

Multidisciplinary Optimization Method for Designing Boundary-Layer-Ingesting Inlets

David L. Rodriguez*

Stanford University, Stanford, California 94305

DOI: 10.2514/1.38755

The blended wing body is a revolutionary conceptual aircraft design with rear-mounted, overwing engines. Two types of engine installations are being considered for this aircraft. The more conventional design has engines installed within nacelles that are mounted on pylons above the wing surface. The engines on the second design are partially buried in the wing and use boundary-layer-ingesting inlets mounted on the wing upper surface. For both designs, and especially for the latter, the tight coupling between the aircraft aerodynamics and the propulsion system presents a challenging design integration problem. This paper presents a design method that approaches the problem using multidisciplinary optimization. A Navier–Stokes flow solver, an engine analysis method, and a nonlinear optimizer are combined to create a design tool that properly addresses the tight coupling of the problem. The appropriate objective function and constraints are identified and an effective set of design variables is selected. The optimization method is then applied to both aircraft designs as sample problems to demonstrate the capabilities of the method and to provide some preliminary insight into advantages and disadvantages of both engine installations. Potential paths for future research in this field are also discussed.

Nomenclature

C_L	=	aircraft lift coefficient
$C_{L_{\text{cruise}}}$	=	aircraft cruise lift coefficient
C_M	=	aircraft pitching moment coefficient
DC_{60}	=	fan-face circumferential distortion descriptor
\dot{m}_a	=	required engine airflow rate computed by propulsion analysis method
\dot{m}_a^*	=	predicted engine airflow rate computed by aerodynamic analysis method
\dot{m}_f	=	fuel-burn rate
\dot{m}_{fan}	=	airflow rate through the engine fan face
p_{back}	=	inlet backpressure for controlling airflow rate
p_t	=	total pressure
p_{t_∞}	=	freestream total pressure
$p_{t_{360}}$	=	area-averaged total pressure over the entire engine fan face
$p_{t_{60,\text{min}}}$	=	minimum area-averaged total pressure over any 60 deg sector of the engine fan face
T	=	required engine thrust
T_t	=	total temperature
T_{t_∞}	=	freestream total temperature
X	=	geometric design variables used in the optimization problem
α	=	aircraft angle of attack
η_r	=	inlet pressure recovery

Introduction

BOUNDARY layer ingestion (BLI) by a propulsion system is a design concept that can improve propulsive efficiency. Ingesting the low-momentum flow in a boundary layer essentially reduces ram drag. This concept has been researched intermittently for several decades and even applied successfully to several

configurations, though never on a commercial aircraft design. Perhaps the first documented investigation of boundary layer ingestion was conducted by Smith [1]. In 1947, Smith completed an analysis of two commercial turbojet-powered aircraft designs: one with standard inlets and one with boundary-layer-ingesting inlets. In this case, the BLI inlets were actually slots that were installed over the entire span of the wing. Smith showed that the aircraft that implemented boundary layer suction on the wing had about a 30% improvement in fuel efficiency and a 7% higher optimum cruise speed. Of course, this example takes advantage of not only having engines that ingest boundary layer, but also the reduced drag from sucking boundary layer flow over the extent of the wing. In 1960, Lynch [2] of the Douglas Aircraft Company performed a wake momentum analysis on a turbofan installation that ingested a typical fuselage boundary layer flow and obtained very different results. Lynch predicted only a 3% improvement in propulsive efficiency, accompanied by a 6–10% decrease in maximum engine thrust. Lynch's analysis also depended on the assumption of minimal inlet losses, which may not necessarily be the case with a BLI inlet. In 1970, Douglass [3] performed an energy wake analysis on a commercial aircraft with aft fuselage-mounted engines. The inlets were mounted close enough to the fuselage to ingest the boundary layer flow. Douglass's analysis suggested up to a 10% improvement in propulsive efficiency on such a configuration when compared with the standard pylon-mounted engine installation. However, Douglass also stated that a boundary-layer-ingesting inlet installation would require a larger (and heavier) engine, and so some of this benefit may be mitigated. Smith [4] analyzed an unducted propeller ingesting wake flow off an axisymmetric body. Smith extended simple actuator disk theory to account for losses due to the radially varying velocity distributions ingested by a propulsor behind a body's wake. Smith reported that unpublished studies showed that the unducted propulsors on cruise missiles exhibit a 7% improvement in propulsive efficiency when the wake of the missile is ingested. Smith's analysis also suggested that improvements of up to 20% are possible if the wake is ingested in an ideal manner. All of these investigations dealt with different configurations using different methods of analysis and hence produced differing results. Though not a consensus, these previous analyses suggest that a properly designed BLI inlet on certain aircraft configurations can significantly improve performance. On the other hand, the concept has not been successfully applied to a commercial aircraft design because of the high risk and limited design capabilities in this field.

However, the BLI concept is being considered for the blended wing body (BWB), a passenger transport configuration [5] currently

Received 27 May 2008; revision received 16 February 2009; accepted for publication 1 March 2009. Copyright © 2009 by David L. Rodriguez. Published by the American Institute of Aeronautics and Astronautics, Inc., with permission. Copies of this paper may be made for personal or internal use, on condition that the copier pay the \$10.00 per-copy fee to the Copyright Clearance Center, Inc., 222 Rosewood Drive, Danvers, MA 01923; include the code 0021-8669/09 \$10.00 in correspondence with the CCC.

*Ph.D. Graduate, Department of Aeronautics and Astronautics; currently Senior Engineer, Desktop Aeronautics, Inc., 1900 Embarcadero Road, Suite 101, Palo Alto, CA 94303-3310; dlr@desktopaero.com. Senior AIAA Member.




Fig. 1 Artist's rendition of a blended wing body commercial aircraft design.

being researched and developed by The Boeing Company, NASA, and several universities and depicted in Fig. 1. Reference [6] presented results from three theoretical analyses of the benefits of boundary layer ingestion for this BWB aircraft configuration. All the analyses predicted an increase in propulsive efficiency of nearly 2% when BLI was used. This would be a substantial improvement for this aircraft and therefore was the impetus for the research presented here.

Because of the engine location, the BWB airframe aerodynamics and propulsion system are more tightly coupled than the same systems on a conventional aircraft with nacelles mounted on pylons, especially when BLI inlets are used. In other words, this aircraft exhibits an unusually more substantial tradeoff between propulsive and aerodynamic efficiencies. Hence, conventional design methods are unable to properly integrate these systems and fail to realize and fully exploit the advantages of BLI. State-of-the-art aerodynamic design tools (such as Navier–Stokes-based shape optimization schemes) can reduce drag, but cannot predict the tradeoff on engine performance. Likewise, engine simulators can predict the effects of ingesting boundary layer on propulsive efficiency, but cannot address the tradeoff on aerodynamic performance. Even these modern design methods ultimately fail because they treat aerodynamics and propulsion as separate disciplinary problems and do not truly address the tight coupling between these systems on this aircraft configuration. This lack of an effective propulsion integration method was the motivation for a research project at Stanford University in the late 1990s. The goal of this work was to develop and demonstrate a new inlet design method that not only provides a means to effectively design this type of propulsion system, but also helps to determine the viability of applying the BLI concept to the BWB aircraft design. The following sections introduce the inlet design method and present results from its application to two sample problems. These results provide some insight into the advantages and disadvantages of the BLI design concept.

