NOTES ON NUMERICAL FLUID MECHANICS
VOLUME 78, pp 135-178

COMPUTATIONAL FLUID DYNAMICS FOR THE 21ST CENTURY

Simulation Based Aerodynamic Design

A. Jameson, L. Martinelli, J. Alonso, J. Vassberg and J. Reuther

Proceedings of a Symposium Honoring Prof. Satofuka on the Occasion of his 60th Birthday

Kyoto, Japan

July 15 - 17, 2000

Simulation Based Aerodynamic Design

A. Jameson
Dept of Mechanical and Aerospace Engineering
Stanford University
Stanford, California 94305 USA
650-725-6802
jameson@baboon.stanford.edu

L. Martinelli
Dept of Mechanical and Aerospace Engineering
Princeton University
Princeton N.J. 08544 U.S.A.
609-258-6652
gigi@princeton.edu

J.J. Alonso
Stanford University
Stanford, California 94305 USA

J. C. Vassberg
Boeing Commercial Airplane Group
Long Beach, CA 90846

J. Reuther
NASA Ames Research Center
Moffet Field, CA 96035

Abstract– This paper reviews the status of advanced computational simulation techniques in the aerodynamic design of modern aircraft. An outline of the aircraft design process is provided, and the most relevant trade-offs between disciplines are presented to justify the leading role that aerodynamic design plays in this truly multidisciplinary process. A control theory based adjoint approach which considerably reduces the computational cost of the calculation of design sensitivities is presented. Our experience with this method is described with the help of several computational design examples that cover a substantial range of objective functions, computational models, and geometric complexity. Using the adjoint method, entire design optimization calculations can be completed with a computational cost equivalent to only a few (typically less than ten) analysis runs. In comparison with current practice in an industrial setting, the design approach presented in this paper can yield large computational savings as well as reduced turn around times that can be used to either decrease the time to market or to increase the number of design iterations within a given time frame.

TABLE OF CONTENTS

1. INTRODUCTION

The past 25 years have seen a revolution in the entire engineering design process as computational simulation has come to play an increasingly dominant role. Nowadays engineers spend most of their time at workstations rather than at their drafting tables.

Most notably, computer aided design (CAD) methods have essentially replaced the drawing board as the basic tool for definition and control of the configuration. Software systems such as CATIA and Unigraphics provide solid modeling capabilities that enable designers to prepare complex layouts

without the need to build mockups. Computer visualization techniques enable the designer to verify that there will be no interference between different parts in the layout, and greatly facilitate decisions on the routing of all the electrical wiring and hydraulic piping.

Similarly, structural analysis is now almost entirely carried out by computational methods typically based on the finite element method. Commercially available software systems such as NASTRAN, ANSYS, or ELFINI have been progressively developed and augmented by new features, and can treat the full range of requirements for aeronautical structures, including analysis of stressed skin structures into the nonlinear range. They are also very carefully validated before each new release against a comprehensive suite of test cases, and engineers can place complete confidence in the outcome. Accordingly, the structural design is routinely committed on the basis of computational analysis, while structural testing is limited to the role of verification that the design truly meets its specified requirements of ultimate strength and fatigue life.

Computational simulation of fluid flow has not yet reached the same level of maturity. While commercial software for the simulation of fluid flow is offered by numerous vendors, aircraft companies continue to make substantial investments in the in-house development of their own methods, such as Boeing's TRANAIR, or Lockheed's TEAM programs. At the same time there are major ongoing efforts to develop the science of computational fluid dynamics (CFD) in government research agencies such as NASA, Japan's ARL, or in Europe, France's ONERA, Germany's DLR, Holland's NLR and Sweden's FFA, all of which are a source of industrially used computer programs. This reflects the fact that fluid flow is generally more complex and harder to predict than the behavior of structures. The complexity and range of phenomena of fluid flow is well illustrated in Van Dyke's Album of Fluid Motion [1].

The concept of a numerical wind tunnel, which might eventually allow computers "to supplant wind tunnels in the aerodynamic design and testing process", was already a topic of discussion in the decade of 1970-1980. In their celebrated paper of 1975 [2], Chapman, Mark and Pirtle listed three main objectives of computational aerodynamics:

1. To provide flow simulations that are either impractical or impossible to obtain in wind tunnels or other ground based experimental test facilities.

2. To lower the time and cost required to obtain aerodynamic flow simulations necessary for the design of new aerospace vehicles.

3. Eventually, to provide more accurate simulations of flight aerodynamics than wind tunnels can.

Chapman, Mark, and Pirtle also noted that the inherent limitations of computational and wind tunnel simulations are complementary. Wind tunnels are limited by the size of the models that can be placed in them, and by the density, temperature and velocity of the flow that they can sustain, with the consequence that flight-Reynolds numbers cannot be realized with complete models. Their accuracy is also limited by wall and support interference, and by aeroelastic distortion. Computers are not limited in any of these ways, but they are limited in speed and memory, which in turn limit the attainable complexity and resolution of the simulations.

Computational Fluid Dynamics (CFD) has now matured to the point at which it is widely accepted as a key tool for aerodynamic design. Algorithms have been the subject of intensive development for the past two decades. The principles underlying the design and implementation of robust schemes which can accurately resolve shock waves and contact discontinuities in compressible flows are now quite well established. It is also quite well understood how to design high order schemes for viscous flow, including compact schemes and spectral methods. Adaptive refinement of the mesh interval (h) and the order of approximations (p) has been successfully exploited both separately and in combination in the h-p method [3].

Despite these advances, CFD is still not being exploited as effectively as one would like in the design process. This is partially due to the long set-up times and high costs, both human and computational, of complex flow simulations. A continuing obstacle to the treatment of configurations with complex geometry has been the problem of mesh generation. Several general techniques have been developed, including algebraic transformations and methods based on the solution of elliptic and hyperbolic equations. In the last few years methods using unstructured meshes have also begun to gain more general acceptance.

The fidelity of mathematical modeling of high Reynolds number flows continues to be limited by computational costs. Thus, accurate and cost-effective simulation of viscous flow at high Reynolds numbers associated with full scale flight remains a challenge. Several routes are available toward the reduction of computational costs, including the reduction of mesh requirements by the use of higher order schemes, improved convergence to steady state by sophisticated acceleration methods, and the exploitation of massively parallel computers. With the present state of the art, however, it is still cheaper to obtain massive quantities of data (such as the loads data across the flight envelope) by wind tunnel testing, since the incremental cost of obtaining additional data is very small once a wind tunnel model has been built. Using computational simulation, the cumulative cost of generating data for the full flight envelope becomes very large since a separate run is required for each data point. In addition, as the flight conditions vary substantially, the mesh

may even need to be regenerated or adaptively refined in order to accurately capture the changing flow physics present in different areas of the aircraft envelope. Computational simulation has the key advantage, on the other hand, that it allows the rapid exploration of numerous alternative designs. Thus CFD and wind tunnel testing can be effectively used in complementary roles, with CFD the prime tool for the initial design studies, and wind tunnel testing the prime tool for final verification of the design concept and acquisition of the full aerodynamic data required for completion of the detailed design.

This paper examines ways to exploit computational simulation more effectively in the overall design process, with the primary focus on aerodynamic design, while recognizing that this should be part of an integrated multi-disciplinary process. The design process itself is surveyed in the next section. The following two sections examine the industrial requirements for effective and trustworthy CFD software, and the way in which optimization techniques can be integrated with CFD. Section 5 discusses recent industrial experiences in the application of CFD and optimization techniques to a major project for a commercial aircraft. Finally Section 6 discusses ways in which the design process might be re-engineered to exploit computational simulation more effectively.

2. THE DESIGN PROCESS

The design process can generally be divided into three phases: conceptual design, preliminary design, and final detailed design, as illustrated in Figure 1. The conceptual design stage defines the mission in the light of anticipated market requirements, and determines a general preliminary configuration capable of performing this mission, together with first estimates of size, weight and performance. In the preliminary design stage the aerodynamic shape and structural skeleton progress to the point where detailed performance estimates can be made and guaranteed to potential customers, which can then, in turn, formally sign binding contracts for the purchase of a certain number of aircraft. At this stage the development costs are still fairly moderate, in the range of 50 - 100 million dollars. In the final design stage the structure must be defined in complete detail, together with complete systems, including the flight deck, control systems (involving major software development for fly-by-wire systems), electrical and hydraulic systems, landing gear, weapon systems for military aircraft, and cabin layout for commercial aircraft. Major costs are incurred at this stage, during which it is also necessary to prepare a detailed manufacturing plan, together with appropriate facilities and tooling. The development costs to reach the point of initial production are in the range of 3 - 10 billion dollars. Thus, the final design would normally be carried out only if sufficient orders have already been received to indicate a reasonably high probability of at least achieving a significant portion of return on the investment.

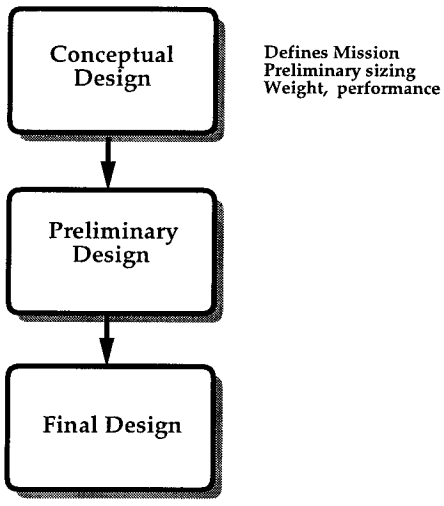

Figure 1: The Overall Design Process

Figure 2 provides a closer look at the conceptual design stage. In the case of commercial aircraft the mission is defined on the basis of airline requirements. Desired payload-range characteristics follow from route analysis between representative city pairs such as Los Angeles - Tokyo, including data on expected traffic volume, desired frequency, and prevailing weather patterns. At the same time it is necessary to consider issues of airport compatibility, including constraints on gate size and noise regulations, as well as the growth possibilities of the aircraft and the existing and upcoming competition in similar markets. A preliminary synthesis using simplified aerodynamic and structural models and statistical data bases provides an initial configuration and sizing, together with performance estimates, taking into account requirements for stability and control. Software for aircraft synthesis such as NASA Ames' ACSYNT program is available to assist in this process. For commercial aircraft it is necessary to estimate both the operating cost and the cost of ownership, while for military aircraft the life time cycle cost may be a determining factor. In either case it is generally assumed that the selling price is likely to be proportional to the gross weight of the aircraft.

The result of the initial synthesis may confirm the feasibility of the proposed mission. On the other hand it may suggest that it is too ambitious, requiring an excessively large and expensive aircraft, or, alternatively, that a more testing mission could be accomplished with an aircraft of acceptable size. Thus the process will generally be iterated until it arrives at a mission and corresponding design that can be

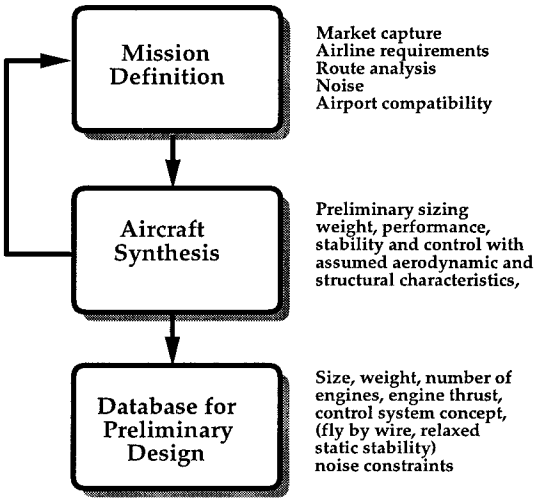

Figure 2: The Conceptual Design Process

expected to attain the desired market capture and return on investment. Concurrently discussions will proceed both with potential customers to verify market interest, and with major vendors such as the engine manufacturers to assure the availability of appropriate power plants and systems. These discussions may well lead to further iteration of the mission and design concept in an ongoing process. Vendors may also be approached to share in the development costs as risk sharing partners, or to undertake substantial development costs of their own to provide components which meet the design requirements.

In the development of commercial aircraft, aerodynamic design plays a leading role in the preliminary design stage. The definition of the external aerodynamic shape may actually be finalized in the preliminary design. The aerodynamic lines of the Boeing 777 were frozen, for example, when initial orders were accepted before the initiation of the detailed design of the structure. Figure 3 illustrates the way in which the aerodynamic design process is embedded in the overall preliminary design. The starting point is an initial CAD definition resulting from the conceptual design. The inner loop of aerodynamic analysis is contained in an outer multidisciplinary loop, which is in turn contained in a major design cycle involving wind tunnel testing. In recent Boeing practice three major design cycles, each requiring about 4 - 6 months, are used to finalize the wing design. Improvements in CFD which would allow the elimination of a major cycle would significantly shorten the overall design process and therefore reduce costs. In the development of the MDXX, McDonnell Douglas planned to rely on high level CFD together with the experimental database which had been developed for the MD12 and expected to eliminate the need for a sequence of major design cycles.

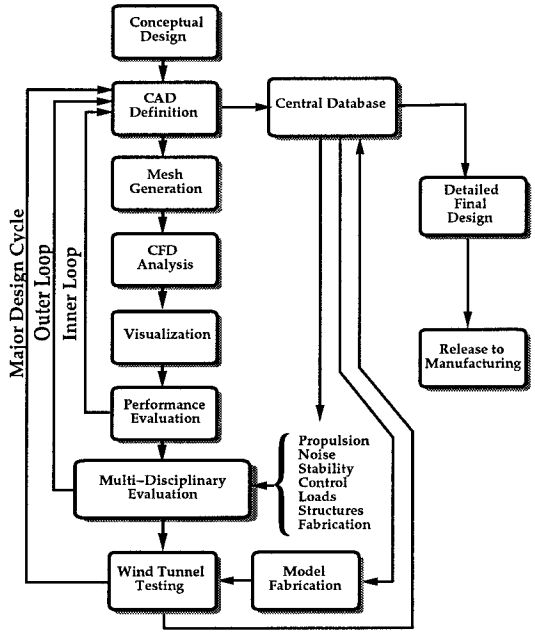

Figure 3: The Aerodynamic Design Process

The inner aerodynamic design loop is used to evaluate numerous variations in the wing definition. In each iteration it is necessary to generate a mesh for the new configuration prior to performing the CFD analysis. Computer graphics software is then used to visualize the results, and the performance is evaluated. The first studies may be confined to partial configurations such as wing-body or wing-body-nacelle combinations. At this stage the focus is on the design of the clean wing. Key points of the flight envelope include the nominal cruise point, cruise at high lift and low lift to allow for the weight variation between the initial and final cruise as the fuel is burnt off, and a long range cruise point at lower Mach number, where it is important to make sure there is no significant drag creep. Other defining points are the climb condition, which requires a good lift to drag ratio at low Mach number and high lift coefficient with a clean wing, and the buffet condition. This is typically taken as the high lift cruise point increased to a load of 1.3 g to allow for maneuvering and gust loads. Both wing section modifications such as the thickness to chord ratio and camber distributions, and planform variations such as the sweepback angle or aspect ratio may be considered. While the detailed design of the high lift system and control surfaces may be deferred to a later stage, the planform must provide the necessary space

for both high lift systems and control surfaces outside the main structural box, and it must also accommodate the landing gear. This generally requires an extension of the in board trailing edge to form an area known as a "yehudi".

The aerodynamic analysis interacts with the other disciplines in the next outer loop. These disciplines have their own inner loops, not shown in Figure 3. For an efficient design process the fully updated aero-design data base must be accessible to other disciplines without loss of information. For example, the thrust requirements in the power plant design will depend on the drag estimates for take-off, climb and cruise. In order to meet airport noise constraints a rapid climb may be required while the thrust may also be limited. Initial estimates of the lift and moments allow preliminary sizing of the horizontal and vertical tail. This interacts with the design of the control system, where the use of a fly-by-wire system may allow relaxed static stability and tail surfaces of reduced size.

In fact, the interaction between disciplines, as well as the effect of disciplinary constraints on the other participating disciplines is quite strong. For example, in the case of the interaction between the aerodynamics and the structures, it often pays off to consider both disciplines jointly rather than as separate modules. This interaction between disciplines is not reserved only to the analysis of the coupled system, but should extend into the computation of coupled sensitivities to be used during the design process. As this multidisciplinary process becomes more closely integrated, the second loop in Figure 3 will slowly disappear. Aero-structural design is only an example of the interaction between disciplines. Additional disciplines may have as strong an impact on the design of the coupled system and will have to be considered in a way which will be much different from sequential interaction. In fact, this may be also true in the design of systems different from aircraft (spacecraft, automobiles, ships, chip manufacturing and layout, etc.).

