AIAA 2002-2838

Aerodynamic Shape Optimization for Supersonic Aircraft

James Reuther
NASA Ames Research Center
Moffett Field, CA 94035

32nd AIAA Fluid Dynamics Conference

June 24-27 2002, St. Louis, MO

Aerodynamic Shape Optimization for Supersonic Aircraft

James Reuther*
NASA Ames Research Center, Moffett Field, CA 94035

This paper presents a historical perspective on the development and application of aerodynamic shape optimization (ASO) methods for supersonic aircraft design. The motivation for using ASO for supersonic transport design is discussed. The accelerated research and development that several ASO methods underwent during the NASA High Speed Research Program is presented. The various ASO algorithms are compared to facilitate an understanding of the proper application of these methods. The continued development of ASO for supersonic aircraft on several fronts is described, including aero-structural optimization, sonic-boom minimization and laminar flow wing design. The paper closes with a discussion on future research activities and possible pathways to the creation of the next supersonic commercial transport.

Introduction

Considering the remarkable developments in many high technology fields during the past few decades, it is surprising that virtually no change in the speed of commercial air transportation has occurred since the introduction of the Boeing 707 in the late 1950s. Instead, the aging Concorde continues to fly at supersonic speeds for an extremely limited market and all other commercial aircraft have kept well below Mach 1. Several attempts have been mounted by U. S., European and Japanese organizations to develop the technologies needed for the next generation supersonic commercial transport. All of these endeavors, while succeeding in developing or improving key technologies, have still failed to produce a viable aircraft. Nevertheless, some of the technologies are noteworthy because they also have applications to other flight regimes and disciplines.

One of the most rewarding technologies fostered during the supersonic aircraft programs of the 1990s (NASA's High Speed Research (HSR) program, the European Supersonic Research Program (ESRP), and the SST Program of NAL in Japan) was aerodynamic shape optimization (ASO). In contrast to the evolutionary development process that progressed slowly over the last four decades to produce the current generation of transonic commercial transports, the next supersonic transport will have to achieve a level of efficiency that somehow allows it to become economically competitive with its transonic rivals. Achieving this level of efficiency in all aspects of the aircraft's engineering can be addressed mainly with advanced approaches to aircraft design. A critical engineering element of any commercial transport, and particularly for a supersonic transport, is its aerodynamic efficiency. The NASA HSR program therefore nurtured the rapid technology development of ASO methods, which were proven through wind tunnel validation to far surpass the "cut-and-try" approaches used in the past. Both during and subsequent to the HSR program, ASO methods have been adopted as a workhorse by many aerodynamic design teams. Their success has, in turn, facilitated the advent of even more sophisticated design strategies that attempt true high-fidelity multidisciplinary optimization (MDO).

In many respects, the real hurdles that must be overcome in achieving cost-effective commercial supersonic transportation are not the daunting engineering challenges that still need years of work, but instead the environmental and economical hurdles. With current technologies, large commercial supersonic transports will produce significant sonic booms, present a hazard to the ozone layer and require design, development, test and evaluation (DDT&E) costs that exceed any single corporation's ability to finance. It is no wonder that the two most serious U. S. efforts to develop affordable supersonic transportation—the SST program (circa 1970) and the NASA HSR program (1990-1999)—stalled well before serious corporate money was invested. Each of these efforts took the direct approach to attaining supersonic commercial travel by proposing aircraft that carried 300 passengers, flew at least twice the speed of sound and had the range for trans-Pacific routes. The "make or break" risks and challenges for such a vehicle were simply too grave.

* Senior Research Scientist, Reacting Flow Environments Branch, Associate Fellow, AIAA

However, there are other stepping-stone approaches to achieving the goal of affordable supersonic travel that may not entail navigating similar hazards. Two such examples are Boeing's Sonic Cruiser aircraft development program and the supersonic business jet studies under way at Lockheed Martin and Northrop Grumman under DARPA's Quite Supersonic Platform (QSP) program.[1] In the case of the Sonic Cruiser, it is suspected by some[2] (including the author) that the aircraft, while marginally viable in its stated mission of operation at Mach 0.95, is really intended to be a Mach 1.4 aircraft over water. By first certifying the airplane at Mach 0.95, Boeing strategically opens the door for eventually obtaining approval for regular commercial flight operations above Mach 1. Similarly, it may be far less risky and more financially palatable first to undertake a supersonic business-jet development as a pathfinder for larger commercial aircraft.

In either of these more likely paths to affordable supersonic commercial travel, the role played by ASO methods will be significant. Any future supersonic business jet or Sonic Cruiser will constitute a radical departure from existing aircraft configurations. Just as in the HSR program, it will not be enough for these new aircraft to out-perform existing aircraft; they must do so while being economically competitive and hence aerodynamically efficient. Use of ASO and future MDO methods seems the only viable approach to rapidly achieving the necessary levels of efficiency on such radically different concepts that have no legacy of previous designs upon which to base improvements.

This paper presents a historical account of NASA's HSR program, focused upon the development of aerodynamic design optimization strategies. Attention then turns to developments in ASO methods that have occurred since the HSR program and steps that are currently under way on multiple fronts to implement MDO strategies. Finally, an assessment of future work toward an affordable supersonic commercial transport is presented.

Motivation for Aerodynamic Shape Optimization for Supersonic Transports

The most significant difference between the HSR program and previous attempts to make the leap toward a supersonic transport was its stronger emphasis on targeting an economically viable aircraft. The target aircraft that was specified to validate technology developments was to fly at Mach 2.4 and carry 300 passengers 5,500 nautical miles with a seat-mile cost only 15% above that of a 747 or 777. Developing a second-generation supersonic transport that could compete on such near-equal footing with contemporary transonic transports implied that entirely new design and engineering practices had to be employed.

Perhaps the most important element of this newer design and engineering philosophy is achieving highly refined aerodynamic efficiency. Prior to HSR, design of SST configurations was performed through the application of linear theory-type methods. The best documented of these approaches are those used during the U. S. SST Program of the 1970s.[3-9] Unfortunately, these approaches fail to provide the requisite level of aerodynamic efficiency. Conceptual design studies explored during the 1980s by Kulfan[10] and others showed that even with anticipated improvements in propulsion and structures it would still be necessary to improve the *M(L/D)* (Mach times lift-to-drag ratio) beyond that achievable by linear theory by 10%. Considering that even a 0.1% improvement in aerodynamic efficiency for transonic transports is difficult to achieve, the required improvement in supersonic efficiency seemed unattainable with traditional methods.

Counterbalancing the situation, during the late 1980s and early 90s, was the advent of more powerful computer systems and more mature computational fluid dynamics (CFD). By the beginning of the HSR program in 1990, many aerodynamicists were already routinely computing solutions in three dimensions for complex configurations. Thus, the stage had been set for a serious attempt at improving aerodynamic efficiency through the use of nonlinear CFD methods.

Beyond Linear Methods: NASA's High Speed Research Program

The development of accurate and efficient CFD methods for arbitrary configurations meant that aircraft design teams could repeatedly test subtle design changes computationally and resort to wind tunnel tests only for validation purposes. These modern techniques could account for nonlinear aerodynamic effects including compressibility, viscosity and geometric complexity. Indeed, by the time the HSR program had commenced, design teams working on improved transonic configurations were actively applying CFD. Unlike subsonic applications, the aerodynamic efficiency of transonic and supersonic designs can be extremely sensitive to subtle changes in the aircraft shape. As a result, many years can be spent developing the final wing shape for a new transonic transport. This reality is despite the fact that new configurations already have highly refined starting points—witness the fleet of existing aircraft. A future supersonic aircraft will not have the luxury of a four-decade long aerodynamic maturity effort. It must achieve a level of aerodynamic maturity comparable with that of contemporary transonic transports much more rapidly. However, using CFD methods in the "cut-and-try" approach of the past can be painfully slow when a design problem must be parameterized by hundreds of shape variables. The obvious answer was to couple mature CFD algorithms with numerical optimization procedures.

During the 1970s and 80s, efforts by Drela,[11] Hicks,[12,13] Kennelly[14] and others had successfully led to airfoil design methods based upon numerical optimization techniques. However, it was not until the beginning of the HSR program that a critical need for advanced, complex-configuration design optimization methods materialized. By the conclusion of the HSR program in 1999, aerodynamic shape optimization methods had matured to the extent that they are now considered a mainstay of aerodynamic design.

Aerodynamic shape optimization strategies based upon CFD fall into two basic categories: gradient-based approaches and non-gradient-based approaches. The choice between these two strategies is often dependent upon the application. Problems more suited to non-gradient approaches are characterized by a design space that may be parameterized by a limited number of variables or a space that does not vary smoothly with respect to the design parameters. An example of this type is the MDO problem of determining the correct planform shape for a supersonic aircraft. Structures, integrated propulsion and aerodynamics all play together to produce a highly nonlinear and discontinuous design space that cannot be tackled by a gradient-based approach. On the other hand, the final refinements to determine the best wing loft for a fixed planform is a problem that is best tackled by a gradient-based approach.

An interesting subtype of gradient-based ASO is the design method referred to as the inverse method. The idea behind inverse methods is to reverse the boundary conditions of the flow solution method. Instead of calculating the surface pressure distribution for given geometry, the technique attempts to calculate the shape that produces a desired pressure distribution—hence the term inverse. The approach can be thought of as a special case of a gradient-based optimization strategy where the objective is obtaining the desired pressure distribution and the changes to the surface shape are the design parameters. Of course, many inverse methods solve the underlying optimization problem in a tightly coupled process where the distinction between what is considered the CFD tool and what is considered the design tool becomes blurred. The important point that distinguishes them as gradient methods is the fact that the objective (meeting a target pressure distribution) varies smoothly with respect to the design parameters (changes in the shape). In any case, the use of inverse methods for the design of transonic wings was perhaps the first big breakthrough for ASO during the late 1970s.[15]

Many years of experience had shown skilled aerodynamicists what the shapes of near-optimal pressure distributions should be at transonic speeds. Inverse methods allowed these engineers to design the subtle shapes that would produce exactly the desired pressures. Unfortunately, the situation in supersonic design was different. Without the experience base, there simply was not a clear understanding of what an optimal pressure distribution should be. In addition, for supersonic aircraft with shocks emanating from wing, fuselage and nacelles/pylons/diverters, designing all components together is indispensable for achieving optimum configurations. Consequently, engineers were focused on improving the preeminent figure of merit, $M(L/D)$, rather than target pressures. To improve this parameter, researchers pursued the application of gradient methods.

