45th AIAA Aerospace Sciences Meeting and Exhibit
8 - 11 January 2007, Reno, Nevada

AIAA 2007-361

Surrogate-Based Design Optimization of a Large Asymmetric Launch Vehicle Payload Fairing

Patrick H. Reisenthel[1] and Robert E. Childs[2]
Nielsen Engineering & Research, Inc., Mountain View, California 94043

and

John E. Higgins[3]
United States Air Force, Kirtland AFB, New Mexico 87117

The problem of designing a payload fairing large enough to encompass a spacecraft with an optical mirror up to twice the diameter of an expendable launch vehicle (EELV) is considered. Optimization and computational fluid dynamics (CFD) methods are used to design an asymmetric exterior shape for the fairing that meets specific aerodynamic design goals. These aerodynamic goals include a balance of low lateral force on the fairing and smooth variations in that force with respect to angle of attack across a range of Mach numbers near Mach 1. Multiple design objectives are met by means of subsystem metamodels which are combined in a single performance index. Finally, a discussion is given of the combination of individual subsystem uncertainties to assess global uncertainty on the objective, and how these might affect the optimization strategy.

Nomenclature

C_m	=	pitching moment coefficient
M	=	Mach number
OBJV	=	objective function value
ref	=	reference quantity
RSS	=	root sum of squares
U	=	uncertainty on the objective function
α	=	angle of attack
δ	=	uncertainty on the response surface
ϕ	=	roll angle

I. Introduction

In their review of the status of MDO, Lewis and Mistree[1] noted that the design and optimization of complex multidisciplinary systems was experiencing a shift in which multidisciplinary tradeoffs and detailed analysis were carried out at progressively earlier stages of the design cycle. Because the knowledge of the system is imprecise at this stage, a corresponding, parallel, shift was described in which fuzzy heuristics and precise mathematical models needed to be used together. The design optimization method used in this paper can be viewed, in a sense, as the product of these shifts. High-fidelity computer simulations are brought in early in the design cycle through the use of fast-running metamodel approximations of each subsystem. These dynamically evolving metamodels are used as surrogates for the more expensive computer analyses. Because of the efficiency of the metamodel evaluations, this approach additionally lends itself to the comprehensive characterization of uncertainty. The resulting method can be described as a "probabilistic metamodel" allowing the assessment of design uncertainties and risk between design alternatives.

Global metamodels and response surface technology are increasingly used in a variety of fields, including structural reliability, instrument calibration, and aerodynamic and trajectory optimization, to name a few.[2-11] Because of their analytical nature, these models can be used for automated searches and are naturally well-suited to

[1] Chief Scientist, 605 Ellis Street, Suite 200, Mountain View, CA 94043, Senior Member.
[2] Vice President, 605 Ellis Street, Suite 200, Mountain View, CA 94043, Senior Member.
[3] R&D Engineer, 3550 AFRL/VF, Aberdeen Avenue SE, Kirtland AFB, NM 87117, Member.

the acceleration of optimization tasks[12-14] and rapid strategy evaluation. A central issue to constructing appropriate response surface models is the so-called curse of dimensionality, in which the number of data points required to characterize/support the surface increases exponentially with the number of independent variables. This difficulty is well-known and, in effect, precludes the use of conventional schemes, such as polynomial and piecewise-polynomial (finite-element) approximations. Neural networks,[15,16] support vector machines,[17] kernel methods,[18,19] and multidimensional splines[20] all have proven capable of multidimensional data generalization and of (partially) circumventing this difficulty. The particular approach used in this study is based on self-training radial basis function networks which form the basis of the NEAR-RS (response surface) technology,[21] a software system consisting of two modules: a metamodel identification module, and a metamodel evaluation/interrogation module. A graphical user interface included in this second module serves as a multidimensional viewer, which facilitates the visualization of trends in high-dimensional data. A key aspect of the technology is the ability to estimate further sampling needs and model quality, based on uncertainty estimation. The uncertainty estimation in NEAR-RS is based on propagating statistical descriptions of uncertainty in measurements or input data to estimates of uncertainty in the response surface coefficients themselves. This approach (described in Ref. 21) makes use of the covariance of the output measurements and is based on the theory of best linear unbiased estimation. This technology provides a rational basis for propagating uncertainty estimates throughout the design space of each subsystem. The subsystem metamodels are then combined in a single performance index, or objective function, which is used to guide the design. The last portion of the paper describes how the individual subsystem uncertainties are combined to assess the global uncertainty on the objective, and how these might affect the optimization strategy.

II. Large Asymmetric Launch Vehicle Payload Fairing Design

This paper concerns the aerodynamic and structural design of a payload fairing large enough to encompass a spacecraft with an optical mirror up to twice the diameter on an evolved expendable launch vehicle (EELV).[22]

A.Objective and Methods

The main design objectives at the early stages of the work were related to stability and control, as well as the mass of the fairing. Optimization and computational fluid dynamics (CFD) methods were used to design an asymmetric exterior shape for the fairing that meets specific aerodynamic design goals. These aerodynamic goals include a balance of a low lateral force on the fairing and smooth variations in that force with respect to angle of attack, across a range of Mach numbers near Mach 1.0.

The fairing surface was defined principally in terms of analytical functions. The cost to perform design optimization increases significantly with the number of independent design variables; therefore, it was important to minimize the number of control variables needed to parameterize the surface shape. This can be accomplished with an analytical surface shape definition, rather than a surface defined by discrete points. Also, the approach sought to impose the designer's understanding of aerodynamics on the optimization process, and to focus the design on shapes that would trigger flow separation gradually, to alleviate excessive aerodynamic forces on the fairing.

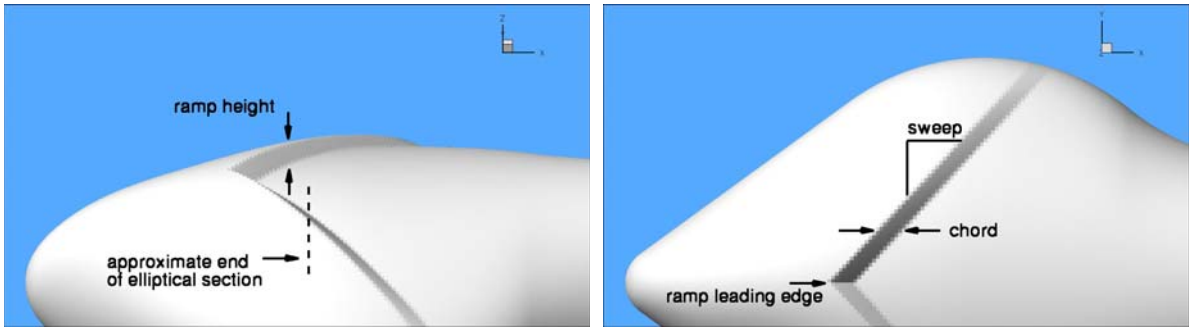

Figure 1. Side and top views of payload fairing geometry.