With this in mind, first estimates of the aerodynamic loads allow the design of an initial structural skeleton, which in turn provides an estimate of the structure weight. One of the main trade-offs is between aerodynamic performance and wing structure weight. The requirement for fuel volume may also be an important consideration. An increase in the thickness to chord ratio both increases fuel volume, and allows the same bending moment to be carried with reduced skin thickness, with an accompanying reduction in weight. On the other hand it will lead to a decrease in the drag rise Mach number. The induced drag, which typically contributes around 40 percent of the cruising drag, varies inversely as the square of the span. Thus a 5 percent increase in the wing span could produce a total drag reduction of the order of 4 percent, but would lead to an increase in wing weight because of the increase in the root bending moment. The wing span may in fact be limited by airport gate constraints.

The taper ratio and span load distribution also affect the trade-off between aerodynamic performance and wing weight. While an elliptic span load distribution minimizes the induced drag for a given span, a more triangular load distribution reduces the root bending moment. A large root chord may be dictated by the need to accommodate the landing gear and flaps, but it also has the advantage of increasing the root thickness for a fixed thickness to chord ratio, yielding a weight reduction. For example, the root chord of the MDXX was increased at a late stage in the design to accommodate larger flaps, and this contributed a significant weight reduction. In order to maintain a moderately efficient span load distribution with a highly tapered planform the outboard wing must operate with higher local section lift coefficient than the inboard wing. This can have an adverse effect on the behavior near buffet, as the outboard wing will incur a shock stall before the inboard wing, leading to a reduction of lift behind the center of gravity, and consequently a high speed pitch-up. This is unacceptable for certification if it is too severe.

An increase in the wing sweepback angle may be used to increase the drag rise Mach number. Alternatively it allows an increase in the thickness to chord ratio for the same drag rise Mach number, with a resulting weight reduction. This is partially offset by the increase in the length of the wing. Moreover, an increase in the sweep back angle will aggravate the problem of high speed pitch-up. Most modern highly loaded wings have sweep back angles no greater than 35 degrees at the $\frac{1}{4}$ chord line.

Manufacturing constraints must also be considered in the final definition of the aerodynamic shape. For example, the curvature changes in the span wise direction must be limited. This avoids the need for shot peaning which might otherwise be required to force curvature in both the span wise and chord wise directions.

From the complexity of these trade-offs it can be seen that a crucial requirement for aerodynamic analysis is to make trustworthy predictions with fast enough turn around so as not to delay the outer multidisciplinary cycle. In order to allow the completion of the major design cycle in 4 - 6 months, the cycle time for the multidisciplinary loop should not be greater than about 2 weeks. Considering the need to examine the performance of design variations at all the key points of the flight envelope, this implies the need to turn around aerodynamic analyses in a few hours. The computational costs are also important because the cumulative costs of large numbers of calculations can become a limiting factor.

It is also evident that the number of possible design variations is too large to permit their exhaustive evaluation, and thus it is very unlikely that a truly optimum solution can be found without the assistance of automatic optimization procedures. Ultimately there is a need for multi-disciplinary

optimization (MDO), but this can only be effective if it is based on sufficiently high fidelity modeling of the separate disciplines. As a step in this direction there could be significant pay-offs from the application of optimization techniques within the disciplines, where the interactions with other disciplines are taken into account through the introduction of constraints. For example the wing drag can be minimized at a given Mach number and lift coefficient with a fixed planform, and constraints on minimum thickness to meet requirements for fuel volume and structure weight.

3. INDUSTRIAL CFD

In order to carry out the inner loop of the aerodynamic design process the main requirements for effective CFD software are:

1. Sufficient and known level of accuracy
2. Acceptable computational and manpower costs
3. Fast turn around time

Performance estimation in the cruise condition is crucial to the design of transport aircraft. The error should be in the range of $\pm\frac{1}{2}$ percent. The drag coefficient of a long range transport aircraft such as the Boeing 747 is in the range of .0275 (275 counts), depending on the lift coefficient, which is in approximately .5. The drag coefficient of proposed supersonic transport designs is in the range of .0120 to .0150 at much lower lift coefficients in the range of .1 - .12. Thus one should aim to predict drag with an accuracy of the order of $\pm$ .0001 ($\pm$1 count). Manufacturers have to guarantee performance, and errors can be very expensive through the costs of redesign, penalty payments and lost orders.

A first consideration is the choice of appropriate mathematical models of fluid flow which are adequate for trustworthy flow predictions. Many critical phenomena of fluid flow, such as shock waves and turbulence, are essentially non-linear. They also exhibit extreme disparities of scales. While the actual thickness of a shock wave is of the order of a mean free path of the gas particles, on a macroscopic scale its thickness is essentially zero. In turbulent flow energy is transferred from large scale motions to progressively smaller eddies until the scale becomes so small that the motion is dissipated by viscosity. The ratio of the length scale of the global flow to that of the smallest persisting eddies is of the order $\mathrm{Re}^{\frac{3}{4}}$, where Re is the Reynolds number, typically in the range of 30 - 50 million for an aircraft. In order to resolve such scales in all three space directions a computational grid with the order of $\mathrm{Re}^{\frac{9}{4}}$ cells would be required. This is beyond the range of any current or foreseeable computer. Consequently mathematical models with varying degrees of simplification have to be introduced in order to make computational simulation of flow flow feasible and produce viable and cost-effective methods.

Figure 4 (supplied by Pradeep Raj) indicates a hierarchy of models at different levels of simplification which have proved useful in practice. Efficient flight is generally achieved by the use of smooth and streamlined shapes which avoid flow separation and minimize viscous effects, with the consequence that useful predictions can be made using inviscid models. Inviscid calculations with boundary layer corrections can provide quite accurate predictions of lift and drag when the flow remains attached, but iteration between the inviscid outer solution and the inner boundary layer solution becomes increasingly difficult with the onset of separation. Procedures for solving the full viscous equations are likely to be needed for the simulation of arbitrary complex separated flows, which may occur at high angles of attack or with bluff bodies. In order to treat flows at high Reynolds numbers, one is generally forced to estimate turbulent effects by Reynolds averaging of the fluctuating components. This requires the introduction of a turbulence model. As the available computing power increases one may also aspire to large eddy simulation (LES) in which the larger scale eddies are directly calculated, while the influence of turbulence at scales smaller than the mesh interval is represented by a subgrid scale model.

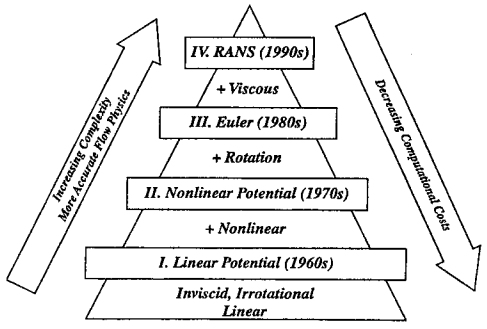

Figure 4: Hierarchy of Fluid Flow Models

Computational costs vary drastically with the choice of mathematical model. Panel methods can be effectively used to solve the linear potential flow equation with higher-end personal computers. Studies of the dependency of the result on mesh refinement, performed by this author and others, have demonstrated that inviscid transonic potential flow or Euler solutions for an airfoil can be accurately calculated on a mesh with 160 cells around the section, and 32 cells normal to the section. Using multigrid techniques 10 to 25 cycles are enough to obtain a converged result. Consequently airfoil calculations can be performed in seconds on Pentium-class personal computers. Correspondingly accurate three-dimensional inviscid calculations can be per-

formed for a wing on a mesh, say with $192\times32\times48 = 294,912$ cells, in minutes on a single processor Cray T90 or in one and a half hours on a workstation.

Viscous simulations at high Reynolds numbers require vastly greater resources. Careful two-dimensional studies of mesh requirements have been carried out at Princeton by Martinelli [4]. He found that on the order of 32 mesh intervals were needed to resolve a turbulent boundary layer, in addition to 32 intervals between the boundary layer and the far field, leading to a total of 64 intervals. In order to prevent degradations in accuracy and convergence due to excessively large aspect ratios (in excess of 1,000) in the surface mesh cells, the chordwise resolution must also be increased to 512 intervals. Reasonably accurate solutions can be obtained in a 512×64 mesh in 100 multigrid cycles. Translated to three dimensions, this would imply the need for meshes with 5–10 million cells (for example, $512\times64\times256 = 8{,}388{,}608$ cells as shown in Figure 5). When simulations are performed on less fine meshes with, say, 500,000 to 1 million cells, it is very hard to avoid mesh dependency in the solutions as well as sensitivity to the turbulence model.

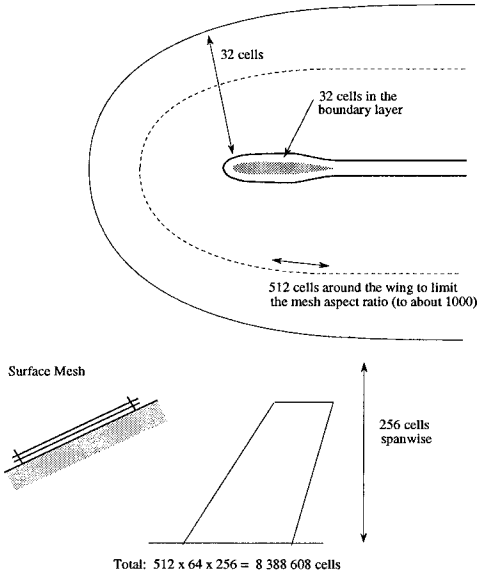

Figure 5: Mesh Requirements for a Viscous Simulation

A typical algorithm requires of the order of 5,000 floating point operations per mesh point in one multigrid iteration. With 10 million mesh points, the operation count is of the order of 0.5×10^{11} per cycle. Given a computer capable of sustaining 10^{11} operations per second (100 gigaflops), 200 cycles could then be performed in 100 seconds. Simulations of unsteady viscous flows (flutter, buffet) would be likely to require 1,000–10,000 time steps. A further progression to large eddy simulation of complex configurations would require even greater resources. The following estimate is due to W.H. Jou [5]. Suppose that a conservative estimate of the size of eddies in a boundary layer that ought to be resolved is $1/5$ of the boundary layer thickness. Assuming that 10 points are needed to resolve a single eddy, the mesh interval should then be $1/50$ of the boundary layer thickness. Moreover, since the eddies are three-dimensional, the same mesh interval should be used in all three directions. Now, if the boundary layer thickness is of the order of 0.01 of the chord length, 5,000 intervals will be needed in the chordwise direction, and for a wing with an aspect ratio of 10, 50,000 intervals will be needed in the spanwise direction. Thus, of the order of $50 \times 5,000 \times 50,000$ or 12.5 billion mesh points would be needed in the boundary layer. If the time dependent behavior of the eddies is to be fully resolved using time steps on the order of the time for a wave to pass through a mesh interval, and one allows for a total time equal to the time required for waves to travel three times the length of the chord, of the order of 15,000 time steps would be needed. Performance beyond the teraflop (10^{12} operations per second) will be needed to attempt calculations of this nature, which also have an information content far beyond what is needed for engineering analysis and design. However, the designer does not need to know the details of the eddies in the boundary layer. The primary purpose of such calculations is to improve the calculation of averaged quantities such as skin friction, and the prediction of global behavior such as the onset of separation. The current use of Navier-Stokes and large eddy simulations is to try to gain an improved insight into the physics of turbulent flow, which may in turn lead to the development of more comprehensive and reliable turbulence models.

Turbulence Modeling

It is doubtful whether a universally valid turbulence model, capable of describing all complex flows, could be devised [6]. Reference [7] provides a detailed assessment of the current status of turbulence modeling in aeronautical applications. Algebraic models [8, 9] have proved fairly satisfactory for the calculation of attached and slightly separated wing flows. These models rely on the boundary layer concept, usually incorporating separate formulas for the inner and outer layers, and they require an estimate of a length scale which depends on the thickness of the boundary layer. The estimation of this quantity by a search for a maximum of the vorticity times a distance to the wall, as in the Baldwin-Lomax model, can lead to ambiguities in internal flows, and also in complex vortical flows over slender bodies and highly swept or delta wings [10, 11]. The Johnson-King model [12], which allows for non-equilibrium effects through the introduction of an ordinary differential equation for the maximum shear stress, has improved the prediction of flows with shock induced separation [13, 14].

Closure models depending on the solution of transport equations are widely accepted for industrial applications. These models eliminate the need to estimate a length scale by detecting the edge of the boundary layer. Eddy viscosity models typically use two equations for the turbulent kinetic energy k and the dissipation rate ϵ, or a pair of equivalent quantities [15, 16, 17, 18, 19, 20, 21]. Models of this type generally tend to present difficulties in the region very close to the wall. They also tend to be badly conditioned for numerical solution. The $k-\omega$ model has so far been the most consistently accurate model from this family. The $k-l$ model [22] is designed to alleviate this problem by taking advantage of the linear behavior of the length scale l near the wall. Similarly, the $k-\tau$ model was designed to alleviate the numerical stiffness problem that the $k-\omega$ model has exhibited near the wall [7]. In an alternative approach to the design of models which are more amenable to numerical solution, new models requiring the solution of one transport equation have recently been introduced [23, 24]. In particular, the Spalart-Allmaras one-equation turbulence model has gained wide acceptance in the aerospace community and is currently used with most of the existing flow solvers. The performance of the algebraic models remains competitive for wing flows, but the one- and two-equation models show promise for broader classes of flows. In order to achieve greater universality, research is also being pursued on more complex Reynolds stress transport models, which require the solution of a larger number of transport equations.

The selection of sufficiently accurate mathematical models and a judgment of their cost effectiveness ultimately rests with industry. As the design progresses through the three phases of conceptual design, preliminary design, and detailed design. The appropriate CFD models will vary in complexity. In the conceptual and preliminary design phases, the emphasis will be on relatively simple models which can give results with very rapid turn-around and low computer costs, in order to evaluate alternative configurations and perform quick parametric studies. The detailed design stage requires the most complete simulation that can be achieved with acceptable cost.

Algorithms and Mesh Generation

The computational simulation of fluid flow presents a number of severe challenges for algorithm design. At the level of inviscid modeling, the inherent nonlinearity of the fluid flow equations leads to the formation of singularities such as shock waves and contact discontinuities. Moreover, the geometric configurations of interest are extremely complex, and often contain sharp edges which lead to the shedding of vortex sheets. Extreme gradients near stagnation points or wing tips may also lead to numerical errors that can have global influence. Numerically generated entropy may be convected from the leading edge for example, causing the formation of a numerically induced boundary layer which can lead to separation. The need to treat exterior domains of infinite extent is also a source of difficulty. Boundary conditions imposed at artificial outer boundaries may cause reflected waves which significantly interfere with the flow. When viscous effects are also included in the simulation, the extreme difference of the scales in the viscous boundary layer and the outer flow, which is essentially inviscid, is another source of difficulty, forcing the use of meshes with extreme variations in mesh interval. For these reasons CFD, has been a driving force for the development of numerical algorithms.

An essential requirement for industrial CFD is the capability to treat extremely complex geometric configurations. A key choice that must be made is the nature of the mesh used to divide the flow field into discrete subdomains. The discretization procedure must allow for the treatment of complex configurations. The principal alternatives are Cartesian meshes, body-fitted curvilinear meshes, and unstructured tetrahedral/hybrid meshes. Each of these approaches has advantages which have led to their use. The Cartesian mesh minimizes the complexity of the algorithm at interior points and facilitates the use of high order discretization procedures, at the expense of greater complexity, and possibly a loss of accuracy, in the treatment of boundary conditions at curved surfaces. This difficulty may be alleviated by using mesh refinement procedures near the surface. With their aid, schemes which use Cartesian meshes have recently been developed to treat very complex configurations [25, 26, 27, 28].

Body-fitted meshes have been widely used and are particularly well suited to the treatment of viscous flow because they readily allow the mesh to be compressed near the body surface. With this approach, the problem of mesh generation itself has proved to be a major pacing item. In order to treat very complex configurations it generally proves expedient to use a multiblock [29, 30] procedure, with separately generated meshes in each block, which may then be patched at block faces, or allowed to overlap, as in the Chimera scheme [31, 32]. While a number of interactive software systems for grid generation have been developed, such as EAGLE, GRIDGEN, GRAPE, and ICEM, the generation of a satisfactory grid for a very complex configuration may require man-months of effort. However, new techniques which alleviate the grid-generation bottleneck are continually being developed. For example, Vassberg has recently introduced the concept of Globally Elliptic Meshing (GEM) [33] which automates the generation of a multi-block field mesh for semi-complex configurations.

