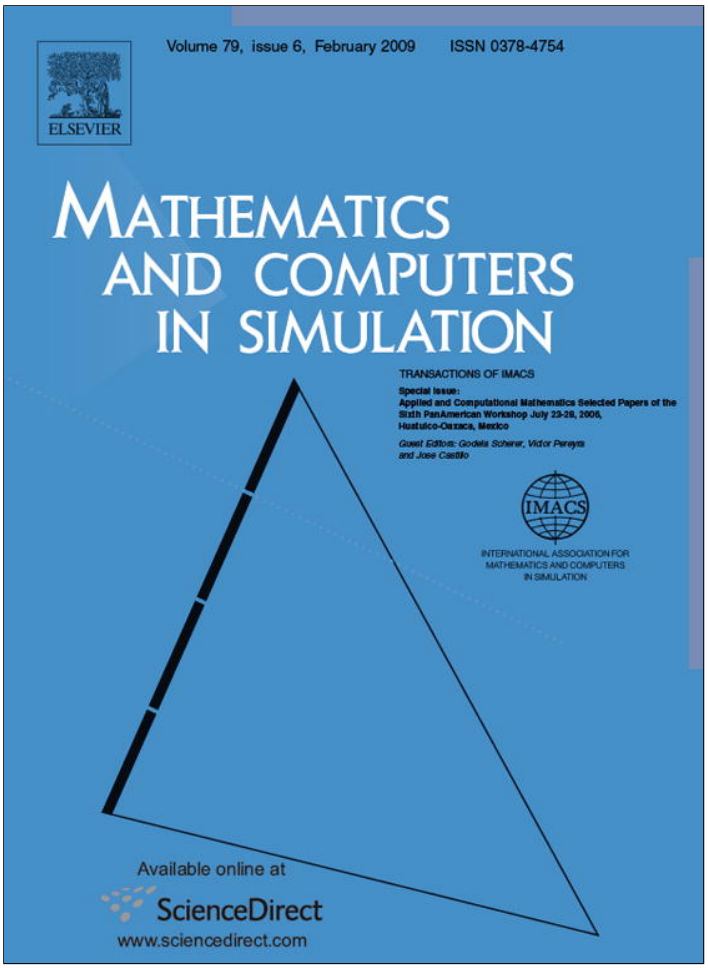

Available online at www.sciencedirect.com
ScienceDirect
Mathematics and Computers in Simulation 79 (2009) 1948–1958
MATHEMATICS AND COMPUTERS IN SIMULATION
www.elsevier.com/locate/matcom

Aircraft design optimization

J.J. Alonso [a,b], P. LeGresley [c], V. Pereyra [d,*]

[a] *Department of Aeronautics and Astronautics, Stanford University, Stanford, CA 94305, United States*
[b] *Fundamental Aeronautics, NASA, Washington, DC, United States*
[c] *Advanced Numerical Design, United States*
[d] *Weidlinger Associates Inc., 399 El Camino Real, #200, Mountain View, CA 94040, United States*

Received 28 August 2006; received in revised form 24 April 2007; accepted 2 July 2007
Available online 6 July 2007

Abstract

In this paper we describe briefly a set of procedures for the optimal design of full mission aerospace systems. This involves multi-physics simulations at various fidelity levels, surrogates, distributed computing and multi-objective optimization. Low-fidelity analysis is used to populate a database of inputs and outputs of the system simulation and Neural Networks are then designed to generate inexpensive surrogates. Higher fidelity is used only where is warranted and also to do a local exploration after global optimization techniques have been used on the surrogates in order to provide plausible initial values. The ideas are exemplified on a generic supersonic aircraft configuration, where one of the main goals is to reduce the ground sonic boom.

Keywords: Optimal design; Surrogates; Aerospace systems; Multi-objective optimization; Neural Networks

1. Introduction

Optimization problems in many industrial applications are extremely hard to solve in a general manner. Good examples of such problems can be found in the design of aerospace systems. Due to the high level integration of today's systems and the increase in complexity of analysis and design methods for the evaluation of system performance, such problems are characterized by multi-disciplinary simulations, goal functionals and constraints that are expensive to evaluate and, in addition, they often have large-dimensional design parameter spaces that further complicate the solution of the problem.

Moreover, the resulting optimization problems are frequently non-convex, i.e., multi-modal and ill-conditioned, making complete optimizations of complex aircraft configurations prohibitively expensive. Finally, the simulations may require using legacy or commercial codes that have to be used as black boxes that, in particular, may not produce the derivative information required by some optimization techniques.

Multi-modal, ill-conditioned problems require global optimization techniques and regularization [13,15] and present some of the most challenging problems for robust initialization and ulterior accurate solution. In high-dimensional spaces the available techniques are problematic at their best and one often must resort to surrogate models, divide and

* Corresponding author.
E-mail addresses: jjalonso@stanford.edu (J.J. Alonso), plegresl@advnumdes.com (P. LeGresley), pereyra@wai.com (V. Pereyra).

0378-4754/$32.00 © 2007 IMACS. Published by Elsevier B.V. All rights reserved.
doi:10.1016/j.matcom.2007.07.001

conquer techniques and parallel computing, in order to even have a chance to solve the problem in a reasonable time [2,7,14,15].

The most common surrogate models consist of low degree polynomial approximations. This is inadequate to surrogate highly discontinuous functions such as supersonic boom. In this paper we use Neural Networks to produce surrogates that more faithfully and successfully reproduce this type of function. We also use a fast training tool for generating the NN using the Variable Projection method of ref. [12].

