

**Using Automatic Differentiation
with the Quasi-Procedural
Method for Multidisciplinary
Design Optimization**

*Steve Altus Christian Bischof
Paul Hovland Ilan Kroo*

**CRPC-TR95634-S
October 1995**

Center for Research on Parallel Computation
Rice University
6100 South Main Street
CRPC - MS 41
Houston, TX 77005

Using Automatic Differentiation with the Quasi-Procedural Method for Multidisciplinary Design Optimization

Paul Hovland*

Steve Altus†

Christian Bischof‡

Ilan Kroo†

AIAA 96-0090

Abstract

As computers have become increasingly powerful, the field of design optimization has moved toward higher fidelity models in the early stages of design. One way in which this movement is manifested is in the increasing popularity of multidisciplinary design optimization (MDO). Because the models used in MDO are large and complicated, a modular design is desirable. There are many design parameters to optimize, and the robustness of the method requires that derivatives be computed accurately and efficiently. This paper describes how the quasi-procedural program architecture developed by Takai and Kroo and the technique of automatic differentiation can be combined to address these needs effectively. The two techniques are explained, the manner in which they were integrated into a single framework is described, and the result of using this framework for an optimization problem in airplane design is presented.

1 Introduction

Over the past several years, there has been a movement in the field of optimization toward multidisciplinary design optimization (MDO), moving away from the traditional method of iteratively executing local optimizations of disciplinary components. For example, an airplane designer may incorporate fluid dynamics and structural analysis into a single model. For optimization using this multidisciplinary model to be practical, the model should

have a modular design and derivatives must be computed efficiently. Modularity is important because the multidisciplinary model may be created by several development teams and also because it simplifies the integration of new code when a better model for any of the disciplines becomes available. The quasi-procedural method is a modular framework for optimization that promotes efficiency by avoiding redundant computations. Automatic differentiation improves performance through fast, accurate computation of derivatives and supports the integration of new code by automatically creating modules that compute derivatives with minimal user intervention.

This paper describes how automatic differentiation and the quasi-procedural method were applied to the optimization of an airplane design. Sections 2 and 3 introduce the quasi-procedural method and automatic differentiation, respectively. Section 4 describes how these complementary techniques were incorporated into a single environment and Section 5 presents our experimental results using this environment. Section 6 concludes with an analysis of how the combined approach might facilitate MDO, and also an assessment of when this approach might not be appropriate.

2 The quasi-procedural method

The quasi-procedural method (QPM) is a form of non-procedural programming.^{6, 9} Unlike conventional procedural programs, in which computation proceeds from inputs to outputs according to a rigid structure, non-procedural systems are free to reorganize computations as necessary to compute the desired outputs. Thus, it is the outputs that drive the computation, rather than the inputs. However, non-procedural programming is of limited value at fine granularity, as programmers can often utilize knowledge of a computation to develop extremely efficient small procedural subprograms.

*Computer Science Department, University of Illinois at Urbana-Champaign, 1304 W. Springfield Ave., Urbana, IL 61801, hovland@uiuc.edu. This work was performed in part while the author was a student research associate at Argonne National Laboratory.

†Department of Aeronautics and Astronautics, Stanford University, Stanford, CA 94305,
{altus, kroo}@leland.stanford.edu.

‡Mathematics and Computer Science Division, Argonne National Laboratory, 9700 S. Cass Ave., Argonne, IL 60439, bischof@mcs.anl.gov.

Copyright ©1995 by the American Institute of Aeronautics and Astronautics, Inc. All rights reserved.

2.1 A composite system

The quasi-procedural method attempts to exploit the best of both methods by allowing the programmer to develop efficient subprograms and providing a system for linking these subprograms so that they can be executed non-procedurally. This linkage system is demand-driven. When the value of a variable is requested, the executive system determines which subprogram is responsible for computing that value and runs the appropriate subroutine. If that routine requires inputs, it informs the executive; the executive provides the desired inputs either by looking them up in a database or by executing additional routines.

This type of request-driven execution is depicted in Figure 1. For example, if the value of I is desired, the executive invokes subprogram 5. Subprogram 5 requires values for G and H, which causes the executive to invoke subprograms 3 and 4. This sequence of requests continues until values for A, B, C, and D are provided, at which point execution of the subprograms commences, and a value for I is produced. This approach is coupled with a database that stores intermediate values and maintains information about their validity, enabling the executive to determine whether to use a previously computed value or to recompute the variable; this feature is called consistency maintenance.