Methodology

A BLI inlet exhibits a strong coupling between the airframe aerodynamics and the propulsion system. The more low-momentum flow the inlet ingests, the lower the aerodynamic drag on the airframe and hence the lower the thrust requirement on the propulsion system. On the other hand, ingesting this boundary layer flow also reduces the inlet pressure recovery and thereby reduces propulsive efficiency. An effective design method must be able to address this tradeoff between aerodynamic and propulsive efficiencies and therefore must be inherently multidisciplinary in approach. The inlet design tool developed for this work addresses this coupling by integrating three basic tools: an aerodynamics analysis method, a propulsion system simulator, and a nonlinear optimizer. Brief descriptions of these tools and the integration techniques follow. In later sections, the basic optimization problem, which consists of design variables, an objective, and a set of constraints, is discussed.

Methods for computing body forces, fan-face flow properties, and gradients of these quantities with respect to the design variables are also presented.

Design Tool Components

The flow through a transonic boundary-layer-ingesting inlet is very complex and can feature vortical flow, thick separating viscous layers, and strong shocks. The aerodynamic analysis for the design tool must be able to capture these complex flow features, and therefore a Navier–Stokes method is necessary. The solver currently used for this work is CFL3D [7], a cell-centered finite volume code developed at NASA Langley Research Center. CFL3D uses a third-order upwind-biased spatial discretization for the convective and pressure terms and a central-difference scheme for the shear stress and heat transfer terms. The time-advancement scheme is implicit. CFL3D is extremely versatile with its large selection of turbulence models and ability to handle multiblock, patched, embedded, and even overlapping computational meshes. This versatility along with its popularity in the aerospace industry and validated accuracy in similar applications were the primary reasons for selecting CFL3D for the inlet design tool. The turbulence model that was used for all computations is the one-equation Spalart–Allmaras [8] model. This model was chosen for its simplicity and applicability to the multiblock grid topology necessary to analyze the full BWB configuration. CFL3D with the Spalart model has also been validated for BLI inlets in low-speed flow with wind-tunnel experiments as detailed in [9]. The task of CFL3D within the inlet design tool is to provide predictions for aerodynamic forces (lift, drag, and pitching moment) along with the necessary fan-face properties (predicted airflow rate, inlet pressure recovery, and fan-face distortion levels).

The propulsion system simulator is an engine analysis code developed at NASA John H. Glenn Research Center at Lewis Field. The NEPP [10] code allows for one-dimensional, steady-state, thermodynamic analysis of gas turbine engines. The user specifies the design specifications of each engine component (such as the fan, turbine, burner, and inlet), allowing the simulator to predict the design and even offdesign performance of the engine. The NEPP code was chosen for its speed, ability to handle offdesign conditions, and wide usage in industry and government research. For the inlet design tool, the engine simulator takes the required thrust and inlet pressure recovery provided by the aerodynamics analysis and computes the fuel-burn rate and required inlet airflow rate. The BWB engines studied in this work are advanced ducted propellers. Though the bypass ratio of the engine has not been finalized, the range of bypass ratios being considered is wide; ratios as high as 22 have been considered. However, for the work presented here, the bypass ratio is fixed at 12 to keep the engine diameters reasonable on the current BWB geometry.

Developed at Stanford University, NPSOL [11] is a gradient-based nonlinear optimizer capable of handling both linear and nonlinear constraints. An advanced quasi-Newton algorithm is implemented in which each major iteration involves the solution of a quadratic subproblem based on the current objective function value, gradient, and Hessian approximation. NPSOL was chosen for its robustness and efficiency with highly nonlinear problems. The design of a BLI inlet is indeed a very difficult nonlinear problem with many constraints, as will be discussed in later sections.

The three tools discussed previously were integrated into a multidisciplinary design tool for which the architecture is depicted in Fig. 2. Taking computed information from the aerodynamics and propulsion analyses, the optimizer modifies a selected set of design variables to minimize an objective function. For all the work discussed in this paper, this objective function is the engine fuel-burn rate. This is an appropriate objective function for design problems in which the aircraft weight is assumed to remain constant; the goal is to maximize the performance of the inlet design and essentially maximize the range of the aircraft with a specified weight and engine design. For a typical application of the inlet design tool, the set of design variables includes geometric variables to modify the inlet and airframe shape, the angle of attack to control the total lift of the

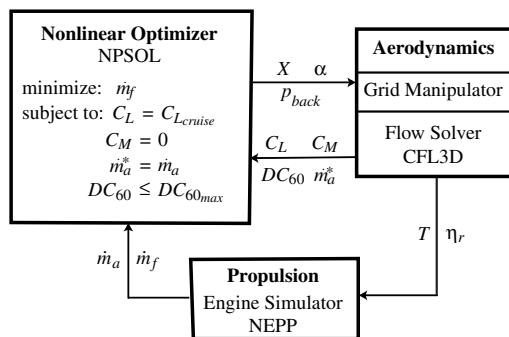


Fig. 2 Architecture of the inlet design tool.

airframe, and an inlet backpressure level to ensure airflow compatibility between the engine simulator and the flow solver.

This airflow compatibility constraint is critical to the inlet design methodology. The inlet airflow rate predicted by the flow solver must be the same as the required airflow rate that is computed by the engine analysis to make the problem consistent. Using the predicted airflow rate as a design variable can satisfy this constraint. As it minimizes the objective function, the optimizer modifies the value of the airflow rate predicted by the flow solver to ensure that it is identical to the required airflow rate computed by the engine simulator. Of course, in practice, the primary design variable used to satisfy this constraint is not the predicted airflow rate but rather a backpressure level. The boundary condition used in the flow solver to control the airflow rate through the inlet is the standard engine-outflow type; at the boundary, the static pressure is specified and the values of the other flow variables are extrapolated from inside the flowfield. To obtain a specific flow rate, this backpressure must be determined iteratively, which is extremely costly in terms of computation time and resources. By using the backpressure value itself as the design variable, this iterative process is avoided and the optimizer has direct control of the boundary condition. The optimizer then modifies the backpressure level to ensure that the airflow compatibility constraint is satisfied as the optimal design (with minimum fuel-burn rate) is sought.

Several additional nonlinear constraints are necessary to complete the problem. Because the aircraft inlet is optimized at cruise conditions, the computed lift of the aircraft must be equal to the cruise lift coefficient to ensure a level-flight condition. To trim the aircraft, a pitching moment constraint must also be enforced. Finally, because engine manufacturers set limits on fan-face distortion levels, yet another constraint must be added. For this work, the distortion figure of merit that is used is the DC_{60} descriptor, which will be further discussed in the next section. Geometric constraints on the wing and nacelle (such as those necessary due to the aircraft structural design) manifest themselves as a set of linear constraints on the geometric design variables used in the problem.

During each iteration of the optimization process, the design variables are passed from the optimizer to the flow solver, as shown in Fig. 2. After modifying the inlet/airframe shape and the computational grid according to these design variables, the flow solver computes values for required aircraft thrust and inlet pressure recovery. These values are passed to the engine simulator in which the fuel-burn rate and required airflow rate are computed and passed back to the optimizer. The lift, pitching moment, and circumferential distortion coefficients along with the predicted airflow rate are also passed back to the optimizer by the flow solver. The optimizer uses all of these computed values to determine a new set of design variables that attempt to minimize the objective function while satisfying the constraints. The process repeats until an optimal solution is found.

