AIAA

AIAA 2000-4886

First-Order Model Management with Variable-Fidelity Physics Applied to Multi-Element Airfoil Optimization

N. M. Alexandrov, E. J. Nielsen,
R. M. Lewis, and W. K. Anderson

NASA Langley Research Center
Hampton, VA 23681

**8th AIAA/USAF/NASA/ISSMO Symposium
on Multidisciplinary Analysis & Optimization
6-8 September 2000 / Long Beach, CA**
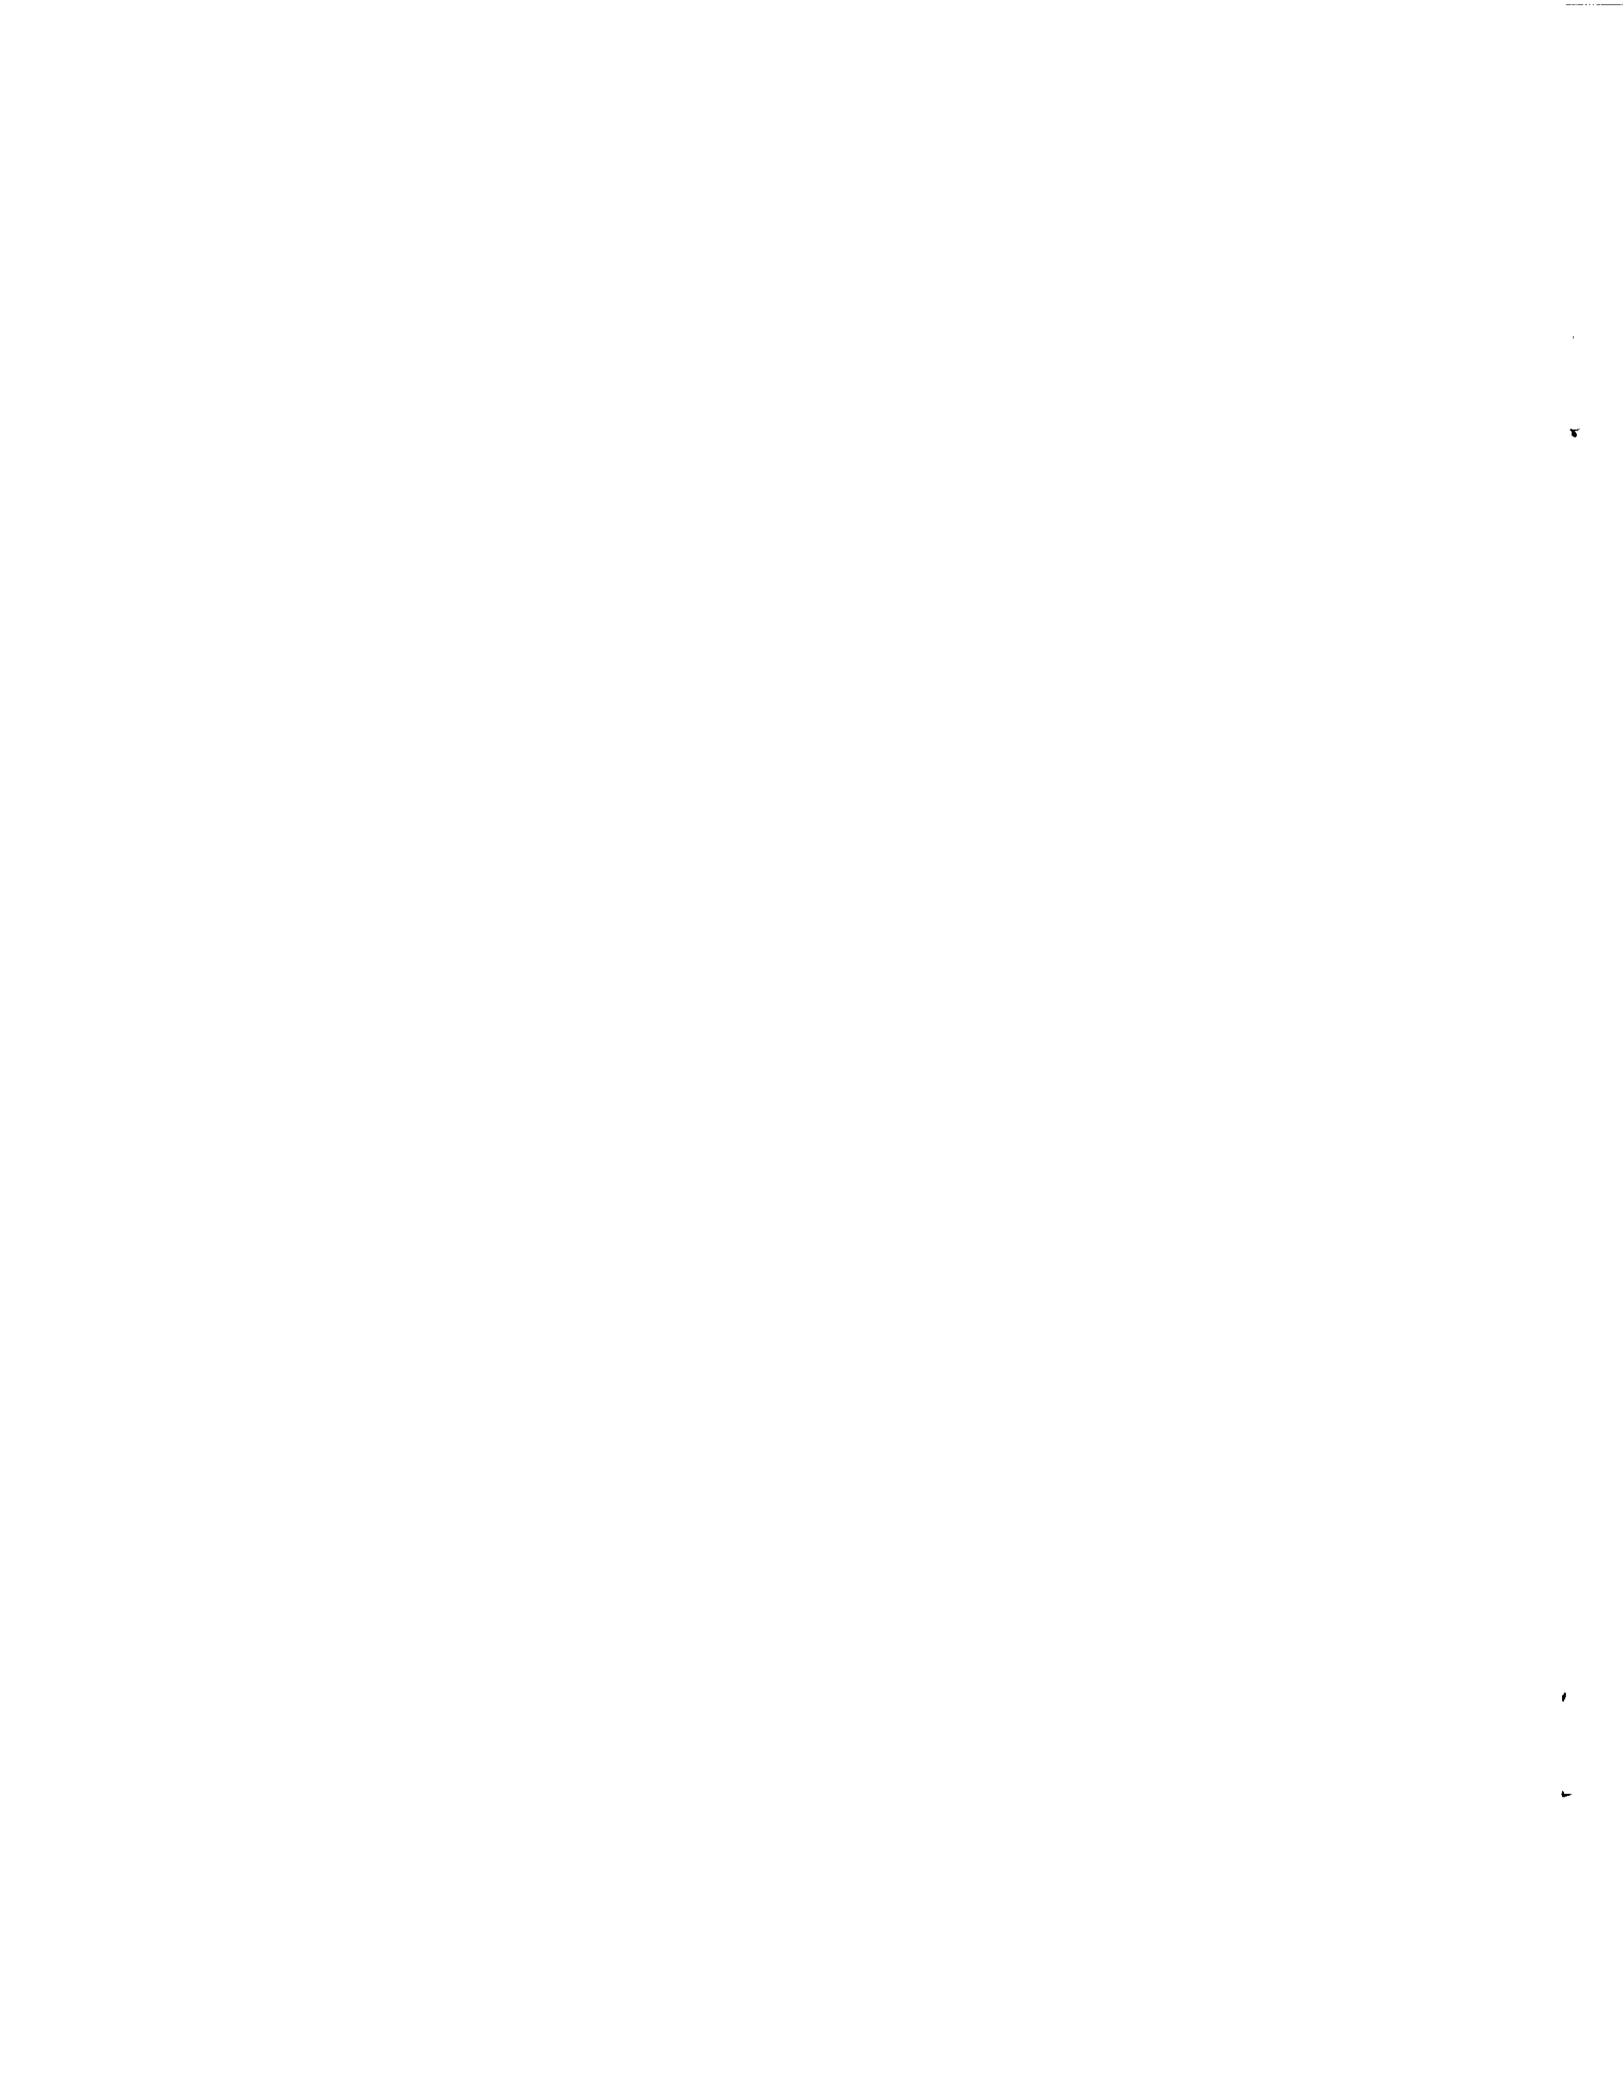
FIRST-ORDER MODEL MANAGEMENT WITH VARIABLE-FIDELITY PHYSICS APPLIED TO MULTI-ELEMENT AIRFOIL OPTIMIZATION