Three generations of the aerodynamic design have been considered to date. Only the first two are discussed in this paper. In the first-generation family of designs, the payload fairing geometry was initially parameterized using nine free design variables, corresponding to a two-dimensional (unswept) version of the ramp shown in **Fig. 1**. The nine control variables were the longitudinal minor axis and focus for the nose, the depths of the minor and major axes, three variables controlling the height, location, and slope of the separation trigger, and two geometry transition parameters. A main outcome of optimizing this first-generation family of designs was to confirm that flow separation was needed in order to achieve the aerodynamic goals. The geometry of the second-generation family of

designs is illustrated in Fig. 1, which shows side and top views of the three-dimensional payload fairing geometry, along with some of the geometric design parameters. This fairing geometry was controlled by five independent variables, which were (1) the length of the elliptical section of the nose (in side view), (2) the location of the apex of the swept rearward-facing ramp (denoted "ramp leading edge" in the top view), (3) the ramp sweep angle (parameterized by its tangent), (4) the length of the backward-facing ramp ("chord" in top view), and (5) the height of the ramp.

Because the flow around the payload fairing involves flow separation, a Navier-Stokes CFD methodology (Overflow[23]) was used to predict the steady-state forces and moments on the launch vehicle. The turbulence model used was the Spalart-Allmaras model.[24] The active independent variables were used to automatically remesh the vehicle's surface and recompute a new volume grid using Hypgen.[25] Grid convergence studies were conducted to determine the optimal grid resolution needed. Using this grid resolution, individual CFD calculations took between two and six hours, depending on the rate of convergence, which was affected by flow physics.

B.Approach

A global optimization scheme (described below) was applied to a continuous, analytical surrogate, rather than the CFD calculations themselves. This allows a thorough search of the solution space without the cost of performing actual CFD solutions. Naturally, the approach relies on the response surface being representative of the performance of the designs, which may not be the case, particularly at the early stages of the design, when the uncertainty is large (see discussion below). Nevertheless, our sampling or "database populating" strategy based on minimizing the uncertainty only in promising areas of the design space was found to yield solutions satisfying the problem constraints in a relatively small number of CFD calculations. Specifically, the response surfaces corresponding to multiple dependent variables were combined into a single objective function on which the search was performed. The individual dependent variables were the overall pitching moment, C_m, at multiple aerodynamic conditions, each a specified combination of Mach number, angle of attack, and, possibly, roll angle. At Mach 1.0, it was found that the load C_m could be minimized relatively easily, to the point of zero load or even load reversal. Such a condition, however, was the result of massive flow separation over the surface of the fairing and, consequently, tended to exhibit poor aerodynamic characteristics (i.e., a nonsmooth response with respect to angle of attack). Another problem (not shown) was that minimizing the load at Mach 1.0 (the nominal specification) frequently resulted in configurations that exhibited excessive loads at slightly supersonic Mach numbers, due to the reduced extent of flow separation in supersonic flow. Both difficulties were addressed by carrying out the optimization based on a measure of global aerodynamic performance, rather than on a single load. This measure of global aerodynamic performance was specified as a multiobjective function designed to penalize the undesired behaviors. The objective function was a linear combination of multiple "bucket" functions, one per explicitly optimized aerodynamic condition. For each aerodynamic condition, the bucket function was designed to be minimal near a specified target value $C_{m,target}$, to return high values above a constraint $C_{m,max}$, and to rise again if the C_m value was too low.

For example, for the first-generation family of designs a two-objective function was defined initially, based on loads simultaneously optimized at the aerodynamic conditions $(M, \alpha, \phi)_{aero1}=(1.0, 5^o, 0^o)$ and $(M, \alpha, \phi)_{aero2}=(1.2, 5^o, 0^o)$. Note that, of the initially nine design variables of the first-generation design, three were found to have minimal effect on the aerodynamic performance and were, consequently, set to fixed values. The remaining active independent variables defined a six-dimensional parameter space, which was initially seeded using Latin Hypercube sampling (LHS) design-of-experiments methods.[26] The results of these seed calculations were used to form the initial response surfaces (one for each aerodynamic condition). A global search of the objective function over the response surfaces would then "suggest" the next set of CFD runs based on the minimization of the global performance index. This new batch of CFD calculations would then be used to populate and enrich the design space, resulting in a new "design iteration" in this process. The global search performed at each design iteration was the result of an exhaustive "scatter and poll" search strategy[27] on the global performance metamodel. The particular implementation of the scatter and poll strategy used here involved parallel searches, each started from random locations on the surface (the "scatter" phase) and, from each of these random locations, uses a local nongradient-based direct search[28] employing an interrogation stencil that was either translated or focused (contracted), depending on the results of the evaluations on the stencil (the "poll" phase).

A mixed strategy was used to enrich the data set (i.e., compute further CFD solutions). It was found in the early stages of design that the suggested optima were often on the boundary of the domain. Since some of the range limits in the control variables were not "hard" limits, their ranges were expanded, allowing the optimizer to "explore" beyond the initial boundaries by populating the data set in these areas. Each new batch of CFD calculations used

LHS randomization for the independent variables, running 10 to 16 solutions at a time on a cluster of Linux workstations. By contrast, the intermediate and later stages of the design were characterized by a *focusing* ("exploitation") of the parameter space on the most promising regions, at a given iteration of the design. Note that these regions of interest did not always remain the most promising ones, as further CFD calculations would either confirm or invalidate the trend predicted by the response surface at the previous design iteration.