The advantage of gradient methods is that the number of function evaluations (CFD solutions) required ranges from $O(n)$ to $O(n^2)$, where n is the number of design variables. In contrast, for non-gradient methods, the number of function evaluations normally needed scales by a very high power of n. Thus, non-gradient methods are usually attempted for only small numbers of design parameters while gradient methods have been employed for problems with several hundred design parameters. This paper focuses on the challenges of applying gradient-based design methods because they constitute the bulk of the modern CFD-based ASO techniques. One simple reason is that most aerodynamic objective functions vary smoothly with respect to their design parameters. Even modest drag reductions represent substantial savings over the service life of an airliner. Thus, there are significant potential benefits in applying gradient-based strategies.

The HSR program was a joint NASA and industry effort. The primary participants in the area of cruise aerodynamics were Boeing, McDonnell Douglas, NASA Ames and NASA Langley. Each group developed its own CFD-based design strategies and capabilities. By 1992, the NASA Ames team had combined three-dimensional Euler CFD methods (FLO57[16] and FLO67) with an unconstrained quasi-Newton numerical optimization scheme.[17] At around the same time, the Boeing team had implemented a novel, tightly-coupled constrained optimization strategy using the TRANAIR[18] CFD code and NPSOL[19] optimization algorithm. Meanwhile, the Langley team was investigating a more heuristic approach to supersonic wing design termed "Natural Flow Wing."[20] The McDonnell Douglas team joined the fray by initially adopting the NASA Ames strategies then pursuing a CFL3D-based approach.[21]

Both the Boeing and Ames strategies relied on solving a nonlinear optimization problem often characterized by a classical Lagrangian function where the goal is to minimize a nonlinear cost function,

$$I(w,u), \tag{1}$$

subject to a nonlinear constraint functional,

$$R(w,u) = 0 . \tag{2}$$

Here, R is the governing flow equation, w is a vector of flow solver unknowns and u is a vector of design parameters. The problem is then formulated as a Lagrangian:

$$F(w,u,\lambda) = I(w,u) + \lambda^T R(w,u) \tag{3}$$

The necessary optimality conditions can then be written as:

$$\begin{aligned}
\frac{\partial F}{\partial u} &= \frac{\partial I}{\partial u} + \lambda^T \frac{\partial R}{\partial u} = 0 \\
\frac{\partial F}{\partial w} &= \frac{\partial I}{\partial w} + \lambda^T \frac{\partial R}{\partial w} = 0 \\
\frac{\partial F}{\partial \lambda} &= R = 0
\end{aligned} \tag{4}$$

By rearranging the above equations we can arrive at the adjoint equation:

$$\lambda = -\left(\frac{\partial R}{\partial w}\right)^{-T} \left(\frac{\partial I}{\partial w}\right)^{T}, \tag{5}$$

which when substituted back into (4) yields:

$$\begin{aligned}
\frac{dF}{du} &= \frac{\partial I}{\partial u} - \left(\frac{\partial I}{\partial w}\right)\left(\frac{\partial R}{\partial w}\right)^{-1}\left(\frac{\partial R}{\partial u}\right) = 0 = \frac{dI}{du} \\
R &= 0
\end{aligned} \tag{6}$$

So the optimality conditions are simply that the governing equations must be satisfied and that the gradient of the cost function with respect to the optimization variables must vanish:

$$\frac{dI}{du} = 0 . \tag{7}$$

The various aerodynamic shape optimization strategies are distinguished by the method that is employed to solve equations (1)-(7). In most cases an iterative procedure is established which progressively drives the gradient,

$$\frac{dI}{du} = \frac{\partial I}{\partial u} - \left(\frac{\partial I}{\partial w}\right)\left(\frac{\partial R}{\partial w}\right)^{-1}\left(\frac{\partial R}{\partial u}\right) \tag{8}$$

to zero. In many cases, the approach must rely on calculating the gradient (or an approximation of it) at each iteration. For example, the original approach employed at NASA Ames[9] is to solve for the value of the gradient by using finite-differences—a discrete shape change is made for each design variable in succession, with the flow being recomputed to obtain the corresponding changes in the cost function. The approach is extremely simple to implement. However, it is computationally expensive, as the number of flow solutions required for each gradient calculation is proportional to the number of design variables. Furthermore, coupled CFD and numerical optimization algorithms that converge to satisfy (4) require multiple iterations and hence multiple gradient vector evaluations. The coupling of the numerical optimization algorithm with the CFD solver is typically achieved by having the optimizer on the outer loop and treating the flow solver as a black box for both function analysis and finite difference gradient evaluations. For typical nonlinear ASO problems, the number of gradient evaluations required is $O(n)$. When this is combined with the $O(n)$ CFD solutions that are needed for each gradient vector, it is apparent that solving the complete ASO problem using finite difference gradients requires $O(n^2)$ CFD solutions.

The direct finite difference approach implemented at NASA Ames also has several other demanding requirements. First, a completely automatic method of mesh generation is needed. For structured meshes on simple wing and wing/body geometries, several iterative implicit (elliptic) solvers that could automate the grid generation were already available.[9] Second, the flow solver must be robust in terms of obtaining solutions for a variety of shapes and achieving a high degree of convergence. Finite difference gradient accuracy demands that flow solutions vary smoothly with the design parameters, which in turn depend upon meshes that vary smoothly with design parameters and thorough convergence of the iterative flow solutions. By 1992 the necessary $O(n^2)$ wing/body Euler CFD solutions, with n in the neighborhood of 50 variables, could be computed with a dedicated 2-6 month effort. The first wind tunnel-verified results from the NASA Ames capability were obtained in 1994[22], and by 1996 several designs had been validated.[23]

Working in parallel and in conjunction with the HSR program, the Boeing team developed its own ASO methodology based upon the TRANAIR CFD code. The Boeing ASO method relied upon a more complex strategy of satisfying (4). Unlike the NASA Ames approach where the numerical optimization algorithm was on the outer loop, the Boeing approach put the optimization into a sub-problem buried within the larger design strategy. The motivation for this approach has two parts. First, as will be discussed, the TRANAIR flow solution strategy offers a natural way of assisting in the solution of the design problem. Second, the advantage of a tighter coupling of the optimization method implies a more computationally efficient design algorithm. TRANAIR solves for w in (2) through a Newton iteration:

$$\frac{\partial R}{\partial w}\delta w = -R \tag{9}$$

The δw solution is used to update w and the process is repeated. Of course, the solution strategy is not that simple. Complicating matters is the fact that TRANAIR uses a solution-adaptive Cartesian mesh. The structure of the Jacobian matrix in (9) as well the number of unknowns changes with each adaptation cycle. Finally, for any three-dimensional problem worth solving, the size of the sparse Jacobian is such that a direct solution of (9) through factorization would be impractical in terms of CPU time and storage. Thus, an iterative method known as GMRES[24] is employed.

Integrating the design problem involves first applying the chain rule to (1):

$$\frac{dI}{du} = \frac{\partial I}{\partial u} + \left(\frac{\partial I}{\partial w}\right)\left(\frac{\partial w}{\partial u}\right). \tag{10}$$

Now, by combining (8) and (10),

$$\left(\frac{\partial R}{\partial w}\right)\left(\frac{\partial w}{\partial u}\right) = -\left(\frac{\partial R}{\partial u}\right). \tag{11}$$

Solving for $(\partial w/\partial u)$ allows substitution back into (10) and direct calculation of the gradient vector. Unlike the corresponding Newton flow solution problem, $(\partial w/\partial u)$ and $(\partial R/\partial u)$ are matrices (i.e,. they have n columns that each have the same length as the number of flow solution unknowns). However, the GMRES iterative flow solution strategy at the core of TRANAIR can be applied to a problem with multiple right-hand sides with significant computational savings over just applying the iterative strategy to each column of (11) in turn (which would be equivalent to calculating the flow solution n times just as in the finite difference approach).

Another aspect of the TRANAIR-based approach is the need for a closed water-tight surface mesh as a starting point for the Cartesian volume mesh process. This surface meshing is not easily automated. Furthermore, any motion of the surface produces a topologically different Cartesian mesh. Dropping or adding computational nodes and changing their connectivity through the flow solution stencil thus changes the number of unknowns and the structure of the Jacobian matrix $(\partial R/\partial w)$. The main difficulty of using TRANAIR in combination with numerical optimization arises from its use of a Cartesian mesh approach. Cartesian meshes simply cannot be moved smoothly to account for surface changes. Thus, perturbing a design variable and reconstructing the surface and volume mesh for each term of a single gradient evaluation is not a viable option—the signal-to-noise ratio tends to be poor because of the discrete, as opposed to smooth, changes in the mesh.

Realizing these mesh related limitations, the group at Boeing chose an alternative approach. Instead of attempting to move the surface and volume meshes directly during optimization, they simulated mesh movement by imposing a transpiration boundary condition (i.e., a boundary condition where mass was permitted to pass through the surface). The transpiration boundary condition has been used to develop an approximation to the right-hand side vectors of (11). Furthermore, the substitution of transpiration in place of surface motion means the flow solutions vary smoothly with the design parameters. Once (11) is solved for $(\partial w/\partial u)$, a substitution back into (10) yields the needed gradient, and the optimization algorithm can proceed. At each iteration of NPSOL, (11) and (10) can be recomputed with updates to both the flow unknowns and the transpiration forcing function. Then, after NPSOL converges on an answer to the approximate Lagrangian problem, the surface is relofted as dictated by the transpiration condition and the solution recomputed. If a significant difference is found between the solution of the approximate problem using transpiration and solution obtained after relofting, the design process could be repeated on the relofted shape.

This Boeing TRANAIR-OPT[25,26] approach capitalized on both the way the flow governing equations were solved and a tighter integration of the numerical optimization algorithm. By 1994, the Boeing group began designing wing/body configurations with TRANAIR-OPT.[27] In 1995, the first wind tunnel models of designs developed using the technique had been built and tested.[23,28]

The Adjoint Approach and Complex Configurations

By the beginning of 1994, the various aerodynamics teams within the HSR program had become convinced of the potential benefits of ASO methods. However, it was also apparent that significant improvements to each of the new ASO techniques were needed. Properly treating constraints, improving the computational efficiency and treating more complex geometries were all deemed vital. The McDonnell Douglas (MDA) team moved on from the initial Ames method and developed their own approaches that pushed the limits of what could be addressed using finite differencing. By 1996, MDA had successfully tested a model after employing techniques that included nacelle/diverter integration.[29-31] Developing a finite difference-based design method accommodating nacelles and diverters proved to be a significant hurdle. The structured-mesh CFD codes were forced to use multiblock topologies with many more grid points, increasing the computational cost of each flow solution by a factor of about 4. Furthermore, the automated methods for mesh generation proved to be very difficult to develop.