The alternative is to use an unstructured mesh in which the domain is subdivided into tetrahedra. This in turn requires the development of solution algorithms capable of yielding the required accuracy on unstructured meshes. This approach has been gaining acceptance, as it is becoming apparent that it can lead to a speed-up and reduction in the

cost of mesh generation that more than offsets the increased complexity and cost of the flow simulations. Two competing procedures for generating triangulations which have both proved successful are Delaunay triangulation [34, 35], based on concepts introduced at the beginning of the century by Voronoi [36], and the moving front method [37].

For a detailed review of CFD algorithms in current use the reader is referred to reference [38]. Another key issue is the validation of CFD software for industrial use. For a better understanding of this issue it is important to distinguish the different sources of error. These include modeling errors because the mathematical model does not adequately represent the true physics of the flow, numerical errors and programming errors. Numerical errors include discretization errors, and errors in the numerical solution of the discrete model, if for example, an iterative procedure is not fully converged. The asymptotic behavior of discretization errors may be estimated by numerical analysis, and their magnitude in practice can be estimated by mesh refinement studies. It is hard to guarantee the elimination of programming errors, but their likelihood can be reduced by the use of modular programming. Then it should be possible to obtain the same result when alternative implementations are substituted for each module. Mesh refinement studies may also help the detection of programming errors by exposing discrepancies from the predicted asymptotic behavior as the mesh spacing is reduced, or discrepancies from known results for special cases, such as the fact that the drag should be zero in two dimensional subsonic inviscid flow. It is only after the correctness of the program and the accuracy of the numerical solution procedure have been independently verified that it is possible to assess the modeling errors which may arise, for example, from the use of an inappropriate turbulence model. For a more detailed discussion of validation procedures the reader is referred to reference [39] .

4. Flow Solvers And Parallel Implementations

In our most recent papers [40, 41, 42, 43], the Euler adjoint-based design formulation was extended to treat complete aircraft configurations using a new parallel multiblock implementation. This extension of the method required the replacement of our single-block flow and adjoint solvers [44, 45, 46] with their multiblock counterparts.

In order to use CFD in an automated design environment, the flow solver must meet fundamental requirements of accuracy, efficiency, robustness, and fast convergence. High accuracy is required since the improvements predicted by the method for the design in question can only be as good as the accuracy of the flow analysis. Efficiency of the flow solver is also critical since the optimization of the design will generally require the computation of many flow and adjoint solutions. The robustness of the method (i.e. its ability to obtain a flow solution for a variety of configuration shapes and flow conditions) is of great importance in order to guarantee the continuity of the design process. The last aspect, rapid convergence, is also of significant importance; in highly refined aerodynamic design applications, the benefit of aerodynamic optimization lies in obtaining the last few percentage points in aerodynamic efficiency. In such cases, the solutions must be highly converged such that the noise in the figure of merit is well below the level of realizable improvement.

In our three-dimensional single-block applications, the first author's FLO87 and the second author's FLO107 codes easily met all of the above criteria [47]. Both FLO87 and FLO107 achieve fast convergence with the aid of multigridding and implicit residual smoothing, typically obtaining solutions that converge to machine accuracy. The challenge addressed in reference [41] was to meet these strict convergence requirements within the framework of a cell-centered multiblock flow solver. Thus, the general strategy employed in developing the multiblock scheme consisted of using the same numerical algorithms used in the single-block code. This goal can be achieved by constructing and maintaining a halo of cells surrounding each block in the mesh. Then, updates to the internal cells within each block can be performed independently using single-block techniques. This strategy requires the identification of the halo cells and their corresponding "donor" cells (from which they inherit their values) in the interior of the neighboring blocks. The values of these halo cells need to be loaded with appropriate flow field data at regular intervals in the main solution algorithm. A double halo is required so that the flow field values of the complete stencil necessary to calculate the fluxes of all the internal cells in a block are available, even at the block boundaries. This halo approach to the solution of the flow in a multiblock domain can be easily mapped into a distributed memory parallel computer when one realizes that the "donor" cells may now simply reside in a different processor.

The main strategies that are used to accomplish the parallelization of the design procedure are a domain decomposition model, a SPMD (Single Program Multiple Data) strategy, and the MPI (Message Passing Interface) library for message passing. The choice of MPI was determined by the requirement that the resulting code be easily portable to different parallel computing platforms as well as to homogeneous and heterogeneous networks of workstations while still achieving high computational efficiency.

As an example, Figure 6 shows the parallel scalability of the multiblock design method for a mesh containing 1.8 million cells using up to 32 processors of an SGI Origin2000 parallel computer. Despite the fact that the multigrid technique is used in both the flow and adjoint solvers, the demonstrated parallel speedups are quite good.

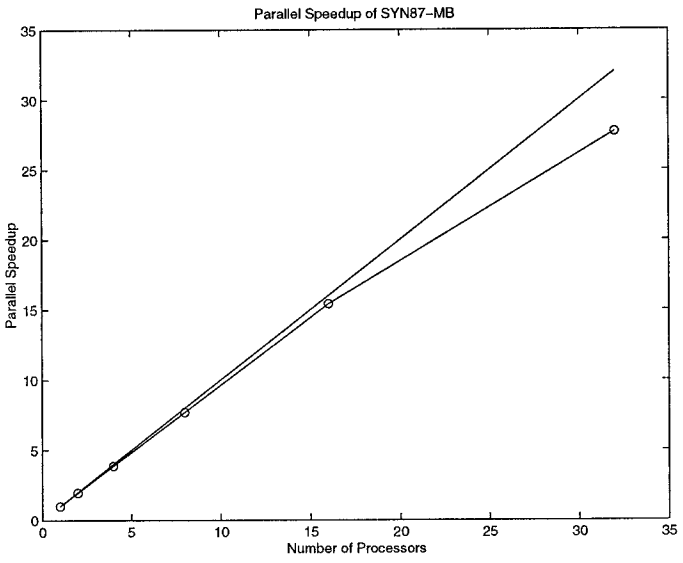

Figure 6: Scalability Study for Multiblock Design Method.

5. AERODYNAMIC SHAPE OPTIMIZATION

Traditionally the process of selecting design variations has been carried out by trial and error, relying on the intuition and experience of the designer. It is not at all likely that repeated trials in an interactive design and analysis procedure can lead to a truly optimum design. In order to take full advantage of the possibility of examining a large design space the numerical simulations need to be combined with automatic search and optimization procedures. This can lead to automatic design methods which will fully realize the potential improvements in aerodynamic efficiency.

The simplest approach to optimization is to define the geometry through a set of design parameters, which may, for example, be the weights α_i applied to a set of shape functions $b_i(x)$ so that the shape is represented as

$$f(x) = \sum \alpha_i b_i(x).$$

Then a cost function I is selected which might, for example, be the drag coefficient or the lift to drag ratio, and I is regarded as a function of the parameters α_i. The sensitivities $\frac{\partial I}{\partial \alpha_i}$ may now be estimated by making a small variation $\delta\alpha_i$ in each design parameter in turn and recalculating the flow to obtain the change in I. Then

$$\frac{\partial I}{\partial \alpha_i} \approx \frac{I(\alpha_i + \delta\alpha_i) - I(\alpha_i)}{\delta\alpha_i}.$$

The gradient vector $\frac{\partial I}{\partial \alpha}$ may now be used to determine a direction of improvement. The simplest procedure is to make a step in the negative gradient direction by setting

$$\alpha^{n+1} = \alpha^n - \lambda \frac{\partial I}{\partial \alpha},$$

so that to first order

$$I + \delta I = I + \frac{\partial I^T}{\partial \alpha}\delta\alpha = I - \lambda \frac{\partial I^T}{\partial \alpha}\frac{\partial I}{\partial \alpha}.$$

More sophisticated search procedures may be used such as quasi-Newton methods, which attempt to estimate the second derivative $\frac{\partial^2 I}{\partial \alpha_i \partial \alpha_j}$ of the cost function from changes in the gradient $\frac{\partial I}{\partial \alpha}$ in successive optimization steps. These methods also generally introduce line searches to find the minimum in the search direction which is defined at each step. The main disadvantage of this approach is the need for a number of flow calculations proportional to the number of design variables to estimate the gradient. The computational costs can thus become prohibitive as the number of design variables is increased. Reference [48] compares the performance of several search methods, including three quasi-Newton schemes, on a nonlinear test problem. This study showed that an explicit search method can be developed which exhibits a convergence rate that is independent of the number of design variables.

An alternative approach is to cast the design problem as a search for the shape that will generate the desired pressure distribution. This approach recognizes that the designer usually has an idea of the kind of pressure distribution that will lead to the desired performance. Thus, it is useful to consider the inverse problem of calculating the shape that will lead to a given pressure distribution. The method has the advantage that only one flow solution is required to obtain the desired design. Unfortunately, a physically realizable shape may not necessarily exist, unless the pressure distribution satisfies certain constraints. The difficulty that the target pressure may be unattainable may be circumvented by treating the inverse problem as a special case of the optimization problem, with a cost function which measures the error in the solution of the inverse problem. For example, if p_d is the desired surface pressure, one may take the cost function to be an integral over the body surface of the square of the pressure error,

$$I = \frac{1}{2}\int_{\mathcal{B}} (p - p_d)^2 d\mathcal{B},$$

or possibly a more general Sobolev norm of the pressure error. This has the advantage of converting a possibly ill posed problem into a well posed one. It has the disadvantage that it incurs the computational costs associated with optimization procedures.

Application of Control Theory

In order to reduce the computational costs, it turns out that there are advantages in formulating both the inverse problem and more general aerodynamic problems within the framework of the mathematical theory for the control of systems

governed by partial differential equations [49]. A wing, for example, is a device to produce lift by controlling the flow, and its design can be regarded as a problem in the optimal control of the flow equations by variation of the shape of the boundary. If the boundary shape is regarded as arbitrary within some requirements of smoothness, then the full generality of shapes cannot be defined with a finite number of parameters, and one must use the concept of the Frechet derivative of the cost with respect to a function. Clearly, such a derivative cannot be determined directly by finite differences of the design parameters because there are now an infinite number of these. Using techniques of control theory, however, the gradient can be determined indirectly by solving an adjoint equation which has coefficients defined by the solution of the flow equations. The cost of solving the adjoint equation is comparable to that of solving the flow equations. Thus the gradient can be determined with roughly the computational costs of two flow solutions, independently of the number of design variables, which may be infinite if the boundary is regarded as a free surface.

For flow about an airfoil or wing, the aerodynamic properties which define the cost function are functions of the flowfield variables (w) and the physical location of the boundary, which may be represented by the function $\mathcal{F}$, say. Then

$$I = I\left(w, \mathcal{F}\right),$$

and a change in $\mathcal{F}$ results in a change

$$\delta I = \frac{\partial I^T}{\partial w}\delta w + \frac{\partial I^T}{\partial \mathcal{F}}\delta \mathcal{F}, \tag{1}$$

in the cost function. Using control theory, the governing equations of the flowfield are introduced as a constraint in such a way that the final expression for the gradient does not require reevaluation of the flowfield. In order to achieve this δw must be eliminated from (1). Suppose that the governing equation R which expresses the dependence of w and $\mathcal{F}$ within the flowfield domain D can be written as

$$R\left(w, \mathcal{F}\right) = 0. \tag{2}$$

Then δw is determined from the equation

$$\delta R = \left[\frac{\partial R}{\partial w}\right]\delta w + \left[\frac{\partial R}{\partial \mathcal{F}}\right]\delta \mathcal{F} = 0. \tag{3}$$

Next, introducing a Lagrange Multiplier ψ, we have

$$\begin{aligned}\delta I &= \frac{\partial I^T}{\partial w}\delta w + \frac{\partial I^T}{\partial \mathcal{F}}\delta \mathcal{F} - \psi^T\left(\left[\frac{\partial R}{\partial w}\right]\delta w + \left[\frac{\partial R}{\partial \mathcal{F}}\right]\delta \mathcal{F}\right)\\ &= \left\{\frac{\partial I^T}{\partial w} - \psi^T\left[\frac{\partial R}{\partial w}\right]\right\}\delta w + \left\{\frac{\partial I^T}{\partial \mathcal{F}} - \psi^T\left[\frac{\partial R}{\partial \mathcal{F}}\right]\right\}\delta \mathcal{F}.\end{aligned}$$

Choosing ψ to satisfy the adjoint equation

$$\left[\frac{\partial R}{\partial w}\right]^T \psi = \frac{\partial I}{\partial w} \tag{4}$$

the first term is eliminated, and we find that

$$\delta I = \mathcal{G}\delta \mathcal{F}, \tag{5}$$

where

$$\mathcal{G} = \frac{\partial I^T}{\partial \mathcal{F}} - \psi^T\left[\frac{\partial R}{\partial \mathcal{F}}\right].$$

The advantage is that (5) is independent of δw, with the result that the gradient of I with respect to an arbitrary number of design variables can be determined without the need for additional flow-field evaluations. In the case that (2) is a partial differential equation, the adjoint equation (4) is also a partial differential equation and appropriate boundary conditions must be determined.

After making a step in the negative gradient direction, the gradient can be recalculated and the process repeated to follow a path of steepest descent until a minimum is reached. In order to avoid violating constraints, such as a minimum acceptable wing thickness, the gradient may be projected onto the allowable subspace within which the constraints are satisfied. In this way one can devise procedures which must necessarily converge at least to a local minimum, and which can be accelerated by the use of more sophisticated descent methods such as conjugate gradient or quasi-Newton algorithms.

In order to make sure that each new shape in the optimization sequence remains smooth, it proves essential to smooth the gradient and to replace $\mathcal{G}$ by its smoothed value $\bar{\mathcal{G}}$ in the descent process. This also acts as a preconditioner which allows the use of much larger steps. If in fact the optimum is not a smooth shape, it should be noted that the optimization process based on an implicitly-smoothed gradient will still achieve the optimum state in the limit, approaching it with a set of decreasingly smooth shapes. To apply smoothing in the ξ_1 direction, for example, the smoothed gradient $\bar{\mathcal{G}}$ may be calculated from a discrete approximation to

$$\bar{\mathcal{G}} - \frac{\partial}{\partial \xi_1}\epsilon\frac{\partial}{\partial \xi_1}\bar{\mathcal{G}} = \mathcal{G} \tag{6}$$

where ϵ is the smoothing parameter. If one sets $\delta \mathcal{F} = -\lambda\bar{\mathcal{G}}$, then, assuming the modification is applied on the surface $\xi_2 =$ constant, the first order change in the cost function is

$$\begin{aligned}\delta I &= -\iint \mathcal{G}\delta \mathcal{F}\, d\xi_1 d\xi_3\\ &= -\lambda\iint\left(\bar{\mathcal{G}} - \frac{\partial}{\partial \xi_1}\epsilon\frac{\partial \bar{\mathcal{G}}}{\partial \xi_1}\right)\bar{\mathcal{G}}\, d\xi_1 d\xi_3\\ &= -\lambda\iint\left(\bar{\mathcal{G}}^2 + \epsilon\left(\frac{\partial \bar{\mathcal{G}}}{\partial \xi_1}\right)^2\right)d\xi_1 d\xi_3\\ &< 0,\end{aligned}$$

assuring an improvement if λ is sufficiently small and positive, unless the process has already reached a stationary point

at which $\bar{\mathcal{G}} = 0$ (and therefore, according to Equation 6, $\mathcal{G} = 0$).

It turns out that this approach is tolerant to the use of approximate values of the gradient, so that neither the flow solution nor the adjoint solution need be fully converged before making a shape change. This results in very large savings in the computational cost so that a full optimization can be obtained with a cost equivalent to that of 2-10 flow solutions.

Quasi-Newton methods accelerate the search procedure by recursively estimating the Hessian matrix, A, of second derivatives from the changes of the gradient during the search steps, and then using a Newton step $\delta\mathcal{F} = -A^{-1}\mathcal{G}$. With N design variables, N steps are required to obtain a complete estimate of the Hessian, after which there is the possibility of obtaining very rapid convergence. Most quasi-Newton methods also require line searches to find the minimum along the search direction at each step. The performance of a variety of search methods has been evaluated in Reference [48] using a trajectory optimization problem (the brachistochrone) as a representative model. The study verified that the serach cost of a simple steepest descent method applied to this problem scales as N^2, where N is the number of design variables, while the cost of quasi-Newton methods scaled linearly with N as expected. On the other hand, with the correct amount of smoothing, the smooth descent method requires a fixed number of steps, independent of N. Considering that the evaluation of the gradient by a finite difference method requires $N + 1$ flow calculations, while the cost of its evaluation by the adjoint method is roughly that of two flow calculations, one arrives at the estimates of total computational cost given in Tables 1-2.

Steepest Descent	N^2 steps
Quasi-Newton	N steps
Smoothed Gradient	K steps (independent of N)

Table 1: Computational Cost of Search Algorithm as a Function of the Number of Design Variables, N.