C.Results

Figure 2 illustrates the result of the optimization procedure for the first-generation family of designs, in the form of a cumulative plot of the calculated loads $C_{m,aero1}$ and $C_{m,aero2}$ after several design iterations. The plot insets show the front and side views of the initial and final geometric designs. In this example, the various symbols are associated with the various design iterations, in which successive waves of CFD data points were harvested in the process. The results improve at successive design iterations because the metamodel is cumulatively enhanced each time by the new knowledge provided by the latest CFD runs. In this case, the objective was to reduce each pitching moment coefficient C_m to less than twice the corresponding load for a conventional fairing used on the Boeing Delta 4 Heavy. Thus, any point such that $(C_m / C_{m,ref})_{aero1} < 2.0$ and $(C_m / C_{m,ref})_{aero2} < 2.0$ (i.e., inside the dashed line in Fig. 2) constitutes an acceptable aerodynamic design. The main point of Fig. 2 is that, after just 10 design iterations (less than 230 geometric configurations), a large number of acceptable designs were found, guided by the cumulatively enhanced metamodel. This rather small number of CFD runs (for a six-dimensional parameter space and a highly nonlinear response) is to be contrasted with the total number of metamodel interrogations used to drive the optimization, which is on the order of three million.

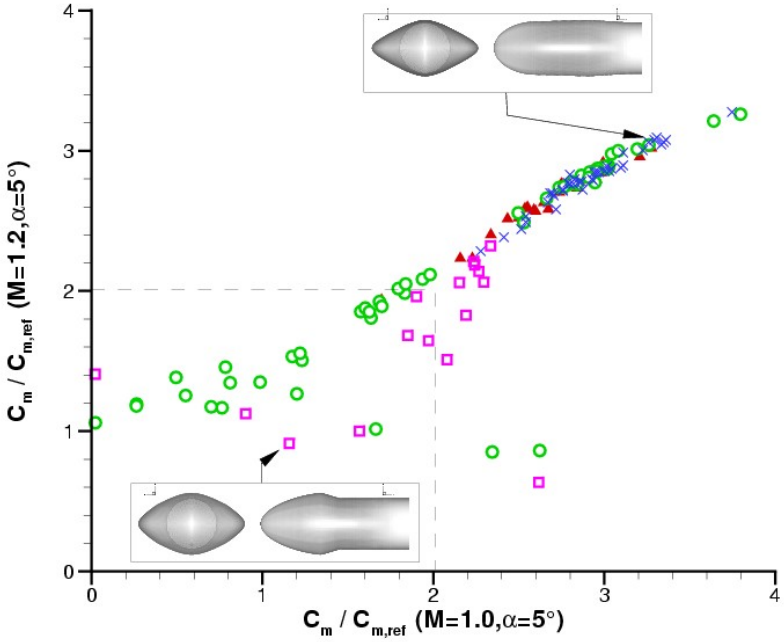

Figure 2. Multiobjective optimization using cumulatively enhanced aerodynamic metamodels.

While the optimization exercise shown in Fig. 2 was largely successful at the explicitly optimized aerodynamic control points, the presence of (and, indeed, reliance on) flow separation to reduce the moment imply that the aerodynamic characteristics of the vehicle may be nonlinear. A consequence of this nonlinearity is that one cannot rely on load reduction at one angle of attack to force a corresponding load reduction at all other angles of attack, as would be the case for a configuration with linear aerodynamics. This was verified for the nominally "acceptable" solutions identified in Fig. 2, the majority of which exhibited highly nonlinear behavior, such as load sign reversals at intermediate and nearby angles of attack.

A somewhat brute-force solution to this problem is to include additional aerodynamic constraints as part of the objective function. However, the addition of such constraints, when spread over a large range of Mach numbers and angles of attack, would be computationally prohibitive, since the number of CFD runs per geometric configuration would grow accordingly. Parsimonious use of aerodynamic control points in the objective function is a critical concern. To help choose candidate control points, aerodynamic sweeps were calculated for a number of geometries.

These led to the following observations: (1) relaxing the target load from $C_m / C_{m,ref} \approx 1.0$ (50% safety margin) to $C_m / C_{m,ref} \approx 1.8$ (10% safety margin) reduces angle-of-attack nonlinearity, and (2) at supersonic speeds, the loads behave approximately linearly. These observations led to the definition of a new multiobjective bucket function seeking to proportionately constrain the Mach 1.0 aerodynamic loads at two angles ($\alpha = 5^o$ and $\alpha = 1^o$), to specify the $\alpha = 5^o$ target loads at Mach 1.2 and 1.0 to 15% and 45% margin levels, respectively, relative to the reference fairing at the same conditions, and to add a stability (effective slope) constraint at Mach 0.9 and $\alpha = 1^o$. In essence, the new objective function sought to achieve load characteristics which mimicked those of the reference vehicle, multiplied by 1.1 at $(M, \alpha, \phi)_{aero1}=(1.0, 5^o, 0^o)$, and 1.7 at $(M, \alpha, \phi)_{aero2}=(1.2, 5^o, 0^o)$, plus two linearizing conditions $(M, \alpha, \phi)_{aero3}=(1.0, 1^o, 0^o)$ and $(M, \alpha, \phi)_{aero4}=(0.9, 1^o, 0^o)$. The optimization procedure described above was repeated for the case of four simultaneous objectives, and successfully resulted in several candidate solutions, starting shortly after two hundred geometric configurations. The angle-of-attack characteristics of two of these solutions are illustrated in **Fig. 3**.

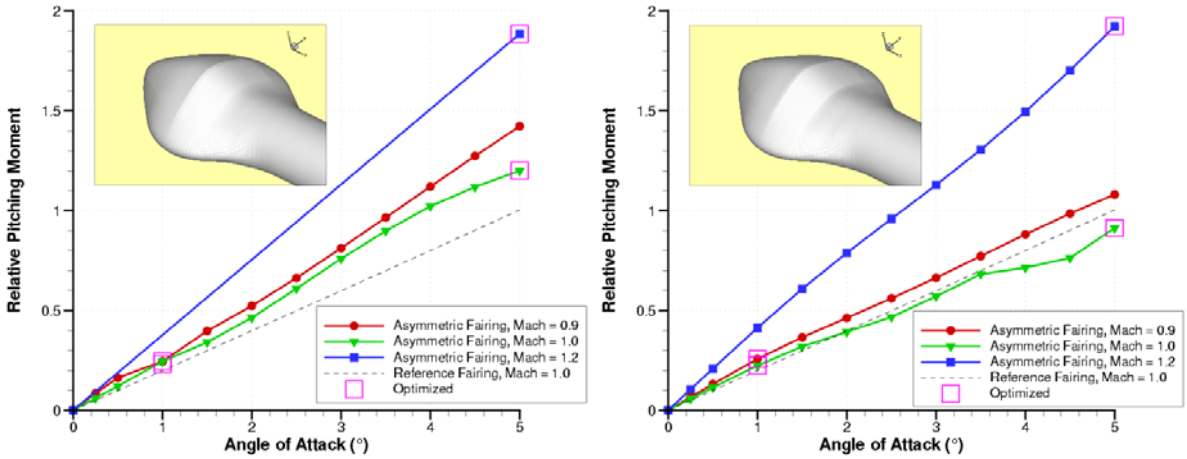

Figure 3. Two examples of aerodynamic characteristics achieved as a result of a 4-point optimization.