In addition to all of these complexities, the available optimization strategies are such that, given initial conditions that are not in the vicinity of the true global optimum they may fail to converge or even produce a feasible sub-optimal design. This is particularly true of system vehicle designs (aircraft and spacecraft), where many disciplines interact in a complex multi-disciplinary analysis, and where the models of each of the disciplines may be noisy and even discontinuous.

However, if a reasonable starting point can be provided for the design, the likelihood of convergence to the optimum is greatly improved: the design space is much better behaved in the neighborhood of the optimum and since the range of variations of the design parameters is more restricted, multi-modal design spaces are typically avoided.

This view of the initialization problem requires a flexible environment that can be tailored to specific applications in order to have the best possibility of success in the design of complex multi-disciplinary systems. For that reason, we have chosen to develop an advanced toolbox for design that will contain the necessary modules (approximation, optimization, discipline-specific analysis) to tackle a wide variety of aerospace-related problems. These tools can be combined by advanced users and developers to create new design applications with relatively low investment. We present here some preliminary results of what has been achieved so far.

2. The problem considered and the various modules employed for its solution

We start from a high plateau with an existing prototype of a general Toolbox for Optimization, which has been successfully applied to other multi-disciplinary realms [16], and a number of proven tools for aircraft design [1,2,7–9]. In this work we sketch some of the principal components necessary to carry on this type of task. We focus on the generation of sample data for the aerodynamic and boom characteristics of a generic supersonic aircraft configuration, the fitting of this data using Neural Networks, and the use of the resulting surrogate models in a representative design optimization problem.

For this test example, a small number of the most relevant design parameters were selected, as described below. Although this is a small number of design variables relative to a realistic aerospace vehicle design, it has all of the elements necessary to exercise and evaluate the proposed initialization and optimization techniques. The main goals are to maximize the total range of the aircraft and to minimize the perceived loudness of the ground boom signature (measured in dBA), while satisfying a number of mission and buildability constraints, such as:

- Structural integrity of the aircraft for a $N = 2.5$ g pull-up maneuver.
- Take-off field length < 6000 ft.
- Landing field length < 6000 ft.

For aerodynamic modeling we are currently using the A502 solver, also known as PanAir [5,6], a flow solver developed at Boeing to compute the aerodynamic properties of arbitrary aircraft configurations flying at either subsonic or supersonic speeds. This code uses a higher order (quadratic doublet, linear source) panel method, based on the solution of the linearized potential flow boundary-value problem. Results are generally valid for cases that satisfy the assumptions of linearized potential flow theory—small disturbance, not transonic, irrotational flow and negligible viscous effects. Once the solution is found for the aerodynamic properties on the surface of the aircraft, A502 can then easily calculate the flow properties at any location in the flow field, hence obtaining the near-field pressure signature needed for sonic boom prediction. In keeping with the axisymmetric assumption of sonic boom theory, the near-field pressure can be obtained at arbitrary distances below the aircraft [3].

To provide a suitable description of the geometry to A502 a well-defined surface in three-dimensional space, which constitutes the outer mold line (OML) of the vehicle, needs to be generated as a function of the selected design variables. Because of the importance of the OML, a separate utility, Aerosurf, is used to generate and manage it. Aerosurf was specifically created for the analysis and design of aircraft configurations. The baseline geometry of an

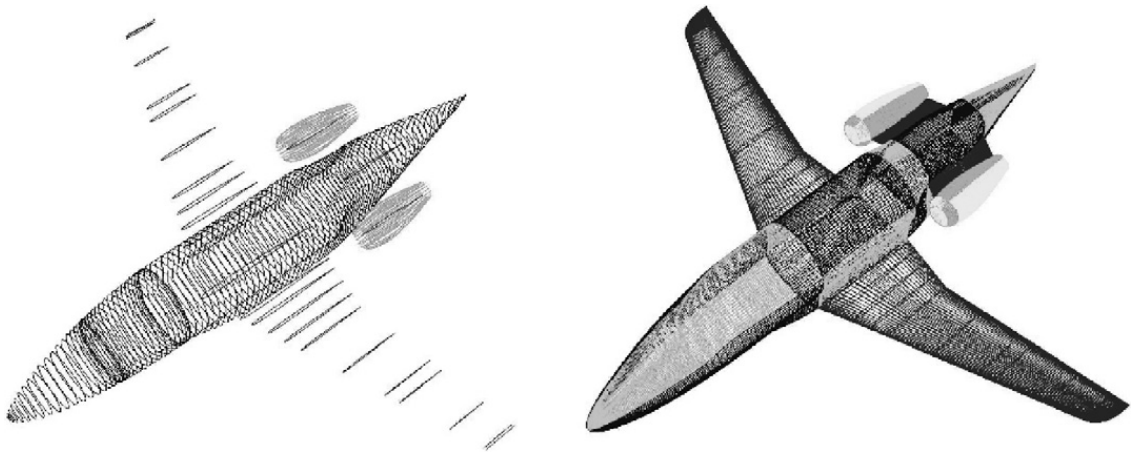

Fig. 1. Un-intersected components of a transonic business jet configuration (left) and intersected components forming the OML (right).

aircraft configuration is given to Aerosurf in the form of separate components, each one being described by a series of point-wise cross sections. These components can be fuselages, pylons, nacelles, and wing-like surfaces. After lofting the sections that define each component using a bi-cubic spline method, Aerosurf intersects these components and divides the resulting surface into a series of patches. At this stage, Aerosurf creates a parametric description of each patch and then distributes points on their surface, forming a fine structured watertight mesh. Thus, the set of points formed by the grids of all patches represents a watertight discretization of the OML within Aerosurf. For Euler or Navier–Stokes computations the surface mesh can be used for the generation of the volume meshes as well as the deformation of an existing mesh for design purposes (Fig. 1).