The adjoint method can be applied to a variety of measures of performance. It should be remembered, however, that gradient search methods depend on the assumption that the cost-function depends continuously on the design parameters. This can be violated, if, for example, one attempts to calculate the sensitivity of the pressure at a fixed location, because there is the possibility that a shape modification could result in a shock moving over that location. The movement of the shock, however, is continuous as the shape changes, with the consequence that integrated quantities such as the drag coefficient also depend continuously on the shape. The adjoint equation allows the sensitivity of the drag coefficient to be calculated without the explicit evaluation of pressure sensitivities.

Finite Difference Gradients + Steepest Descent	$\mathcal{O}(N^3)$
Finite Difference Gradients + Quasi-Newton Search	$\mathcal{O}(N^2)$
Adjoint Gradients + Quasi-Newton Search	$\mathcal{O}(N)$
Adjoint Gradients + Smoothed Gradient Search	$\mathcal{O}(K)$ (independent of N)

Table 2: Total Computational Cost of Design as a Function of the Number of Design Variables, N.

In reference [50] the first author derived the adjoint equations for transonic flows modeled by both the potential flow equation and the Euler equations. The theory was developed in terms of partial differential equations, leading to an adjoint partial differential equation. In order to obtain numerical solutions, both the flow and the adjoint equations must be discretized. The control theory might be applied directly to the discrete flow equations which result from the numerical approximation of the flow equations by finite element, finite volume or finite difference procedures. This leads directly to a set of discrete adjoint equations with a matrix which is the transpose of the Jacobian matrix of the full set of discrete nonlinear flow equations. On a three-dimensional mesh with indices i, j, k the individual adjoint equations may be derived by collecting together all the terms multiplied by the variation $\delta w_{i,j,k}$ of the discrete flow variable $w_{i,j,k}$. The resulting discrete adjoint equations represent a possible discretization of the adjoint partial differential equation. If these equations are solved exactly they can provide the exact gradient of the cost function which results from the discretization of the flow equations, which is itself inexact. This may facilitate the asymptotic convergence of the search procedure. On the other hand any consistent discretization of the adjoint partial differential equation will yield the exact gradient in the limit as the mesh is refined.

There are a number of benefits to be gained from developing the theory for the partial differential equations of the flow. First, the true optimum shape belongs to an infinitely dimensional space of design parameters, and the theory provides an indication, in principle, of how such a solution could be approached if sufficient computational resources were available. Second, it provides insight into the nature of the adjoint equations, and the connection between the formulation of the cost function and the boundary conditions needed to assure a well-posed problem. Third, in certain circumstances the discrete solution may lose the property of continuous depen-

dence of the design parameters. It may, for example, contain non-differentiable flux limiters. Also, if adaptive mesh refinement is used, there will be a discontinuous change in the solution whenever a mesh point is added or deleted. Finally, the differential equation theory provides a guideline for the design of iterative solution methods for the adjoint equation, both in the case when the adjoint equation is separately discretized and in the case when the discrete adjoint equations are derived directly from the discrete flow equations. The theory for standard multigrid methods, for example, depends on the property that the discrete equations on a sequence of meshes all represent the same differential equation. It turns out that the same multigrid solution method can readily be used for both the flow and the adjoint equation.

The adjoint method has also been extended to treat the compressible Navier Stokes equations [51, 52]. The details of the mathematical derivation can be found in Reference [53]. In addition, in some of our applications where the satisfaction of constraints was important, the adjoint method was used to provide sensitivity information to an external optimization method such as NPSOL [54]. NPSOL implements a sequential quadratic programming algorithm and allows for the consideration of both linear and non-linear constraints. The overall cost of the optimization procedure was greatly increased, however, because of the need to conduct line searches every time a design change direction is chosen.

6. INDUSTRIAL EXPERIENCE AND RESULTS

The methods described in this paper have been quite thoroughly tested in industrial applications in which they were used as a tool for aerodynamic design. They have proved useful both in inverse mode to find shapes that would produce desired pressure distributions, and for direct minimization of the drag. They have been applied both to well understood configurations that have gradually evolved through incremental improvements guided by wind tunnel tests and computational simulation, and to new concepts for which there is a limited knowledge base. In either case they have enabled engineers to produce improved designs.

Substantial improvements are usually obtained with 20 – 200 design cycles, depending on the difficulty of the case. One concern is the possibility of getting trapped in a local minimum. In practice this has not proved to be a source of difficulty. In inverse mode, it often proves possible to come very close to realizing the target pressure distribution, thus effectively demonstrating convergence. In drag minimization, the result of the optimization is usually a shock-free wing. If one considers drag minimization of airfoils in two-dimensional inviscid transonic flow, it can be seen that every shock-free airfoil produces zero drag, and thus optimization based solely on drag has a highly non-unique solution, since a theorem is available that states that a shock-free solution in transonic flow is an isolated point. Different shock-free airfoils can be obtained by starting from different initial profiles. One may also influence the character of the final design by blending a target pressure distribution with the drag in the definition of the cost function.

Similar considerations apply to three-dimensional wing design. Since the vortex drag can be reduced simply by reducing the lift, the lift coefficient must be fixed for a meaningful drag minimization. In order to do this the angle of attack α is adjusted during the flow solution. It has proved most effective to make a small change $\delta\alpha$ proportional to the difference between the actual and the desired lift coefficient every few iterations in the flow calculation. A typical wing of a transport aircraft is designed for a lift coefficient in the range of 0.4 to 0.6. The total wing drag may be broken down into vortex drag, drag due to viscous effects, and shock drag. The vortex drag coefficient is typically in the range of 0.0100 (100 counts), while the friction drag coefficient is in the range of 45 counts, and the shock drag at a Mach number just before the onset of severe drag rise is of the order of 15 counts. With a fixed span, typically dictated by structural limits or a constraint imposed by airport gates, the vortex drag is entirely a function of span loading, and is minimized by an elliptic loading unless winglets are added. Transport aircraft usually have highly tapered wings with very large root chords to accommodate retraction of the undercarriage. An elliptic loading may lead to excessively large section lift coefficients on the outboard wing, leading to premature shock stall or buffet when the load is increased. The structure weight is also reduced by a more inboard loading which reduces the root bending moment. Thus the choice of span loading is influenced by other considerations. The skin friction of transport aircraft is typically very close to flat plate skin friction in turbulent flow, and is very insensitive to section variations. An exception to this is the case of smaller executive jet aircraft, for which the Reynolds number may be small enough to allow a significant run of laminar flow if the suction peak of the pressure distribution is moved back on the section. This leaves the shock drag as the primary target for wing section optimization. This is reduced to zero if the wing is shock-free, leaving no room for further improvement. Thus the attainment of a shock-free flow is a demonstration of a successful drag minimization. In practice range is maximized by maximizing $M\frac{L}{D}$, and this is likely to be increased by increasing the lift coefficient to the point where a weak shock appears. One may also use optimization to find the maximum Mach number at which the shock drag can be eliminated or significantly reduced for a wing with a given sweepback angle and thickness. Alternatively one may try to find the largest wing thickness or the minimum sweepback angle for which the shock drag can be eliminated at a given Mach number. This can yield both savings in structure weight and increased fuel volume. If there is no fixed limit for the wing span, such as a gate constraint,

increased thickness can be used to allow an increase in aspect ratio for a wing of equal weight, in turn leading to a reduction in vortex drag. Since the vortex drag is usually the largest component of the total wing drag, this is probably the most effective design strategy, and it may pay to increase the wing thickness to the point where the optimized section produces a weak shock wave rather than a shock-free flow [55].

The first major industrial application of an adjoint based aerodynamic optimization method was the wing design of the Beech Premier [56] in 1995. The method was successfully used in inverse mode as a tool to obtain pressure distributions favorable to the maintenance of natural laminar flow over a range of cruise Mach numbers. Wing contours were obtained which yielded the desired pressure distribution in the presence of closely coupled engine nacelles on the fuselage above the wing trailing edge.

During 1996 some preliminary studies indicated that the wings of both the McDonnell Douglas MD-11 and the Boeing 747-200 could be made shock-free in a representative cruise condition by using very small shape modifications, with consequent drag savings which could amount to several percent of the total drag. This led to a decision to evaluate adjoint-based design methods in the design of the McDonnell Douglas MDXX during the summer and fall of 1996. In initial studies wing redesigns were carried out for inviscid transonic flow modeled by the Euler equations. A redesign to minimize the drag at a specified lift and Mach number required about 40 design cycles, which could be completed overnight on a workstation.

Three main lessons were drawn from these initial studies: (i) the fuselage effect is too large to be ignored and must be included in the optimization, (ii) single-point designs could be too sensitive to small variations in the flight condition, typically producing a shock-free flow at the design point with a tendency to break up into a severe double shock pattern below the design point, and (iii) the shape changes necessary to optimize a wing in transonic flow are smaller than the boundary layer displacement thickness, with the consequence that viscous effects must be included in the final design. The reason that these items are critical to the aerodynamic shape optimization of a commercial aircraft is that a one percent improvement to the aircraft performance is a significant accomplishment. In the case of the MDXX, the Douglas engineers had done an exceptional job before our optimization capabilities were brought to bear on this effort. During their wing design, the man-in-the-loop considered the fuselage, nacelle and viscous effects over a wide range of flow conditions yielding a near-optimum design. While a manual design process requires man-years of effort to complete, the performance of the final design cannot be compromised. Hence, the simulations on which the optimization process is built must accurately represent the essential elements of the geometry as well as the flow physics.

In order to meet the first two of these considerations, the second phase of the study was concentrated on the optimization of wing-body combinations with multiple design points. These were still performed with inviscid flow to reduce computational cost and allow for fast turnaround. It was found that comparatively insensitive designs could be obtained by minimizing the drag at a fixed Mach number for three fairly closely spaced lift coefficients such as 0.5, 0.525, and 0.55, or alternatively three nearby Mach numbers with a fixed lift coefficient.

The third phase of the project was focused on the design with viscous effects using as a starting point wings which resulted from multipoint inviscid optimization. While the full viscous adjoint method was still under development, it was found that useful improvements could be realized, particularly in inverse mode, using the inviscid result to provide the target pressure, by coupling an inviscid adjoint solver to a viscous flow solver. Computer costs are many times larger, both because finer meshes are needed to resolve the boundary layer, and because more iterations are needed in the flow and adjoint solutions. In order to force the specified lift coefficient the number of iterations in each flow solution had to be increased from 15 to 100. To achieve overnight turnaround a fully parallel implementation of the software had to be developed. Finally it was found that in order to produce sufficiently accurate results, the number of mesh points had to be increased to about 1.8 million. In the final phase of this project it was planned to carry out a propulsion integration study using the multiblock versions of the software. This study, well underway, was not completed due to the cancelation of the entire MDXX project.

In the next subsections we present examples of the use of the adjoint method for viscous inverse and drag minimization in two dimensional flow. We then show a three-dimensional wing design using the Euler equations and a wing design using the full viscous adjoint method in its current form, implemented in the computer program SYN107. These calculations were all performed using the simple descent optimization method with smoothing of the gradient described in Section 5.. This has proved to be very efficient: in all cases the final optimum design was achieved with a total computational cost equivalent to the cost of from 2 to 10 converged flow solutions. The remaining subsections present results of optimizations for complete configurations in inviscid transonic and supersonic flow using the multiblock parallel design program, SYN107-MB.

Inverse design of an airfoil in transonic viscous flow

Our first example shows an inverse design in two dimensional viscous transonic flow obtained using the design code SYN103. The target pressure is that of the section of the ONERA M6 wing at Mach .75 and a lift coefficient of .50 and is denoted by the circles in the Figures. It was calculated using

SYN103 in analysis mode, thus it should be exactly realizable. A C-type mesh was used which contained 256 intervals in the chordwise and 96 cells in the normal direction for a total of 24,576 cells. The design calculation was started with the NACA 0012 airfoil as the initial profile, and the ONERA M6 pressure distribution was almost exactly recovered in 25 design cycles. In the first cycle 120 iterations were used in both the flow and the adjoint solutions. In the subsequent cycles only 30 iterations were used in both the flow and adjoint solutions. Figure 7 shows the initial profile and pressure distribution with the pressure coefficient plotted vertically in the negative direction. It then shows the results after one, five and twenty five design cycles, with the target represented by circles. It also superposes on each redesigned profile the smoothed gradient plotted in the direction of the shape modification. A fixed scale is used so that it is possible to observe the decrease in the magnitude of the gradient as the calculation converges enough to ensure that they were fairly close to convergence. The root mean square error between the target and actual pressure was reduced from .0530 to .0016 in the course of the entire calculation which took 3569 seconds using a single Silicon Graphics R10000 processor. A fully converged flow solution using 500 iterations on the same mesh took 936 seconds, so the cost of the entire design calculation was about that of three and one half flow solutions.

Drag reduction of an airfoil in viscous flow

The next example shows a redesign of the RAE2822 airfoil to reduce the drag at a fixed lift coefficient of .65 in transonic flow at Mach .75. In this case a shock free flow was obtained after 10 design cycles, in each of which both the flow and the adjoint solutions were calculated with 25 multigrid cycles. A grid with 512×64 cells was used. The pressure drag was reduced from .0091 to .0041, while the viscous drag remained essentially constant. The constraint was imposed that the thickness of the profile could not be reduced by only permitting outward movement from the initial profile. Figure 8 displays the sequence of pressure distributions, showing the elimination of the shock wave. It also shows the initial profile, and the smoothed gradient superposed on the subsequent profiles. It can be seen that the gradient continues to have an inward component, indicating that the drag might be further reduced if a thickness reduction were permitted. Furthermore, it should be noted that the sensitivities in the vicinity of the trailing edge want to cross the upper and lower surfaces; this corresponds to inserting a sink at the trailing edge which reduces the drag on the airfoil. In order to prevent this non-physical shape from forming, special boundary conditions are imposed on the smoothed-gradient equation (6 that freeze the position of the upper and lower trailing-edge points. In this case the entire design calculation took 1549 seconds on an R10000 processor, while a fully converged flow solution took 1160 seconds, so the cost of the design was equivalent to less than two flow solutions.

Three point inviscid redesign of the Boeing 747 wing

The third example shows a redesign of the wing of the Boeing 747 to reduce its drag in a typical cruising condition. It has been our experience that drag minimization at a single point tends to produce a wing which is shock free at its design point, but tends to display undesirable characteristics off its design point. Typically, a double shock pattern forms below the design lift coefficient and Mach number, and a single, but fairly strong shock above the design point. To alleviate this tendency the calculation was performed with three design points. In carrying out multipoint designs of this kind a composite gradient is calculated as a weighted average of the gradients calculated for each design point separately. In this case the design points were selected as lift coefficients of .38, .42 and .46 for the exposed wing at Mach .85. Because the fuselage has a significant effect on the flow over the wing, the calculations were performed for the wing body combination, but the shape modifications were restricted to the wing alone. The fuselage also contributes to the lift, so that the total lift coefficient at the mid design point was estimated to be .50.

The results are displayed in Figures 9 - 11 and in Table 3 which shows the pressure drag at three design points of the initial wing, and the final wing after 30 design cycles. It can been seen that a drag reduction was obtained over the entire range of lift coefficients, and at the mid design point the redesigned wing is almost shock free. Figure 12 shows the modification in the wing section about half way out the span. It can be seen that a useful drag reduction can be obtained by a very small change in the wing shape. This is because of the extreme sensitivity of the transonic flow. Also, it is clear that without a tool of this kind it would be almost impossible to find an optimum shape.

Transonic Viscous Wing-Body Design

A typical result for drag minimization of a wing body combination in transonic viscous flow is presented next. The viscous adjoint optimization method was used with a Baldwin-Lomax turbulence model. The initial wing is similar to one produced during the MDXX design studies. Figures 13-15 show the result of the wing-body redesign on a C-H mesh with $288 \times 96 \times 64$ cells. The wing has sweep back of about 38 degrees at the 1/4 chord. A total of 44 iterations of the viscous optimization procedure resulted in a shock-free wing at a cruise design point of Mach 0.86, with a lift coefficient of 0.61 for the wing-body combination at a Reynolds number of 101 million based on the root chord. Using 48 processors of an SGI Origin2000 parallel computer, each design iteration takes about 22 minutes so that overnight turnaround

for such a calculation is possible. Figure 13 compares the pressure distribution of the final design with that of the initial wing. The final wing is quite thick, with a thickness to chord ratio of about 14 percent at the root and 9 percent at the tip. The optimization was performed with a constraint that the section modifications were not allowed to decrease the thickness anywhere. The design offers excellent performance at the nominal cruise point. A drag reduction of 2.2 counts was achieved from the initial wing which had itself been derived by inviscid optimization. Figures 14 and 15 show the results of a Mach number sweep to determine the drag rise. The drag coefficients shown in the figures represent the total wing drag including shock, vortex, and skin friction contributions. It can be seen that a double shock pattern forms below the design point, while there is actually a slight increase in the drag coefficient at Mach 0.85. This wing has a low drag coefficient, however, over a wide range of conditions. Above the design point a single shock forms and strengthens as the Mach number increases.