Force Computation

Normally, the force computation process within a computational fluid dynamics (CFD) method is straightforward: integrate the pressure and skin-friction forces on all solid walls. However, for the




Fig. 3 Representative control volume used to compute forces on a BWB aircraft.

case with an engine, the correct method is not so clear. Accurately analyzing the flow within the engine and all its parts is obviously not an option due to its complexity and therefore the engine must be modeled. The typical way of modeling the engine is to introduce a flow exit plane in the inlet and a flow entrance plane within the engine nozzle. This engine modeling is represented in Fig. 3. Special boundary conditions are applied at the engine entrance (fan face) and exit (exhaust plane) to properly set the flow properties through these boundaries. Computing forces on this configuration is not straightforward due to the fact that some flow exits the computational domain at the fan face and some flow enters at the engine exit plane, but it is still possible. Consider the control volume shown in Fig. 3 that includes all the solid surfaces and the engine fan face and exhaust planes. The pressure, shear forces, and momentum flux at the control volume boundaries must all be integrated to compute the total force acting on this control volume. Note that the momentum flux through the solid walls is zero and the shear forces on the fan face and engine exhaust planes are negligible.

For steady, unaccelerated cruise flight, the total integration of these forces on this control volume should be zero. However, in the inlet design method, the engine simulator requires a gross thrust value that balances the airframe and inlet ram drag in the design problem. To compute the gross thrust of the engine using the CFD method, the force due solely to the flow entering the computational domain at the engine exhaust plane must be determined. This means the momentum flux and pressure must be integrated on this boundary face. However, because the total force integration is zero, the force acting on all solid surfaces and the fan face is equal (except in sign) to the gross thrust and is the total drag (including any ram drag) of the aircraft. Therefore, the gross thrust requirement of the engine can also be computed by integrating the pressure and skin friction on all solid surfaces and the pressure and momentum on the fan face.

But why not just compute the gross thrust by integrating the forces on the exhaust plane? The reason is reduce the complexity of the problem. To estimate the gross thrust of the aircraft in this manner requires that the flow properties at the engine exit plane are indeed that which would provide a total integrated force of zero. Normally, this would mean iteratively determining the exit flow total pressure and total temperature until the total integrated force is zero. This iteration could be eliminated by introducing more design variables (probably one for each engine if the total temperature were set by engine operating conditions) and introducing yet another constraint. This of course makes the problem more complex.

To determine if this added complexity is really necessary, a simple numerical experiment was conducted on a two-dimensional version of this problem. The total pressure of the exhaust flow was set to a wide range of values of flight and engine operating conditions. The various drag components were computed and compared to determine if the drag (which equals the required gross thrust) of the aircraft was significantly affected by varying the engine exhaust conditions. The results of this experiment are given in Table 1 and suggest that it is indeed possible to leave the exit-plane conditions constant and still

Table 1 Comparison of BWB drag for varying exhaust-plane boundary conditions

p_t/p_{t_∞}	T_t/T_{t_∞}	Total force coefficient	Drag coefficient
1.8	1.5	0.0381	0.0279
2.3	1.5	0.0066	0.0280
3.0	1.5	-0.0338	0.0282

effectively perform the optimization, especially because only small changes in the aircraft drag are expected.

Computation of Fan-Face Total Pressure Recovery

The method of computing pressure recovery is always an issue of much debate. The propulsion analysis method requires a single value that represents the total pressure losses at the fan face. Because the total pressure loss can vary over the fan-face area, the only possibility is some sort of integrated quantity. When the total pressure distribution is not constant over the fan face, the correct integration scheme is not obvious. Many methods exist, including area-weighted integration, mass-flow-weighted integration, momentum-weighted integration, and even various forms of energy-weighted integration. All these methods can give very different results for a fan face with a widely varying total pressure distribution, such as that found in a BLI inlet, for example. The method chosen for this inlet design method is based on the work by Livesey [12]. This integration method computes pressure recovery by assuming a constant entropy flux and weighting the integration appropriately:

$$\eta_r = \frac{1}{p_{t_\infty}} \exp \left[\frac{1}{\dot{m}_{\text{fan}}} \int_{\text{fan}} \{ \ell_n(p_t) \} d\dot{m} \right] \quad (1)$$

Note that momentum is not conserved with this method. Although other methods that do conserve momentum exist, these methods either do not conserve other important properties (such as mass flow rate) or involve some sort of conservative mixing process that reduces the total energy of the flow. Though it is still unclear whether the method of Eq. (1) is the correct method (if there even is a correct method), Livesey [12] provides some valid arguments for selecting this formula. This integrated quantity is therefore the value of pressure recovery that is sent to the engine simulator in the inlet design method.

Distortion Computation

Although there are many types of descriptors for rating levels of fan-face distortion, the BWB program adopted the DC_{60} definition in preliminary design work. This descriptor gives levels of circumferential distortion and is given by

$$DC_{60} = \frac{p_{t_{360}} - p_{t_{60,\min}}}{p_{t_{360}}} \quad (2)$$

In Eq. (2), $p_{t_{360}}$ is the area-averaged total pressure on the entire fan face. The other area-averaged total pressure term, $p_{t_{60,\min}}$, is computed on the 60 deg sector of the fan face with the minimum area-averaged total pressure. For this definition, pools of low total pressure far from the center of the fan face induce high distortion levels.

Gradient Computation

Two methods of computing gradients (sensitivities) for the optimizer are presented in this section: finite differencing and the complex variable method. Although complex variables were used for all optimization work in this paper, the finite difference method is presented for comparison.

Finite Differencing

When analytic gradient computation is not possible or practical, the simplest and most common method for computing sensitivities is the method of finite differences. The method is straightforward and easy to implement. Perturbations in each design variable are made and the perturbed objective and constraint values computed. A forward-difference gradient computation (based on a truncated Taylor series expansion) is given by

$$f'(x) \approx \frac{f(x + \Delta x) - f(x)}{\Delta x} \quad (3)$$

which is well known to be first-order-accurate in Δx . For N design variables, the forward finite difference method requires $N + 1$

function evaluations to compute all the necessary gradients. Ideally, the step size is chosen to be very small to provide the most accurate gradient possible. However, this is not always possible, due to noise from the analysis method or other inaccuracies.

Whereas the scheme is easy to implement, the finite difference method can be problematic when the accuracy of the objective function is not satisfactory. For the gradient computation to be accurate, the objective function computation must be even more accurate, because the subtraction between two large numbers in Eq. (3) gives a relatively small number usually comparable with Δx . Unfortunately, CFD codes are notorious for not converging to highly accurate solutions, especially when viscous flow is present. Even if the CFD code could be converged to high accuracy, such convergence is extremely costly in terms of computation time, which is highly undesirable in an optimization scheme. To be effective and efficient, the step size must be chosen to filter out the inherent error of the CFD analysis and yet still produce accurate gradients. A too-small step size allows noise in the CFD solution to corrupt the gradient, whereas a too-large step size will produce inaccurate results for functions for which the gradients vary greatly with the design variable.