Because of the expense of the multiblock flow solutions and the difficulty of multiblock automated mesh generation, the Ames team next developed an alternative way to incorporate nacelle and diverter effects.[32] Instead of directly meshing the nacelles and diverters, their influence was imposed during the single-block CFD calculations in an *a posteriori* manner. Two corrections were made to the wing/body CFD solutions during optimization. First, the interference pressure distribution created by the nacelle/diverter combination (as determined from separate nacelles-on and nacelles-off analyses by another solver) was added to the wing/body solution. Figure 1 shows the lower surface pressures as calculated by AIRPLANE[33] for a complete configuration compared with the same effects interpolated onto the single-block optimization mesh (containing no nacelles or diverters).

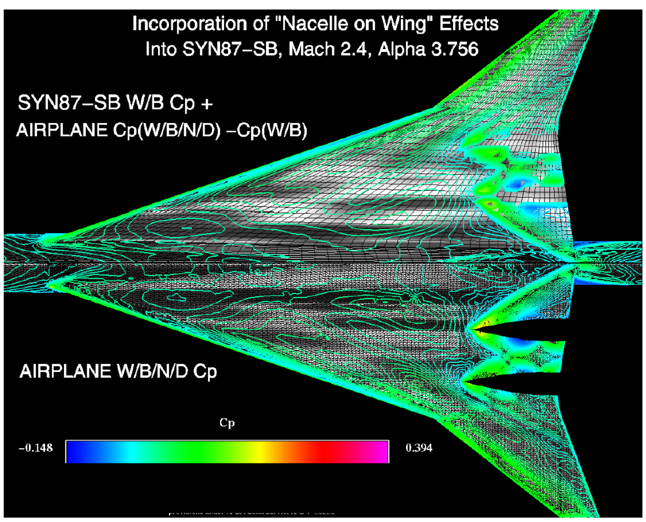

Fig. 1 Lower surface pressure deviation due to nacelles interpolated onto single-block mesh compared with wing/body/nacelle/diverter solution

Second, the buoyancy effects, defined by the changes in the local flow-field about the nacelles due to surface changes in the rest of the configuration, were taken into account. Details of accounting for these two effects are presented in Reference 32. This "Pseudo Nacelle" approach allowed shape optimization with relatively simple wing/body meshes and highly-vectorized single-block flow solutions, resulting in significant conservation of computer resources. Periodically, the input interference pressure distribution was updated as the wing/body shape evolved.

Realizing the limitations in performance of finite difference-based ASO, the Ames team started in 1993 to pursue a novel alternative design strategy based upon work that Jameson had been investigating for some years in two dimensions.[33,34] Instead of using finite differences to calculate the gradients, the new approach employed the "adjoint" equation, as is now outlined. Equations (8) and (5) are rewritten as:

$$\frac{dI}{du} = \frac{\partial I}{\partial u} + \lambda^T \frac{\partial R}{\partial u} \tag{12}$$

and

$$\left(\frac{\partial R}{\partial w}\right)^T \lambda = -\left(\frac{\partial I}{\partial w}\right)^T . \tag{13}$$

Note that solving (13) for λ does not depend in any way on the choice of design parameters, *u*. Also, it turns out that the rest of the right-hand side quantities in (12) can be readily calculated without significant computational expense. It is important to examine the difference between the adjoint equation and the key design equation (11) solved in the Boeing approach. First, the transpose of the Jacobian matrix instead of the Jacobian is needed for (13). Second, (13) requires the solution for only a single column vector, as opposed to the multiple ones needed for (11). These differences between equations (11) and (13) are critical. In particular, forming the transpose of the "true" Jacobian matrix can often prove difficult. Most CFD methods, and especially those involving three dimensions, never form the complete Jacobian matrix. It is simply too complicated and would provide no additional value. While obtaining the true Jacobian would allow one to use (9) to perform a Newton solution to the governing equations, it is usually the case that the cost of performing a true Newton iteration is too expensive in terms of CPU time and memory. Recall that this is the very reason that TRANAIR uses GMRES to solve (9). Unfortunately, the transpose of the Jacobian seen in (13) is no less difficult to manage than the original Jacobian. Thus, even in cases where a Jacobian or its transpose is readily available, the difficulties in solving (13) by direct means has discouraged many investigators.

In fact, the explicit multigrid iterative flow solution methods favored by Jameson[35] did not even construct an approximate Jacobian. Instead of constructing and solving a discrete version of (13), Jameson's original idea was to return to the governing partial differential equations for the flow and the integral form of the cost functional. Then the steps needed to construct (12) and (13) were repeated on these continuous operators through the method of calculus of variations. Once the process had been completed, a differential form of the adjoint equation was obtained. This differential equation could then be discretized and solved using the same strategy as that used for the flow equations. This approach bypassed the need to form the transposed Jacobian matrix and allowed all the carefully developed techniques used to solve the flow equations to be applied to the adjoint equations. By 1995, the Ames team in close collaboration with Jameson and others had managed to implement the adjoint-based ASO technique for the potential equation as well as the Euler equations in two and three dimensions.[36-38]

The new continuous adjoint methods had a tremendous advantage over the finite difference methods. With the solution of one flow and one adjoint system, the gradient with respect to any number of design parameters could be obtained very inexpensively. Suddenly, it was possible to perform optimization studies using hundreds of design variables at acceptable turn around times. Equations (12) and (13) could be solved for the gradient while the same numerical optimization method could be used to drive the overall design towards a minimum. Moreover, the solutions did not have to be so highly converged because they were no longer being differenced.[39]

There were some significant challenges to the continuous adjoint approach. Unlike the extremely simple finite difference methods that had been used up until that point, the adjoint approach demanded not only a significant degree of mathematical understanding of the optimization methods and calculus of variations, but also a thorough understanding of the numerical methods used to discretize and solve the governing equations.

Substantial work was devoted to investigating different discretizations of the adjoint equations.[40] The adjoint boundary conditions proved to be especially challenging. However, with some persistence, even seemingly complicated situations like treating the Pseudo Nacelle method could be modeled as an adjoint boundary condition.[32] Figure 2 shows a cut through the adjoint field solution at a wing station where a nacelle is supposed to be. The adjoint field solution must be thought of as a companion to the flow solution. It has just the same number of unknowns, at the same grid point locations, with roughly the same discrete stencil. However, when compared with the discrete flow equations, the boundary conditions are driven by the form of the cost function instead of, say, the no-penetration condition. Figure 2 illustrates how the adjoint solution includes forcing terms in the volume mesh to account for the presence of the pseudo nacelles. Note the reverse supersonic appearance of the adjoint field solution. The reversed Mach cone lines emanating from a given point, say on the surface, indicate the upstream domain within which a disturbance in the flow may influence the solution, and hence the design objective, at the given point.

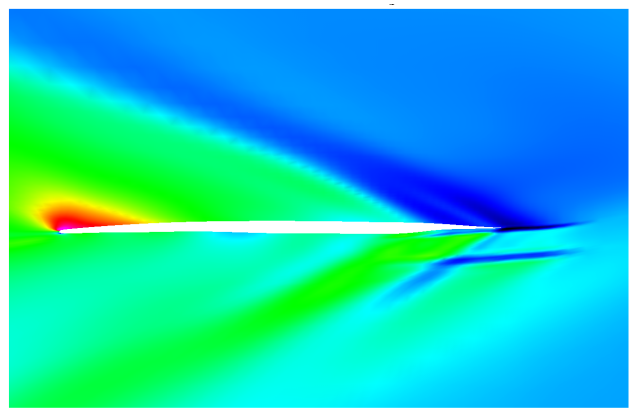

Fig. 2 A cut though an adjoint field solution around a SST wing with pseudo nacelles apparent on the lower surface trailing edge; the flow is left to right

By 1997, the single-block adjoint approach had become the primary method of performing numerical optimization at NASA Ames.[40,41] However, even with this rapid ability to calculate gradients and similar progress toward improving the efficiency of ASO methods at both McDonnell Douglas and Boeing, the detailed aerodynamic design of a given configuration remained a laborious process. It was rarely the case that a single optimization run yielded an acceptable result. Multiple design runs were typically made with different design variables and different objective functions. Many design attempts were quickly discarded because they yielded configurations with undesirable characteristics (e.g., spanwise waviness of the wing) that had not been constrained in the problem formation. The complexity of the design problem combined with the lack of existing comparative designs meant that problem formulation had to be learned as part of the design process itself. All of the different design attempts could be thought of a computational exercise in learning what constitutes a good supersonic aircraft design. Transonic aircraft went through the same exercise via countless wind tunnel tests and many generations of progressively better aircraft culminating in an understanding of desirable pressure distributions. By the use of ASO, the HSR program compressed the design process such that, gradually over the course of a few months, through repeatedly reformulating the design problem and changing the design variables, a new design was obtained that had significantly improved cruise performance.

