 96-07-01, Virginia Polytechnic Institute and State University, Blacksburg, VA, 1996.

- P. E. MacMillin, O. B. Golovidov, W. H. Mason, B. Grossman, and R. T. Haftka, "An MDO investigation of the impact of practical constraints on an HSCT optimization," in *AIAA 35th Aerospace Sciences Meeting and Exhibit*, AIAA Paper 97-0098, Reno, NV, 1997.
- L. A. McCullers, "Aircraft configuration optimization including optimized flight profiles," in *Proc. Symp. on Recent Experiences in Multidisciplinary Analysis and Optimization*, NASA CP-2327, J. Sobieski, ed., pp. 396–412, 1984.
- W. D. McGrory, D. C. Slack, M. P. Applebaum, and R. W. Walters, *GASP Version 2.2 Users' Manual*, Aerosoft, Inc., Blacksburg, VA, 1993.
- M. Papila and R. T. Haftka, "Uncertainty and wing structural weight approximations," in *40th AIAA/ASME/ASCE/ASC Structures, Structural Dynamics, and Material Conf.*, AIAA Paper 99-1312, St. Louis, MO, 1999.
- P. J. Rousseuw and A. M. Leroy, *Robust Regression and Outlier Detection*, John Wiley: New York, NY, 1987.
- M. Snir, S. Otto, S. Huss-Lederman, D. W. Walker, and J. Dongarra, *MPI: The Complete Reference*, vol. 1: The MPI Core, MIT Press: Cambridge, MA, 1998.
- J. Sobieszczanski-Sobieski and R. T. Haftka, "Multidisciplinary aerospace design optimization: Survey of recent developments," *Structural Optim.* vol. 14, no. 1, pp. 1–23, 1997.
- J. A. Snyman, "An improved version of the original Leap-Frog dynamic method for unconstrained minimization: LFOP1(b)," *Appl. Math. Modelling* vol. 7, pp. 216–218, 1983.
- G. Tel, *Introduction to Distributed Algorithms*, Cambridge University Press: Cambridge, UK, 1994.
- Vanderplaats, Miura, and Associates, Inc., *GENESIS User Manual*, Version 1.3, Goleta, CA, 1993.
- Vanderplaats Research & Development, Inc., *DOT: Design Optimization Tools*, Version 4.20, Goleta, CA, 1995.
- L. T. Watson and C. A. Baker, "A fully-distributed parallel global search algorithm," *Engrg. Comput.* vol. 18, no. 1, pp. 155–169, 2001.