This strong dependence on step size is illustrated in the following example. The clean BWB wing (no engines) with a relatively coarse computational grid was used to compute the gradient of the lift coefficient with respect to angle of attack. The convergence of the gradient is computed by applying the finite difference after each iteration of both solutions of the CFD code and hence in lockstep. Three different step sizes were used to fully illustrate the importance of the selection of step size, and the results are shown in Figs. 4a–4c. Note that as the step size gets smaller to try to improve accuracy, the inaccuracies in the CFD code cause the gradient computation to fail. Therefore, to make finite differencing a productive method, the step size for each design variable in a problem must be chosen appropriately, which can be very difficult and tedious.

Complex Variables

Computing gradients with complex variables is not a new concept. Lyness [13] and Lyness and Moler [14] proposed this concept in 1967. The method is based on a Taylor series expansion of a function that is perturbed in the imaginary plane:

$$f(x + i\Delta x) = f(x) + i\Delta x f'(x) - (\Delta x)^2 f''(x) - i(\Delta x)^3 f'''(x) + \dots \quad (4)$$




Fig. 4 Lift gradient convergence on BWB wing using a–c) finite differences and d) complex step method.

Rearranging the terms of Eq. (4) to group the real and imaginary parts and factoring out the imaginary step $i\Delta x$ gives

$$f(x + i\Delta x) = [f(x) - (\Delta x)^2 f''(x) + \dots] + i\Delta x[f'(x) - (\Delta x)^2 f'''(x) + \dots] \quad (5)$$

The real part of Eq. (5) is the value of the function to order $(\Delta x)^2$ and the imaginary part is the first derivative of the function also to order $(\Delta x)^2$. What this allows is the computation of the objective function and a gradient at the same time and to the same accuracy. Because there is no large number subtraction required to compute the gradient, the step size can be as small as possible without needing the extreme accuracy required in the finite difference method. This property makes the complex step method very attractive because there are virtually no step size issues that must be considered as with the finite difference method.

Applying this method to a flow solver is rather straightforward. The solver code is modified to use complex instead of real arithmetic. For demonstration, a complex variable version of CFL3D was used on the clean BWB wing computational grid of the preceding section, and the gradient of lift with respect to angle of attack was once again computed. The convergence history of the gradient using the complex variable method is shown in Fig. 4d. Note that the convergence of the gradient is virtually identical to that of finite differences with a proper step size (Fig. 4a) for this case.

On the programming side, the complex variable method is very attractive because changes can be made to the code directly. Other methods, such as automatic differentiation [15], require that a code be postprocessed each time a change is made. The resulting differentiated code is usually very difficult to read and can be highly inefficient computationally.

On the other hand, the gradient computation using complex variables does come at the price of significantly more computation time. A CFD code running with complex arithmetic requires about 3 times as many floating-point operations as the standard code and it still has to be run once for each design variable. Although computing the objective and gradients only takes N solutions, each solution takes 3 times longer than the solutions in the finite difference scheme. Then again, the solutions in the finite difference scheme must be very tightly converged to obtain gradients with any kind of accuracy, due to the subtraction errors. With no subtraction errors present, the complex variable solutions often in practice do not have to be as tightly converged to produce gradients of the same accuracy as the finite difference method, though this does not seem to be the case in Fig. 4. This is because the lift coefficient is easily converged (even with finite differences) to high accuracy in the clean BWB case. Based on the author's experience with more complicated flows, however, the complex variable method can converge gradients more quickly in terms of iterations (and hence computation time) than the finite difference method.

There are other tricks that can hasten the computation of gradients using the complex step method. Although these techniques have not been proven rigorously, the general experience of the author during the course of the development and application of the inlet design method has shown them to be effective. First of all, the complex part of the solution will usually converge in fewer time steps if the real part of the solution is already converged. Therefore, one useful technique is to first run the standard CFD code to obtain the real part, and then run the complex step method to compute gradients. This has been found to speed up the entire process tremendously. Also, the imaginary parts of the solutions for each design variable can be saved for subsequent design iterations, reducing convergence times even more, particularly when the optimizer makes only small changes in the design variables, such as when the method has nearly reached the optimal solution.

The Design Problems

To demonstrate the inlet optimization method, two design problems were identified on two unique engine installations on the BWB. The first configuration had its engines mounted inside nacelles

on vertical pylons above the upper surface of the aircraft. The second configuration had engines that were partially buried in the upper wing surface and used boundary-layer-ingesting inlets. The purpose for analyzing and optimizing both configurations was twofold. The first and foremost objective was to demonstrate the capability and versatility of the inlet design method. The second was to provide some preliminary insight into the advantages and disadvantages of boundary layer ingestion over the more conventional engine installation.

To keep computation times manageable in a research environment with limited computational resources, the problems were kept as simple as possible and the computational grids for the aerodynamic analyses were relatively coarse. Naturally, the coarseness of the meshes can and does affect the accuracy of the analyses. However, to reiterate, the primary purpose of these problems was to demonstrate the capabilities and effectiveness of the inlet design method, not to perform a proper redesign of the BWB engine installation. To at least address the accuracy of the final results from a coarse-grid optimization, the baseline and optimized inlet designs were also analyzed with much finer meshes to verify the relative performance improvement between these designs.

Another restriction that was necessary to keep computation times manageable was limiting the number of design variables. For gradient-based quasi-Newton optimization schemes, the number of design cycles necessary to converge to an optimum is approximately the number of design variables minus the number of constraints. Using the complex variable method to compute gradients requires one analysis for each design variable. Therefore, the total number of analyses needed to complete one optimization problem with few constraints is on the order of the number of design variables squared. As a result, the number of design variables used for the optimization problems presented here ranged from 16 to 22.

Ideally, for either problem, both the centerline and outboard nacelles on the BWB would be optimized simultaneously to obtain the optimal engine installation. Because the number of design variables was limited, however, this was not always possible for the problems presented here. Another unfortunate side effect of limiting the number of design variables is that the design space is consequently limited and even an approximation of the true optimal design may not have been attainable. Of course, there is no way to verify that the final result was even close to the true optimal design without performing an optimization with a large number of design variables. This could indeed be done with either a vast amount of computational resources (and patience) or by significant improvements in the efficiency of the gradient calculation method. Such tasks are left for future work and are discussed in a later section.

Some minor variations to the ideal constraints were also necessary for these optimization problems. As in the ideal problem, the fuel-burn rate is minimized while maintaining the cruise lift. Although the aircraft would normally need to satisfy a trim constraint ($C_M = 0$),




Fig. 5 The 3-view of the podded BWB configuration.