The method for computing the ground boom signature is shown in Fig. 2. At the near-field plane location, the pressure signature created by the aircraft is extracted and propagated down to the ground using extrapolation methods based on geometric acoustics.

The location of the near-field must be far enough from the aircraft so that its flow field is nearly axisymmetric and there are no remaining diffraction effects, which cannot be handled by the extrapolation scheme. Since A502/Panair

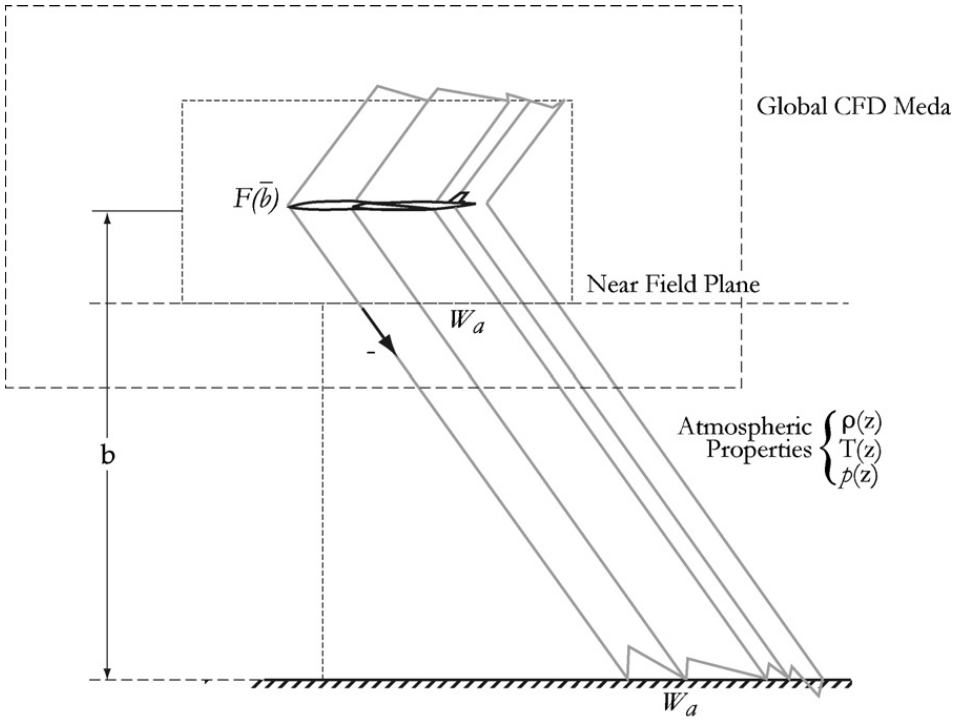

Fig. 2. Sonic boom propagation procedure.

only uses a surface mesh for all of its calculations, it is able to obtain near-field pressures at arbitrary distances without changes in the computational cost.

In this work we are using the Sboom [19] extrapolation method to propagate near-field signatures into ground booms. The sonic boom extrapolation method accounts for vertical gradients of atmospheric properties and for stratified winds (although the winds have been set to zero in this work). The method relies on results from geometric acoustics for the evolution of the wave amplitude, and utilizes isentropic wave theory to account for non-linear waveform distortion due to atmospheric density gradients and stratified winds.

Our past research on low-boom aircraft design focused on reducing the magnitude of only the initial peak of the ground boom signature [1,9]. This requirement, which had been suggested as the goal of the DARPA-sponsored Quiet Supersonic Platform (QSP) program ($\Delta p_0 < 0.3$ psf), neglects the importance of the full signature, which depends on the more geometrically complex aft portion of the aircraft where empennage, engine nacelles and diverters create more complicated flow patterns. Moreover, such designs often have two shock waves very closely following each other in the front portion of the signature [6,8], a behavior that is not robust and is therefore undesirable.

For these reasons, we are computing the perceived loudness of the complete signature (dBA). Frequency weighting methods are used to account for the fact that humans do not have an equal response to sounds of different frequencies. In these calculations, less weighting is given to the frequencies to which the ear is less sensitive. In addition, all signatures computed are post-processed to add a physical rise-time across the shock waves that yield loudness numbers that are more representative of those perceived in reality.

3. Generation of response surfaces (surrogates)

Using the tools described above, 300 configurations obtained via Latin Hypercube Sampling (LHS) were generated. These sample data were fitted using Neural Networks (NN). The prototype NN was a single hidden layer perceptron with sigmoid activation functions that provided a general non-linear model. We use for its fast training (i.e., determination of the NN parameters) a Variable Projection algorithm (VARPRO [12,18]) to solve the resulting non-linear least-squares separable problem in order to generate a reduced cost approximation of the objective space. This is combined with a global optimization algorithm since the resulting problems are generally multi-modal [4].