N. M. ALEXANDROV*, E. J. NIELSEN[†], R. M. LEWIS[‡] AND W. K. ANDERSON[§]

Abstract

First-order approximation and model management is a methodology for a systematic use of variable-fidelity models or approximations in optimization. The intent of model management is to attain convergence to high-fidelity solutions with minimal expense in high-fidelity computations. The savings in terms of computationally intensive evaluations depends on the ability of the available lower-fidelity model or a suite of models to predict the improvement trends for the high-fidelity problem. Variable-fidelity models can be represented by data-fitting approximations, variable-resolution models, variable-convergence models, or variable physical fidelity models. The present work considers the use of variable-fidelity physics models. We demonstrate the performance of model management on an aerodynamic optimization of a multi-element airfoil designed to operate in the transonic regime. Reynolds-averaged Navier-Stokes equations represent the high-fidelity model, while the Euler equations represent the low-fidelity model. An unstructured mesh-based analysis code FUN2D evaluates functions and sensitivity derivatives for both models. Model management for the present demonstration problem yields fivefold savings in terms of high-fidelity evaluations compared to optimization done with high-fidelity computations alone.

Key Words: Aerodynamic optimization, airfoil design, approximation concepts, approximation management, model management, nonlinear programming, surrogate optimization, variable-fidelity

Background

Approximations and low-fidelity models have long been used in engineering design to reduce the cost of optimization (e.g., [1-3]). An overview of approximations in structural optimization, for instance, can be found in [4]. Accounts of recent efforts in developing methodologies for variable-complexity modeling are relayed in [5,6].

The present work concerns an approach, the Approximation and Model Management Framework (AMMF) [7-10], designed to enable rapid and early integration of high fidelity nonlinear analyses and experimental results into the multidisciplinary optimization process. This is accomplished by reducing the frequency of performing high-fidelity computations within a single optimization procedure.

Until recently, procedures for the use of variable-fidelity models and approximations in design had relied on heuristics or engineering intuition. In addition, with a few exceptions (e.g., [11], [12]), the analysis of algorithms had focused on convergence to a solution of the approximate or surrogate problem ([13], [14]). The AMMF methodology discussed here and in related papers is distinguished by a systematic approach to alternating the use of variable-fidelity models that results in procedures that are provably globally convergent to critical points or solutions of the high-fidelity problem.

Model management can be, in principle, imposed on any optimization algorithm and used with any models. In [15], we considered AMMF schemes based on three nonlinear programming methods and demonstrated them on a 3D aerodynamic wing optimization problem and a 2D airfoil optimization problem. In both cases, Euler analysis performed on meshes of varying degree of refinement formed a suite of variable-resolution models. Results indicated approximately threefold savings (similar across the three schemes) in terms of high-fidelity function evaluations. The AMMF based on the sequential quadratic programming (SQP) approach was judged to be the most promising for single-discipline problems with a modest

---

number of design variables, as well as for certain formulations of the MDO problem.

The study in [15] has served as a proof of concept for AMMF in the case where low-fidelity models are represented by data-fitting approximations (kriging, splines and polynomial response surfaces) or variable-resolution models. The present work considers, arguably, the most challenging combination of the high and low-fidelity models within a single optimization procedure – that of variable-fidelity physics models. The performance of the first-order model management is demonstrated on an aerodynamic optimization of a multi-element airfoil. Variable-fidelity models are represented by an unstructured mesh-based analysis run in viscous and inviscid modes.

In the next section, we describe the AMMF under investigation and discuss the points of interest for the current study. We then present the demonstration problem, followed by a discussion of the numerical experiments and results.

AMMF under investigation

For the present demonstration, the optimal design problem is represented by the bound constrained nonlinear programming problem:

$$\begin{aligned} \min \quad & f(x) \\ \text{s. t.} \quad & l \leq x \leq u, \end{aligned} \tag{1}$$

where x is the vector of design variables, the objective f is a continuously differentiable nonlinear function, and $l \leq x \leq u$ denotes bound constraints on the design variables.