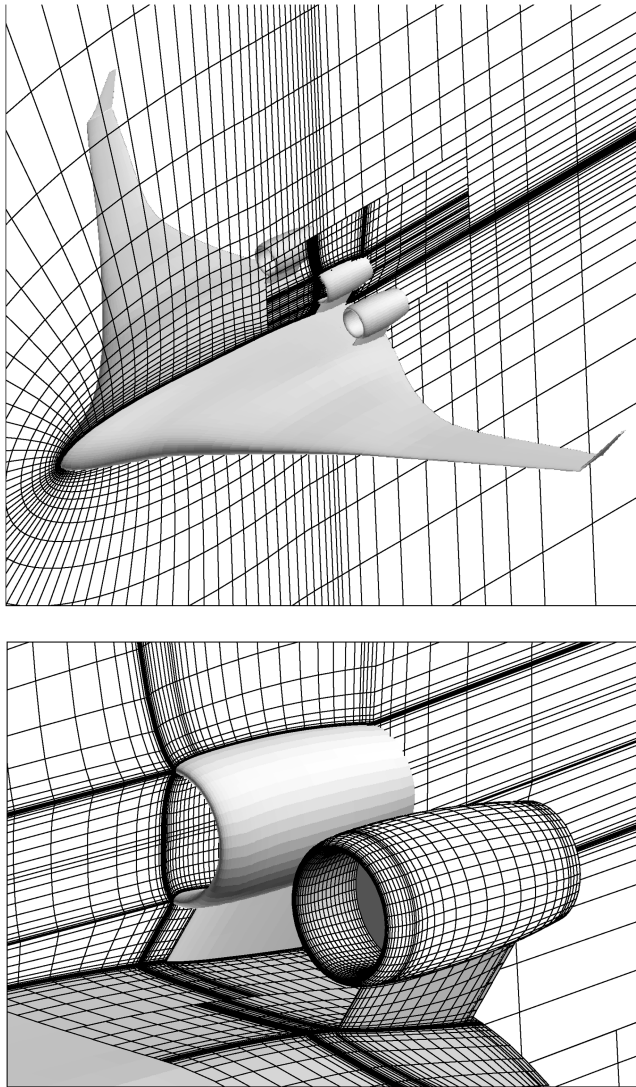


Fig. 6 Typical coarse computational grid on the podded BWB configuration.

for these problems, the pitching moment is constrained to not become worse (further from trimmed) than the wing-alone configuration. The reason for this constraint modification is that the BWB wing design used for these optimization problems is not a trimmed design. Therefore, the moment is not allowed to become worse so that no large trim drag penalty is introduced. Likewise, the fan-face distortion is constrained to be no worse than a given value.

Podded Nacelle Optimization Problem

The baseline podded configuration is shown in Fig. 5 in 3-view format. The pylons were designed to have a thickness reasonable for structural considerations and to have symmetric airfoil sections. The engines are canted at a 5 deg angle of attack from the wing reference plane and have no cant in the yaw direction. The nacelle wall thickness and the leading-edge radius of curvature of the inlets were selected using conventional preliminary design methods used by McDonnell Douglas Corporation.

A sample view of the computational grid used for the optimization work on this configuration is given in Fig. 6. The computational grid was created by removing a contiguous block of a baseline wing grid and then inserting nacelle grids for which the outer boundaries conform to the removed grid block. The interfaces between the wing and nacelle grids are not point-matched, and therefore the patching algorithm included in CFL3D was used for boundary conditions. The total number of grid points is just over 476,000. Obviously, this is a

very coarse grid for such a complex configuration, but the small grid size was necessary for computational efficiency.

Objective Function and Constraints

The objective function for this optimization case is the fuel-burn rate. The lift and airflow compatibility constraints are necessarily enforced in this problem. Of course, in this case, there are two airflow compatibility constraints because there are two engines (centerline and outboard). Also enforced is a constraint that ensures that the pitching moment coefficient does not become lower than that of the baseline wing alone (-0.17). No constraint on fan-face distortion is enforced, because it is simply not necessary with this inlet design.

Design Variables

To keep computation times manageable with the available resources, only 18 design variables (listed in Table 2) were selected. The angle-of-attack and inlet backpressure variables are necessary to help enforce the lift and airflow compatibility constraints. A wing twist variable (washout) was also added to help enforce the pitching moment constraint. This variable adds a linear twist variation to the outboard section of the wing, which is effective in controlling the pitching moment of a swept wing. Only the outboard wing twist is affected so that the cabin section is not warped.

The remaining variables control the shape and orientation of the nacelles themselves. These include the inlet length, inlet camber (which controls the inlet area distribution and outer shape of the nacelle), the nacelle pitch and yaw, the pylon height, and the pylon camber. Note that no thickness variables were included. Because the baseline nacelles were designed to have the minimum thickness and inlet lip radius, thickness variables were deemed mostly ineffective for this case. This may not be entirely true, because thickness variables would allow for the independent control of inlet area distribution and the outer nacelle contours, which is important for controlling the shock strength and location on the nacelle outer surface. However, computational efficiency dictated the omission of this type of variable. A more detailed description of each of these design variables is given in [16], but is omitted here for conciseness.

Results and Discussion

Two optimization runs were completed on the podded BWB design. For the first case, only the angle-of-attack and inlet backpressure design variables are used along with the lift and airflow compatibility constraints. This optimization determined the performance of the initial podded BWB configuration at cruise conditions, thus providing a baseline for comparison with the optimization result that used all of the design variables. The constraints were satisfied very quickly (3 design steps), as expected in such a low-dimensional design problem. The performance of this baseline configuration is summarized in Table 3.

For the second case, all 18 design variables of Table 2 were used to minimize the fuel-burn rate of the baseline podded BWB

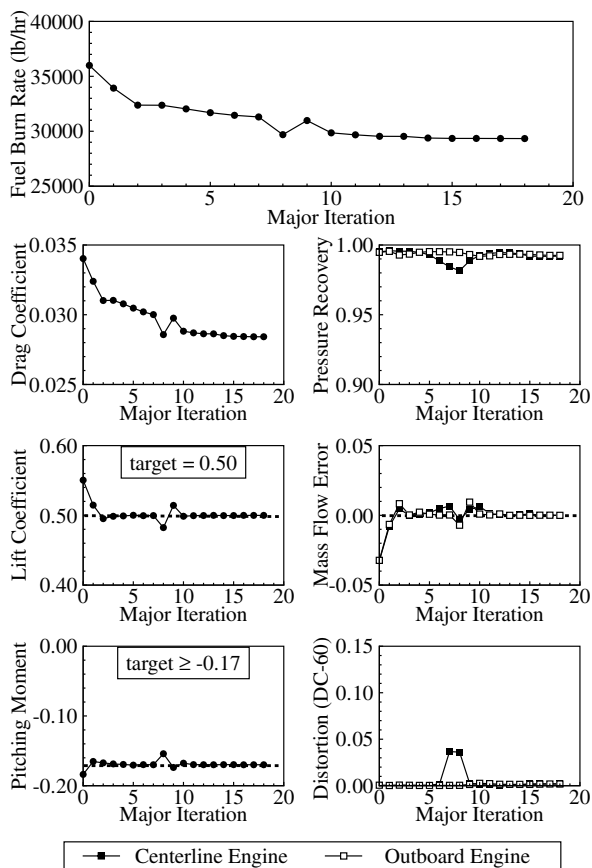

Table 2 Design variables for podded BWB optimization

Variable	Type	Location
1	Angle of attack	Entire aircraft
2	Inlet backpressure	Centerline engine fan face
3	Inlet backpressure	Outboard engine fan face
4	Linear wing twist	Outboard wing section
5	Inlet length	Centerline nacelle
6–7	Inlet/nacelle camber	Centerline nacelle
8	Pylon height	Centerline nacelle/pylon
9	Nacelle pitch	Centerline nacelle
10	Inlet length	Outboard nacelle
11–14	Inlet/nacelle camber	Outboard nacelle
15	Pylon height	Outboard nacelle/pylon
16	Nacelle pitch	Outboard nacelle
17	Nacelle yaw	Outboard nacelle/pylon
18	Pylon camber	Outboard pylon