The training of a single hidden layer Neural Network, using sigmoid activation functions leads to a separable non-linear least-squares problem in which the unknown parameters in the network are determined to best fit in the L_2 sense the training data (i.e., the input/output data contained in the simulation database). This results in a surrogate function that is a linear combination of sigmoids:

$$S(\boldsymbol{x};\boldsymbol{\alpha}) = \sum_j a_j \left(\frac{1}{1 + \exp - (\langle \boldsymbol{\alpha^j}, x\rangle + \alpha_0^j)} \right), \tag{1}$$

where $\boldsymbol{x} \in R^k$ corresponds to the input or design parameters, a_j are the linear parameters, $\boldsymbol{\alpha^j} \in R^k$, α_0^j, are the non-linear parameters, and $\langle , \rangle$ stands for vector inner product. The VARPRO algorithm was designed specially for this situation. It consists of first eliminating the linear weights a_j and then solving the resulting non-linear least-squares problem by a Marquardt type algorithm. As shown in ref. [12,18], this algorithm has been very successful in solving many problems of this type and in particular ill-conditioned problems involving linear combinations of exponentials.

The data used corresponds to a representative supersonic business jet configuration. For each output, the training files contain 300 data points, while a test set of an additional 150 data points are used for evaluation of the fit. There are eight independent or input variables, namely:

- wing reference area,
- wing aspect ratio,
- longitudinal position of the wing,
- wing sweep angle,
- lift coefficient at initial cruise,
- lift coefficient at final cruise,
- altitude at initial cruise,
- altitude at final cruise.

Table 1
Neural Network results for surrogate functionals corresponding to the 10 output variables

Output	Max. value	#Sigmoids	RMS training	RMS test	Time (s)	Max. Res.
1		8	0.0000865	0.00033	872	0.0009
2		8	0.0033	0.206	831	0.15
3		10	0.406	1.2	1872	1.5
4	95	10	0.848	2.0	2332	3.0
5	95	10	0.541	2.01	2393	1.9
6	0.017	10	0.0000721	0.000229	2782	0.00023
7	3	8	0.0364	2.66	1393	0.26
8	99	10	0.414	1.14	1773	1.9
9	94	10	0.747	4.46	2090	3.2
10	94	10	0.545	1.71	1917	2.2

There are 10 dependent variables and we want to approximate each with a different Neural Network of the form (1):

- drag coefficient at initial cruise,
- sonic boom initial rise at initial cruise,
- sonic boom sound level at initial cruise w/out rise time modification,
- sonic boom sound level at initial cruise w/ first type rise time modification,
- sonic boom sound level at initial cruise w/ third type rise time modification,
- drag coefficient at final cruise,
- sonic boom initial rise at final cruise,
- sonic boom sound level at final cruise w/out rise time modification,
- sonic boom sound level at final cruise w/ first type rise time modification,
- sonic boom sound level at final cruise w/ third type rise time modification.

We scale the input variables so that they have zero mean and variance equal to 1. In Table 1, we show the results of training and testing for each one of the 10 outputs.

The approximate aerodynamic and boom model generated using VARPRO will be used in conjunction with low-order, first-principles based weight estimation and performance analysis tools. These include component structural weight estimation, take-off/landing distance, mission range and/or loiter, climb performance, etc.

For a given mission profile, the approximation model developed using VARPRO and the additional weight and performance tools will be used to run representative optimizations.

4. Design optimization

The optimization is accomplished by varying the values of the parameters highlighted in Table 2, where the baseline values were those of the initial design.

For this design test a small number of the most relevant design parameters were selected: wing reference area, wing aspect ratio, longitudinal position of the wing, wing sweep angle, lift coefficient at the initial and final cruise conditions, and altitude at the initial and final cruise conditions. Although this is a small number of design variables relative to a complete aerospace vehicle design, we have all of the elements necessary to exercise and evaluate the initialization and optimization techniques proposed.

The ranges of the variables used in the optimization are chosen to match those of the multiple (Design of Experiments) analyses conducted to fit the VARPRO surrogates. It is expected that the accuracy of the surrogate models is quite reasonable within these bounds, so that meaningful optimizations can be obtained and so that unreasonable results are not arrived at because of inaccuracies in the fits. The performance of the baseline design is also included in this table (as given by the range, boom loudness and take-off and landing field lengths). Note that the baseline design does not meet the take-off field length requirement.

Table 2
Design variables for the optimization problem and performance of the baseline design

Variable name	Baseline value	Allowable range
Maximum take-off weight (MTOW) (lbs)	100,000	80,000–120,000
Wing reference area (ft^2)	1150	1000–1300
Wing aspect ratio	4.0	3–5
Wing sweep angle (°)	52.5	45–60
Wing longitudinal position (percent of fuselage length) (%)	37.5	25–50
Initial cruise altitude (ft)	48,000	42,000–54,000
Final cruise altitude (ft)	52,000	42,000–54,000
Range (nmi)	3261	Output
Boom perceived loudness (dBA)	75.3	Output
Take-off field length (ft)	9633	Output
Landing field length (ft)	5322	Output

Design requirement	Value
Payload weight (lbs)	8000
Cruise mach number	1.8
Ultimate load	2.5
Number of engines	2

Table 3
Summary of design requirements and additional parameters

Design parameter	Value
Wing taper ratio	0.3
Average wing t/c ratio	0.03
Fuselage length (ft)	80
Fuselage width (ft)	6
Fuselage height (ft)	6
Sea level thrust (per engine) (lbs)	15,000
Max. CL at take-off	1.1
Max. CL on approach	1.4
Thrust-specific fuel consumption (kg/h/N)	0.25

Several other parameters are fixed in the optimization. These design requirements and/or parameters are listed in Table 3.