The impact of maintaining information about the validity of intermediate values can be seen by again considering the program in Figure 1. Suppose we compute the value of I as before, then decide we want to change the value of design parameter C and recompute I. Since the values of E, F, and G do not depend on C, there is no need to recompute these values by executing subprograms 1, 2, and 3. Instead, it is sufficient to recompute H and I by executing subprograms 4 and 5.

An important characteristic of the quasi-procedural method in addition to its efficiency is its modularity. Each of the subprograms represents a separate module, and it is easy to replace one of these subprograms without affecting the rest of the computation. As an example of how the quasi-procedural method breaks the tight coupling between subprograms in a procedural program, consider the procedural and quasi-procedural versions of the PASS airplane model depicted in Figures 2 and 3 (see Section 5 for a description of this model). The ability to incorporate new code without having to rewrite, or even recompile, modules corresponding to other parts of the computation is very important to MDO. Thus, support for modular program-

ming and large reductions in computational cost are the two main benefits of using the quasi-procedural method.

2.2 GENIE

GENIE is a generic framework for engineering computations.^{6, 9} GENIE provides a set of routines for linking analysis routines through the central database, which is updated under the control of the quasi-procedural method. This interface is provided via GET and PUT routines, and the computation proceeds in the following manner:

1. The optimizer issues a GET operation, signaling a request for a valid value for the specified objective or constraint.
2. If no valid value is currently available, the GET routine calls the appropriate analysis routine to compute one.
3. The analysis routine issues one or more calls to GET to load the required input variables. In the event that the values of these variables are not known or are invalid, the GENIE executive calls the appropriate analysis routine, and the process repeats.

Thus, GET is called recursively until the values of all required input variables are known. This method enables quantities to be computed only as needed and without a fixed execution path. An example of a simple analysis routine is provided in Figure 4.

A major function of the GENIE framework is maintaining information about the validity of the values of variables. Whenever one of the input parameters for a routine is modified this parameter is marked ‘invalid.’ The quasi-procedural executive then marks all outputs of that routine, as well as variables directly or indirectly dependent on these outputs, as ‘invalid.’

3 Automatic Differentiation and ADIFOR

In general, multidisciplinary design optimization requires the derivatives of an objective function and several constraints with respect to many design parameters. Since the function is typically described by a complicated computer program, using a symbolic manipulator, such as Maple,⁵ is usually not an option. Similarly, developing derivative code by hand is unattractive, because it is complicated, tedious, and prone to errors. This approach is also ill-suited for rapid prototyping, where parts of the system model may change several times, requiring ad-

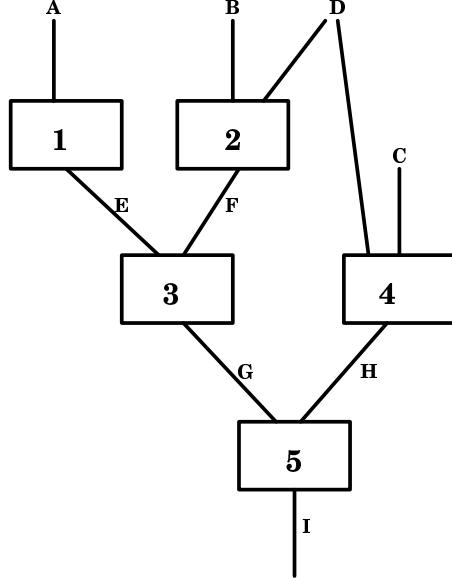


Figure 1: Execution outline for the quasi-procedural method

ditional code development for each new part. Consequently, optimization often relies on divided difference approximations to the desired derivatives. However, if an inappropriate step size is used, these approximations can be grossly inaccurate. This hinders rapid prototyping, because finding a good step size can be difficult and time-consuming, and a new step size must be determined each time the system model changes. Divided difference approximations may also take a long time to compute.