Figure 3 shows the relative pitching moment, defined as the pitching moment divided by that of the reference vehicle at $(M, \alpha, \phi)_{aero1}=(1.0, 5^o, 0^o)$. The Mach 1.0 aerodynamics of the reference vehicle are indicated by the dashed line. The symbols in Fig. 3 represent the results of converged CFD calculations, with the roll angle ϕ kept equal to zero. Each plot depicts four square boxes which indicate the conditions at which the vehicle response was explicitly optimized by way of the objective function. The insets provide a side-rear view of the two fairing geometries, which look similar to the naked eye but clearly result in different aerodynamics, particularly in the transonic range. Figure 3 shows that the near-linearization objective for the aerodynamics was accomplished, and that this objective could be repeated on more than one vehicle.

While the steady aerodynamics objectives were achieved, the presence of a large two-dimensional rearward-facing ramp in the optimized design was deemed problematic in terms of its potential for buffet. Consequently, a second family of fairing designs was considered, in which the flow separation trigger consisted of a V-shaped swept rearward facing ramp, as shown in Fig. 1. Also, stricter limitations were set of the height of the separation ramp, and more stringent convergence requirements were imposed on the CFD calculations. The optimization procedure described for Fig. 2 was repeated, with the exception that the initially set ranges for the control variables were not expanded, i.e., they were considered as "hard" limits. The results of the optimization procedure are summarized in **Fig. 4**.

Figure 4 provides a cumulative plot of the aerodynamic performance, as measured by the pitching moments at two target aerodynamic conditions: $(M, \alpha, \phi)_{aero1}=(0.9, 5^o, 0^o)$ and $(M, \alpha, \phi)_{aero2}=(1.2, 5^o, 0^o)$. These two conditions were included as part of the objective function, with the minimum in the objective function being indicated in the plot by the + symbol. A flooded contour plot of the objective function is shown in **Fig. 5**, with contours saturated at a value of 5, for ease of interpretation. By design, the objective function ranges from 0 to 20 in value, saturates at 2 when the loads are too low, saturates at 20 when the loads are high, and exhibits a steep rise (strong penalty) a short distance above the well (the region near the zero value minimum), located approximately at the 10% and 15% margin conditions *aero1* and *aero2*, respectively. The white contour in Fig. 5 indicates the region

of acceptable performance, ideally defined as an objective function value less than 0.5. In practice, regions with objective function values up to 1.0 were considered worthy of further investigation.

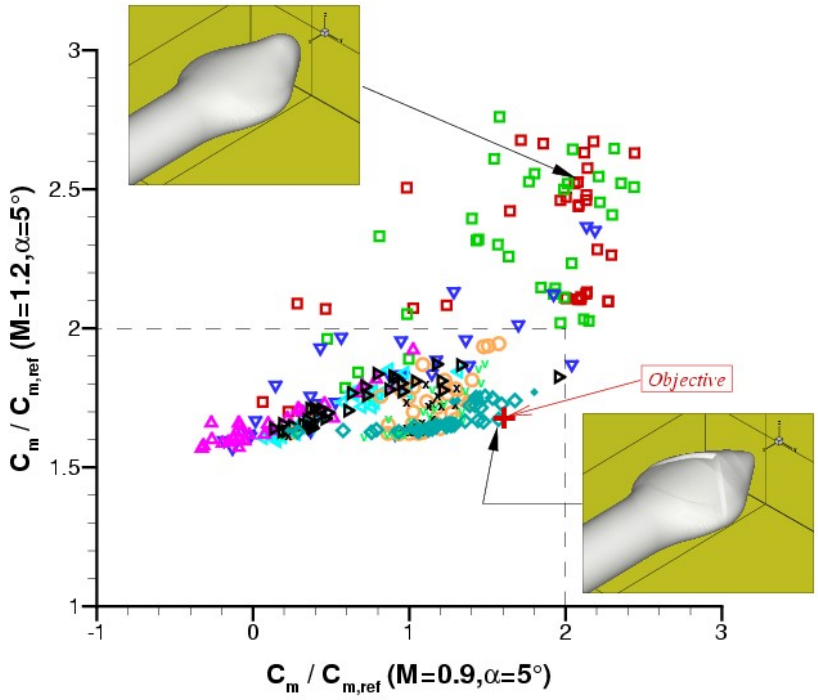

Figure 4. Cumulative plot of two-point performance optimization for chevron-shaped design.