Transonic Multipoint Constrained Aircraft Design

As a first example of the automatic design capability for complex configurations, we show drag reduction for a typical business jet configuration. The objective of the design is to alter the geometry of the wing in order to minimize the configuration inviscid drag at three different flight conditions simultaneously. Realistic geometric spar thickness constraints are enforced. The geometry chosen for this analysis is a full configuration business jet composed of wing, fuselage, pylon, nacelle, and empennage. The inviscid multiblock mesh around this configuration follows a general C-O topology with special blocking to capture the geometric details of the nacelles, pylons and empennage. A total of 240 point-to-point matched blocks with 4, 157, 440 cells (including halos) are used to grid the complete configuration. This mesh allows the use of 4 multigrid levels obtained through recursive coarsening of the initial fine mesh. The upstream, downstream, upper and lower far field boundaries are located at an approximate distance of 15 wing semispans, while the far field boundary beyond the wing tip is located at a distance approximately equal to 5 semispans. An engineering-accuracy solution (with a decrease of 4 orders of magnitude in the average density residual) can be obtained in 100 multigrid cycles. This kind of solution can be routinely accomplished in under 20 minutes of wall clock time using 32 processors of an SGI Origin2000 computer.

In this case the adjoint method was used to supply gradients to the well-known NPSOL optimization software which implements a sequential quadratic programming algorithm with provisions for handling of constraints. The initial configuration was designed for Mach = 0.8 and C_L = 0.3. The three operating points chosen for this design are Mach = 0.81 with C_L = 0.35, Mach = 0.82 with C_L = 0.30, and Mach = 0.83 with C_L = 0.25. In order to demonstrate the advantage of a multipoint design approach, the final solution at the middle design point will be compared with a single point design at the same conditions. As the geometry of the wing is modified, the design algorithm computes new wing-fuselage intersections. The wing component is made up of six airfoil defining sections. Eighteen Hicks-Henne design variables are applied to five of these sections for a total of 90 design variables. The sixth section at the symmetry plane is not modified. Spar thickness constraints were also enforced on each defining station at the $x/c = 0.2$ and $x/c = 0.8$ locations. Maximum thickness was forced to be preserved at $x/c = 0.4$ for all six defining sections. To ensure an adequate included angle at the trailing edge, each section was also constrained to preserve thickness at $x/c = 0.95$. Finally, to preserve leading edge bluntness, the upper surface of each section was forced to maintain its height above the camber line at $x/c = 0.02$. Combined, a total of 30 linear geometric constraints were imposed on the configuration.

Figure 16 shows the surface of the configuration colored by the local coefficient of pressure, C_p, before and after redesign for the middle design point. One can clearly observe that the strength of the shock wave on the upper surface of the configuration has been considerably reduced.

Figure 17 shows the initial and final airfoil geometries and C_p distributions for the middle design point after 5 NPSOL design iterations. It is evident that the new design has significantly reduced the shock strengths on both upper and lower wing surfaces. Upon examination of the results, this is also found to be true for the other two design points. The transitions between design points are also quite smooth. For comparison purposes, a single point drag minimization study (Mach = 0.81 and C_L = 0.25) is carried out starting from the same initial configuration and using the same design variables and geometric constraints.

Figure 18 shows comparisons of the solution from the three-point design with that of the single point design at the third design point. Interestingly, the upper surface shapes for both final designs are very similar for the first two design points. However, in the case of the single point design, a strong lower surface shock appears at the Mach = 0.83, C_L = 0.25 design point. The three-point design is able to suppress the formation of this lower surface shock and achieves a 9 count drag benefit over the single point design at this condition. However, it has a 1 count penalty at the single point design condition. Examination of the reusults shows that the three-point design features a weak single shock for one of the three design points and a very weak double shock at another design point. Table 4 summarizes the drag results for the two designs. The C_D values have been normalized by the drag of the initial configuration at the second design point.

Supersonic Constrained Aircraft Design

For supersonic design, provided that turbulent flow is assumed over the entire configuration, the inviscid Euler equations suffice for aerodynamic design since the pressure drag is not greatly affected by the inclusion of viscous effects. Moreover, flat plate skin friction estimates of viscous drag are often very good approximations. In this study we show drag reduction of a generic supersonic transport configuration used in reference [40] using the inviscid Euler equations to model the flow.

The baseline supersonic transport configuration was sized to accommodate 300 passengers with a gross take-off weight of 750,000 lbs. The supersonic cruise point is Mach 2.2 with a C_L of 0.105. Figure 19 shows that the planform is a cranked-delta configuration with a break in the leading edge sweep. The inboard leading edge sweep is 68.5 degrees while the outboard is 49.5 degrees. Since the Mach angle at $M = 2.2$ is 63 degrees it is clear that some leading edge bluntness may be used inboard without a significant wave drag penalty. Blunt leading edge airfoils were created with thickness ranging from 4% at the root to 2.5% at the leading edge break point. These symmetric airfoils were chosen to accommodate thick spars at roughly the 5% and 80% chord locations over the span up to the leading edge break. Outboard of the leading edge break where the wing sweep is ahead of the Mach cone, a sharp leading edge was used to avoid unnecessary wave drag. The airfoils were chosen to be symmetric, biconvex shapes modified to have a region of constant thickness over the mid-chord. The four-engine configuration features axisymmetric nacelles tucked close to the wing lower surface. This layout favors reduced wave drag by minimizing the exposed boundary layer diverter area. However, in practice it may be problematic because of the channel flows occurring in the juncture region of the diverter, wing, and nacelle at the wing trailing edge.

The computational mesh on which the design is run has 180 blocks and 1,500,000 mesh cells (including halos), while the underlying geometry entities define the wing with 16 sectional cuts and the body with 200 sectional cuts. In this case, where we hope to optimize the shape of the wing, care must be taken to ensure that the nacelles remain properly attached with diverter heights being maintained.

The objective of the design is to reduce the total drag of the configuration at a single design point (Mach = 2.2, C_L = 0.105) by modifying the wing shape. Just as in the transonic case, 18 design variables of the Hicks-Henne type are chosen for each wing defining section. Similarly, instead of applying them to all 16 sections, they are applied to 8 of the sections and then lofted linearly to the neighboring sections. Spar thickness constraints are imposed for all wing defining sections at $x/c = 0.05$ and $x/c = 0.8$. An additional maximum thickness constraint is specified along the span at $x/c = 0.5$. A final thickness constraint is enforced at $x/c = 0.95$ to ensure a reasonable trailing edge included angle. An iso-C_p representation of the initial and final designs is depicted in Figure 19 for both the upper and lower surfaces.

It is noted that the strong oblique shock evident near the leading edge of the upper surface on the initial configuration is largely eliminated in the final design after 5 NPSOL design iterations. Also, it is seen that the upper surface pressure distribution in the vicinity of the nacelles has formed an unexpected pattern. However, recalling that thickness constraints abound in this design, these upper surface pressure patterns are assumed to be the result of sculpting of the lower surface near the nacelles which affects the upper surface shape via the thickness constraints. For the lower surface, the leading edge has developed a suction region while the shocks and expansions around the nacelles have been somewhat reduced. Figure 20 shows the pressure coefficients and (scaled) airfoil sections for four sectional cuts along the wing. These cuts further demonstrate the removal of the oblique shock on the upper surface and the addition of a suction region on the leading edge of the lower surface. The airfoil sections have been scaled by a factor of 2 so that shape changes may be seen more easily. Most notably, the section at 38.7% span has had the lower surface drastically modified such that a large region of the aft airfoil has a forward-facing portion near where the pressure spike from the nacelle shock impinges on the surface. The final overall pressure drag was reduced by 8%, from $C_D = 0.0088$ to $C_D = 0.0081$.

7. CONCLUSIONS

A procedure for the aerodynamic shape optimization of modern aircraft has been developed and presented. The method is based on ideas from control theory and relies on the solution of an adjoint equation for the inexpensive computation of design sensitivities that would typically involve recalculation of the flow. The procedure has been implemented in several computer programs that can deal with varying levels of sophistication in the flow model as well as in the geometric complexity of the configuration of interest. The method has been used in several occasions in an industrial setting, and in all cases it has yielded significant performance improvements over the baseline configuration. Furthermore, it has been verified that the computational cost of a complete design calculation is generally a small multiple of a typical analysis run (ranging from 2 to 10), consistent with the estimates of Table 2. Together with the use of large scale parallel computers, this has allowed overnight turnaround for design calculations on complete aircraft configurations.

8. ACKNOWLEDGEMENTS

This work has benefited from the generous support of the Air Force Office of Scientific Research under Grant No. AF F49620-98-1-002, as well as the contributions of the Boeing Company and the Raytheon Aircraft Company.

REFERENCES

[1] M. Van Dyke. *An Album of Fluid Motion.* The Parabolic Press, Stanford, 1982.

[2] D.R. Chapman, H. Mark, and M.W. Pirtle. Computers vs. wind tunnels in aerodynamic flow simulations. *Astronautics and Aeronautics*, 13(4):22–30, 35, 1975.

[3] J.T. Oden, L. Demkowicz, T. Liszka, and W. Rachowicz. *h-p* adaptive finite element methods for compressible and incompressible flows. In S. L. Venneri A. K. Noor, editor, *Proceedings of the Symposium on Computational Technology on Flight Vehicles*, pages 523–534, Washington, D.C., November 1990. Pergamon.

[4] L. Martinelli and A. Jameson. Validation of a multigrid method for the Reynolds averaged equations. *AIAA paper 88-0414*, 1988.

[5] W.H. Jou. Boeing Memorandum AERO-B113B-L92-018, September 1992. To Joseph Shang.

[6] M.H. Ha. The impact of turbulence modelling on the numerical prediction of flows. In M. Napolitano and F. Solbetta, editors, *Proc. of the 13th International Conference on Numerical Methods in Fluid Dynamics*, pages 27–46, Rome, Italy, July 1992. Springer Verlag, 1993.

[7] ECARP. *European Computational Aerodynamics Research Project, Validation of CFD Codes and Assessment of Turbulence Models.* Vieweg.

[8] T. Cebeci and A.M.O. Smith. *Analysis of Turbulent Boundary Layers.* Academic Press, 1974.

[9] B. Baldwin and H. Lomax. Thin layer approximation and algebraic model for separated turbulent flow. AIAA Paper 78-257, 1978.

[10] D. Degani and L. Schiff. Computation of turbulent supersonic flows around pointed bodies having crossflow separation. *J. Comp. Phys.*, 66:173–196, 1986.

[11] L. Martinelli, A. Jameson, and E. Malfa. Numerical simulation of three-dimensional vortex flows over delta wing configurations. In M. Napolitano and F. Solbetta, editors, *Proc. 13th International Confrence on Numerical Methods in Fluid Dynamics*, pages 534–538, Rome, Italy, July 1992. Springer Verlag, 1993.

[12] D. Johnson and L. King. A mathematically simple turbulence closure model for attached and separated turbulent boundary layers. *AIAA Journal*, 23:1684–1692, 1985.

[13] C.L. Rumsey and V.N. Vatsa. A comparison of the predictive capabilities of several turbulence models using upwind and centered - difference computer codes. AIAA Paper 93-0192, AIAA 31st Aerospace Sciences Meeting, Reno, NV, January 1993.

[14] T.J. Kao, T.Y. Su, and N.J. Yu. Navier-Stokes calculations for transport wing-body configurations with nacelles and struts. AIAA Paper 93-2945, AIAA 24th Fluid Dynamics Conference, Orlando, July 1993.

[15] W.P. Jones and B.E. Launder. The calculation of low-Reynolds-number phenomena with a two-equation model of turbulence. *Int. J. of Heat Tran.*, 16:1119–1130, 1973.

[16] D.C. Wilcox. A half a century historical review of the k-ω model. AIAA Paper 91-0615, AIAA 29th Aerospace Sciences Meeting, Reno, NV, January 1991.

[17] C.G. Speziale, E.C. Anderson, and R. Abid. A critical evaluation of two-equation models for near wall turbulence. AIAA Paper 90-1481, June 1990.

[18] R. Abid, C.G. Speziale, and S. Thangam. Application of a new k-τ model to near wall turbulent flows. AIAA Paper 91-0614, AIAA 29th Aerospace Sciences Meeting, Reno, NV, January 1991.

[19] F. Menter. Zonal two-equation k-ω turbulence models for aerodynamic flows. AIAA Paper 93-2906, AIAA 24th Fluid Dynamics Meeting, Orlando, July 1993.

[20] T.J. Coakley. Numerical simulation of viscous transonic airfoil flows. AIAA Paper 87-0416, AIAA 25th Aerospace Sciences Meeting, Reno, NV, January 1987.

[21] F. Liu and X. Zheng. A strongly coupled time-marching method for solving the Navier-Stokes and $\kappa - \omega$ turbulence model equations with multigrid. *J. Comp. Phys.*, 128:289–300, 1996.

[22] B. R. Smith. A near wall model for the $k - l$ two equation turbulence model. *AIAA paper 94-2386*, AIAA 25th Fluid Dynamics Conference, Colorado Springs, CO, June 1994.

[23] B.S. Baldwin and T.J. Barth. A one-equation turbulence transport model for high Reynolds number wall-bounded flows. AIAA Paper 91-0610, AIAA 29th Aerospace Sciences Meeting, Reno, NV, January 1991.

[24] P. Spalart and S. Allmaras. A one-equation turbulent model for aerodynamic flows. AIAA Paper 92-0439, AIAA 30th Aerospace Sciences Meeting, Reno, NV, January 1992.

[25] J.E. Melton, S.A. Pandya, and J.L. Steger. 3D Euler flow solutions using unstructured Cartesian and prismatic grids. AIAA Paper 93-0331, Reno, NV, January 1993.

[26] S.S. Samant, J.E. Bussoletti, F.T. Johnson, R.H. Burkhart, B.L. Everson, R.G. Melvin, D.P. Young, L.L. Erickson, and M.D. Madson. TRANAIR: A computer code for transonic analyses of arbitrary configurations. AIAA Paper 87-0034, 1987.

[27] M. Berger and R.J. LeVeque. An adaptive Cartesian mesh algorithm for the Euler equations in arbitrary geometries. AIAA Paper 89-1930, 1989.

[28] A.M. Landsberg, J.P. Boris, W. Sandberg, and T.R. Young. Naval ship superstructure design: Complex three-dimensional flows using an efficient, parallel method. *High Performance Computing 1993: Grand Challenges in Computer Simulation*, 1993.

[29] N.P. Weatherill and C.A. Forsey. Grid generation and flow calculations for aircraft geometries. *J. Aircraft*, 22:855–860, 1985.

[30] K. Sawada and S. Takanashi. A numerical investigation on wing/nacelle interferences of USB configuration. AIAA paper 87-0455, AIAA 25th Aerospace Sciences Meeting, Reno, NV, January 1987.

[31] J.A. Benek, P.G. Buning, and J.L. Steger. A 3-D Chimera grid embedding technique. AIAA Paper 85-1523, AIAA 7th Computational Fluid Dynamics Conference, Cincinnati, OH, 1985.

[32] J.A. Benek, T.L. Donegan, and N.E. Suhs. Extended Chimera grid embedding scheme with applications to viscous flows. AIAA Paper 87-1126, AIAA 8th Computational Fluid Dynamics Conference, Honolulu, HI, 1987.

[33] J. Vassberg. Multi-block mesh extrusion driven by a globally elliptic system. In *5th U.S. National Congress on Computational Mechanics, 2nd Symposium on Trends in Unstructured Mesh Generation*, University of Colorado, Boulder, CO, August 1999.

[34] B. Delaunay. Sur la sphere vide. *Bull. Acad. Science USSR VII: Class Scil, Mat. Nat.*, pages 793–800, 1934.

[35] T. J. Barth. Aspects of unstructured grids and finite volume solvers for the Euler and Navier Stokes equations. In *von Karman Institute for Fluid Dynamics Lecture Series Notes 1994-05*, Brussels, 1994.

[36] G. Voronoi. Nouvelles applications des parametres continus a la theorie des formes quadratiques. Deuxieme memoire: Recherches sur les parallelloedres primitifs. *J. Reine Angew. Math.*, 134:198–287, 1908.

[37] R. Lohner and P. Parikh. Generation of three-dimensional unstructured grids by the advancing front method. AIAA Paper 88-0515, Reno, NV, January 1988.

[38] A. Jameson. The present status, challenges, and future developments in Computational Fluid Dynamics. Technical report, 77th AGARD Fluid Dynamics Panel Symposium, Seville, October 1995.