An alternative to all of these techniques is automatic differentiation. Automatic differentiation is a technique for computing the derivatives of a function expressed in the form of a computer program.⁸ The execution of a computer program consists of the composition of many elementary functions (such as multiplication, square root, and hyperbolic cosine), for each of which an analytic expression for derivatives is well known. So, by simply applying the chain rule

$$\frac{\partial}{\partial t} f(g(t)) \Big|_{t=t_0} = \left(\frac{\partial}{\partial s} f(s) \Big|_{s=g(t_0)} \right) \left(\frac{\partial}{\partial t} g(t) \Big|_{t=t_0} \right)$$

repeatedly, it is possible to compute the derivatives of the function. For example, the code segment:

```

y = 2*x*x + 3*x + 7
z = 4 * sin(x)
f = sqrt(y*y + z*z)

```

may be converted into:

```

y = 2.0*x*x + 3.0*x + 7.0
g_y = 4.0*x*g_x + 3.0*g_x
z = 4.0 * sin(x)
g_z = 4.0*g_x*cos(x)
f = sqrt(y*y + z*z)
g_f = -0.5 * (2.0*y*g_y + 2.0*z*g_z) / f

```

Here, g_{-var} represents the derivatives of var with respect to the independent variable(s), in this case x . Thus, if g_x is initialized to 1.0 (since $\frac{\partial x}{\partial x} = 1.0$), upon exit the value of g_f will be $\frac{\partial f}{\partial x}$. If x, y, z , and f were vectors, then the appropriate initial value for g_x would be an identity matrix and g_f would still represent $\frac{\partial f}{\partial x}$. For this problem, the value of g_x is propagated through the derivative computation; hence, g_x is termed the *seed matrix* for this computation.⁴

While symbolic differentiation uses the rules of calculus in a more or less mechanical way, automatic differentiation is intimately related to the program for the computation of the function to be differentiated. By applying the chain-rule step by step to the elementary operations executed in the course of computing the “function,” automatic differentiation computes exact derivatives (within the limits of finite precision arithmetic) and avoids the potential pitfalls of divided differences. The technique of automatic differentiation is directly applicable to complicated functions with branches and loops.

Automatic differentiation is amenable to modular program design. The ability to control which derivatives are computed through an appropriate