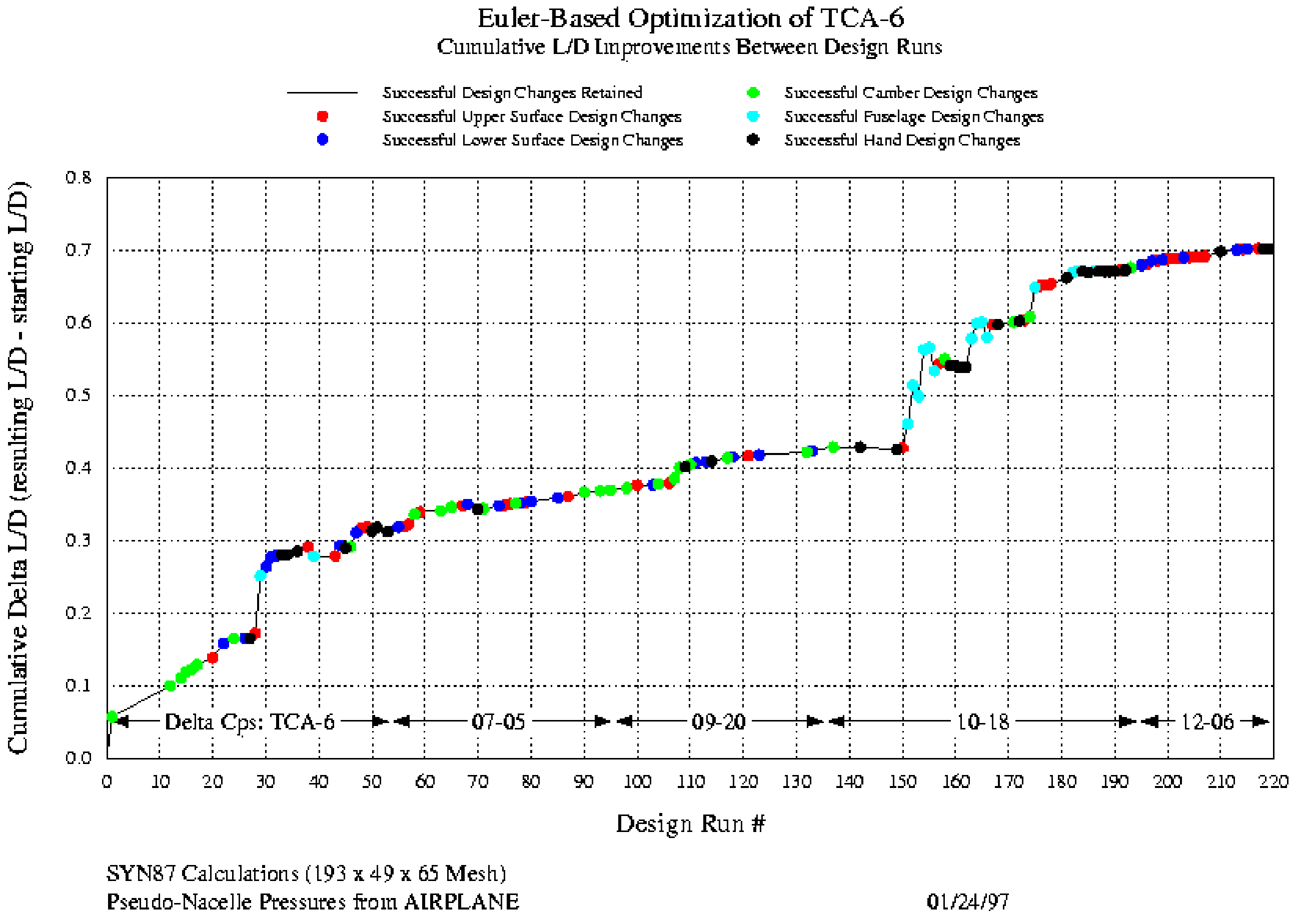

Fig. 3 TCA optimization history using pseudo nacelle method of SYN87-SB, with delta Cp updates indicated

Figure 3 illustrates the extent to which this process was pushed. The Technology Configuration Aircraft (TCA) configuration, which was one of the later HSR study concepts, underwent hundreds of optimization attempts. As the figure shows, progress toward the final design improvement occurred gradually. Periodic updates to the nacelle/diverter interference pressure distribution are indicated by dates above the horizontal axis. Figure 4 depicts some of the changes that resulted in and around the nacelle region as the refinement progressed. Reductions in variation of the pressures impinging onto the lower wing surface correlate with improved aerodynamic efficiency.

By the end of the HSR program, the Ames team had extended the adjoint technique to address complete aircraft configurations via a generalized multiblock implementation.[42,43] The jump to a multiblock grid took some time owing to difficulties that were not related to the adjoint solver. A key issue for applying ASO to complex configurations turns out to be how the aircraft surfaces and volume meshes are moved in response to design changes. Recall that these issues were at the heart of Boeing's decision to resort to transpiration boundary conditions early in the HSR program.

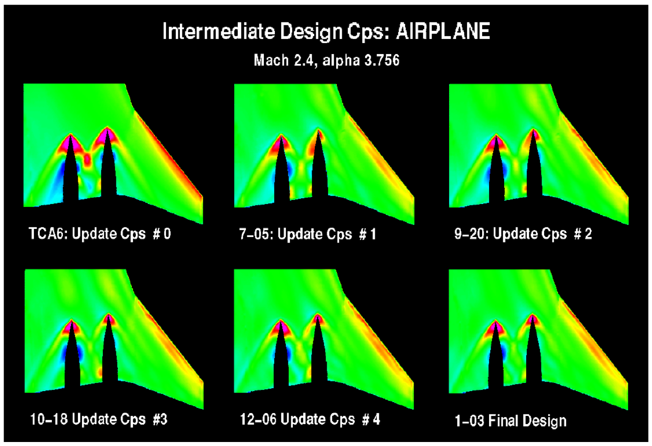

Fig. 4 Lower wing surface pressures from initial, intermediate and final TCA designs achieved with pseudo-nacelle method, calculated by AIRPLANE

Consider the example of applying shape changes to individual aircraft components such as the wing, the fuselage, the nacelles, etc. Inevitably, the shape changes will result in changes that affect the way the components intersect with each other. A multiblock mesh that conforms to the outer mold line (OML) of the aircraft must follow these changes and warp appropriately. Figure 5 illustrates the multiblock mesh constructed for a generic supersonic transport. The black lines denote some of the block boundaries while the colored meshes illustrate the general topology of the meshing. Note the complexity of the mesh, making its movement between design changes a challenge. Between 1997 and 1999, much of the work at Ames was devoted to the implementation of a multiblock mesh perturbation scheme that was completely automated and computationally efficient.[44-48]

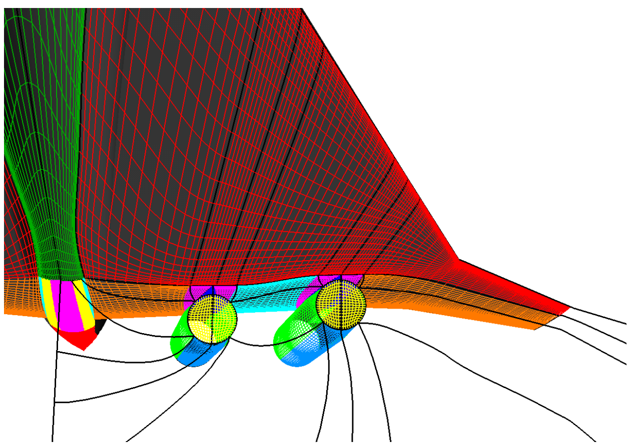

Fig. 5 Multiblock surface grid and some block outlines for a generic supersonic transport

Examining, the simplified equations (1)-(13) shows no dependence on the mesh used to solve the governing equations. In reality, the complete derivation of the adjoint equations results in terms called mesh sensitivities,

$$\frac{\partial X}{\partial u}, \tag{14}$$

where X represents the mesh point locations. Even with the adjoint approach, to eliminate the dependence of the gradient upon each design variable, we must also be able to calculate efficiently (or eliminate) the mesh sensitivities (14). For an approach like TRANAIR where a Cartesian mesh is used, (14) cannot be properly defined and the transpiration boundary condition must be used in its place. In the case of body-fitted structured mesh approaches such as those used at Ames and MDA, it was possible to construct algorithms that would smoothly move the mesh to account for the changing surface shapes. At Ames, the multiblock ASO code included routines that would reintersect components and reconstruct a water-tight structured surface representation of the OML for every design perturbation. A multiblock algebraic perturbation scheme would then adjust the block volume grids for each surface change, preserving as well as possible the qualities of the initial volume grid blocks with no need for further elliptic smoothing. The required mesh sensitivities were thus obtained by finite differencing, where the mesh perturbations were completely explicit. Even for highly complex meshes involving hundreds of blocks and very complicated OML surfacing, these explicit procedures were robust and efficient, requiring only a very small fraction of the total computation time used during optimization.

A more detailed presentation of results obtained with the single-block and multiblock ASO methods at NASA Ames during HSR program is available in Reference 49.

The McDonnell Douglas team not only implemented their own mesh deformation and complex geometry strategies,[50] but also pursued several alternative approaches to calculating gradients more efficiently than with finite differences.[51-56] One of these employed a discrete adjoint technique.[51,53-56] The discrete adjoint differs from the continuous adjoint in that the discrete form of the flow equations (not the continuous differential form) is used to construct the corresponding adjoint equations. In some respects, this has an advantage over the continuous approach in that there are no difficulties in determining the discretization of the adjoint system. However, as discussed before, great care must be taken to develop a solution strategy for (13) that avoids excessive CPU or memory costs. The first implementations of the discrete adjoint did not eliminate these perceived weaknesses.

There are two basic approaches to constructing an efficient discrete adjoint. Either it is either carefully hand-derived[53] or it is developed through the application of automatic differentiation.[57,58] Automatic differentiation is a process by which the actual source code for the state equations is algorithmically (step-by-step) differentiated. The process can be applied to any existing source code the result being new source code that, when compiled and executed, produces sensitivities. Two forms of automatic differentiation have been developed: the forward mode that is analogous to solving equation (11), and the reverse mode that is analogous to solving the adjoint or equation (13). Although the application of automatic differentiation usually results in methods for calculating gradients that are less efficient than a continuous adjoint, it does have the advantage that it is automated. Thus, adjoint solvers can be quickly derived for new CFD methods such as those employing the Navier-Stokes equations with complex turbulence models, but there is a trade-off between ease of implementation and the level of computational efficiency. During the HSR program, it was usually the case that the work involved to develop an analytically differentiated adjoint solver did pay off.

The Boeing ASO approach remained TRANAIR-based throughout the HSR program. The difficult work performed early in the program to develop a design capability using TRANAIR with transpiration boundary conditions paid off. The method was improved at several intervals, but primarily it was simply extended to treat more complex configurations and multipoint design cases[59] with correspondingly more constraints. In particular, wing curvature constraints implemented in TRANAIR-OPT helped control the spanwise waviness that tended to hamper other approaches relying on off-line surface smoothing.

During the life of the HSR program, a concerted effort by all participants produced several very capable and mature ASO methodologies. The developmental support that HSR provided was critical. Furthermore, the program gave the various teams a forum through which they could frequently exchange ideas. The program also presented a constructive way to compete the various ASO approaches. Design test cases were defined that progressively became more challenging in terms of geometric complexity and the number of design constraints that had to be met. Figure 6 shows the baseline TCA configuration that was used as the final test case during 1998. Optimizing the TCA required the integrated treatment of wing, fuselage, canard, horizontal tail, diverters and nacelles. The three teams (Boeing-Long Beach, Boeing-Seattle and NASA Ames) each completed an integrated multipoint design (weighted optimization at three Mach and C_L conditions) involving hundreds of design variables and constraints. Figure 7 depicts some of the fuselage constraints. Figure 8 shows results typical of how the pressures on the lower surface of a wing were altered by ASO. Promising designs obtained from the ASO methods were verified through a battery of wind tunnel tests.