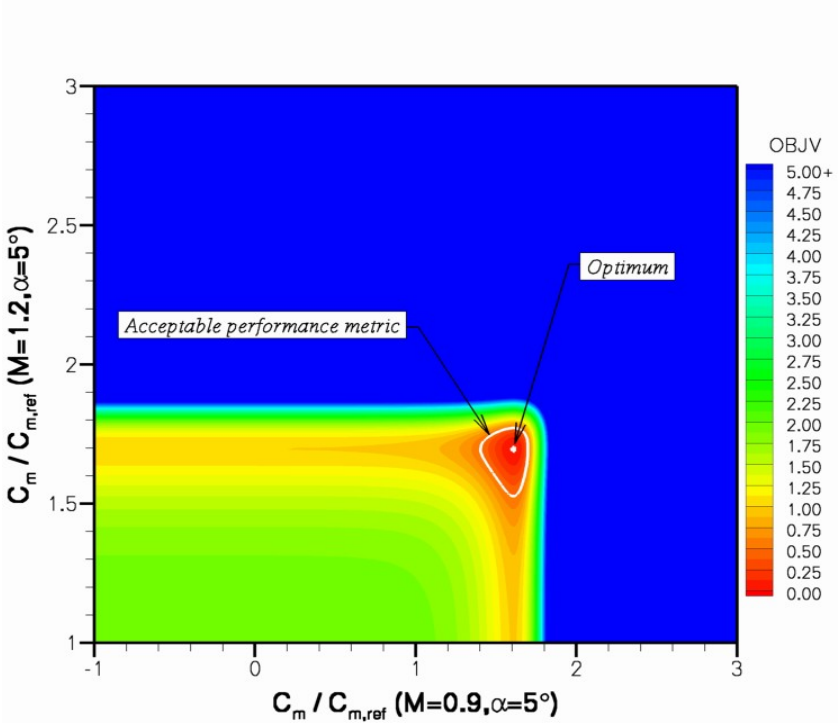

Figure 5. Plot of objective function value for optimization results shown in Fig. 4.

The results shown in Fig. 4 represent 320 fairing configurations, where different symbols indicate successive design iterations, and the insets provide a perspective view of the initial and final geometries. An initial 64 seed configurations were used to populate the five-dimensional design space, and form an initial response surface for

global searching and direction discovery. These initial calculations represented in this case an initial investment of 20% of the computational resources used. This number is, of course, problem dependent, but compares favorably with the Sóbester et al's[13] rule-of-thumb recommendation of 35% of the available computational budget as a safe choice of initial Latin Hypercube sample size.

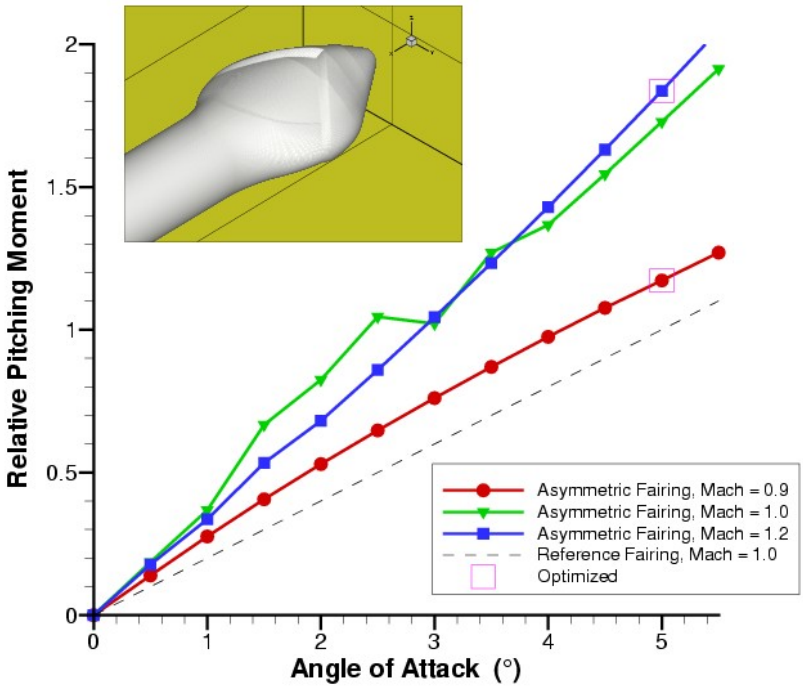

Figure 6. Angle-of-attack aerodynamic characteristics for two-point optimization.

Although a near-linearity objective for the angle-of-attack aerodynamics was not explicitly formulated in this case, owing to the computational burden this would have entailed, it is interesting to note (see **Figure 6**) that near-linearity did result at the two Mach numbers (0.9 and 1.2) that were explicitly specified in the objective function. Similarly to Fig. 3, the load curves depicted are nondimensionalized by the reference vehicle pitching moment at $(M, \alpha, \phi) = (1.0, 5^{\circ}, 0^{\circ})$. For consistency with previous results, the optimized vehicle response at Mach 1.0 is also shown in the figure, indicating a lack of smoothness that was later found to be characteristic of other Mach numbers.

D.Optimization Strategy

The criteria or strategies used to enrich the design space in successive iterations of the optimization were based on aerodynamics expertise, but also, to some extent, on heuristics, and on decisions made at each iteration of the response surface metamodels. Therefore, the results shown in this paper are not those of a fully automated system but, rather, of a system involving an expert-in-the-loop. Future versions of this system will be directed toward greater automation, for faster turnaround and, perhaps more importantly, more precise documentation and archiving of the processes leading to a given design. The asymmetric payload fairing design example described in this paper is but one illustration of surrogate-based optimization, for which there are many examples in the literature. In particular, the reader is referred to the seminal work of Jones et al[12] concerning efficient global optimization when dealing with limited quantities of expensive computational data, in which a closed-form expected improvement function is used to balance local and global searches of the design space. This balance between the need to continue searching the overall space (exploration) and the desire to focus in on a local optimum (exploitation) has also been recognized in more recent studies[13] emphasizing the possibility of automating these procedures for large problems.

The three optimization examples shown in this paper differ from each other by the objectives specified, the initial seeding of the calculations, the dynamic parameter bound settings, and considerations of global *vs.* semilocal metamodels for exploitation. What they have in common is a strategy primarily focused on exploitation, rather than

a continued balance between exploitation and exploration. Aside from an initial design-of-experiments investment in calculations aimed at a coarse characterization of the design space, the successive waves of design space enrichment were based on exploiting the optima suggested by the response surfaces obtained at any given iteration of the design. Typically, only one or two regions (three at the very most) of global optimality could be "followed," namely, focused-on for further CFD calculations. As will be seen below, adding computational points in promising regions of the design space amounts to an uncertainty reduction strategy in those areas. If the results improve, i.e., if the objective function is reduced, then the most promising subregions of *that* surface are followed, and so on. This results in a focusing action, equivalent to a "numerical zooming" procedure in a complicated multidimensional space. If, on the other hand, the new CFD results prove to degrade the objective function performance, then backtracking to the range covered by the previous response surface iterate becomes necessary. Unfortunately, this results in a complicated decision tree structure which (a) is potentially difficult to manage, and (b) becomes quickly unaffordable. It is suggested that one of the keys to unlock this difficulty is the use of the metamodel uncertainty.