[39] A. Jameson and L. Martinell. Mesh refinement and modelling errors in flow simulation. *AIAA paper 96-2050*, AIAA 27th Fluid Dynamics Conference, New Orleans, LA, June 1996.

[40] J. Reuther, J.J. Alonso, M.J. Rimlinger, and A. Jameson. Aerodynamic shape optimization of supersonic aircraft configurations via an adjoint formulation on parallel computers. *AIAA paper 96-4045*, 6th AIAA/NASA/ISSMO Symposium on Multidisciplinary Analysis and Optimization, Bellevue, WA, September 1996.

[41] J. Reuther, A. Jameson, J. Farmer, L. Martinelli, and D. Saunders. Aerodynamic shape optimization of complex aircraft configurations via an adjoint formulation. *AIAA paper 96-0094*, 34th Aerospace Sciences Meeting and Exhibit, Reno, Nevada, January 1996.

[42] J. Reuther, A. Jameson, J. J. Alonso, M. J. Rimlinger, and D. Saunders. Constrained multipoint aerodynamic shape optimization using an adjoint formulation and parallel computers. *AIAA paper 97-0103*, 35th Aerospace Sciences Meeting and Exhibit, Reno, Nevada, January 1997.

[43] J. Reuther, J. J. Alonso, J. C. Vassberg, A. Jameson, and L. Martinelli. An efficient multiblock method for aerodynamic analysis and design on distributed memory systems. *AIAA paper 97-1893*, June 1997.

[44] A. Jameson. Optimum aerodynamic design via boundary control. In *AGARD-VKI Lecture Series, Optimum Design Methods in Aerodynamics*. von Karman Institute for Fluid Dynamics, 1994.

[45] J. Reuther and A. Jameson. Aerodynamic shape optimization of wing and wing-body configurations using control theory. *AIAA paper 95-0123*, 33rd Aerospace Sciences Meeting and Exhibit, Reno, Nevada, January 1995.

[46] J. Reuther and A. Jameson. Supersonic wing and wing-body shape optimization using an adjoint formulation. Technical report, The Forum on CFD for Design and Optimization, (IMECE 95), San Francisco, California, November 1995.

[47] A. Jameson. Multigrid algorithms for compressible flow calculations. In W. Hackbusch and U. Trottenberg, editors, *Lecture Notes in Mathematics, Vol. 1228*, pages 166–201. Proceedings of the 2nd European Conference on Multigrid Methods, Cologne, 1985, Springer-Verlag, 1986.

[48] Jameson A. and J. Vassberg. Studies of alternative numerical optimization methods applied to the brachistochrone problem. In *Proceedings of OptiCON'99*, Newport Beach, CA.

[49] J. L. Lions. *Optimal Control of Systems Governed by Partial Differential Equations*. Springer-Verlag, New York, 1971. Translated by S.K. Mitter.

[50] A. Jameson. Aerodynamic design via control theory. *Journal of Scientific Computing*, 3:233–260, 1988.

[51] A. Jameson, L. Martinelli, and N. A. Pierce. Optimum aerodynamic design using the Navier-Stokes equations. *Theoretical and Computational Fluid Dynamics*, 10:213–237, 1998.

[52] A. Jameson, N. Pierce, and L. Martinelli. Optimum aerodynamic design using the Navier-Stokes equations. *AIAA paper 97-0101*, January 1997.

[53] A. Jameson and L. Martinelli. Aerodynamic shape optimization techniques based on control theory. In *Course on "Computational Mathematics Driven by Industrial Applications"*, Martina Franca, Italy, 1999. Fondazione CIME International Mathematical Summer Center.

[54] 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, Jan. 1986.

[55] A. Jameson. Optimum aerodynamic design using control theory. *Computational Fluid Dynamics Review*, pages 495–528, 1995.

[56] J. Gallman, J. Reuther, N. Pfeiffer, W. Forrest, and D. Bernstorf. Business jet wing design using aerodynamic shape optimization. *AIAA paper 96-0554*, 34th Aerospace Sciences Meeting and Exhibit, Reno, Nevada, January 1996.

Design Conditions		Initial	Three Point Design	
Mach	C_L	C_D Original	C_D Redesign	C_D Reduction (%)
0.85	0.38	0.0071	0.0064	9.8
0.85	0.42	0.0086	0.0077	10.4
0.85	0.46	0.0106	0.0095	10.3

Table 3: Pressure Drag Reduction for Multipoint Design.

Design	Conditions	Initial	Single Point Design	Three Point Design
Mach	C_L	Relative C_D	Relative C_D	Relative C_D
0.81	0.35	1.00257	0.85003	0.85413
0.82	0.30	1.00000	0.77350	0.77915
0.83	0.25	1.08731	0.81407	0.76836

Table 4: Pressure Drag Reduction for Single and Multipoint Designs.

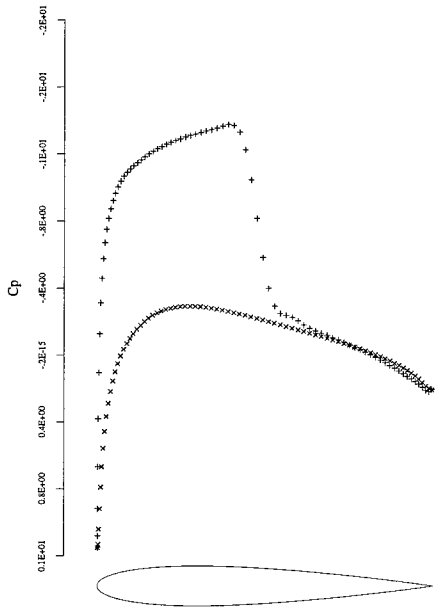

7a: C_p after Zero Design Cycles.
Design Mach 0.75, $C_l = 0.5008$,
$C_d = 0.0225$.

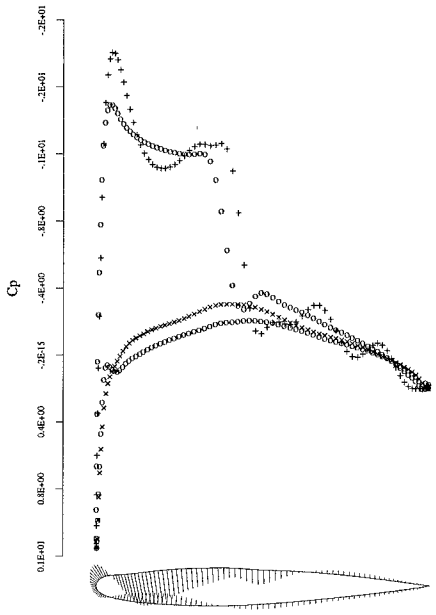

7b: C_p after One Design Cycle.
Design Mach 0.75, $C_l = 0.4841$,
$C_d = 0.0185$.

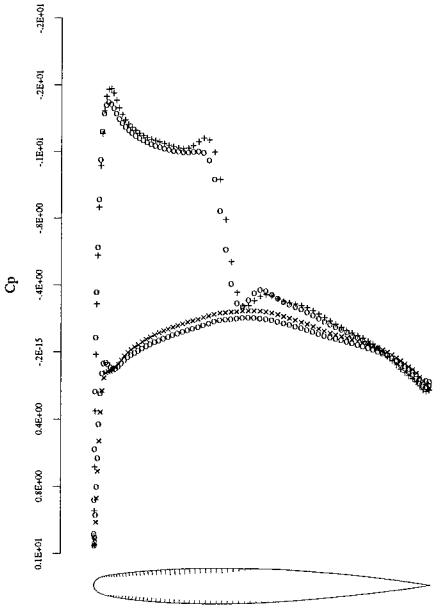

7c: C_p after Five Design Cycles.
Design Mach 0.75, $C_l = 0.4994$,
$C_d = 0.0148$.

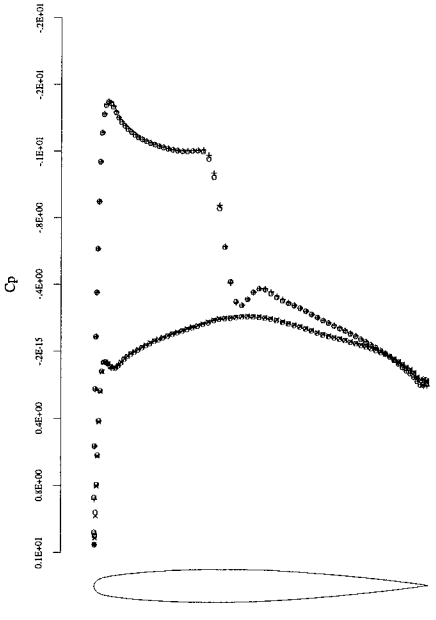

7d: C_p after Twenty five Design Cycles.
Design Mach 0.75, $C_l = 0.5007$,
$C_d = 0.0118$.

Figure 7: Inverse Design of an ONERA Airfoil. The vectors on the airfoil surface represent the direction and magnitude of the gradient.

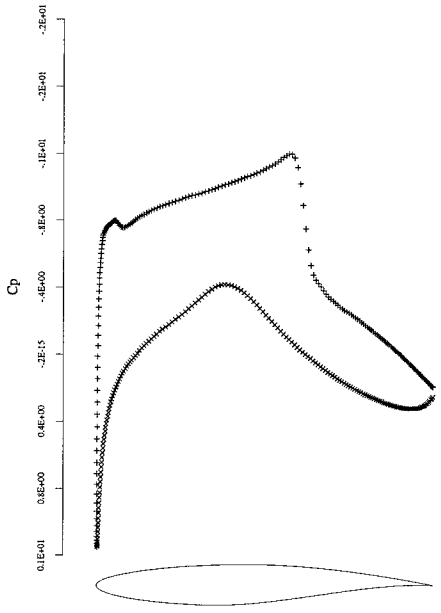

8a: C_p after Zero Design Cycles.
Design Mach 0.75, $C_l = 0.6450$,
$C_d(pressure) = 0.0091$,
$C_d(viscous) = 0.0056$.

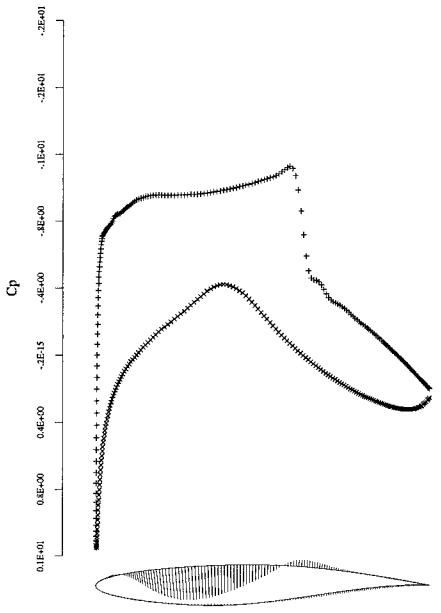

8b: C_p after One Design Cycle.
Design Mach 0.75, $C_l = 0.6512$,
$C_d(pressure) = 0.0066$.
$C_d(viscous) = 0.0057$.

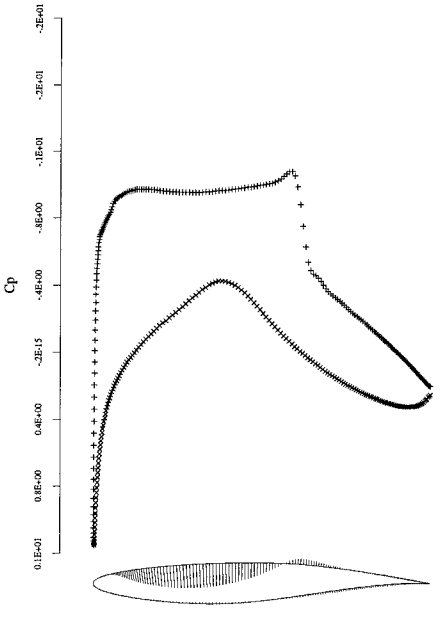

8c: C_p after Two Design Cycles.
Design Mach 0.75, $C_l = 0.6510$,
$C_d(pressure) = 0.0054$,
$C_d(viscous) = 0.0057$.

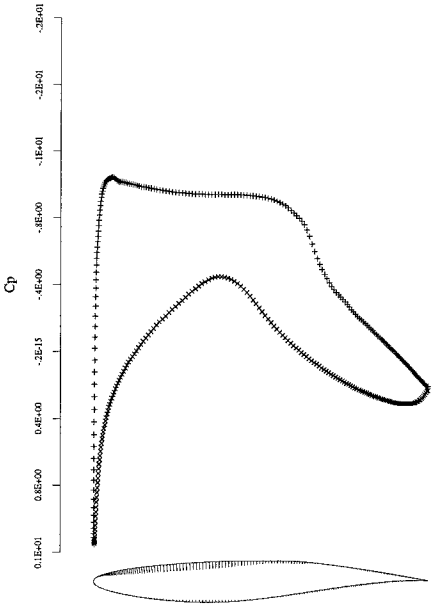

8d: C_p after Ten Design Cycles.
Design Mach 0.75, $C_l = 0.6460$,
$C_d(pressure) = 0.0041$.
$C_d(viscous) = 0.0058$.

Figure 8: Drag Minimization of an RAE2822 Airfoil. The vectors on the airfoil surface represent the direction and magnitude of the gradient.

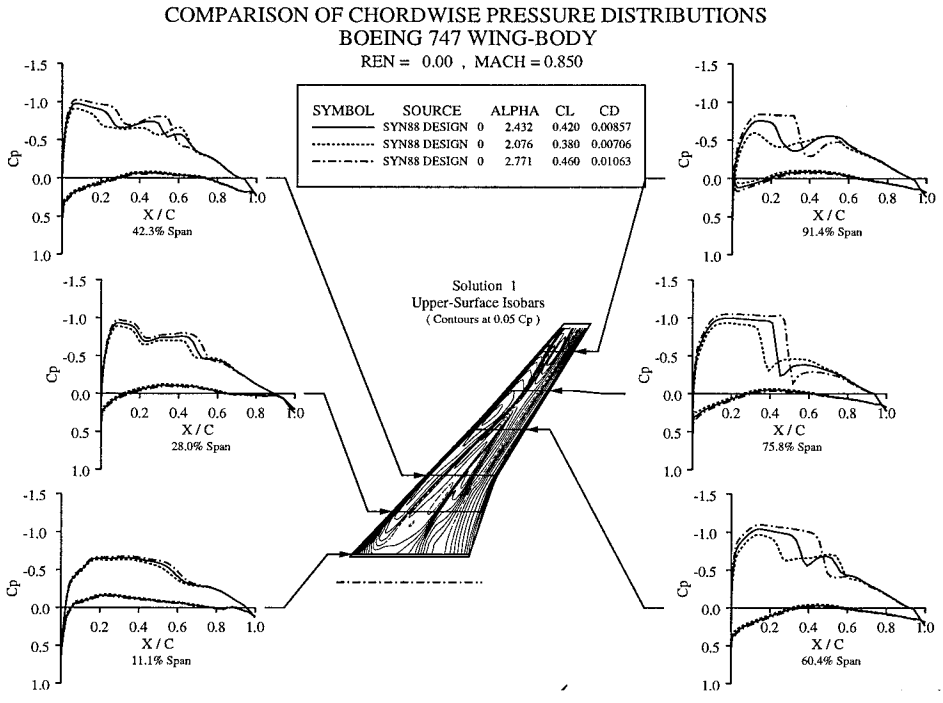

Figure 9: Pressure distribution of the Boeing 747 Wing-Body before optimization.

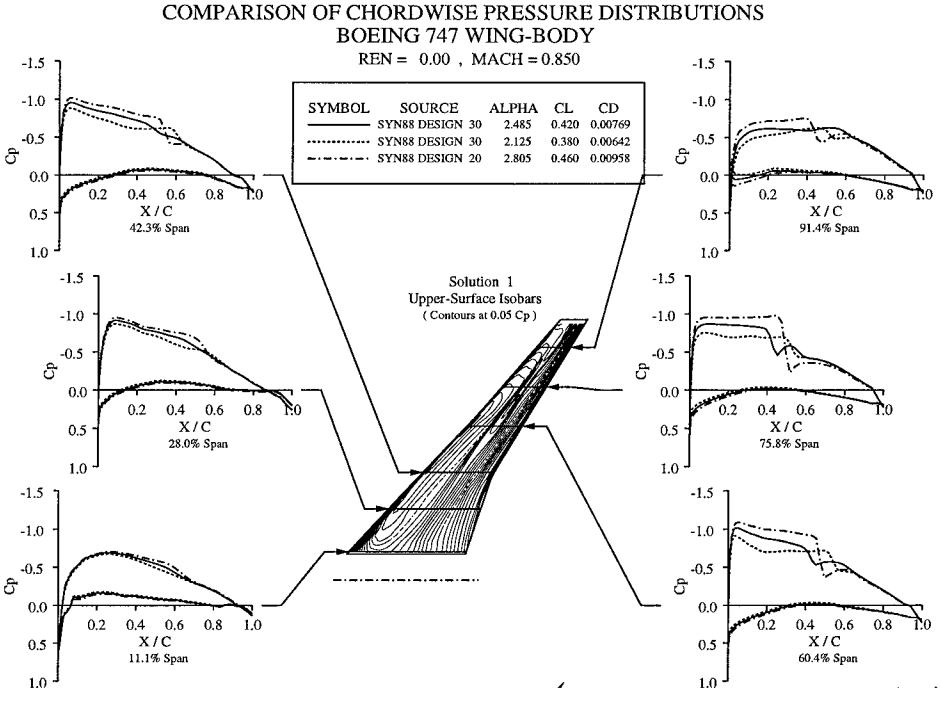

Figure 10: Pressure distribution of the Boeing 747 Wing-Body after a three point optimization.