Table 3 Performance comparison of baseline and optimized podded BWB configurations

Performance parameter	Baseline	Optimized
Fuel-burn rate, lb/h	32467	29336
Lift coefficient	0.5000	0.5000
Pitching moment coefficient	-0.1676	-0.1700
Drag coefficient	0.03074	0.02842
Centerline engine mass flow error	0.000	0.000
Centerline engine inlet pressure recovery	0.996	0.992
Outboard engine mass flow error	0.000	0.000
Outboard engine inlet pressure recovery	0.996	0.993



configuration. The final results are compared with the baseline results in Table 3. The optimization history is given in Fig. 7. With the 18 design variables, the inlet design method was able to reduce the fuel-burn rate by almost 10%. This performance improvement was

**Fig. 7 Optimization history of the podded BWB configuration.****Fig. 8 Comparison of baseline and optimized podded BWB configurations.**

entirely due to the reduction in drag of over 23 counts, because the pressure recovery of both inlets actually decreased slightly.

A comparison of the baseline and optimized geometries is given in Fig. 8. The major difference is the height of the centerline pylon. The optimizer has moved the nacelles as far apart and as far from the wing as possible with the selected design variables to relieve shocks in the channels between the nacelles and the wing. The reduction of shock strength is better illustrated the pressure-contour plots in Fig. 9. The channel between nacelles is clearly a major source of drag, and therefore the optimizer did what it could to improve the flowfield in this region. Because the spanwise location of the nacelles was held fixed due to engine-out performance constraints, the best the optimizer could do was raise the centerline nacelle as high as possible and still not violate the moment constraint. The centerline nacelle was raised higher than the outboard nacelles, because then only one pylon's wetted area would increase, whereas if the outboard nacelles were raised higher, two pylons' wetted areas would increase. Of course, the structural weight of the pylon is probably the true limiting factor in the pylon height, but because structural analysis was not included in the method at this point, the pylon weight penalty was not addressed.

The pressure contours (Fig. 9) also show that the pressure gradients in the vicinity of the inlets have been reduced. The nacelle yaw, pitch, and inlet camber have been tailored to allow the flow to be ingested cleanly and therefore with minimal drag penalties. The yaw of the outboard nacelles has also been increased to further alleviate the shock strengths in the channel. This could cause problems in the engine-out flight condition, because the moment arm about the

**Fig. 9 Comparison of pressure contours on baseline and optimized podded BWB configurations.**




Fig. 10 The 3-view of the BLI BWB configuration.

aircraft center of gravity of the outboard nacelle has been increased. Future design work will have to include a yaw constraint to address this potential problem.

Boundary-Layer-Ingesting Nacelle Optimization Problem

The baseline BLI configuration is shown in Fig. 10 in 3-view format. This configuration provided a very challenging case for the inlet design method. Flow separation is present within both inlets and in the channel between the nacelles on the wing upper surface. This flow separation often degraded the convergence properties of the flow solver, which naturally hindered the computation of accurate gradients. To somewhat alleviate this problem, the trailing-edge camber of the wing section between the nacelles was altered to reduce the adverse pressure gradient. This reduced the amount of channel flow separation to reasonable levels so that more accurate gradients could be computed. This, of course, altered the pitching moment of the aircraft significantly, but fortunately, the change in pitching moment resulted in an aircraft that was closer to being trimmed. The overall nacelle and inlet design is very similar to a design provided by McDonnell Douglas Corporation.

Two views of the coarse computational grid used for this optimization are shown in Fig. 11. This grid was created in the same way as the podded grid: a block of the baseline wing grid was removed and the nacelle grids were inserted with non-point-match interfaces. The total number of points in this grid is just over 627,000, which is quite a bit more than that used on the podded grid. More points were used in this grid to better capture the flow features within the inlet, whereas the podded inlets had relatively featureless flow. Of course, this is still a rather coarse grid, to keep computation times reasonably low during the optimization process.

Objective Function and Constraints

The objective function for this optimization case is again the fuel-burn rate. Constraints are identical to those in the podded design, with a few minor differences. As discussed previously, modifying the trailing edge of the aircraft changed the pitching moment of the aircraft significantly. Consequently, the pitching moment was constrained to be no less than -0.13 , which is the pitching moment of the modified wing-alone configuration. The one other difference is that a distortion constraint is enforced, because this is clearly an active constraint for this inlet design. Note that the distortion is enforced to be no more than 10% for the centerline inlet and 12% for the outboard inlet, not the 5% suggested by Anabtawi [17]. The lower value of distortion is quite aggressive based on the values computed on the baseline inlet design and, in fact, may not be attainable with a limited number of design variables. However, to demonstrate that the inlet design method is capable of reducing or at least constraining distortion, the more attainable level of 10–12% is used. Also included are a few geometric constraints to ensure that the nacelle

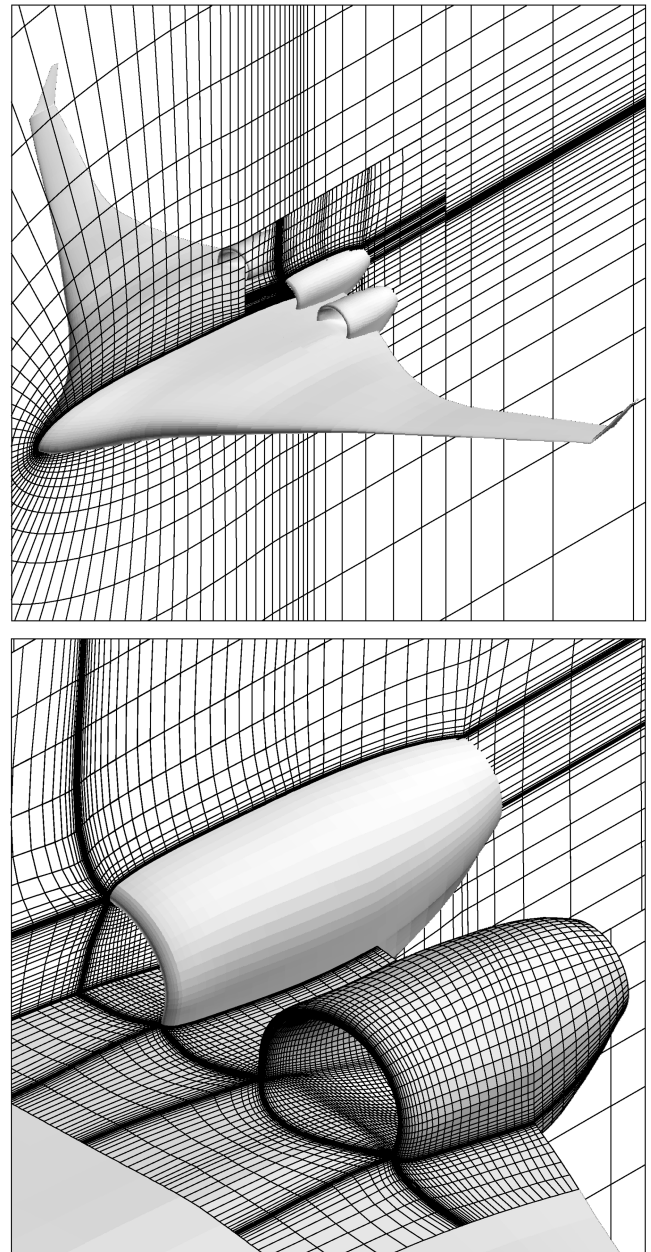


Fig. 11 Typical coarse computational grid on the BLI BWB configuration.

maintains a reasonable minimum thickness for structural considerations. Note that no nacelle leading-edge radius constraints are included, because the design variables selected are incapable of affecting the baseline leading-edge radius.