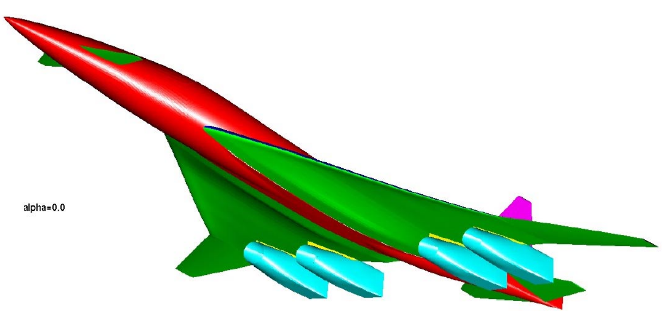

Fig. 6 TCA complete aircraft configuration used for testing ASO methods

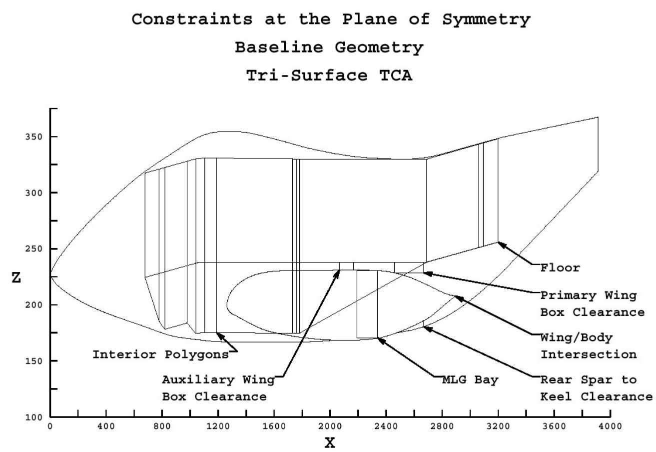

Fig. 7 Some TCA fuselage constraints enforced during optimization

Thus, while the development of key design methods such as the continuous adjoint approach and the transpiration-based TRANAIR approach had started even before the HSR program began, there is no doubt that the requirements and support provided by HSR greatly accelerated their maturation. The series of design competitions in particular, where each of the three teams submitted a candidate configuration, proved to be a highly motivating factor spurring refinement of all of the design methods. Even though the atmosphere was at times combative, the competitions yielded impressive results. By the end of the HSR program, the three design teams had developed ASO methods that could handle a level of complexity previously seen only at the detailed design stage. Furthermore, as illustrated in Figure 9, each team was able to produce designs that improved the *M(L/D)* of the linear-theory baseline by the program goal of 10%.

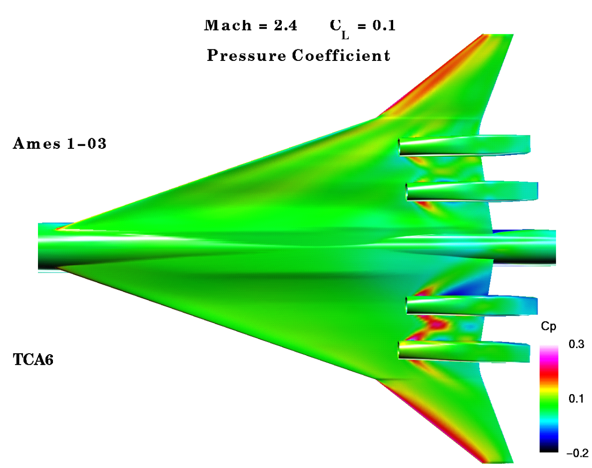

Fig. 8 Comparison of lower surface Cp for baseline and optimized designs

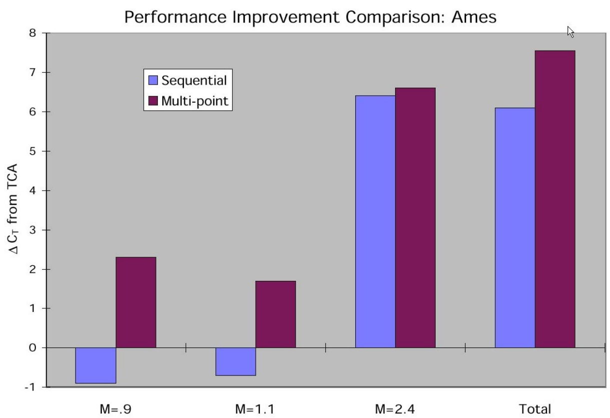

Fig. 9 Improvements achieved at three Mach numbers for two designs: the first design is optimized at M=2.4 only; the second is a multipoint design optimized for all three Mach numbers

Application of ASO After HSR

Following termination of the HSR program, no single forum was available in the United States to foster the continued development of ASO. On the other hand, long before the end of the HSR program, aerodynamic shape optimization had found other customers who valued the capability. Raytheon aircraft used the adjoint approach as early as 1995[60] to design a transonic wing for a new business jet configuration. Figure 10 depicts a before- and after-optimization representation of the Premier I business jet. Note that the upper surface wing shock has been significantly weakened by ASO. In 1996, McDonnell Douglas employed adjoint-based ASO to design a new wing for the planned MD-XX aircraft.[61]

The TRANAIR design code has become an important tool for Boeing's transonic transport design team. It is a small irony that the capabilities that were developed to produce economically competitive supersonic transports are now the same tools being used to enhance their transonic rivals. However, this situation may not continue indefinitely. The desire to make the leap to economical supersonic flight continues to compel new research and development activities. As has already been shown under the HSR program, any new supersonic aircraft will benefit significantly from the latest developments in ASO.

Despite the current situation where a formal program to develop a supersonic transport does not exist, cost-effective supersonic commercial transports are almost certainly inevitable in the long run. If, for example, Boeing's Sonic Cruiser is inherently capable of efficient flight at speeds of up to Mach 1.4, then even if it is initially certified for Mach 0.95, the push for faster travel will eventually lower the barriers to commercial travel at higher speeds. Extending this scenario further, if the Sonic Cruiser does become an efficient supersonic transport, it will not hold the title uniquely for long. Business jet manufacturers and other commercial competitors will surely be obliged to respond to such a challenge.

Given the likelihood that Sonic Cruiser will slip well beyond the introduction of Airbus's A380, another scenario leading to routine supersonic commercial travel would be the development of an efficient supersonic business jet. Successful introduction of such an aircraft would significantly reduce the risk of developing a commercial supersonic counterpart. While the pathway to be followed remains unclear, it is safe to assume that the technological advancement of commercial supersonic travel has not hit a dead end, and that effective application of ASO will be a key to success.

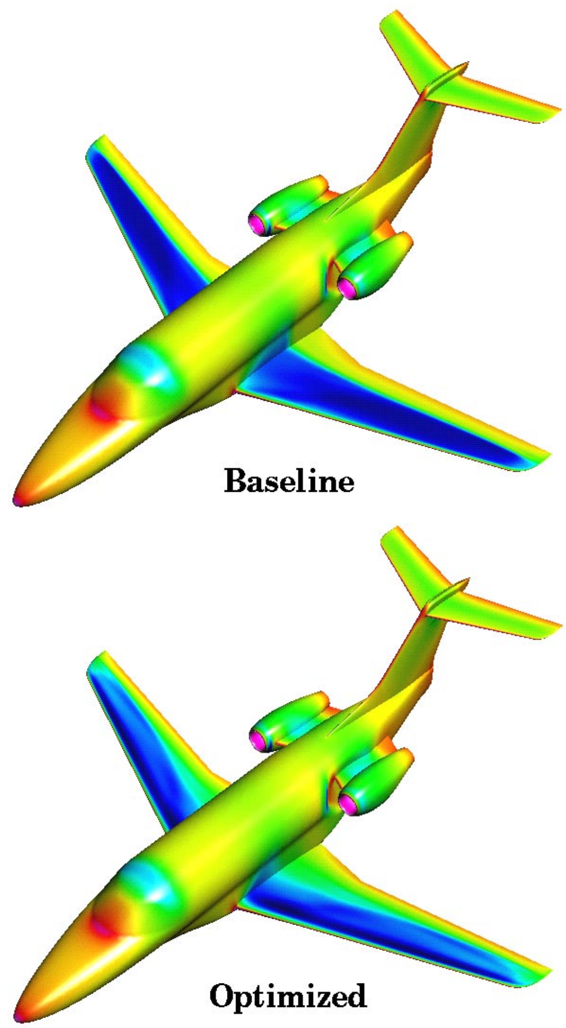

Fig. 10 Comparison between baseline and optimized Premier I business jet representations: note the reduction in wing upper surface shock strength and a general reduction in high Mach number (low pressure) regions; red represents high pressures and blue represents low pressures

New Developments in ASO

While aircraft manufacturers consider their next move toward supersonics, the development of improved ASO methods for supersonic design has not ceased. Rather, important research continues on several fronts.

At NASA Ames, work to enhance design optimization capabilities is proceeding with a new ASO method[62] based upon the unstructured CFD method of AIRPLANE. Compared with the structured mesh CFD methods that were used for ASO during the HSR program, an unstructured mesh approach promises several advantages. Primary among these is the relative ease with which surface and volume meshes can be constructed for complex configurations. The intricate multiblock meshes used during HSR, such as that shown in Figure 5, often took a month or more to prepare. Only once the initial mesh was laboriously completed with interactive grid generation tools could the design optimization process begin. Another advantage of unstructured meshes is the level of complexity that can be treated with these methods. While structured meshes with more than 500 blocks were constructed during the HSR project to cover complete aircraft configurations,[49] these meshes were still limited to cruise point shapes. No ASO for complete high-lift configurations including flap or slat gaps has yet been attempted using structured meshes. Cartesian and overset structured grids are both problematic for gradient-based methods. ASO for high-lift configurations is most likely to benefit from an unstructured implementation.

The switch to unstructured meshes is motivated by hopes of removing the unappealing labor-intensive step of initial grid generation. However, development of a robust and efficient unstructured ASO approach is not trivial. Recall that Boeing resorted to transpiration boundary conditions to build the design version of TRANAIR simply because they were dealing with a Cartesian mesh. Researchers exploring the use of unstructured meshes will have to develop novel mesh movement strategies that not only perform efficiently but also retain the crucial property of varying the mesh smoothly with respect to surface changes.

Meanwhile, the continuous adjoint approach has been extended to perform Navier-Stokes-based design for high-lift systems.[63] The difficulty in accurately analyzing, much less designing, high lift systems using CFD is still severe. Nevertheless, recent results presented in Reference 63 indicate that success in solving these problems is on the horizon.