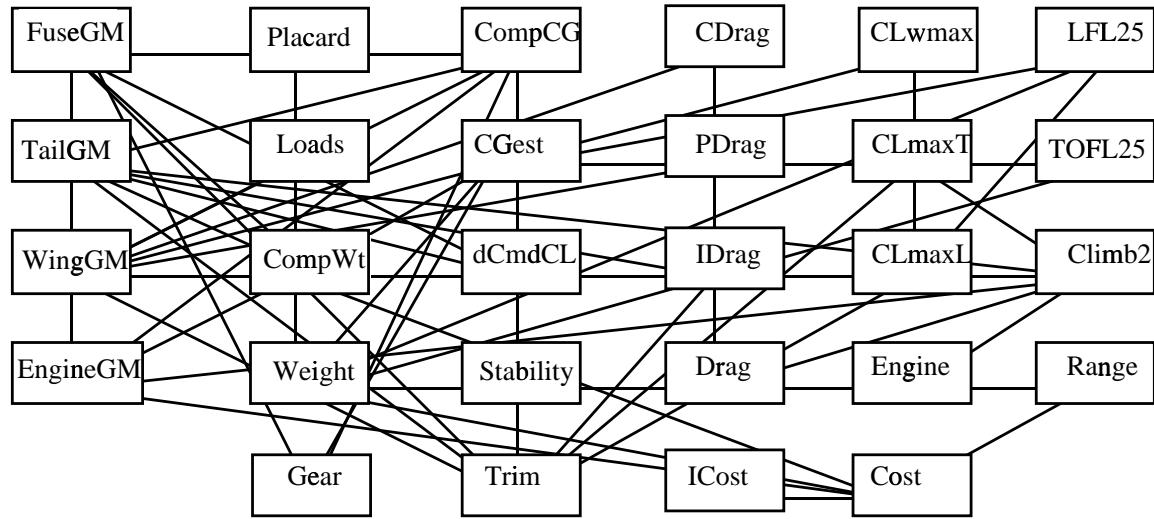


Figure 2: Procedural version of the PASS airplane model

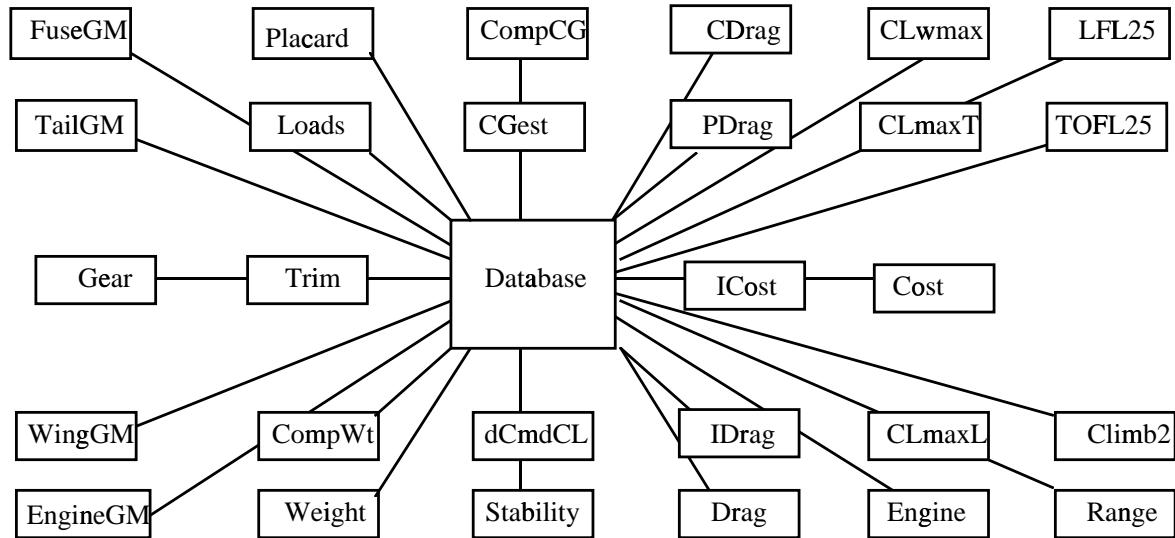


Figure 3: Quasi-procedural version of the PASS airplane model

```

C-----  

      subroutine Qpsf  

C     Profile:  

C-----  

C     This routine computes the dynamic pressure for a given  

C     density and velocity.  

C     Declarations:  

C-----  

      implicit none  

      real Vknobs, rho, Vfps, DynamicPress  

      logical Abort  

C     Required inputs:  

C-----  

      call GET(Vknobs,'Speed'    )  

      call GET(rho , 'Density')  

      if(Abort()) return  

C     Calculations:  

C-----  

      Vfps = Vknobs*1.69  

      DynamicPress = .5 * rho * Vfps*Vfps  

C     Repack database:  

C-----  

      call PUT(DynamicPress,'DynamicPress')  

      return  

      end

```

Figure 4: A simple analysis routine

initialization of the seed matrix means that we do not need to know which design parameters are being optimized at the time the derivative module is created. Furthermore, we can process individual modules separately, then connect the derivative modules using seed matrices. For example, suppose module **A** computes $y(x)$ and module **B** computes $z(y)$, and that we process these modules to yield derivative modules **g_A** and **g_B**. Then, if we initialize **g_x** (the seed matrix for module **g_A**) to an identity matrix, module **g_A** will compute $g_y = \frac{\partial y}{\partial x}$. If we then pass **g_y** to module **g_B** as a seed matrix, this module will compute $g_z = \frac{\partial z}{\partial y} \times g_y = \frac{\partial z}{\partial y} \times \frac{\partial y}{\partial x} = \frac{\partial z}{\partial x}$ directly, exactly as if $z(x)$ were computed in a single module.

Many tools have been developed to support automatic differentiation. We used ADIFOR for our application. ADIFOR is a tool that provides automatic differentiation for programs written in Fortran 77.^{1, 2, 3} Given a Fortran subroutine (or collection of subroutines) for a function f , ADIFOR produces Fortran 77 subroutines for the computation of the derivatives of this function. Additional information on the ADIFOR system can be found on the WWW at <http://www.mcs.anl.gov/adifor> and <http://www.cs.rice.edu/~adifor>.

4 Adding automatic differentiation to the QPM framework

In order to allow GENIE to work together with ADIFOR-generated code, several enhancements were made to the GENIE framework. First, additional storage in the database was allocated, so that the database manager would have additional space to store gradients. This storage is provided by an array that is parallel to the array in which the regular variables are stored. This enables lookup, validity checking, and other database maintenance operations to be performed for both a variable and the associated gradient at the same time. Second, the executive was modified by adding two new routines, **G_GET** and **G_PUT**, which return and store the values of a variable and the associated gradient. Finally, one of the GENIE initialization routines was modified to perform the seed matrix initialization required by ADIFOR-generated code. By zeroing all elements of the array containing gradients, then setting the I^{th} element of the gradient associated with the I^{th} design variable equal to 1.0, an identity seed matrix is created automatically.

Making these changes to the GENIE framework has two important benefits. First, the user does not have to deal with derivative object allocation, seed matrix initialization, or interfaces between modules.