The first-order AMMF used here for solving (1) is based on the trust-region strategy, which is a methodology for the improvement of global behavior of the local model-based optimization algorithms [16]. The following pseudo-code describes the AMMF.

Initialize x_c, Δ_c
Do until convergence:
  Select a model a_c such that
    $a_c(x_c) = f(x_c)$ and $\nabla a_c(x_c) = \nabla f(x_c)$
  Solve approximately for $s_c = x - x_c$:
    $$\min_s \quad a_c(x_c + s)$$
    $$\text{s. t.} \quad l \leq x \leq u$$
    $$\| s \|_\infty \leq \Delta_c$$
  Compute $\rho_c \equiv \frac{f(x_c) - f(x_c + s_c)}{f(x_c) - a_c(x_c + s_c)}$
  Update Δ_c and x_c based on ρ_c
End do

Details of the updating strategy can be found, for instance, in [7]. Briefly, the point x_c is accepted if the step s_c results in a simple decrease in the objective function, i.e., if

$f(x_c) > f(x_c + s_c)$. Otherwise the step is rejected. The l_∞ trust region (chosen because it interacts naturally with the bound constraints) is decreased if ρ_c is small. Experience suggests that "small" be taken as less that 10^{-5}. If ρ_c is close to one or greater than one (this indicates excellent predictive properties of the model), the trust region is doubled. Otherwise, it is left unchanged.

The conditions

$$\begin{aligned} a_c(x_c) &= f(x_c) \tag{2} \\ \nabla a_c(x_c) &= \nabla f(x_c) \tag{3} \end{aligned}$$

are known as the first-order consistency conditions, which we will discuss presently.

In conventional optimization, a_c is usually a linear or quadratic model of the objective f. AMMF replaces this local, Taylor series approximation by an arbitrary model required to satisfy the consistency conditions (2)–(3). Regardless of the properties of the low-fidelity model, the consistency conditions force it to behave as a first-order Taylor series approximation at points where they are satisfied. Solving the subproblem of minimizing a_c is itself an iterative procedure that now requires the function and derivative information from the low-fidelity model.

First-order AMMF methods can be shown to converge to critical points or solutions of the high-fidelity problem under appropriate standard assumptions of continuity and boundedness of the constituent functions and derivatives (see [9], for instance), given that the consistency conditions (2)–(3) are imposed at each major iterate x_c.

Qualitatively, the reason a first-order AMMF converges to an answer of the high-fidelity problem may be summarized as follows. Although a lower-fidelity model may not capture a particular feature of the physical phenomenon to the same degree of accuracy (or at all) as its higher-fidelity counterpart, a lower-fidelity model may still have satisfactory predictive properties for the purposes of finding a good direction of improvement for the higher-fidelity model. By imposing the consistency conditions, AMMF ensures that at least at the major iterates, the lower-fidelity model provides the same direction of descent as the high-fidelity counterpart. Two questions arise. How easy is it to impose the first-order consistency? How does the method perform in practice?

The answer to the first question is that imposing the conditions (2)–(3) is straightforward using a correction technique due to Chang et al. [17] This technique corrects a low-fidelity version f_{lo} of an arbitrary function so that it agrees to first-order with a given high-fidelity version f_{hi}. This is done by defining the correction factor β as

$$\beta(x) = \frac{f_{hi}(x)}{f_{lo}(x)}.$$

American Institute of Aeronautics and Astronautics
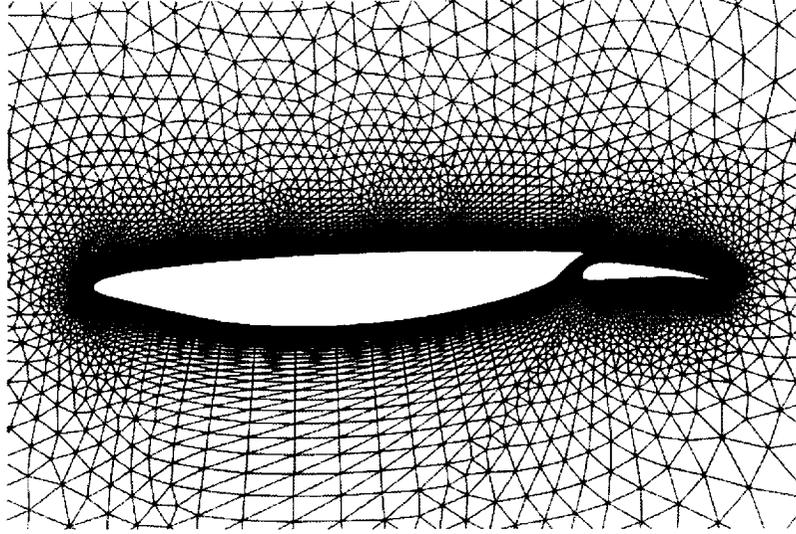
Figure 1: Mesh for the viscous model
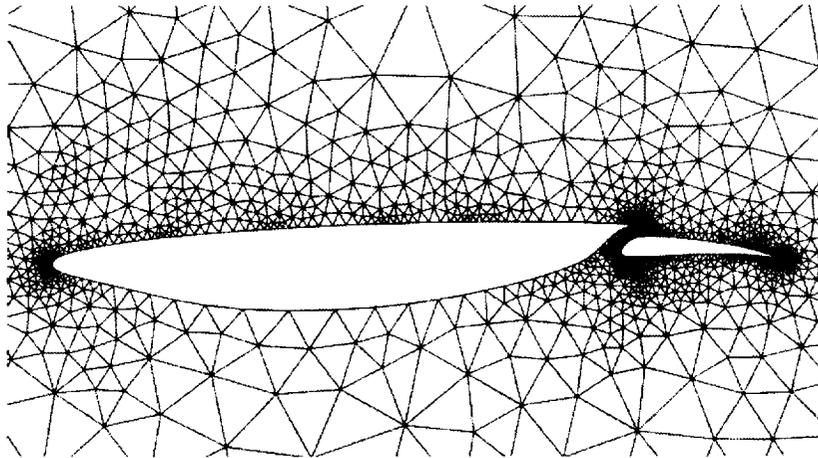
Figure 2: Mesh for the inviscid model