Other techniques for obtaining gradient information efficiently are also undergoing further development. Taylor, et al. have continued to explore the use of automatic differentiation. A recent paper by this team[64] demonstrates the use of automatic differentiation to construct an "Incremental Iterative" adjoint solver by the application of reverse mode automatic differentiation. The approach is interesting because, like the continuous adjoints, it permits the same iterative technique that is used to converge the state equations to be applied to the solution of the co-state equations. The approach not only significantly reduces the CPU requirements for solving the discrete adjoint but also lowers its memory requirements.

Research is also under way to develop discrete sensitivity analysis-based adjoint solvers that do not rely on automatic differentiation. These methods hope to retain the efficiency of continuous adjoint strategies while improving their gradient accuracy by guaranteeing discrete consistency between the state and co-state systems. Recent work, by Nielsen and Anderson, has shown one approach to constructing a discrete adjoint solver and an iterative solver for two- and three-dimensional design using the Navier-Stokes equations on unstructured meshes.[65] They also show methods of treating unstructured mesh motion that are smooth and robust. Hand-derived discrete adjoint solvers have also recently been implemented by Nemec and Zingg,[66] as well as Kim, et al.[67]

Aero-Structural Optimization

Perhaps more interesting than these aerodynamics-only developments of ASO are other recent efforts to extend ASO to address MDO problems. A noteworthy MDO example is the aero-structural optimization work that has recently been pursued by two independent teams.[68,69] These methods solve a coupled adjoint problem for high fidelity computational aerodynamics and structures. As before, an adjoint is used to calculate gradients and the procedure is linked to a numerical optimization algorithm to drive the design process. The new methods overcome an inherent limitation in all the approaches that were applied during HSR, namely their use of fixed geometric constraints. These constraints were imposed to allow a feasible structural design to be developed after the final OML was obtained via ASO. In reality, an optimal aero-structural design would trade thickness and hence aerodynamic efficiency against structural weight. For complex nonlinear problems such as HSCT or supersonic business jet design, the tradeoff between structural weight and aerodynamic efficiency can be quite difficult to determine. The answer to this problem is to solve a tightly coupled aero-structural problem where both the OML and the underlying structure are designed together.

The first step is to develop a coupled aero-structural solver. This coupling is challenging because forces and deflections must be transferred conservatively between the CFD and CSM (computational structural mechanics) modules. Issues such as how often to update quantities through the coupling operator must also be considered.

Once a robust coupled aero-structural solver is developed, a design procedure using the solver must also be developed.

Consider a problem characterized by minimizing

$$I(w,q,u), \tag{15}$$

subject to a nonlinear constraint functionals,

$$R_A(w,q,u) = 0, \tag{16}$$

and

$$R_S(w,q,u) = 0, \tag{17}$$

where R_A is the aerodynamic governing equations, R_S is the structural dynamics governing equations, w is the aerodynamics solution vector, q is the structural dynamics solution vector and u is the design parameters vector. The Lagrangian for the problem becomes

$$\begin{aligned} F(w,q,u,\lambda_A,\lambda_S) &= I(w,q,u) + \\ &\lambda_A{}^T R_A(w,q,u) + \lambda_S{}^T R_S(w,q,u). \end{aligned} \tag{18}$$

The gradient expression can be written as

$$\frac{\partial I}{\partial w} + \lambda_A{}^T \frac{\partial R_A}{\partial w} + \lambda_S{}^T \frac{\partial R_S}{\partial w} = 0, \tag{19}$$

and the adjoint expressions become

$$\begin{aligned} \frac{\partial I}{\partial w} + \lambda_A{}^T \frac{\partial R_A}{\partial w} + \lambda_S{}^T \frac{\partial R_S}{\partial w} = 0 \\ \frac{\partial I}{\partial q} + \lambda_A{}^T \frac{\partial R_A}{\partial q} + \lambda_S{}^T \frac{\partial R_S}{\partial q} = 0. \end{aligned} \tag{20}$$

The two adjoint equations can be rewritten as follows:

$$\left(\frac{\partial R_A}{\partial w}\right)^T \lambda_A = -\left(\frac{\partial I}{\partial w}\right)^T - \left(\frac{\partial R_S}{\partial w}\right)^T \overline{\lambda_S} \tag{21}$$

and

$$\left(\frac{\partial R_S}{\partial q}\right)^T \lambda_S = -\left(\frac{\partial I}{\partial q}\right)^T - \left(\frac{\partial R_A}{\partial q}\right)^T \overline{\lambda_A}. \tag{22}$$

Rearrangement of the adjoint expressions exposes a key fact, namely that except for the final term on the right-hand side of each equation, the rest of the equations are nothing more than the original aerodynamic and structural adjoints independent of any coupling. Therefore, by lagging the final terms in (21) and (22), it is possible to use existing adjoint solvers for the flow and structures. It also turns out that construction of the coupling terms on the right-hand sides of (21) and (22) falls out directly from the implementation of the coupled state solution method discussed earlier. Thus, the same lagging procedure that is used to transfer information between the CFD and CSM solvers is used again to transfer information between their corresponding adjoint solvers. The preliminary results presented in Reference 53 are quite encouraging. Figure 12 below depicts a supersonic business jet wing body design which has been optimized by the coupled aero-structural approach. Note that the baseline side view has no fuselage camber while the optimized design displays a considerable amount—shifting some of the lift burden from the wing to the fuselage. The colors indicate pressures, with red indicating high and blue indicating low. The net reduction in low pressures on the wing upper surface translates into lower wave drag.

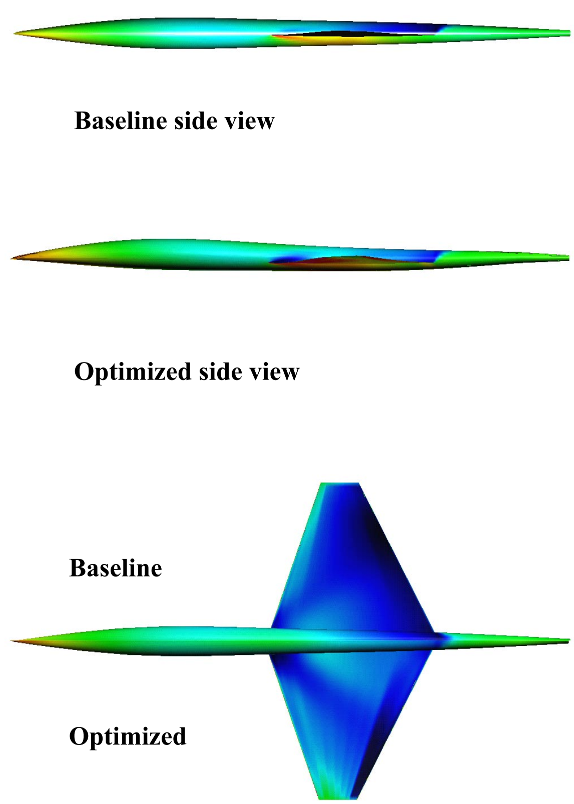

Fig. 12 Comparison of baseline and aero-structural optimization for a supersonic business jet wing/body

Laminar Flow Wings

Another supersonic design idea advocated by Manning and Kroo[70] is drag reduction through the use of the natural laminar flow wing concept. The idea of obtaining natural laminar flow at supersonic speeds is not new. Early in the HSR program, studies were carried out to determine the extent to which laminar flow could be obtained on highly swept wings provided that the pressure distributions are correctly tailored.[71] The HSR program eventually abandoned the concept because it seemed that, for large commercial transports, the wing Reynolds numbers were simply too high for highly swept wings to allow a significant extent of laminar flow. Nevertheless, the Japanese Supersonic Research program has continued to pursue the idea,[72-75] and is currently in the process of testing an unpiloted sub-scale research aircraft[76] to demonstrate the concept. The Japanese have designed their research demonstrator via ASO methods. However, instead of formulating the problem by maximizing *M(L/D)* as was done for much of the HSR program, the Japanese have gone back to the inverse problem of designing to a target pressure distribution. In the case of achieving a high degree of laminar flow for a swept wing, there is only one pressure distribution type that suffices. Their designs only hope to achieve laminar flow over a portion of the upper wing surface. The technique avoids the difficulty in problem formulation that slowed progress during the HSR program. However, it is not clear that the reduction in viscous drag without considering the minimization of wave drag results in optimal designs. Finally, translating the improvements that can be achieved on a sub-scale experimental aircraft to a full-scale commercial transport will be difficult. The significantly larger Reynolds numbers seen on a full-sized aircraft imply that the transition to turbulent flow will occur much earlier.

Manning and Kroo, by contrast, have advocated the use of a sharp leading edge and low wing sweep. Furthermore, they also suggest that the appropriate application would be a business jet and not a commercial transport. The combination of low wing sweep and lower wing Reynolds numbers presents a more attainable degree of laminar flow. In the future, ASO methods will likely be employed to carefully tailor the wing contours of a vehicle with planform similar to that depicted in Figure 12. Trades will have to be made among laminar flow, compressibility drag and structural weight.

Reduced Sonic Boom

As in the case of natural laminar flow, some early efforts within HSR sought to develop configurations with reduced sonic boom signatures. The presence of sonic booms that impact the ground along the ground track of aircraft traveling at supersonic speeds has always made the adoption of widespread commercial high-speed travel problematic. Even for cases where supersonic flight is restricted to flights over water or otherwise unpopulated regions, the environmental impact of these sonic booms is not well understood. Sonic boom studies carried out over the last few decades have determined that a large commercial supersonic transport, designed without sonic boom considerations, would produce prohibitively strong booms, precluding the option of supersonic flight over land. Furthermore, even with vigorous design efforts to moderate the boom, it is doubtful that a large commercial transport could achieve acceptable perceived ground noise levels in the near future. Around 1994, faced with seemingly insurmountable challenges, the HSR program dropped its attempts to reduce sonic boom. The technology for a second-generation SST became focused on efficient travel over water. Like Concorde, a successful aircraft would have flown every flight with loud booms along its over-water routes.

Recently, through work carried out under the DARPA QSP program, researchers at Stanford[1] and elsewhere have been reviewing the possibilities of reduced-boom supersonic travel. The goal of the project is to develop conceptual aircraft designs in the 100,000 lb class, flying at Mach 2.4 with a range of 6,000 nautical miles and producing an initial overpressure below 0.3 psf. The requirements are suited to a supersonic business jet configuration, and because sonic boom design becomes easier as weight and volume are reduced for the same flight altitude, designing a low-boom business jet will be far easier than designing a 750,000 lb, 300-passenger aircraft.