As mentioned previously, the radial basis function networks which form the basis of the metamodeling used in this study permit the propagation of uncertainty, as described in Ref. 21. In the present examples, the metamodel inputs are the results of the CFD calculations, which act as support vectors, and the output is a deterministic response surface based on these vectors. It is anticipated that, in the not-so-distant future, analysis codes, such as CFD, will yield comprehensive uncertainty bars bounding the nominal predictions. In situations where the uncertainty on the input vectors is known, the response surface then becomes a probabilistic response surface, given in the form of a mean (nominal) response surface and its variance. Because of the computational efficiency of the constructed metamodels, Monte Carlo simulations of the combined, multiobjective uncertainty calculations can be carried out in near real time, providing rapid feedback on the objective function and its potential for accuracy.

To illustrate how considerations of global metamodel uncertainty might affect the design strategy, consider the following examples, all derived from the last optimization example. **Figure 7** depicts the objective function along a one-dimensional cut of the design space. For ease of representation, the one-dimensional cut was chosen to parallel the sweep variable axis, and to go through the design optimum (the point at which the objective function is minimized, after 320 configurations). In the figure, this minimum is located around *sweep* = 0.76. Note that the choice of the *sweep* variable is unimportant, as a similar reasoning can be applied to any of the control variables. Figure 7 is shown without uncertainty, and is fairly representative of the optimization work to date. Without any additional information, all regions marked in red (objective function less than 0.5) are regions where possible optima might exist.

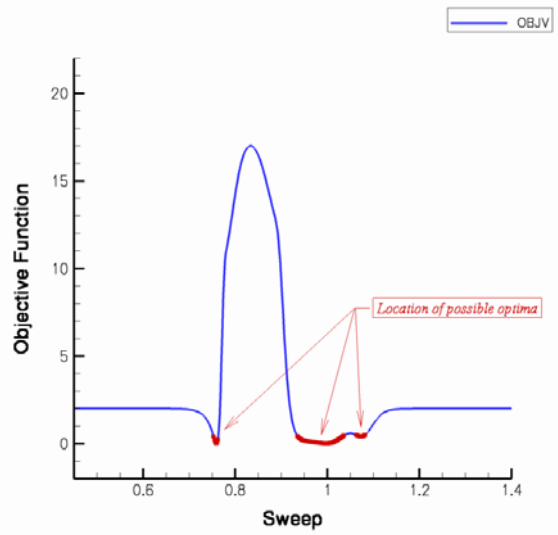

Figure 7. Implications of a search strategy based on response surface alone.

A different picture emerges when one takes into account the uncertainty information. **Figure 8** shows the same cut of the objective function (solid line), with the addition of the plus and minus uncertainty. The latter are depicted as dashed lines in the figure, and were computed from 10,000 Monte Carlo samples, assuming input (CFD) Gaussian probability distributions of uniform variances 0.043 and 0.035 for the $C_m / C_{m,ref}$ loads at Mach 0.9 and Mach 1.2, respectively. It is immediately apparent that the uncertainty on the objective function is substantial. Only around *sweep* = 0.76 is there low uncertainty, owing to the presence of an actual CFD data point there. Since the uncertainty on the response surfaces is strongly correlated with the sampling (the further away from a support vector, the greater the uncertainty), and because of the intentional strongly nonlinear nature of the objective function, any load combination that either exceeds the objective function target loads, or lies far away from any real data point, will produce a large uncertainty on the objective function. Based on this large uncertainty, the message of Fig. 8 is that, if one considers the lower uncertainty bound to indicate the regions where suitable optima may still be found (OBJV-U < 0.5), then the range of sweep (marked in red) is considerably greater than suggested in Fig. 7. This scenario represents one of the extremes, in which no point or region may be excluded from exploration until *proven* unsuitable.

Another extreme is to consider a strategy in which only points or regions with good performance *and* low uncertainty are pursued. This leads to a strategy of focusing, typical of the later stages of optimization in the examples shown in this paper. Such a criterion is illustrated in **Fig. 9**, where the range of points deemed worthy of exploitation (indicated in red) now only represents 1% of the overall range in *sweep*. Figure 9 is evocative of the proverbial "finding a needle in a hay stack," which is what the objective function-driven optimization was designed to do. This is to be contrasted to the strategy of exploration, more appropriately exemplified in Fig. 8, but for which the range of sweeps still remaining to be explored is 60% of the total range. While case dependent, this 60:1 ratio (in a single, one-dimensional cut) certainly gives a feel for the severity of the problem, when confronted with the task of implementing a viable uncertainty-based strategy in multiple dimensions. Yet, both aspects (exploration and exploitation) are necessary. Without exploitation, there is no rationale for the finding of an optimum (except for random chance). And without exploration, there is every chance of being stuck in a local optimum, or finding no solution at all. Additionally, even a pure exploitation strategy cannot be started as such, and needs an initial exploration phase to get the process initialized (as was done in the optimization cases presented in this paper). It is clear from the examples given above that a strategy that primarily emphasizes exploration is a waste of computational resources. It is also evident that one based on exploitation cannot discover anything new, is conditioned by the initial seeding of the design space, and is destined to end with the first promising region found, all at a considerable lost potential for design opportunities. The ideas of Jones et al[12] and Sóbester et al,[13] which call for a proper balance of exploration and exploitation, are most attractive and provide a rational framework for the allocation of computational resources. In the end, time and money limitations may preclude the implementation of an "ideal" global optimization strategy, assuming even that such a strategy can be defined. Nevertheless, the concept of using the response surface (metamodel) uncertainty constitutes a potentially important step in this direction.

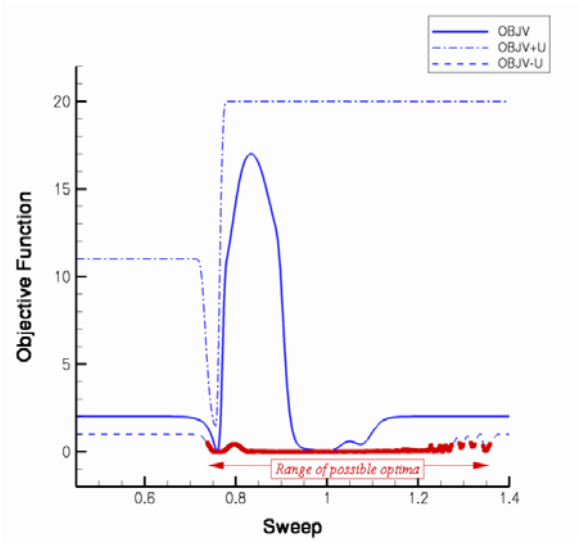

Figure 8. Implications of a search strategy based on rejection of "impossible" areas.