A top view of the baseline configuration can be seen in Fig. 3 and will be used for reference while the results are presented later. The baseline configuration is obtained using the values of the design parameters shown in Table 2. The corresponding performance of the baseline design (range and perceived loudness as well as take-off and landing field lengths) is also presented in the table. Notice that this baseline configuration had been obtained using a slightly lower

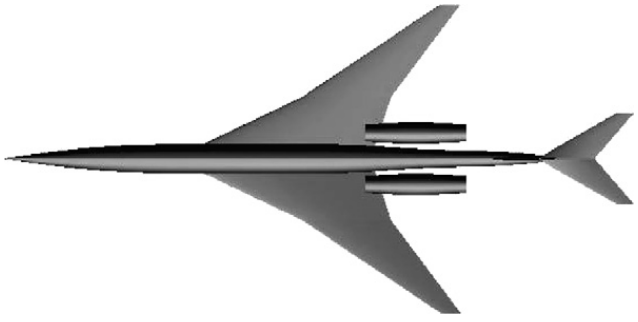

Fig. 3. Top view of the baseline configuration for this design study.

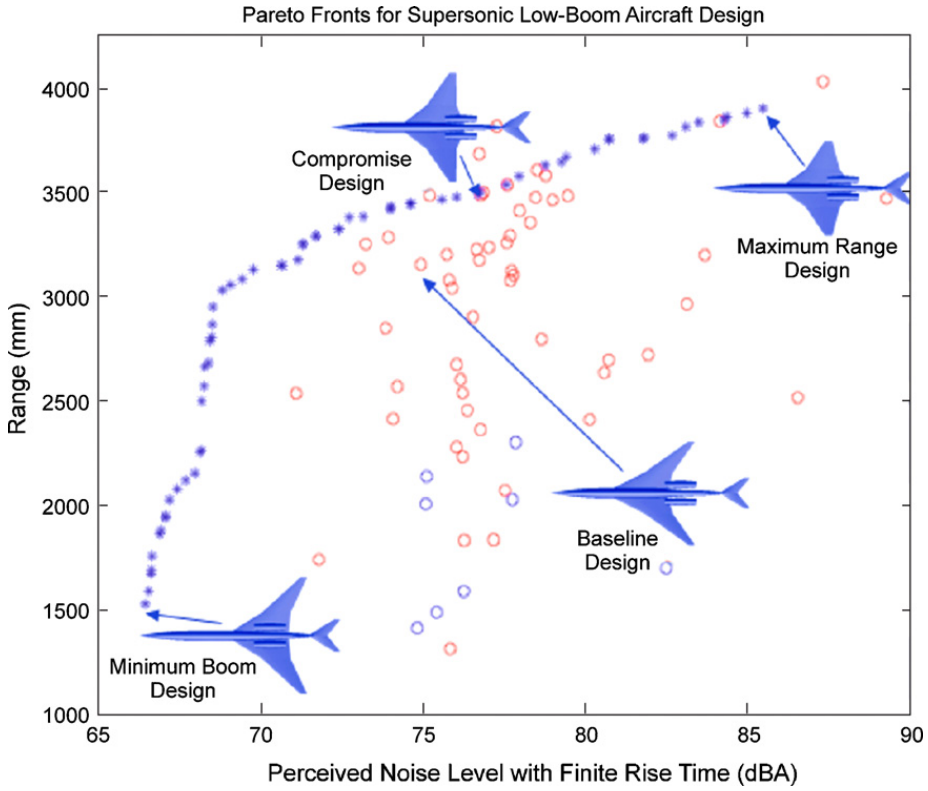

Fig. 4. Results in objective space for the multi-objective genetic algorithm optimization. Initial population: circles; final population: asterisks.

cruise Mach number (1.6) and a single-objective optimization where the function to be minimized was the take-off gross weight.

5. Design results

Using the baseline configuration described in Tables 2 and 3, we generated an initial population by randomly perturbing the parameters around the baseline configuration and calculating the performance, i.e., the perceived noise level and range (these are the open circles in Fig. 4). From this initial population the NSGA-II [10] algorithm, an elitist multi-objective genetic algorithm, evolves the population with the purpose of generating an approximation to the Pareto front, while satisfying all the constraints in Table 3 (the resulting Pareto front is represented by the asterisks). Thus, the Pareto front (the image of the set of Pareto equilibrium points of the problem mapped by the objectives), gives a complete image of all the solutions for this bi-objective optimization, from which we can select a compromise feasible design that satisfies the requirements of the problem.

The initial population is composed of 64 alternative designs that are randomly distributed within the design space. The level of performance of these designs is represented by the circles in Fig. 4. Each of these circles represents a pair of (Range, Loudness) for each of the aircraft in question. Note that the optimization algorithm attempts to drive all designs toward the upper left hand corner of the graph. Of the 64 initial design alternatives some are feasible designs and others are unfeasible. That is the reason why some of the initial designs appear to have extraordinary performance. Notice, again, that the baseline configuration does not meet the constraints (take-off field length in this case) of the optimization problem.