All of these tasks are done automatically in the executive. Once an analysis routine has been run through ADIFOR, the user does not even need to be aware that automatic differentiation is being used. Second, creating a direct association between variables and their derivative objects in the database means that consistency maintenance for the derivative objects can be done with no additional overhead.

5 Experimental Results

To examine the suitability of the QPM-AD combination for multidisciplinary optimization, we applied the techniques to a complete aircraft model. The problem being studied is the synthesis of a twin-engined, 100-passenger, medium-range commercial transport. The objective is to minimize direct operating costs, subject to certain constraints in performance measures such as range and maximum field lengths. The design variables are weights, wing and tail size and shape parameters, engine size, cruise altitudes, and takeoff flap deflection.

The multidisciplinary analysis routines for the PASS aircraft model were processed using ADIFOR, which automatically replaced calls to **GET** and **PUT** with calls to **G_GET** and **G_PUT**, respectively. These routines were compiled and linked with the enhanced GENIE framework and the NPSOL optimizer.⁷

The times required for the optimizer to find a minimum for various problems are reported in Tables 1 and 2. Problem PASS1 involves the optimization of 14 design parameters. This problem was difficult to optimize, and the observation that one design variable was not critical to the design led to its removal, yielding problem PASS2. Problems PASS2 and PASS3 involve the optimization of 13 design parameters, from different starting points. Problem PASS4 is the same as PASS3, but the optimality tolerance is relaxed from 10^{-8} to 10^{-5} . Results are reported for the quasi-procedural method using divided difference approximation, abbreviated QPM-DD, and for the augmented quasi-procedural method, using derivative code generated by ADIFOR, abbreviated QPM-AD. The ratio of QPM-DD to QPM-AD is also reported.

Design optimization using the quasi-procedural method is typically faster than design optimization without the quasi-procedural method.⁶ Because the quasi-procedural method employs consistency maintenance and does not recompute values unless necessary, a great deal of computational cost can be eliminated. Furthermore, the quasi-

	QPM-DD	QPM-AD	ratio
PASS1	46.8	5.0	9.36
PASS2	14.9	4.2	3.55
PASS3	19.9	5.6	3.55
PASS4	19.9	5.5	3.62

Table 1: Runtime (in seconds) for the PASS problem on an IBM RS6000

	QPM/DD	QPM/AD	ratio
PASS1	11.7	1.2	9.75
PASS2	3.7	1.0	3.70
PASS3	4.9	1.3	3.77
PASS4	5.0	1.2	4.17

Table 2: Runtime (in seconds) for the PASS problem on a node of an IBM SP1 parallel computer

procedural method with automatic differentiation often performs much better than the quasi-procedural method using divided difference approximations. There are several reasons for the improvement:

1. The hybrid mode of automatic differentiation implemented by ADIFOR is often more efficient than divided differences. This is especially true of programs with many assignment statements with complicated arithmetic expressions on the right hand side. The derivatives of these expressions are computed using the reverse mode of automatic differentiation, which requires a constant multiple of the time required to evaluate the expression. Divided differences and the forward mode of automatic differentiation require time linear in the number of design variables. This improvement in performance relative to divided differences has been seen in many applications of ADIFOR to engineering codes, and is not specific to the quasi-procedural method (see [1, 2] for details).
2. The perturbing of design variables needed to compute approximate derivatives using divided differences destroys the validity of the values of variables depending on those design variables. Thus, much of the efficiency of the quasi-procedural method is lost. GENIE compensates for this inefficiency by employing a specialized technique⁶ to avoid recomputation as much as

possible. However, this advanced technique is inferior to the ability to compute derivatives without affecting the validity of any values, as is provided by automatic differentiation.

3. The gradient objects used in the automatic differentiation method are vectors. Hence, AD-generated code performs mostly vector updates throughout the computation. Modern superscalar computer architectures can take advantage of the regularity and data locality of these vector updates. The variable perturbations necessary for divided difference approximations must be done one at a time and there is therefore no way to introduce vector operations.

6 Conclusions

Multidisciplinary optimization has three distinguishing features: the system is often modelled by a large, complicated program developed by many different teams; there may be many changes to the system model, due either to rapid prototyping or model refinement; and there is a need for derivatives with respect to many different design parameters, but not necessarily the same set of parameters from iteration to iteration. The quasi-procedural method and automatic differentiation can together provide a framework that is well-suited for optimization problems of this nature.