The Stanford team is taking a two-level design strategy for the project. At the top level, a fast conceptual analysis and optimization system has been assembled to examine the MDO trade space. This MDO tool combines a linear theory drag estimation method with classical boom estimating methods and a genetic optimization algorithm to select interesting concepts quickly. The boom estimation is accomplished via equivalent area ruling combined with propagation using Whitham F-function[77] or Thomas equivalent wave form parameter methods.[78] Once configurations of particular interest have been identified, higher-fidelity nonlinear design methods that employ gradient-based optimization are used to refine the design. Nonlinear sonic boom analysis is performed by using an Euler equation CFD method in the near field (up to about 1 body length) and extrapolation to calculate the ground pressure distribution. The corresponding design method again employs an adjoint solver for the Euler-based flow analysis while the extrapolation part of the analysis can be solved either through finite differences or a special adjoint for the propagation method. Details of the Euler adjoint constructed for sonic boom objective functions, which are similar to those used to simulate the pseudo nacelles during the HSR program, are given in Reference 79.

It is too early to arrive at definitive conclusions, but from the early work in laminar flow wings and low-boom design it appears that designing an efficient, practical and more environmentally palatable supersonic business jet may provide the most realistic incremental step towards a commercial SST. Perhaps it will eventually take an MDO method that brings together laminar flow, low-boom and combined aero-structural optimization to develop a convincing preliminary design of a supersonic business jet.

Conclusions

The clear conclusion that is apparent from work in supersonic design over the last decade is that ASO is an essential element to developing a future commercial capability. The relationship between ASO and supersonic design is one of mutual benefit. Throughout the HSR program, the aerodynamic efficiency gains achieved with ASO were among its greatest success stories. Despite the fact that the program was terminated before any commitment to construct an aircraft, the improvements in aerodynamic efficiency realized by ASO during HSR showed that the aerodynamics goals of the program were attainable. From a design technology development point of view, HSR must be reflected upon as an enormously successful effort. In particular, the adjoint-based and transpiration-based aerodynamic optimization capabilities that have become part of the mainstream of aircraft design owe considerable thanks to HSR. The future of ASO methods seems to be focused on extensions to treat MDO problems. Here again, the pursuit of a supersonic commercial transport, whether through Boeing's Sonic Cruiser or a supersonic business jet, will be instrumental in guiding and challenging the development of high fidelity optimization capabilities. Along the way, many other aircraft design programs will reap the benefits.

References

1. J. R. Wilson "The New Shape of Supersonics," June 2002, Aerospace America.

2. P. Poisson-Quinton and G. Ville, "Speed for Intercontinental Air Transport: From Transonic to Supersonic," Lettre de l'Academie Nationale de l'Air et de l'Espace, (A.N.A.E.), No. 29, Nov. 2001.

3. W. D. Middleton and H. W. Carlson, "Numerical Method of Estimating and Optimizing Supersonic Aerodynamic Characteristics of Arbitrary Planform Wings," Journal of Aircraft, Vol. 2, No. 4, Aug. 1965, pp. 261-265.

4. D. S. Miller, H. W. Carlson, and W. D. Middleton, "A Linearized Theory Method of Constrained Optimization for Supersonic Cruise Wing Design," Proc. of the SCAR Conf., NASA CP-001, Nov. 1976.

5. H. W. Carlson and R. J. Mack, "Estimation of Leading-Edge Thrust for Supersonic Wings of Arbitrary Planform," NASA TP 1270, Oct. 1978.

6. W. D. Middleton, J. L. Lundry, and R. G. Coleman, "A System for Aerodynamic Design and Analysis of Supersonic Aircraft," Dec. 1980, Part 1 - General Description and Theoretical Development, NASA CR-3351.

7. W. D. Middleton, J. L. Lundry, and R. G. Coleman, "A System for Aerodynamic Design and Analysis of Supersonic Aircraft," Dec. 1980, Part 2 - User's Manual, NASA CR-3352.

8. W. D. Middleton, J. L. Lundry, and R. G. Coleman, "A System for Aerodynamic Design and Analysis of Supersonic Aircraft," Dec. 1980, Part 3 - Computer Program Description, NASA CR-3353.

9. W. D. Middleton, J. L. Lundry, and R. G. Coleman, "A System for Aerodynamic Design and Analysis of Supersonic Aircraft," Dec. 1980, Part 4 - Test Cases, NASA CR-3354.

10. R. M. Kulfan, "Projecting and Tracking Advanced Technology Improvement in L/D," July 1996, First NASA/Industry High-Speed Research Configuration Aerodynamics Workshop, Pt. 2, pp. 781-844, also NASA/CP-1999-209690/PT2.

11. M. B. Giles and M. Drela, "Two-Dimensional Transonic Aerodynamic Design Method," Journal of Aircraft, Vol. 25, No. 9, 1987.

12. R. M. Hicks, G. N. Vanderplaats and E. M. Murman, "Airfoil Section Drag Reduction at Transonic Speeds by Numerical Optimization," Society of Automotive Engineers – Business Aircraft Meeting, Paper 760477, 1976.

13. R. M. Hicks and P. A. Henne, "Wing Design by Numerical Optimization," Journal of Aircraft, Vol. 15, No. 7, 1978.

14. R. A. Kennelly, "Improved method for Transonic Airfoil Design-by-Optimization," AIAA Paper 1983-1864.

15. R. M. Hicks and P. A. Henne, "Wing Design by Numerical Optimization," Journal of Aircraft, 15:407-412, 1978.

16. A. Jameson and T. J. Baker, "Solution of the Euler Equations for Complex Configurations," July 1983, AIAA paper 83-1929, AIAA 6th Computational Fluid Dynamics Conference, Danvers, MA.

17. J. Reuther, S. Cliff, R. Hicks, C. P. VanDam, "Practical Design Optimization of Wing/Body Configurations Using the Euler Equations," AIAA Paper 1992-2633, 1992.

18. D. P. Young, R. G. Melvin, M. B. Bieterman, F. T. Johnson, S. S. Samant and J. E. Bussoletti, "A Locally Refined Rectangular Grid Finite Element Method: Application to Computational Fluid Dynamics and Computational Physics," Journal of Computational Physics, 1991, pp. 1-66.

19. P. E. Gill, W. Murray, M. A. Saunders and M. A. Wright, "User's Guide for NPSOL (Version 4.0) A FORTRAN Package Nonlinear Programming, Stanford University Technical Report SOL86-2, Department of Operations Research, 1986.

20. R. M. Wood and S. X. S. Bauer, "The Natural Flow-Design Concept", NASA TP 3193, 1992.

21. R.T. Biedron, "CFL3D Version 6 Home Page," http://cfl3d.larc.nasa.gov/Cfl3dv6/cfl3dv6.html.

22. R. Hicks, S. Agrawal, D. L. Antani, S. Cliff and J. Reuther, "An Experimental/Computational Evaluation of an Arrow-Wing High Speed Civil Transport Designed Using Numerical Optimization and the Euler Equations," April 1996, NASA CDTM-21005.

23. S. Cliff, T. Baker, R. Hicks, J. Reuther, "A Computational/Experimental Study of Two Optimized Supersonic Transport Designs and the Reference H Baseline," July. 1996, First NASA/Industry High-Speed Research Configuration Aerodynamics Workshop, Pt. 3, pp. 845-967, also NASA/CP-1999-209690/PT3.

24. Y. Saad and M. H. Schultz, "GMRES: A Generalized Minimal Residual Method for Solving Non-Symmetric Linear Systems," SIAM Journal on Scientific and Statistical Computing, 7 (1986), pp. 856-869.

25. D. P. Young, W. P. Huffman, R. G. Melvin, M. B. Bieterman, C. L. Hilmes and F. T. Johnson, "Inexactness and Global Convergence in Design Optimization," AIAA Paper 1994-4386, 1994.

26. R. S. Conner, "Overview of HSR Aerodynamic Optimization at Boeing," July. 1996, First NASA/Industry High-Speed Research Configuration Aerodynamics Workshop, Pt. 2, pp. 395-423, also NASA/CP-1999-209690/PT2.

27. K. R. Wittenberg, "Boeing HSR Wing Optimization Using TRANAIR," July. 1996, First NASA/Industry High-Speed Research Configuration Aerodynamics Workshop, Pt. 2, pp. 445-492, also NASA/CP-1999-209690/PT2.

28. S. Yaghmaee and K. M. Mejia, "Update to the Summary of Langley Unitary Test 1649 and Its Implications on Validity of Viscous and Inviscid Analyses," July. 1996, First NASA/Industry High-Speed Research Configuration Aerodynamics Workshop, Pt. 3, pp. 969-1008, also NASA/CP-1999-209690/PT3.

29. P. M. Hartwich, E. R. Unger, A. E. Arslan, and S. Agrawal, "Some Recent Enhancements to Aerodynamic Shape Optimization Methods at McDonnell Douglas Aerospace," July. 1996, First NASA/Industry High-Speed Research Configuration Aerodynamics Workshop, Pt. 2, pp. 423-444, also NASA/CP-1999-209690/PT2.

30. E. R. Unger, J. O. Hager, S. Agrawal, "Supersonic Aerodynamic Design Improvements of an Arrow-Wing HSCT Configuration Using Nonlinear Point Design Methods," July. 1996, First NASA/Industry High-Speed Research Configuration Aerodynamics Workshop, Pt. 3, pp. 1009-1040, also NASA/CP-1999-209690/PT3.

31. R. P. Narducci, P. Sundaram, S. Agrawal, S. Cheung, A. E. Arslan, G. L. Martin, "Experimental Investigation of a Point Design Optimized Arrow-Wing HSCT," July. 1996, First NASA/Industry High-Speed Research Configuration Aerodynamics Workshop, Pt. 3, pp. 1041-1072, also NASA/CP-1999-209690/PT3.

32. J. Reuther and D. Saunders, "Advances in Design Optimization Using Adjoint Methods," July. 1996, First NASA/Industry High-Speed Research Configuration Aerodynamics Workshop, Pt. 2, pp. 493-528, also NASA/CP-1999-209690/PT2.

33. A. Jameson, "Automatic Design via Control Theory," Journal of Scientific Computing, 3:233-260, 1988.

34. A. Jameson, "Automatic Design of Transonic Airfoils to Reduce the Shock Induced Pressure Drag," Proceeding of the 31st Israel Annual Conference on Aviation and Aeronautics, Tel Aviv, pp. 5-17, Feb. 1990.