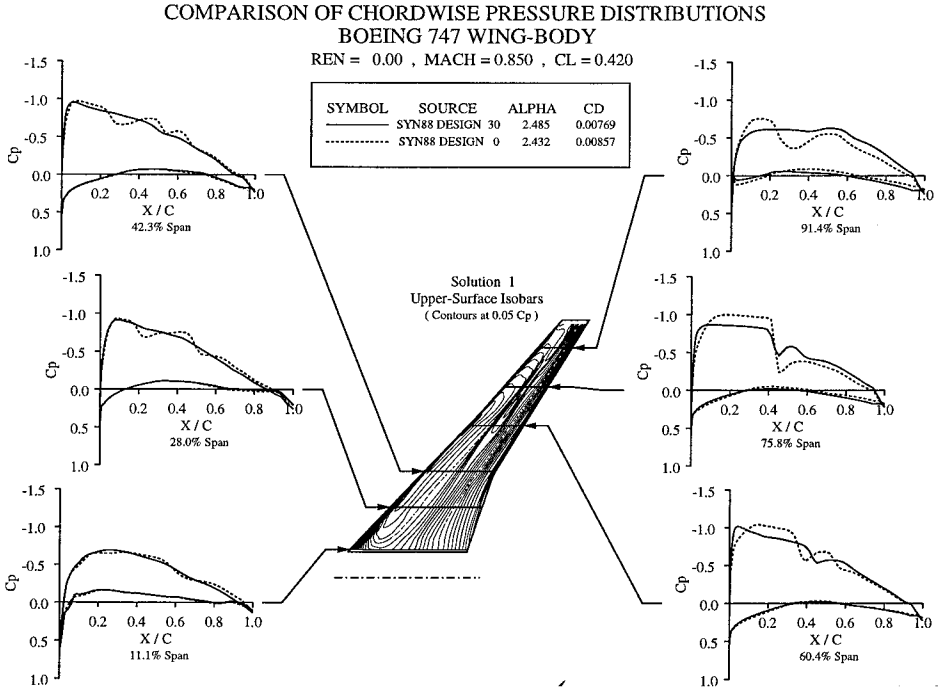

Figure 11: Comparison of Original and Optimized Boeing 747 Wing-Body at the mid design point

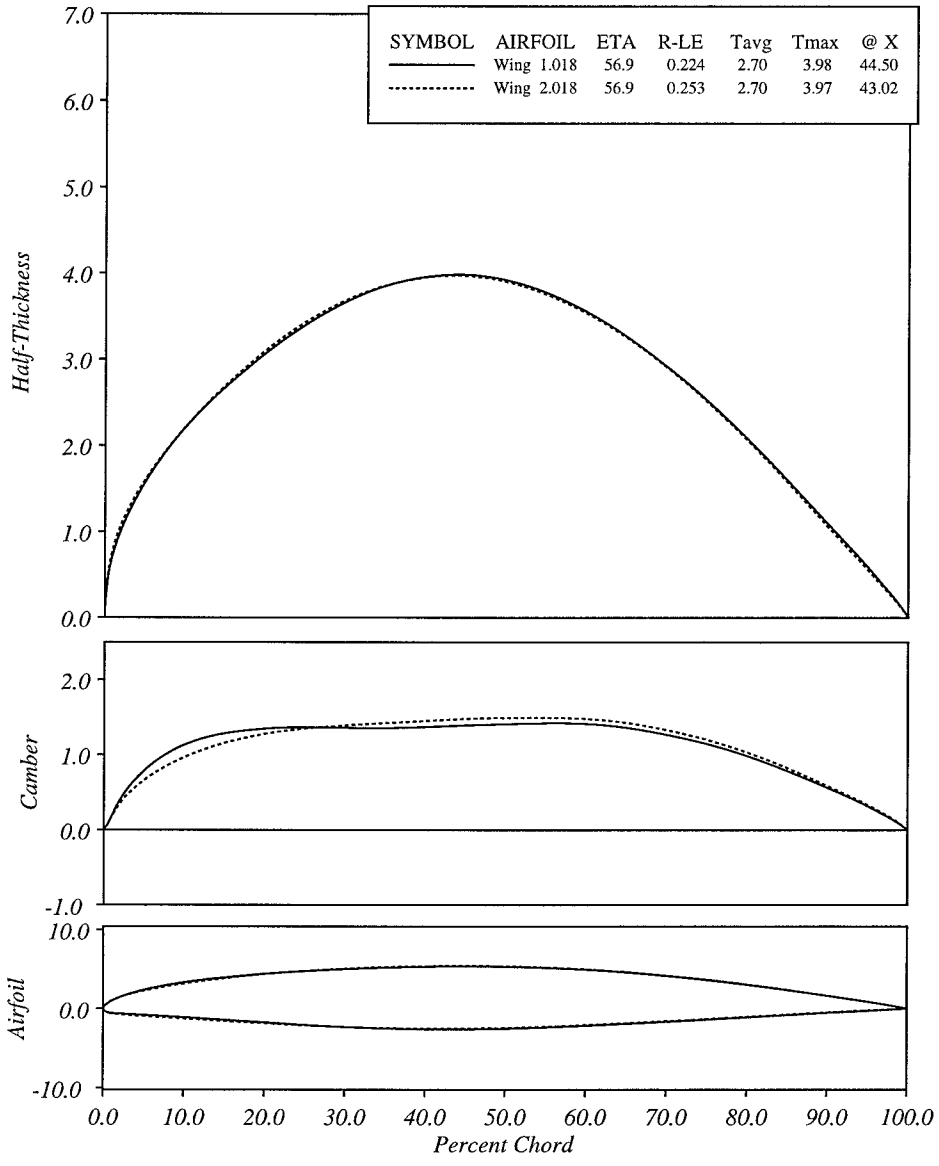

Figure 12: Original and Re-designed Wing section for the Boeing 747 Wing-Body at mid-span.

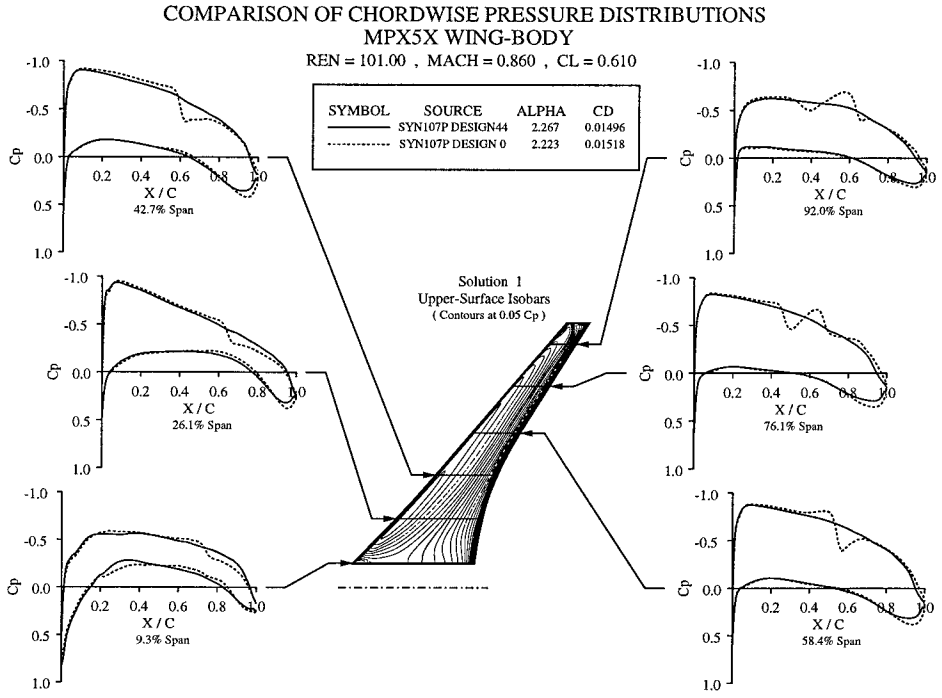

Figure 13: Pressure distribution of the MPX5X before and after optimization.

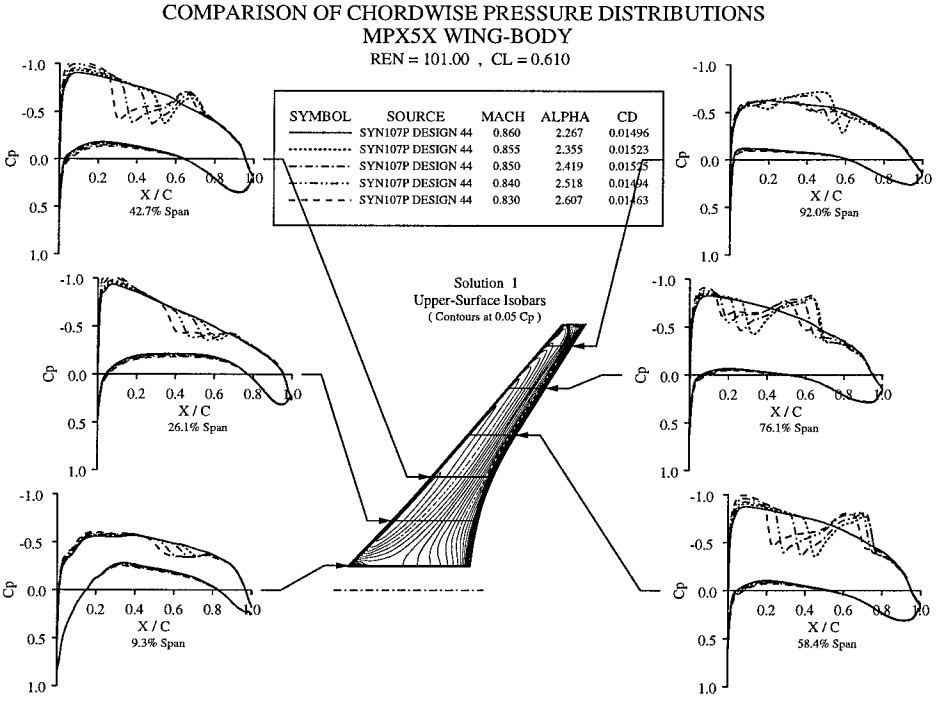

Figure 14: Off design performance of the MPX5X below the design point.

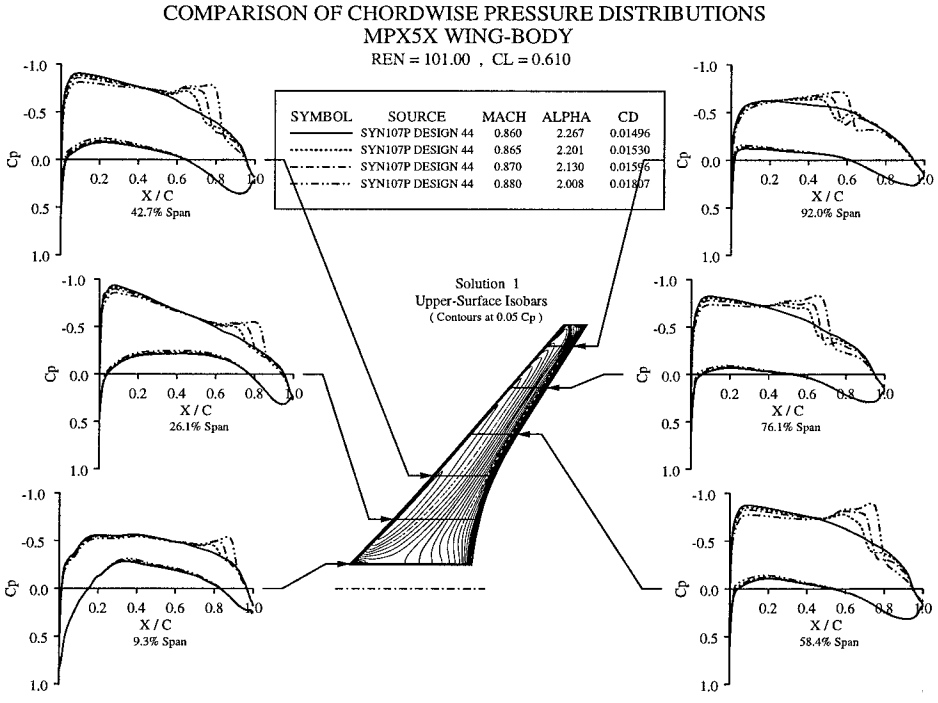

Figure 15: Off design performance of the MPX5X above the design point.

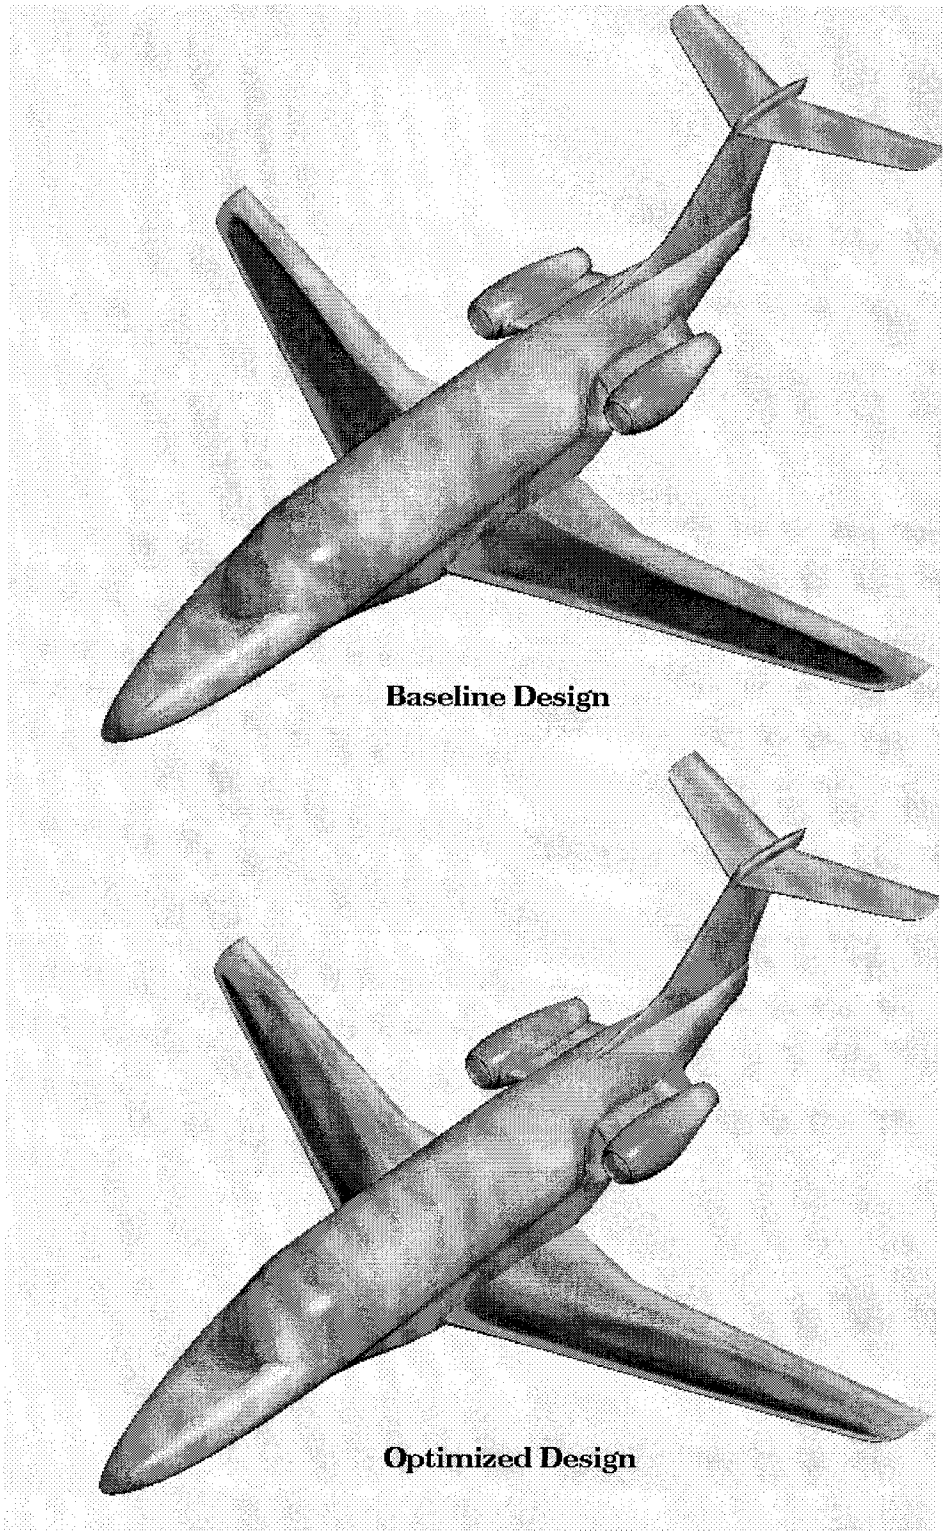

Figure 16: Geometry Surface Colored by local C_p Before and After Redesign.