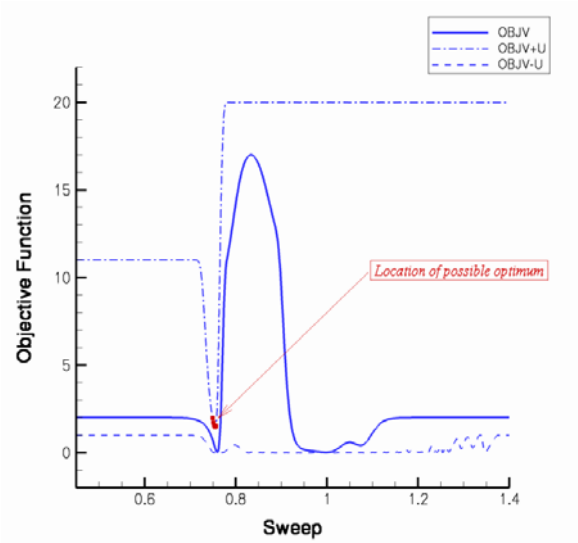

Figure 9. Implications of a search strategy based on exploitation of confirmed "possible" areas only.

As a final observation, we note that the very objective function which successfully forced the global search to focus on an acceptable solution, by virtue of its strong nonlinearity also easily "saturates" the uncertainty, creating a loss of information in the process. This concept is illustrated in **Figure 10**. The lower graph depicts, as before, the nominal objective function, plus and minus its uncertainty, much of it saturated between 0 and 20. The upper graph, on the other hand, depicts the uncertainty on the response surfaces themselves. Recall that two response surfaces feed into the nonlinear objective function. Each is a five-dimensional function of the geometric control variables; one represents the response at $(M, \alpha, \phi)_{aero1}=(0.9, 5^o, 0^o)$, and the other the response at $(M, \alpha, \phi)_{aero2}=(1.2, 5^o, 0^o)$. For ease of representation, these uncertainties are combined in the RSS (root sum of squares) sense and depicted as a single output, one-dimensional cut of

$$RSS(\delta(C_m/C_{m,ref})) \stackrel{\text{def}}{=} \sqrt{(\delta(C_m/C_{m,ref}))^2_{aero1} + (\delta(C_m/C_{m,ref}))^2_{aero2}} \qquad (1)$$

where the δ symbol denotes the total uncertainty. To a large extent, the RSS uncertainty is the result of sampling uncertainty, with a floor level given by the input (CFD) uncertainty. The red lines in both upper and lower graphs indicate the uncertainty that would result in the limit of infinite point coverage. Assuming that the nominal response would not change, one can judge from the lower graph of Fig. 10 which regions of the space are susceptible to uncertainty improvement as a result of further sampling. In all other regions, the only mechanism for reducing the uncertainty further is to reduce the input uncertainty. Since this may not always be an option, for example if the uncertainty is that associated with turbulence modeling rather than truncation error, one should take special care in designing an objective function which not only drives the optimization towards a given goal, but also makes the reduction of its uncertainty feasible. Alternatively, this information can be used to stop adding points in a region where one cannot tell whether the improvement is statistically significant.

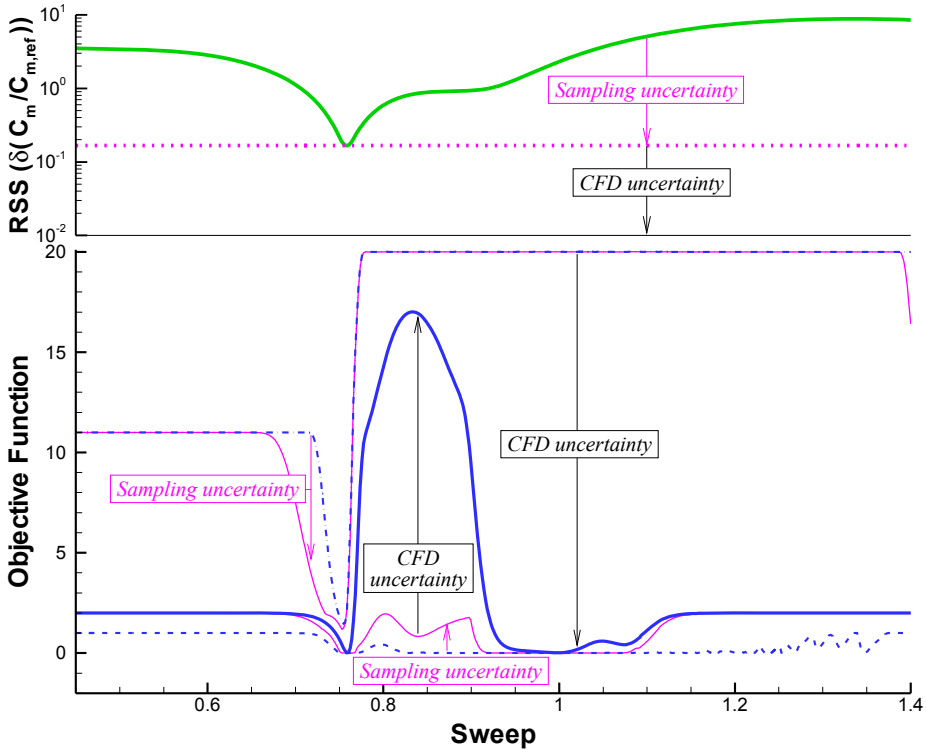

Figure 10. Response surface uncertainty vs. objective function uncertainty.

III. Conclusion

The aerodynamic design of an asymmetric oversized payload fairing subject to stability constraints was used to illustrate how high-fidelity simulations can be incorporated early in the design cycle. The present method uses dynamically evolving metamodels as surrogates for the more expensive computer analyses, in an approach that lends itself to the comprehensive characterization of uncertainty. An objective function formulation was used to specify constraints, and to drive the process towards solutions satisfying the objectives. As a result, candidate solutions were discovered after several hundred CFD runs. While the current design methodology did not make use of the available uncertainty, it is suggested that the latter can be used as a driver in balancing the need for exploration and exploitation in global optimization problems.