Table 4 Design variables used in BWB BLI centerline inlet optimization

Variable	Type	Location
1	Angle of attack	Entire aircraft
2	Inlet backpressure	Centerline engine fan face
3	Inlet backpressure	Outboard engine fan face
4	Linear wing twist	Outboard wing section
5	Inlet length	Outboard nacelle
6	Inlet length	Centerline nacelle
7	Nacelle thickness	Centerline nacelle
8–10	Inlet/nacelle camber	Centerline nacelle
11–12	Inlet lower wall	Centerline nacelle
13–14	Inlet junction/wall	Centerline nacelle
15	Inlet upper/side wall	Centerline nacelle
16	Inlet upper wall	Centerline nacelle

Table 5 Design variables used in BWB BLI outboard inlet optimization

Variable	Type	Location
1	Angle of attack	Entire aircraft
2	Inlet backpressure	Centerline engine fan face
3	Inlet backpressure	Outboard engine fan face
4	Linear wing twist	Outboard wing section
5	Inlet length	Outboard nacelle
6–8	Nacelle thickness	Outboard nacelle
9–13	Inlet/nacelle camber	Outboard nacelle
14–15	Inlet lower wall	Outboard nacelle
16–19	Inlet junction/wall	Outboard nacelle
20–21	Inlet upper/side wall	Outboard nacelle
22	Inlet upper wall	Outboard nacelle

Design Variables

Because the number of design variables in any optimization run had to be limited to conserve computational resources, this optimization was split into two steps. The first optimization concentrated on the centerline nacelle and the second worked primarily on the outboard nacelle. For the first optimization, 16 design variables were used and are listed in Table 4. Most of these variables modify the centerline inlet interior surface by adding or removing smooth bumps, which effectively alter the general shape and area distribution of the inlet. Inlet camber variables control the shape of the inlet lip region, and another design variable controls the length of the inlet. The outboard inlet length is also allowed to vary, because it was expected to have a major effect on the centerline nacelle. As was the case with the podded configuration optimization, the angle of attack, inlet backpressures, and outboard wing twist were allowed to vary to satisfy the lift, airflow compatibility, and pitching moment constraints. Finally, a few variables that modify the exterior shape of the nacelle were included to help reduce airframe drag. A more detailed

description of each of these design variables is given in [6] but is omitted here for conciseness.

Similar design variables, listed in Table 5, were used for the second optimization on the outboard nacelle. Because the outboard nacelle is not symmetric, 22 variables were used, including extra nacelle thickness, camber, and wall shape functions. These did not include any variables on the already-optimized inboard nacelle.

Results and Discussion

As with the podded nacelle, two types of optimization problems were completed. The first problem included only the first three design variables listed in Table 4, with no geometry variables. This run was made to determine the performance of the baseline configuration before the inlet optimization was attempted. The second optimization problem concentrated on the centerline nacelle and included all the design variables in Table 4. The convergence history of the centerline nacelle optimization is shown in Fig. 12. The final optimization concentrated on the outboard nacelle design and included all the variables listed in Table 5. The history of this optimization is shown in Fig. 13. These last two optimizations represent an aggressive application of the inlet design method.

Comparing the performance of the baseline and optimized configurations provides some evidence of the effectiveness of the inlet design method. The performance of these configurations is summarized in Table 6. The first interesting result is that after the centerline and outboard nacelle optimizations, the fuel-burn rate has been reduced over 7% from the baseline value. This result reinforces the importance of proper propulsion/airframe integration if redesigning with a very limited number of design variables can provide such an improvement in aircraft performance. Again referring to Table 6, it is clear that the improvement in performance is due to both the decrease in drag and increase in the centerline inlet pressure recovery. Note that the pressure recovery of the outboard inlet actually decreased, as the method traded inlet performance for

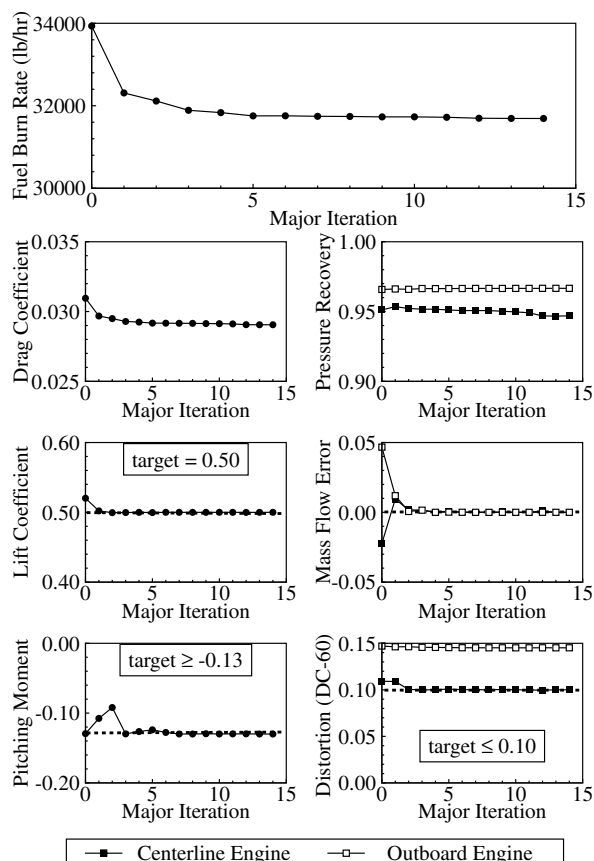


Fig. 12 History of centerline inlet optimization of the BLI BWB configuration.

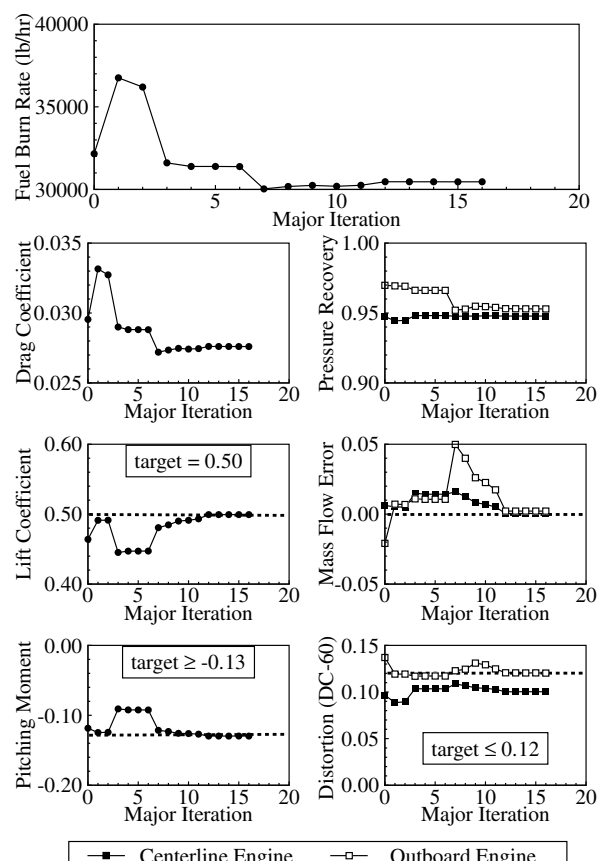


Fig. 13 History of outboard inlet optimization of the BLI BWB configuration.