Given the current design variable vector x_c, one builds a first-order model β_c of β about x_c:

$$\beta_c(x) = \beta(x_c) + \nabla\beta(x_c)^T(x - x_c).$$

The local model of β is then used to correct f_{lo} to obtain a better approximation $a(x)$ of f_{hi}:

$$f_{hi}(x) = \beta(x)f_{lo}(x) \approx a(x) \equiv \beta_c(x)f_{lo}(x).$$

The corrected approximation $a(x)$ has the properties that $a(x_c) = f_{hi}(x_c)$ and $\nabla a(x_c) = \nabla f_{hi}(x_c)$. Zero-order or higher-order corrections are easily constructed as well.

Because the β-correction can make any two unrelated functions match to first order, the framework admits a wide range of models. In the worst case of performance, the subproblem will yield a good predictive step s_c for the

high-fidelity model only at the point x_c. Thus, the high-fidelity information may potentially have to be computed at every step to re-calibrate the low-fidelity information. This would lead AMMF to become, at worst, a conventional optimization algorithm. At its best, the AMMF would be able to take many steps with the corrected low-fidelity model before resorting to re-calibration with expensive evaluations. Which scenario actually takes place depends on the problem at hand.

AMMF has shown promise with low-fidelity models represented by data-fitting approximations and variable-resolution models. In an attempt to evaluate the potential worst-case scenario, we are now considering managing variable physical fidelity models. We view this model combination as the potential worst-case scenario, because
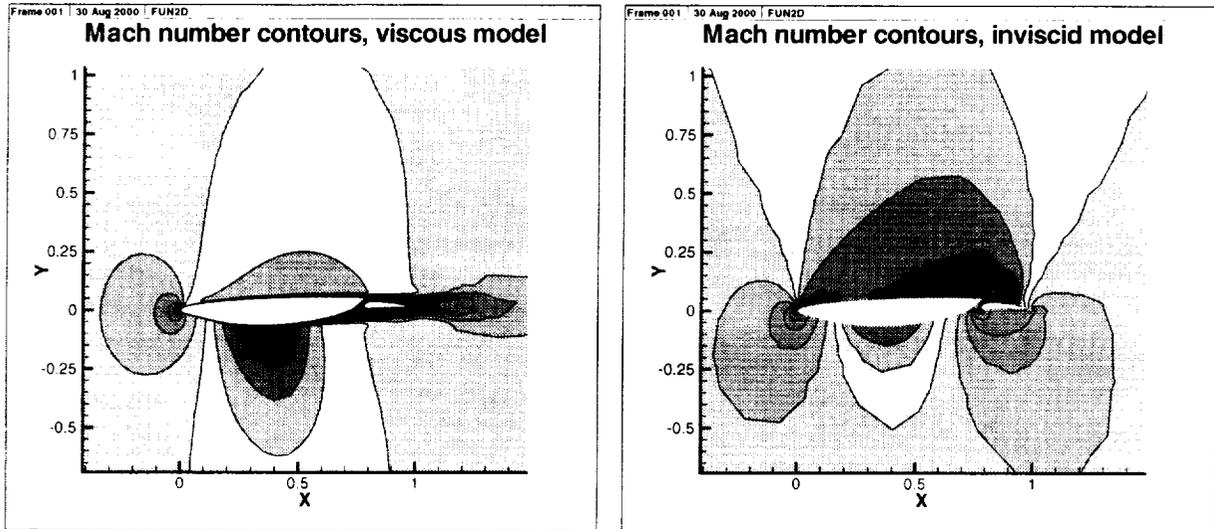
Figure 3: Mach number contours for viscous vs. inviscid model
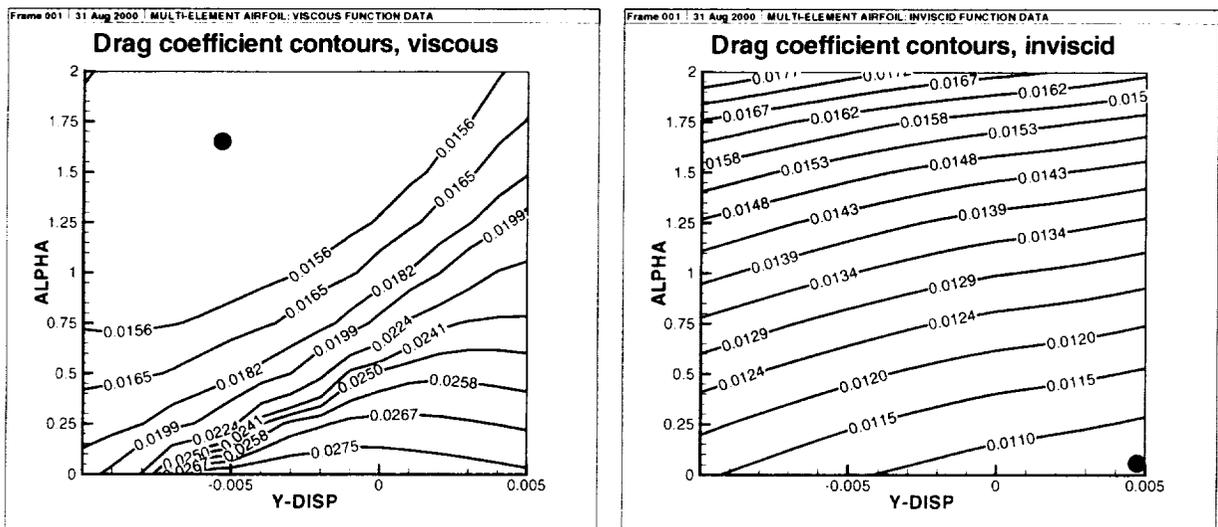
Figure 4: Drag coefficient contours of the viscous and inviscid models