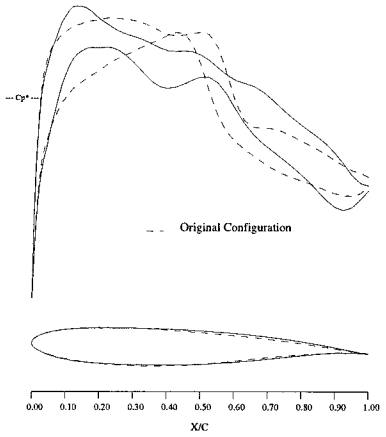

17a: span station $z = 0.190$

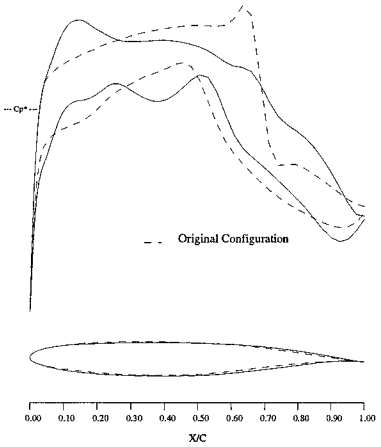

17b: span station $z = 0.475$

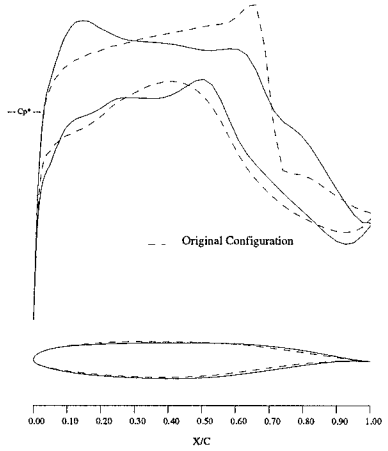

17c: span station $z = 0.665$

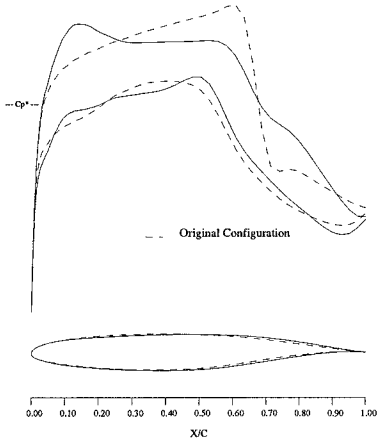

17d: span station $z = 0.856$

Figure 17: Business Jet Configuration. Multipoint Drag Minimization at Fixed Lift.
Design Point 2, $M = 0.82$, $C_L = 0.30$
90 Hicks-Henne variables. Spar Constraints Active.
– – –, Initial Pressures
———, Pressures After 5 Design Cycles.

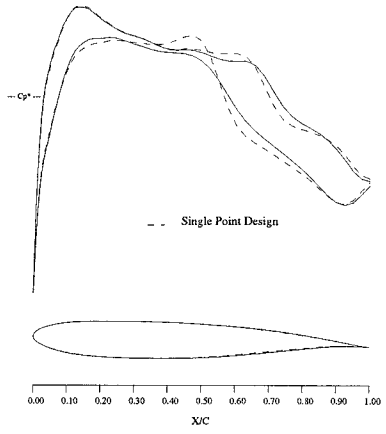

18a: span station $z = 0.190$

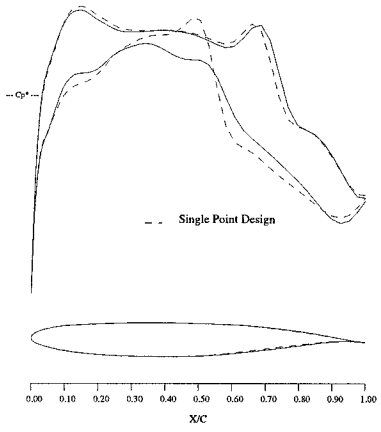

18b: span station $z = 0.475$

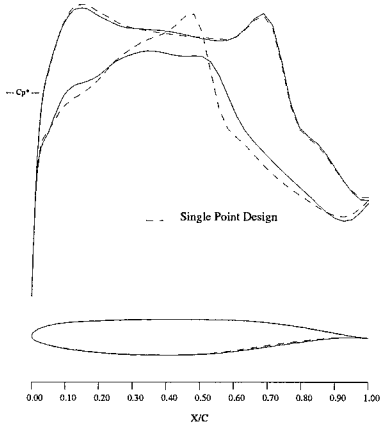

18c: span station $z = 0.665$

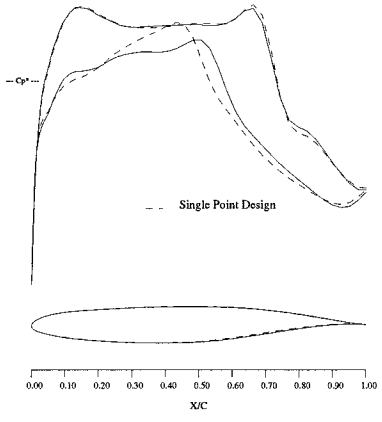

18d: span station $z = 0.856$

Figure 18: Business Jet Configuration. Single Point vs. Multipoint Drag Minimization at Fixed Lift.
Design Point 3, $M = 0.83$, $C_L = 0.25$
90 Hicks-Henne variables. Spar Constraints Active.
– – –, Single Point Design Pressures.
———, Multipoint Design Pressure.

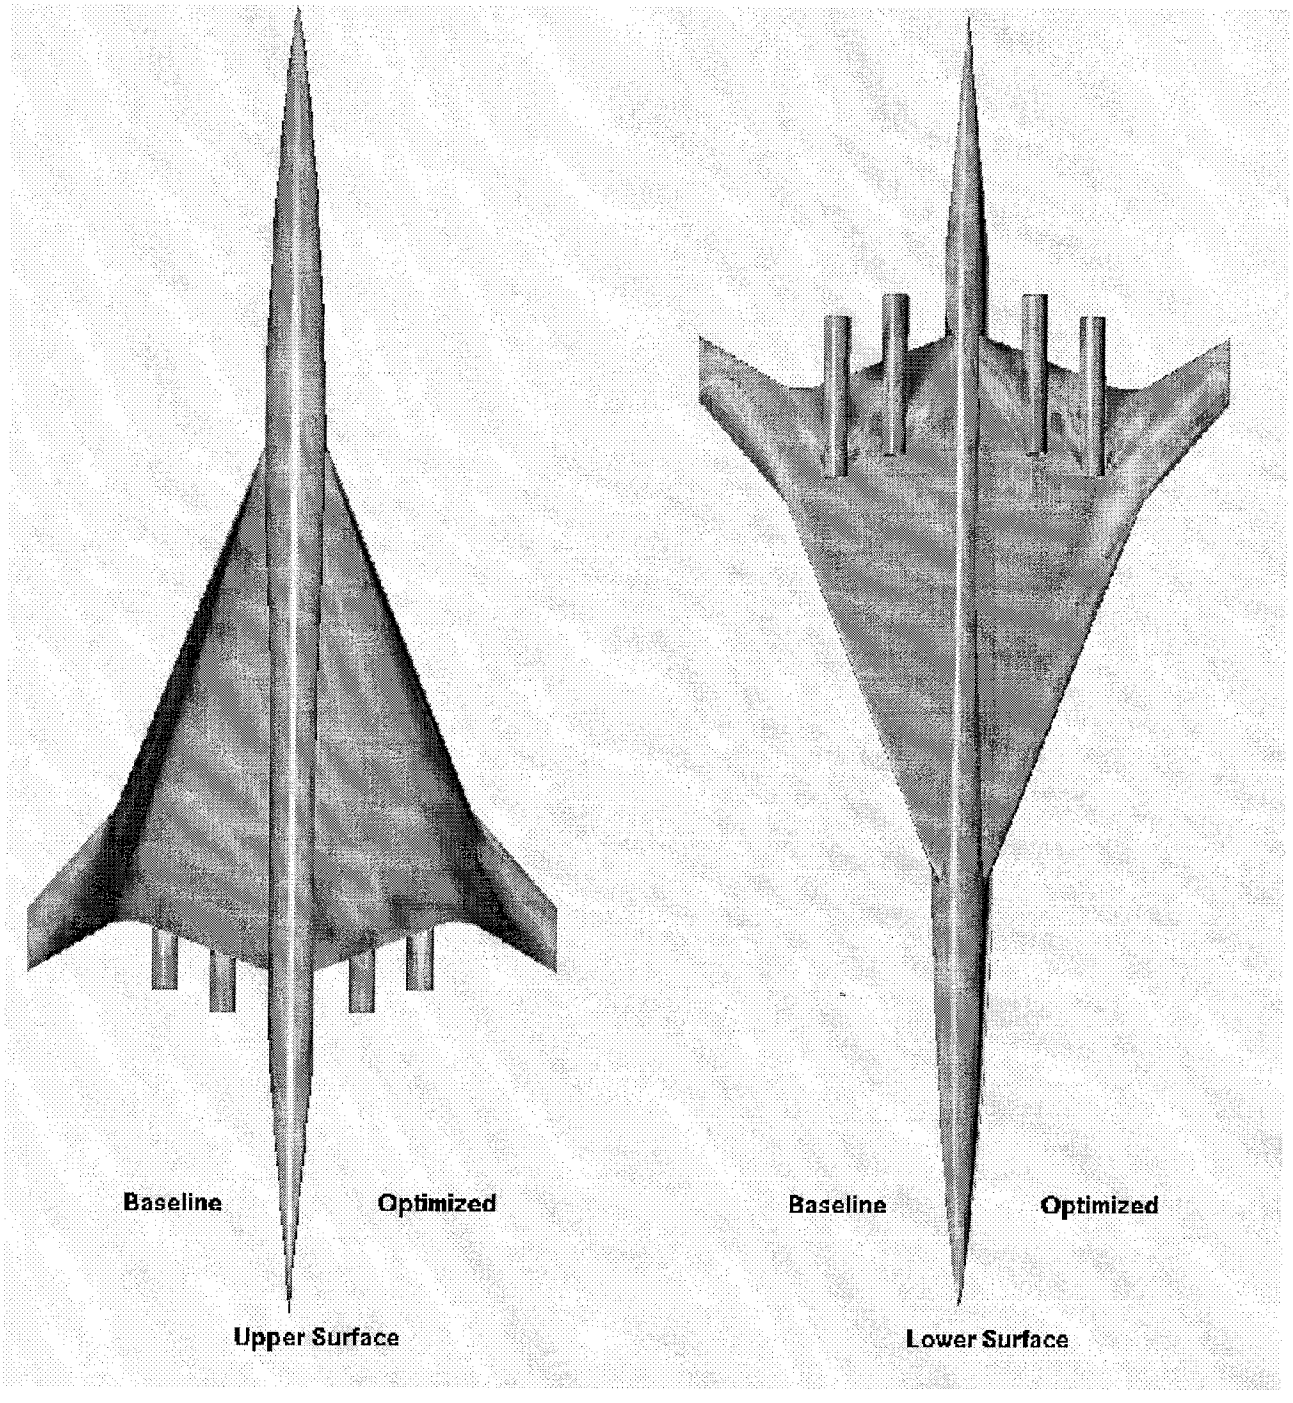

Figure 19: Supersonic Transport Configuration. Iso-C_p Contours on Upper and Lower Surfaces. Baseline and Optimized Designs. $M = 2.2$, $C_L = 0.105$.

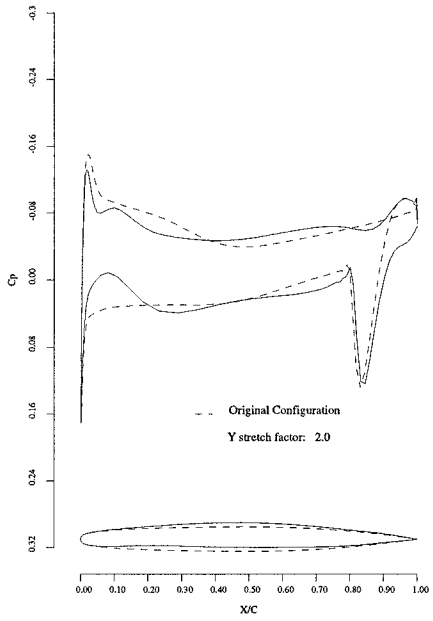

20a: span station $z = 0.194$

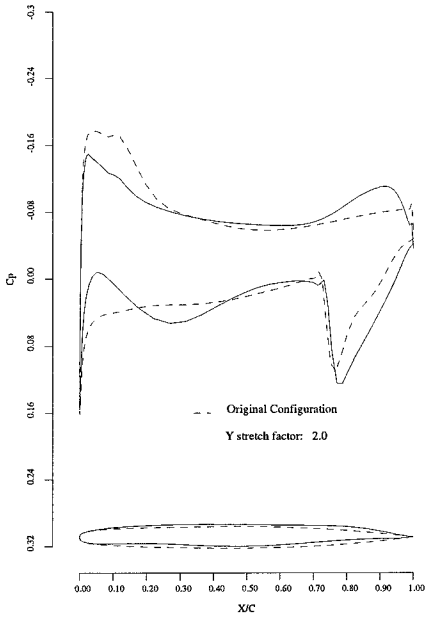

20b: span station $z = 0.387$

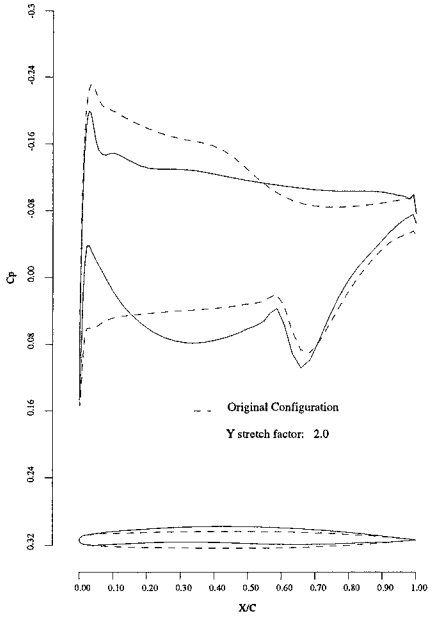

20c: span station $z = 0.581$

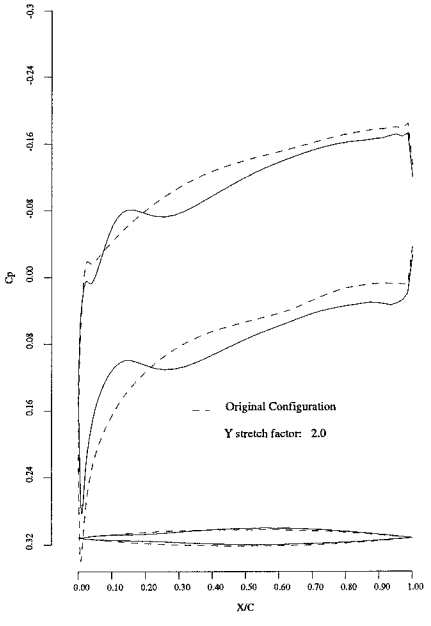

20d: span station $z = 0.775$

Figure 20: Supersonic Transport Configuration. Drag Minimization at Fixed Lift.
$M = 2.20$, $C_L = 0.105$
144 Hicks-Henne variables. Spar Constraints Active.
– – –, Initial Pressures
———, Pressures After 5 Design Cycles.