We choose a crossover probability for the variables of 0.75 and a mutation probability of 0.125. The distribution indices for both the crossover and mutation operations are taken to be 12 and 25%, respectively. The population is allowed to evolve over 1000 generations. As the generations proceed the entire population migrates towards feasibility and attempts to maximize their range and minimize their noise level simultaneously. The result is the Pareto front of solutions given by the asterisks in the graph. This Pareto front represents the set of non-dominated solutions (no other solution is better than any of those solutions in both performance measures simultaneously) of the problem. Notice

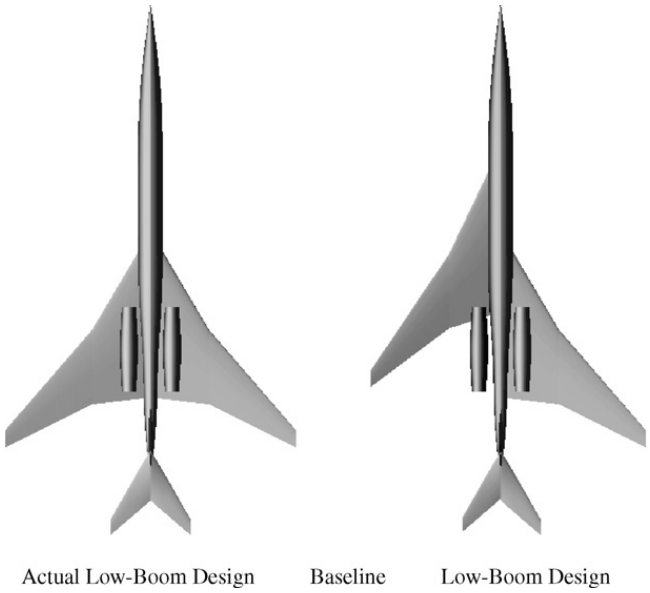

Fig. 5. Result of multi-objective optimization: lowest boom design. Actual low-boom design baseline low-boom design.

that all asterisks on the Pareto front are feasible, indicating that the aircraft that each point represents satisfies the constraints of the problem (which was not the case for the initial population).

The designs/aircrafts on the Pareto front provide the designer with alternatives to trade the relative benefits of each of the two (conflicting) objective functions: one may want to choose an aircraft with very low boom signature but also very low range or an aircraft with extremely high range but rather poor sonic signature or something in between. Notice that the shape of the Pareto front is quite unique: it is initially very steep and then its slope appears to decrease to a smaller value. In fact the Pareto front is well approximated by two straight-line segments with a knee located at a loudness level of approximately 67.5 dBA. The interpretation is clear: it pays off to sacrifice a little bit of the quietness of the aircraft to attain significant improvements in the range of the aircraft. After the loudness increases beyond 67.5 dBA modest improvements in range are obtained at the expense of very large increases in sonic boom loudness.

Figs. 5–7 show representative examples for different aircraft on the Pareto front. In each of these figures we see the resulting configuration on the left, and, on the right, we compare the baseline configuration (left) to the actual configuration (right) side by side. This comparison is provided for reference in order to understand the magnitude of the changes that the optimizer has dictated.

Fig. 5 shows the configuration for the lowest boom aircraft: the asterisk in the lower left corner of Fig. 4. This configuration has a perceived loudness of 66.4 dBA but a very small computed range of 1531 nmi. The optimizer has chosen to move the wing backwards almost as far as it is allowed (to 44.8% of the length of the fuselage). The TOGW ends up at 90,200 lbs and the wing reference area has increased (in order to decrease the wing loading) as far as it is allowed (1300 ft^2). The optimizer also sees benefits in increasing the cruise altitude (in order to lower the boom significantly) to 54,000 ft for the entire cruise segment. One can quickly see that this aircraft is not very realistic since stability and control issues were not included in the formulation. If these considerations were included, the optimizer would be prevented from moving the wing so far back.

Fig. 6 shows, instead, the design with the highest range (and also the highest noise levels). The calculated range for this aircraft is 3902 nmi (close to the 4000 nmi that we had intended to obtain) but the perceived noise level is 85.45 dBA: far too high to be allowed to fly supersonically over land. The changes in the geometry and vehicle characteristics are less dramatic. The TOGW is slightly lower than in the previous example, 89,154 lbs. The wing also has a smaller wing area, a smaller aspect ratio (3.0) and a more forward wing position (25%). The flight altitudes are slightly diminished but close to the upper limit of 54,000 ft. Additional environmental considerations (such as the effect of CO and H_2O emissions on the ozone layer) may actually force a harder constraint on the maximum allowable flight altitude. This design, however, is representative of the type of design that would be obtained if sonic boom considerations were not taken into account at all.

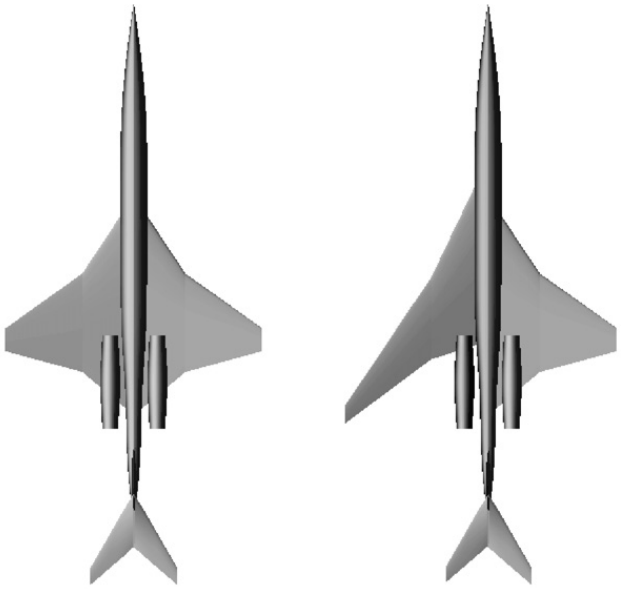

Fig. 6. Result of multi-objective optimization: highest range design. Maximum range design baseline maximum range design.