low-fidelity physics models are expected not to capture the behavior of the high-fidelity counterparts accurately, or at all, over some or all regions of interest in the design space.

Demonstration problem

We consider aerodynamic optimization of a two-element airfoil designed to operate in transonic conditions [18]. The inclusion of viscous effects is very important for obtaining physically correct results. Therefore, the high-fidelity model will be the Reynolds-averaged Navier-Stokes equations and the low-fidelity model will

be the Euler equations. The flow solver, FUN2D, used for this study follows the unstructured mesh methodology [19]. Sensitivity derivatives are provided via a hand-coded adjoint approach [20].

The mesh for the viscous model depicted in Fig. 1 consists of 10449 nodes and 20900 triangles. The mesh for the inviscid model, shown in Fig. 2, comprises 1947 nodes and 3896 triangles. The Mach number is $M_\infty = 0.75$, the Reynolds number is $Re = 9 \times 10^6$, the global angle of attack is $\alpha = 1°$.

Fig. 3 depicts the Mach number contours for the viscous and inviscid model, respectively. The boundary and

Test	hi-fi eval	lo-fi eval	total time	factor
PORT with hi-fi surrogates, 2 var	15/15		negligible	
AMMF with surrogates, 2 var	3/3	18/9	negligible	≈ 5
PORT with hi-fi analyses, 2 var	14/13		≈ 12 hrs	
AMMF with direct analyses, 2 var	3/3	19/9	≈ 2.41 hrs	≈ 4.98

Table 1: AMMF performance vs PORT

shear layers are clearly visible in the viscous case. Because of the importance of the viscous effects in this problems, the use of the inviscid equations for the low-fidelity model should present an important test for the present approach.

The objective of this problem is simply to minimize the drag coefficient by adjusting the global angle of attack and the y-displacement of the flap. In this study, we restrict ourselves to two design variables to enable visualization. The baseline case for both models was constructed at $\alpha = 1°$ and zero y-displacement of the flap.

Fig. 4 depicts the level sets of the drag coefficient for the viscous and inviscid models. The problem appears to support the worse-case scenario: not only is the low-fidelity model not a good representation of the high-fidelity model but, in addition, the descent trends in the two models are reversed. The solution for each problem is marked with a circle. Thus the problem provides a good test of the methodology indeed.

The computational expense necessary to calculate functions and derivatives in the viscous case is considerably greater than that for the inviscid model. We conducted our experiments on an SGI™ Origin™ 2000 workstation with four MIPS RISC R10000 processors. One low-fidelity analysis took approximately 23 seconds and one low-fidelity sensitivity analysis took between 70 and 100 seconds. In contrast, one high-fidelity analysis took approximately 21 minutes and one high-fidelity sensitivity analysis took between 21 and 42 minutes to compute. The measures were taken in CPU time. Thus, the time per low-fidelity evaluation may be considered negligible compared to that required for a high-fidelity evaluation.