Acknowledgments

The support of this work by the Air Force Research Laboratory under SBIR contracts number FA9453-04-M-0132 and FA9453-05-C-0056 is gratefully acknowledged.

References

[1]Lewis, K. and Mistree, F., "The Other Side of Multidisciplinary Design Optimization: Accommodating a Multiobjective, Uncertain and Non-Deterministic World," *Eng. Opt.*, Vol. 31, 1998, pp. 161-189.

[2]Burman, J., Papila, N., Shyy, W., and Gebart, B. R., "Assessment of Response Surface-Based Optimization Techniques for Unsteady Flow Around Bluff Bodies," AIAA 2002-5596.

[3]DeLoach, R. and Erickson, G. E., "Low-Order Response Surface Modeling of Wind Tunnel Data Over Truncated Inference Subspaces," AIAA 2003-0456.

[4]Forrester, A. I. J., Bressloff, N. W., and Keane, A. J., "Response Surface Model Evolution," AIAA 2003-4089.

[5]Hirokawa, N., Fujita, K., and Iwase, T., "Voronoi Diagram Based Blending of Quadratic Response Surfaces for Cumulative Global Optimization," AIAA 2002-5460.

[6]Knill, D. L., Giunta, A., Baker, C. A., Grossman, B., Mason, W. H., Haftka, R. T., and Watson, L. T., "Response Surface Models Combining Linear and Euler Aerodynamics for Supersonic Transport Design," *Journal of Aircraft*, Vol. 36, No. 1, Jan.-Feb. 1999, pp. 75-86.

[7]Krishnamurphy, T., "Response Surface Approximation With Augmented and Compactly Supported Radial Basis Functions," AIAA 2003-1748.

[8]Papila, N., Shyy, W., Griffin, L., and Dorney, D. J., "Shape Optimization of Supersonic Turbines Using Response Surface and Neural Network Methods," AIAA 2001-1065.

[9]Walker, E. L., "Statistical Calibration and Validation of a Homogeneous Ventilated Wall-Interference Correction Method for the National Transonic Facility," NASA TP-2005-213947.

[10]Rais-Rohani, M. and Singh, M. N., "Efficient Response Surface Approach for Reliability Estimation of Composite Structures," AIAA 2002-5604.

[11]Vittal, S. and Hajela, P., "Confidence Intervals for Reliability Estimated Using Response Surface Methods," AIAA 2002-5475.

[12]Jones, D. R., Schonlau, M., and Welch, W. J., "Efficient Global Optimization of Expensive Black-Box Functions," *J. Global Opt.*, Vol. 13, 1998, pp. 455-492.

[13]Sóbester, A., Leary, S. J., and Keane, A. J., "On the Design of Optimization Strategies Based on Global Response Surface Approximation Models, *J. Global Opt.*, Vol. 33, No. 1, 2005, pp. 31-59.

[14]Queipo, N. V., Haftka, R. T., Shyy, W., Goel, T., Vaidyanathan, R., and Tucker, P. K., "Surrogate-Based Analysis and Optimization," *Progress in Aerospace Sciences*, Vol. 41, 2005, pp. 1-28.

[15]Poggio, T. and Girosi, F., "Network for Approximation and Learning," *Proc. IEEE*, Vol. 78, No. 9, 1990, pp. 1481-1497.

[16]Raeth, P. G., Gustafson, S. C., Little, G. R., and Puterbaugh, T. S., "Stretch and Hammer Neural Networks For N-Dimensional Data Generalization," Air Force Wright Laboratory Report WL-TR-91-1146, Jan. 1992.

[17]Vapnik, V. N., *The Nature of Statistical Learning Theory*. Springer, 1995.

[18]Gunn, S. R., "Structural Modelling With Sparse Kernels," *Proceedings of the 13th IFAC Symposium on System Identification*, Rotterdam, Netherlands, 2003.

[19]Gramacy, R. B., Lee H. K. H., and Macready, W., "Parameter Space Exploration With Gaussian Process Trees," Proceedings of the International Conference on Machine Learning, Omnipress and ACM Digital Library, 2004, pp. 353-360.

[20]Friedman, J. H., "Multivariate Adaptive Regression Splines," *The Annals of Statistics*, Vol. 19, No. 1, 1991, pp. 1-67.

[21]Reisenthel, P. H., Love, J. F., Lesieutre, D. J., and Childs, R. E., "Cumulative Global Metamodels with Uncertainty - a Tool for Aerospace Integration," *The Aeronautical Journal*, Vol. 110, No. 1108, Jun. 2006, pp. 375-384.

[22]Childs, R. E., Reisenthel, P. H., Rose, J., and Maly, J., "Large Asymmetric Launch Vehicle Payload Fairing," NEAR TR 611, Nielsen Engineering & Research, Apr. 2005.

[23]Buning, P. G., Jespersen, D. C., Pulliam, T. H., Klopfer, G. H., Chan, W. M., Slotnick, J. P., Krist, S. E., and Renze, K. J., "Overflow User's Manual Version 1.8r," NASA Ames Research Center, 2000.

[24]Spalart, P. R. and Allmaras, S. R., "A One-Equation Turbulence Model for Aerodynamic Flows," AIAA-92-0439, Jan. 1992.

[25]Chan, W. M. and Steger, J. L., "Enhancements of a Three-Dimensional Hyperbolic Grid Generation Scheme," *Appl. Math. and Comput.*, Vol. 51, 1992, pp. 181-205.

[26]Giunta, A. A., Wojtkiewicz, S. F., and Eldred, M. S., "Overview of Modern Design of Experiments Methods for Computational Simulations," AIAA 2003-0649.

[27]Audet, C. and Dennis, J. E., "A Pattern Search Filter Method for Nonlinear Programming Without Derivatives," *SIAM J. Optim.*, Vol. 14, No. 4, 2004, pp. 980-1010.

[28]Kolda, T. G., Lewis, R. M. and Torczon, V., "Optimization by Direct Search: New Perspectives on Some Classical and Modern Methods," *SIAM Review*, Vol. 45, No. 3, 2003, pp. 385-482.