Finally, Fig. 7 shows the design results of an intermediate point in the Pareto front. We have chosen to display the first design that achieves a range greater or equal to 3500 nmi. As it turns out, this design shows a range of 3504 nmi and a perceived noise level of 76.7 dBA. This aircraft can be considered a reasonable compromise between the two competing objectives of boom loudness and aircraft maximum range. The TOGW is found to be 90,708 lbs with a wing reference area of nearly 1300 ft^2. Since no viscous penalty was added for the increasing surface area of the wing, the optimizer chooses to make it as large as possible. The aspect ratio of the wing is 3.93, the sweep angle is nearly 47° and the wing longitudinal position is in between the previous two designs, but closer to the maximum range design (at 29% of the length of the fuselage). As before, the optimizer chooses to maximize the initial and final cruise altitudes mainly for benefits in boom loudness.

Although we have glossed over the complexities of the above task we like to give a hint of it by showing the work flow diagram of Fig. 8.

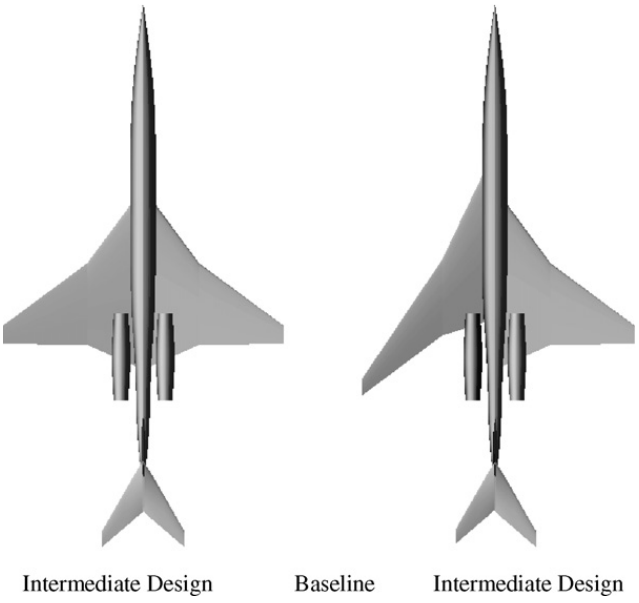

Fig. 7. Result of multi-objective optimization: intermediate design. Intermediate design baseline intermediate design.

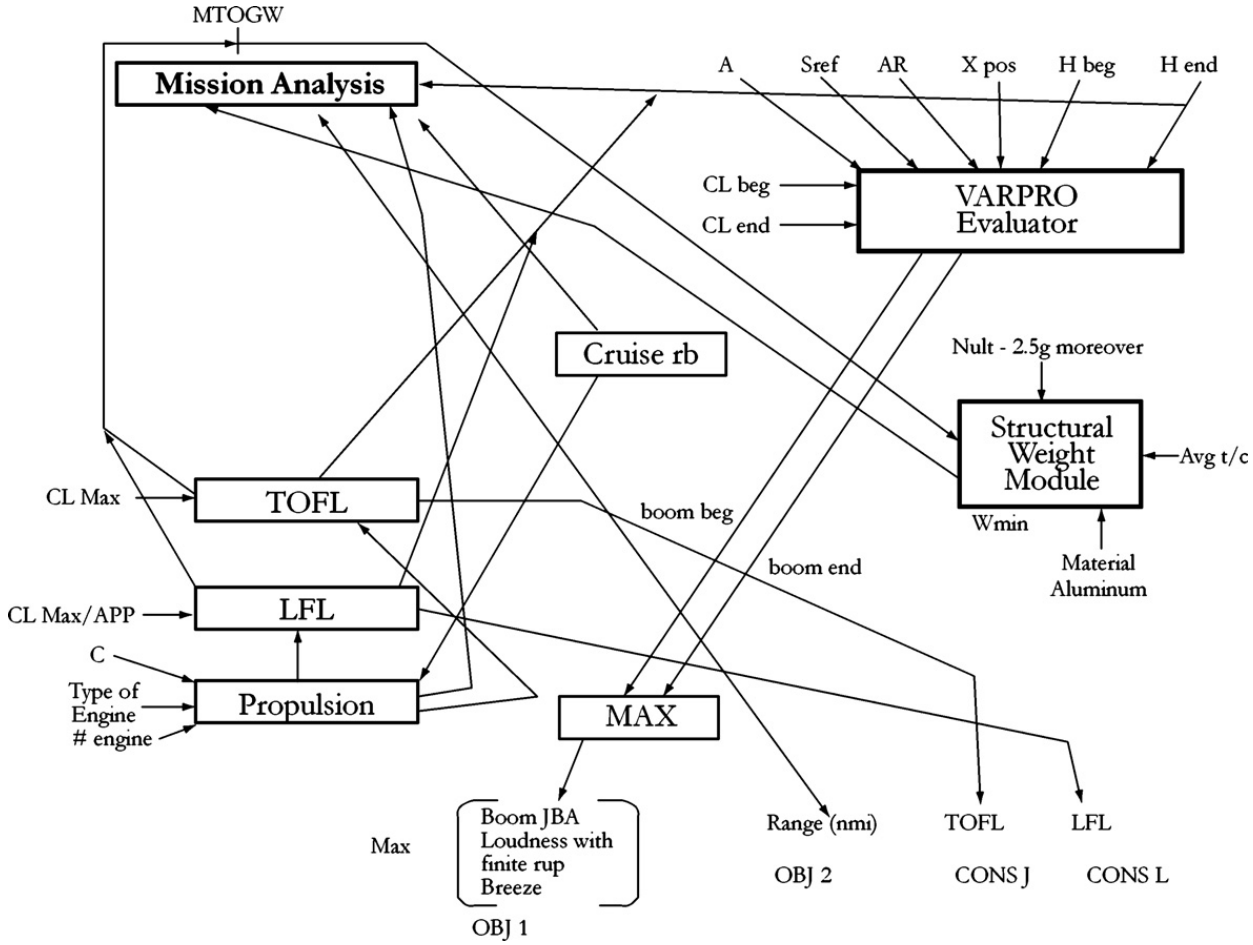

Fig. 8. Computational flow diagram of aircraft design optimization. Inputs on top; outputs on bottom.