Numerical results

We conducted the following computational experiments. Because our test problem has expensive function evaluations, we first built spline substitutes both for the viscous and the inviscid model. Error analysis indicated that the spline fit was highly satisfactory for both models. It should be emphasized that we did not use these

substitutes in the conventional sense, i.e., they were not used to provide lower-fidelity models. Instead, they simply served to provide low-cost substitutes for both models for the problem components in the testing phase. Of course, such a test would never be conducted in a non-research setting, nor would it be considered for a problem with more than a few variables. In our setting, however, it saved us much time by providing an excellent approximation of the actual functions with respect to descent characteristics at a tiny fraction of computational cost. After we ascertained the correctness of our procedures, tests were conducted directly with the flow and adjoint solver, without recourse to substitutes, because the substitutes were expected to smooth out the problem to a certain degree.

The problems were first solved with single-fidelity models alone by using well-known commercial optimization software[¶] PORT [21], in order to obtain a baseline number of function evaluations or iterations to find an optimum. The problems were then solved with AMMF. Identical experiments were conducted with spline substitutes and with the actual flow and adjoint solver.

For each experiment, performance of AMMF was evaluated in terms of the absolute number of calls to the high and low-fidelity function and sensitivity calculations. Because the time for low-fidelity computations was negligible in comparison to the high-fidelity computations, we estimated the savings strictly in terms of high-fidelity evaluations. Table 1 summarizes the number of function (first number) and derivative (second number) computations expended in PORT and in AMMF.

Given the dissimilarity between the high-fidelity and low-fidelity model, we were initially surprised to find that the AMMF performed well: it consistently yielded approximately fivefold savings in terms of high-fidelity computations. The result held both for the spline substitutes and the actual functions. Optimization applied to both cases produced nearly identical iterates.

Following an analysis of the results, we concluded that the savings were not surprising after all. For our combination of models, the β-correction worked extremely well. This is illustrated in Fig. 5. The plot on the left shows the

---

[¶]The use of names of commercial software in this paper is for accurate reporting and does not constitute an official endorsement, either expressed or implied, of such products by the National Aeronautics and Space Administration or Institute for Computer Applications to Science and Engineering.
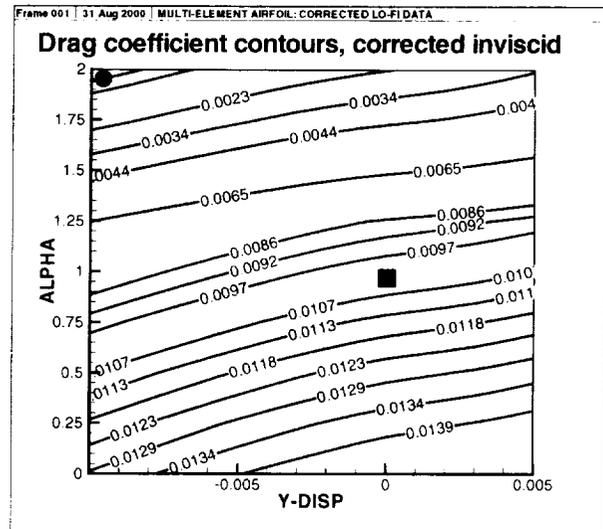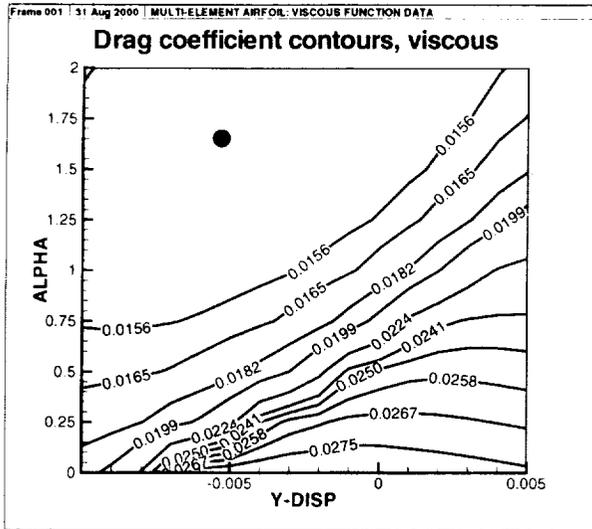
Figure 5: Drag coefficient level sets of the viscous and corrected inviscid models

level sets of the high-fidelity model with the solution. The plot on the right depicts the level sets of the low-fidelity model β-corrected at the initial point. The initial point is marked by a square. We note that the correction is not applied to the entire feasible region during iterations of optimization algorithm. Here, we applied the correction to the entire region to visualize the affect of the correction on the low-fidelity function. The figure clearly shows that the correction, using the function and derivative information at the anchor point (at this iteration - the initial point), reversed the trend of the low-fidelity model, allowing the optimizer to find the next iterate in the left upper corner of the plot, marked by a circle. Similar analysis can be conducted for all iterations. In fact, AMMF located the solution ($\alpha = 1.6305°$, flap y-displacement $= -0.0048$) of the high-fidelity problem already at the next iteration. The high-fidelity drag coefficient at the initial point was $C_D^{initial} = 0.0171$, the high-fidelity drag coefficient at the solution was $C_D^{final} = 0.0148$, a decrease of approximately 13.45%.