Large, complicated programs with multiple authors are most easily expressed in a modular fashion. Both quasi-procedural programming and automatic differentiation support this paradigm. Changes to the system model also create a need for modularity, as well as a mechanism for developing derivative code for new modules as quickly and easily as possible. Automatic differentiation is capable of automatically creating derivative code from function code, with minimal user intervention. Differentiating with respect to many design parameters creates a need for efficient derivative computation. The combination of automatic differentiation and the quasi-procedural method provides derivative code that is much faster than standard divided difference approximations. This speedup can be attributed to the efficiency of the hybrid mode of automatic differentiation, the capability of quasi-procedural programming to avoid redundant computation, the vectorizability of gradient object updates, and the synergistic effects of using the two methods together. Furthermore, derivatives computed using automatic differentiation are more accurate than divided difference approximations, which may

lead to more rapid convergence for some applications. The seed matrix interface of automatic differentiation also provides a convenient mechanism for handling variation in the set of parameters with respect to which we wish to differentiate.⁴

Automatic differentiation and the quasi-procedural method are not without their limitations. In particular, automatic differentiation produces code that computes accurate *local* derivatives. However, if the function being differentiated exhibits high frequency oscillations, this information might not be useful for optimization. Derivative-free optimization techniques such as simulated annealing and genetic algorithms used in combination with the quasi-procedural method have proven effective in handling such functions.¹⁰ However, for problems amenable to classical optimization approaches, the combination of the quasi-procedural method and automatic differentiation promises to be an effective tool for multidisciplinary optimization.

Acknowledgments

This work was supported by the Mathematical, Information, and Computational Sciences Division subprogram of the Office of Computational and Technology Research, U.S. Department of Energy, under Contract W-31-109-Eng-38, by the National Aerospace Agency under Purchase Order L25935D, by the U.S. Department of Defense through an NDSEG fellowship, and by the National Science Foundation, through the Center for Research on Parallel Computation, under Cooperative Agreement No. CCR-9120008. We thank Peter Gage for his many contributions to the development of the combined QPM-AD framework and for his assistance in the refinement of this paper. We also thank Hippokrates Hadjilias for his assistance in the early stages of this research.

References

- [1] C. Bischof, A. Carle, G. Corliss, A. Griewank, and P. Hovland. ADIFOR: Generating derivative codes from Fortran programs. *Scientific Programming*, 1(1):11–29, 1992.
- [2] C. Bischof, A. Carle, P. Khademi, and A. Mauer. The ADIFOR 2.0 system for the automatic differentiation of Fortran 77 programs, 1994. Preprint MCS-P481-1194, Mathematics and Computer Science Division, Argonne National Laboratory, and CRPC-TR94491, Center for Research on Parallel Computation, Rice University.
- [3] C. Bischof, A. Carle, P. Khademi, A. Mauer, and P. Hovland. ADIFOR 2.0 user’s guide. Technical Memorandum ANL/MCS-TM-192, Mathematics and Computer Science Division, Argonne National Laboratory, 1994. CRPC Technical Report CRPC-95516-S.
- [4] C. Bischof and P. Hovland. Using ADIFOR to compute dense and sparse Jacobians. Technical Report ANL/MCS-TM-158, Mathematics and Computer Science Division, Argonne National Laboratory, 1991.
- [5] B. W. Char, K. O. Geddes, G. H. Gonnet, B. L. Leong, M. B. Monagan, and S. M. Watt. *Maple V Language Reference Manual*. Springer Verlag, New York, 1991.
- [6] P. Gage and I. Kroo. Development of the quasi-procedural method for use in aircraft configuration optimization. Technical Report AIAA-92-4693, Fourth AIAA/USAF/NASA/OAI Symposium on Multidisciplinary Analysis and Optimizations, September 1992.
- [7] P. E. Gill, W. Murray, M. A. Saunders, and M. H. Wright. User’s guide for NPSOL (version 4.0): A Fortran package for nonlinear programming. Technical Report SOL 86-2, Department of Operations Research, Stanford University, Stanford, CA, 1986.
- [8] A. Griewank. On automatic differentiation. In *Mathematical Programming: Recent Developments and Applications*, pages 83–108, Amsterdam, 1989. Kluwer Academic Publishers.
- [9] I. Kroo. An interactive system for aircraft design and optimization. Technical Report AIAA-92-1190, AIAA Aircraft Design Conference, February 1992.
- [10] I. Kroo, S. Altus, R. Braun, P. Gage, and I. Sobieski. Multidisciplinary optimization methods for aircraft preliminary design. Technical Report AIAA-94-2543, Fifth AIAA/NASA/USAF/ISSMO Symposium on Multidisciplinary Analysis and Optimization, September 1994.