6. Conclusions and future work

The results presented earlier are meant to serve as an example of the design work that can be done with the technology that has been developed during this work. In fact the designs presented are meant to constitute a valid initial design for continued refinement (during the preliminary design phase). The combination of advanced response surface formulations, the inclusion of the representative analyses in the optimization and the ability to treat multi-objective optimizations with constraints enables the designer to search very large regions of the design space (note the variability in the geometries and the design parameters of the three aircraft presented in Figs. 5–7) which may contain multiple local minima and hone in the areas of the design space that are most promising for further design work.

Areas of improvement for this work include the addition of higher fidelity modules for the analysis of the various disciplines that participate in the design, the integration with more efficient optimization algorithms (such as [11]), and the effective use of multi-fidelity approximations for the response surfaces in the problem.

Also, it would help to have multi-objective optimization methods that do not require as many goal evaluations as the genetic and evolutionary algorithms do. Some initial work in that direction can be found in Pereyra [17] that uses Newton method coupled to continuation techniques and a novel type of constraints that guaranties equispacing of the sampled Pareto front (contrast with the results of Fig. 4 that required thousands of evaluations and provided and adequately sampled Pareto front, although hardly equispaced).

Acknowledgements

This work has been partially sponsored by NSF SBIR Grant DMI-0321420 and AF SBIR Grant 05-250 F051-250-1368.

References

[1] J.J. Alonso, I.M. Kroo, A. Jameson, Advanced algorithms for design and optimization of quiet supersonic platforms (AIAA Paper 02-0114), in: 40th AIAA Aerospace Sciences Meeting and Exhibit, Reno, NV, 2002.

[2] J.J. Alonso, P. LeGresley, E. van der Weide, J. Martins, pyMDO: a framework for high-fidelity multi-disciplinary optimisation (AIAA Paper 2004-4480), in: 10th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference, 2004.

[3] H. Ashley, M. Landahl, Aerodynamics of Wings and Bodies, Dover, New York, 1985, pp. 104–178.

[4] L. Carcione, V. Pereyra, D. Woods, GO: Global Optimization, WA Report, 2005.

[5] R.I. Carmichael, L.I. Erickson, A higher order panel method for predicting subsonic or supersonic linear potential flow about arbitrary configurations, AIAA (1981) 81–1255.

[6] M. Chan, Supersonic aircraft optimization for minimizing drag and sonic boom, Ph.D. Thesis, Stanford University, Stanford, CA 94305, 2003.

[7] S. Choi, J.J. Alonso, H.S. Chung, Design of a low-boom supersonic business jet using evolutionary algorithms and an adaptive unstructured mesh method, in: 45th AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics and Materials Conference, 2004.

[8] H.S. Chung, Multidisciplinary design optimization of supersonic business jets using approximation model-based genetic algorithms, Ph.D. Thesis, Stanford University, Stanford, CA 94305, 2004.

[9] H.S. Chung, S. Choi, J.J. Alonso, Supersonic business jet design using knowledge-based genetic algorithm with adaptive, unstructured grid methodology, in: AIAA 2003-3791, 21st Applied Aerodynamic Conference, Orlando, FL, June 2003.

[10] K. Deb, A. Pratap, S. Agrawal, T. Meyarivan, A fast and elitist multiobjective genetic algorithm: NSGA-II, Technical Report No. 2000001, Indian Institute of Technology Kanpur, Kanpur, India, 2000.

[11] P. Gill, W. Murray, M. Saunders, SNOPT. Stanford University Optimization Laboratory, 2005.

[12] G. Golub, V. Pereyra, Separable nonlinear least squares: the Variable Projection method and its applications, Inverse Problems 19 (2003) R1–R26.

[13] P.C. Hansen, Rank-Deficient and Discrete Ill-Posed Problems, SIAM Publications, Philadelphia, 1998.

[14] V. Pereyra, Asynchronous distributed solution of large scale nonlinear inversion problems, in: V. Pereyra, J. Castillo (Eds.), Special Issue, Applied Numerical Mathematics, vol. 30, 1999, pp. 31–40.

[15] V. Pereyra, Ray tracing methods for inverse problems, Invited Topical Review, Inverse Problems 16 (2000) R1–R35.

[16] V. Pereyra, A toolbox for optimal design, Final Report, NSF SBIR Grant DMI-0321420, 2005.

[17] V. Pereyra, Fast computation of equispaced Pareto manifolds and Pareto fronts for multiobjective optimization problems, Math. Comput. Simul. 79 (2009) 1935–1947.

[18] V. Pereyra, G. Scherer, F. Wong, Variable Projection Neural Network training, Math. Comput. Simul. 73 (2006) 231–243.

[19] C.L. Thomas, Extrapolation of wind-tunnel sonic boom signatures without use of a Whitham F-function, NASA SP-255 (1970) 205–217.