35. A. Jameson, W. Schmidt and E. Turkel, "Numerical Solution of the Euler Equations by Finite Volume Methods with Runge-Kutta Time Stepping Schemes," AIAA paper 1981-1259.

36. J. Reuther and A. Jameson, "Control theory Based Airfoil Design for Potential Flow and a Finite Volume Discretization," AIAA paper 1994-0499.

37. A. Jameson and J. Reuther, "Control Theory Based Airfoil Design Using the Euler Equations," AIAA paper 1994-4272.

38. J. Reuther and A. Jameson, "Aerodynamic Shape Optimization of Wing and Wing-Body Configurations Using Control Theory," AIAA paper 1995-0123.

39. J. Reuther, "Aerodynamic Shape Optimization Using Control Theory," May 1996, *Ph.D. Dissertation*, University of California Davis, CA, also NASA-CR-201064.

40. J. Reuther, D. Saunders and R. Hicks, "Improvements to the Single-Block Adjoint-Based Aerodynamic Shape Design Method: SYN87-SB," 1997 NASA High-Speed Research Program Aerodynamic Performance Workshop, Vol. 1, Part 2, pp. 1199-1256.

41. S. Cliff, J. Reuther and R. Hicks, "Ames Optimized TCA Configuration," 1997 NASA High-Speed Research Program Aerodynamic Performance Workshop, Vol. 1, Part 2, pp. 1257-1347.

42. J Reuther and M. Rimlinger, "Future Advanced in Aerodynamic Shape Optimiation," 1997 NASA High-Speed Research Program Aerodynamic Performance Workshop, Vol. 1, Part 2, pp. 1415-1451.

43. J. Reuther and M. Rimlinger, "Development and Validation of a Multi-Block Adjoint Design Method," 1997 NASA High-Speed Research Program Aerodynamic Performance Workshop, Vol. 1, Part 2, pp. 1348-1414.

44. J. Reuther, M. Rimlinger, D. Saunders and R. Hicks, "SYN107-MB Aerodynamic Shape Optimization Method: Recent Improvements and Current Status," 1998 NASA High-Speed Research Program Aerodynamic Performance Workshop, Vol. 1, Part 1, pp. 693-775.

45. R. Hicks, M. Rimlinger and J. Reuther, "TCA-6 Configuration Optimization," 1998 NASA High-Speed Research Program Aerodynamic Performance Workshop, Vol. 1, Part 1, pp. 837-927.S.

46. Cliff, J. Reuther, D. Saunders and M. Rimlinger, "Nacelle/Diverter Integration into the Design Optimization Process using Pseudo, Warped and Real Nacelles," 1999 NASA High-Speed Research Program Aerodynamic Performance Workshop, Vol. 1, Part 1, pp. 685-746.

47. D. Saunders and J. Reuther, "Automated Euler/Navier-Stokes Grid Generation/Grid Perturbation for Wing/Body Configurations," 1999 NASA High-Speed Research Program Aerodynamic Performance Workshop, Vol. 1, Part 1, pp. 641-684.

48. J. Reuther, M. Rimlinger and D. Saunders, "Geometry Driven Mesh Deformation," 1999 NASA High-Speed Research Program Aerodynamic Performance Workshop, Vol. 1, Part 2, pp. 867-900.

49. S. Cliff, J. Reuther, D. Saunders and R. Hicks, "Single-Point and Multipoint Aerodynamic Shape Optimization of High-Speed Civil Transport," *Journal of Aircraft*, Vol. 38, No. 6, Nov.-Dec. 2001, pp. 997-1005.

50. E. Unger, R. Narducci, J. Hager, P. Hartwich, R. Mendoza and G. Kuruvila, "The AEROSHOP (AEROdynamic Shape Optimization) Toolkit," 1998 NASA High-Speed Research Program Aerodynamic Performance Workshop, Vol. 1, Part 1, pp. 777-836.

51. G. Kuruvila, J. Hager and P. Sundaram, "Aerodynamic Gradients Using Three Methods," 1998 NASA High-Speed Research Program Aerodynamic Performance Workshop, Vol. 1, Part 2, pp. 931-978.

52. R. Narducci and S. Agrawal, "Progress Towards a Multipoint Optimization Procedure," 1998 NASA High-Speed Research Program Aerodynamic Performance Workshop, Vol. 1, Part 2, pp. 1071-1138.

53. G. Kuruvila and R. Narducci, "Design Cycle-Time Reduction Using TLNS3D-Adjoint," 1999 NASA High-Speed Research Program Aerodynamic Performance Workshop, Vol. 1, Part 2, pp. 995-1048.

54. E. Unger, R. Narducci, J. Hager, G. Kuruvila, P. Hartwich and S. Agrawal, "Supersonic Cruise Point Optimization of TCA," 1997 NASA High-Speed Research Program Aerodynamic Performance Workshop, Vol. 1, Part 1, pp. 114-188.

55. J. Hager, P. Hartwich, E. Unger, R. Narducci, G. Kuruvila and S. Agrawal, "Improvements to MDC Nonlinear Aerodynamic Design Tools," 1997 NASA High-Speed Research Program Aerodynamic Performance Workshop, Vol. 1, Part 1, pp. 189-254.

56. R. Narducci, J. Hager, E. Unger, G. Kuruvila, P. Sundaram, P. Hartwich, M. Grant, R. Mendoza, A. Arlson and S. Agrawal, "Technology Development for a Multipoint Optimization Process for HSCT," 1999 NASA High-Speed Research Program Aerodynamic Performance Workshop, Vol. 1, Part 2, pp. 1111-1162.

57. A. Carle and M. Fagan, "Overview of ADIFOR 3.0," Department of Computational and Applied Mathematics, Rice University, CAAM-TR 00-02, Jan. 2000.

58. P. Sundaram and S. Agrawal, "Progress Towards Viscous Design Optimization Using Automatic Differentiation," 1999 NASA High-Speed Research Program Aerodynamic Performance Workshop, Vol. 1, Part 2, pp. 1049-1110.

59. R. Conner, "BCAG Design Optimization Activities," 1998 NASA High-Speed Research Program Aerodynamic Performance Workshop, Vol. 1, Part 1, pp. 625-692.

60. J. Gallman, J. Reuther, N. Pfeiffer, W. Forrest and D. Bernstorf, "Business Jet Design Using Aerodynamic Shape Optimization," AIAA paper 1996-0554.

61. A. Jameson, J. Alonso, J. Reuther, L. Martinelli, and J. Vassberg, "Aerodynamic Shape Optimization techniques Based Upon Control Theory," AIAA paper 1998-2538.

62. S. Cliff, S. Thomas, T. Baker, A. Jameson and R. Hicks, "Aerodynamic Shape Optimization Using Unstructured Grid Methods," AIAA paper 2002-5550.

63. S. Kim, J. Alonso, A. Jameson, "Design Optimization of High-Lift Configurations Using a Viscous Continuous Adjoint," AIAA paper 2002-0844.

64. A. Taylor, L. Green, P. Newman and M. Putko, "Some Advanced Concepts in Discrete Aerodynamic Sensitivity Analysis," AIAA paper 2001-2529.

65. E. J. Nielson and W. K. Anderson, "Recent Improvements in Aerodynamic Design Optimization on Unstructured Meshes," *AIAA Journal*, Vol. 40, No. 6, pp. 1155-1163, June 2002.

66. M. Nemec and D. W. Zingg, "Newton-Krylov Algorithm for Aerodynamic Design Using the Navier-Stokes Equations," *AIAA Journal*, Vol. 40, No. 6, pp. 1146-1154, June 2002.

67. C. Kim, S. Kim and O. H. Rho, "Sensitivity Analysis for the Navier-Stokes Equations with Two-Equations Turbulence Models," AIAA Journal, Vol. 39, No. 5, 2001, pp 838-845.

68. J. R. R. A. Martins, J. J. Alonso and J. Reuther, "High Fidelity Aero-Structural Design Optimization of Supersonic Business Jet," AIAA paper 2002-1483.

69. K. Maute, M. Nikbay and C. Farhat, "Coupled Analytical Sensitivity Analysis and Optimization of Three-Dimensional Nonlinear Aeroelastic Systems," *AIAA Journal*, 39(11):2051-2061, Nov. 2001.

70. V. Manning and I. Kroo, "Multidisciplinary Optimization of a Natural Laminar Flow Supersonic Aircraft," AIAA Paper 1999-2102.

71. B. Bharadvaj, M. Fisher, R. Joslin, L. King and P. Parikh, "Supersonic Laminar Flow Control, An Overview," July 1996, First NASA/Industry High-Speed Research Configuration Aerodynamics Workshop, Pt. 1, pp. 81-98, also NASA/CP-1999-209690/PT1.

72. K. Matsushima, T. Iwamiya and W. Zhang, "Inverse Wing Design for the Scaled Supersonic Experimental Airplane with Ensuring Design Constraints," Jan. 2000, 2nd International Workshop on Numerical Simulation Technology for Design of Next Generation Supersonic Civil Transport, Tokyo, Japan.

73. T. Iwamiya, K. Yoshida, Y. Shimbo, Y. Makino and K. Matsushima, "Aerodynamic Design of Supersonic Experimental Airplane," Jan. 2000, 2nd International Workshop on Numerical Simulation Technology for Design of Next Generation Supersonic Civil Transport, Tokyo, Japan.

74. Y. Yoshikazu, S. Jeong, K. Suzuki, Y. Ueda, K. Matsushima and K. Yoshida, "Optimized Design," Dec. 2001, 3rd International Workshop on Numerical Simulation Technology for Design of Next Generation Supersonic Civil Transport, Tokyo, Japan.

75. D. Sasaki and S. Obayashi, "Automated Aerodynamic Design Optimization System for SST Wing-Body Configuration," Dec. 2001, 3rd International Workshop on Numerical Simulation Technology for Design of Next Generation Supersonic Civil Transport, Tokyo, Japan.

76. K. Sakata, "Supersonic Experimental Airplane (NEXST) for Next Generation SST Technology-Development and Flight Test Plan for the Unmanned Scaled Supersonic Glider," Jan. 2002, AIAA paper 2002-0527.

77. G. B. Whitham, "The flow pattern of a supersonic projectile," 1952, Commun. Pure Applied Math, Vol. 3, pp. 301-348.

78. C. Thomas, "Extrapolation of Sonic Boom Pressure Signatures by the Wave Form Parameter Method," 1972, NASA TN D-6832.

79. S. K. Nadarajah, A. Jameson and J. J. Alonso, "An Adjoint Method for the Calculation of Remote Sensitivities in Supersonic Flow," Jan. 2002, AIAA paper 2002-0261.