Concluding remarks

We believe that the results obtained in this study with AMMF and variable-fidelity physics models are promising. We observed fivefold savings in terms of high-fidelity evaluations compared to conventional optimization. Despite the great dissimilarity between the models, AMMF was able to capture the descent behavior of the high-fidelity model with the assistance of the first-order correction. When the models are greatly dissimilar, first-order information appears indispensable in obtaining reasonable descent directions.

Given the present results, we are cautiously optimistic about several much larger test cases (e.g., 84 variables) that are currently under investigation. Large problems must be tested carefully in AMMF in order to ascertain that its performance is not in some measure an artifact of the problem dimensionality. This does not appear to be the case, because AMMF was previously tested on problems with over ten variables. However, the tests currently conducted with realistic physical models should prove more conclusive, regardless of the outcome.

The performance of AMMF with variable-fidelity physics models raises a number of intriguing questions about the nature of the corrections and an optimal choice of low-fidelity models for a large set of problems. These questions are currently under investigation.

References

[1] L. A. Schmit, Jr. and B. Farshi. Some approximation concepts for structural synthesis. *AIAA Journal*, 12(5):692–699, 1974.

[2] L. A. Schmit, Jr. and C. Fleury. Structural synthesis by combining approximation concepts and dual methods. *AIAA Journal*, 18:1252–1260, 1980.

[3] L. A. Schmit, Jr. and H. Miura. Approximation concepts for efficient structural synthesis. Technical Report CR–2552, NASA, 1976.

[4] J.-F. M. Barthelemy and R. T. Haftka. Approximation concepts for optimum structural design—a review. *Structural Optimization*, 5:129–144, 1993.

[5] 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 problem. *The International Journal of Supercomputing Applications and High Performance Computing*, 10(4):269–299, Winter 1996.

[6] A.A. Giunta. *Aircraft Multidisciplinary Optimization Using Design of Experiments Theory and Response Surface Modeling Methods*. PhD thesis, Department of Aerospace and Ocean Engineering, Virginia Tech, Blacksburg, VA, 1997.

[7] N. M. Alexandrov. A trust region framework for managing approximation models in engineering optimization. In *Proceedings of the Sixth AIAA/NASA/ISSMO Symposium on Multidisciplinary Analysis and Optimization*. AIAA, 1996. AIAA-96-4102.

[8] R. M. Lewis. A trust region framework for managing approximation models in engineering optimization. In *Proceedings of the Sixth AIAA/NASA/ISSMO Symposium on Multidisciplinary Analysis and Design*, September 1996. AIAA Paper 96-4101.

[9] N. M. Alexandrov, J. E. Dennis, Jr., R. M. Lewis, and V. Torczon. A trust region framework for managing the use of approximation models in optimization. *Structural Optimization*, 15(1):16–23, 1998.

[10] N. M. Alexandrov. On managing the use of surrogates in general nonlinear optimization. In *Proceedings of the Seventh AIAA/USAF/NASA/ISSMO Symposium on Multidisciplinary Analysis and Optimization*. AIAA, 1998. AIAA Paper 98-4798.

[11] M. Avriel and A. C. Williams. Complementary geometric programming. *SIAM Journal on Applied Mathematics*, 19(1):125–141, July 1970.

[12] A. J. Morris. Approximation and complementary geometric programming. *SIAM Journal on Applied Mathematics*, 23(4):527–531, December 1972.

[13] V. Braibant and C. Fleury. An approximation-concepts approach to shape optimal design. *Computer Methods in Applied Mechanics and Engineering*, 53:119–148, 1985.

[14] P. Hajela. Geometric programming strategies in large-scale structural synthesis. *AIAA Journal*, 24(7):1173–1178, 1986.

[15] N.M. Alexandrov, R.M. Lewis, C.R. Gumbert, L.L. Green, and P.A. Newman. Optimization with variable-fidelity models applied to wing design. AIAA-2000-0841, January 2000. Proceedings of the 38th Aerospace Sciences Meeting and Exhibit. Also available as ICASE technical report 99-49.

[16] J. J. Moré. Recent developments in algorithms and software for trust region methods. In A. Bachem, M. Grötschel, and B. Korte, editors, *Mathematical Programming. The State of the Art, Bonn 1982*, pages 258–287. Springer-Verlag, Berlin, 1983.

[17] K. J. Chang, R. T. Haftka, G. L. Giles, and P.-J. Kao. Sensitivity-based scaling for approximating structural response. *Journal of Aircraft*, 30(2):283–288, March–April 1993.

[18] Richard T. Whitcomb and Larry R. Clark. An airfoil shape for efficient flight at supercritical mach numbers. Technical Report NASA TM-X-1109, NASA Langley Research Center, 1965.

[19] W. Kyle Anderson and Daryl L.Bonhaus. An implicit upwind algorithm for computing turbulent flows on unstructured grids. *Computers and Fluids*, 23(1):1–21, 1994.

[20] W. Kyle Anderson and Daryl L.Bonhaus. Aerodynamic design on unstructured grids for turbulent flows. Technical Report NASA TM 112867, NASA Langley Research Center, June 1997.

[21] Linda Kaufman and David Gay. *PORT Library: Optimization and Mathematical Programming*. Bell Laboratories, May 1997.