Table 6 Performance comparison of baseline and optimized BLI BWB configurations

Performance parameter	Baseline	Centerline optimized	Outboard optimized
Fuel-burn rate, lb/h	32923	31691	30463
Lift coefficient	0.5001	0.5000	0.4994
Pitching moment coefficient	-0.1181	-0.1300	-0.1297
Drag coefficient	0.02974	0.02905	0.02760
Centerline mass flow error	0.000	0.000	0.001
Centerline inlet pressure recovery	0.927	0.947	0.948
Centerline fan-face distortion (DC_{60})	0.098	0.100	0.100
Outboard mass flow error	0.000	0.000	0.002
Outboard inlet pressure recovery	0.966	0.966	0.953
Outboard fan-face distortion (DC_{60})	0.145	0.145	0.120

lower drag and to satisfy the distortion constraint. The reduction in drag is primarily due to the reduction in shock strength on both nacelles and in the channels, as shown in Fig. 14. The increase in the centerline inlet pressure recovery is also evident in Fig. 14, as the shock within the choked baseline inlet has been eliminated in the optimized inlet.

Referring back to Table 6, note that all constraints were met in the optimized configurations. The lift, pitching moment, and airflow compatibility constraints were not violated by any noticeable amount. However, the outboard optimization did not converge as well on the constraints as the centerline optimization, due to poor flow solution convergence. Also, note that the centerline inlet distortion constraint of 10% was met; then again, the baseline design also satisfied this constraint. This presents the question of whether the distortion constraint was even necessary. Figure 15 compares the total pressure distributions on the centerline engine fan faces of the

baseline and optimized configurations. This comparison shows the obvious improvement in pressure recovery. The pool of low total pressure at the top of the fan face, which is due to flow separation caused by the shock within the choked inlet, has been eliminated in the optimized inlet. This upper pool of low of total pressure, however, was the primary reason that the distortion on the baseline fan face was so low. The average total pressure on the entire fan face was lower, and therefore the normalized value of DC_{60} was also lower [refer back to Eq. (2)]. Therefore, as the optimizer increased inlet pressure recovery, the distortion level also tended to increase. Consequently, the optimizer reshaped the lower wall of the inlet (as seen in Fig. 14) to circumferentially distribute the pool of low total pressure on the fan face. Recall that this low total pressure flow is due to ingesting the wing boundary layer and cannot be eliminated in this design. This redistribution of low-momentum flow is evident on the optimized fan face in Fig. 15. By increasing the value of the average total pressure on the 60 deg sector, which is used to determine the distortion level [Eq. (2)], the distortion constraint is satisfied while still improving the inlet pressure recovery significantly.

On the other hand, the outboard inlet optimization results tell a completely different story. Here, the distortion constraint forced the design to reduce distortion significantly (by at least 2%). Figure 16 gives the total pressure distributions on the outboard fan face before and after the optimization. Note that the baseline flow has only the one low total pressure pool due to the wing boundary layer, because this baseline inlet was not choked. This is the reason for the higher initial distortion levels and better inlet pressure recovery than the baseline centerline inlet. Unfortunately, the optimizer with the limited number of design variables was unable to spread out the low total pressure pool enough to reduce distortion as it did with the centerline inlet. Instead, the optimizer was forced to introduce a new source of low total pressure near the top of the inlet to reduce the DC_{60} level. This new source of low total pressure was flow separation on the upper inlet surface. This flow separation is small but still significant and is the reason why the flow solver was not able to converge well enough for good optimization convergence. It also explains why the pressure recovery in the outboard inlet actually decreased after the optimization. This is a prime example of an optimizer finding the flaw in the definition of the model problem. Because the optimizer could not reduce the distortion in a way

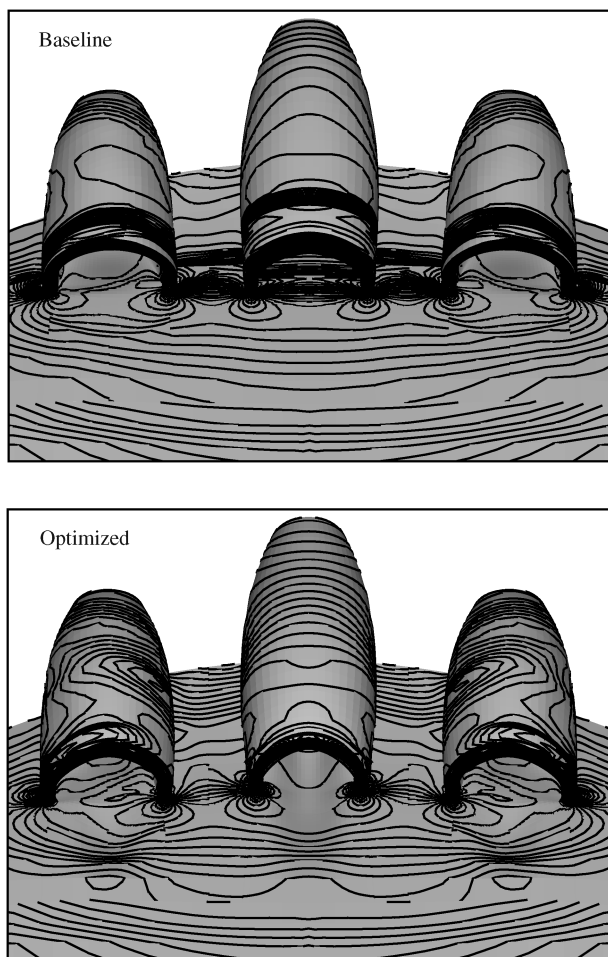


Fig. 14 Comparison of pressure contours on baseline and optimized BLI BWB configurations.

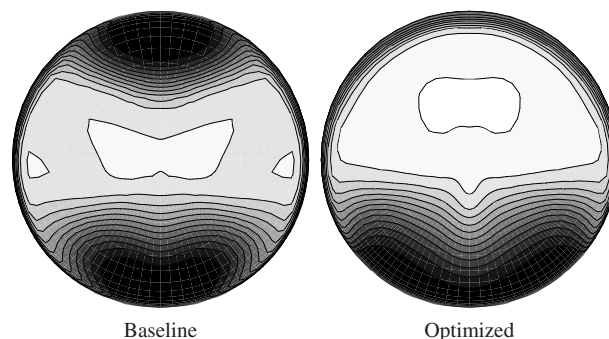


Fig. 15 Total pressure contours (dark is lower) on baseline and optimized centerline fan faces of BLI BWB.

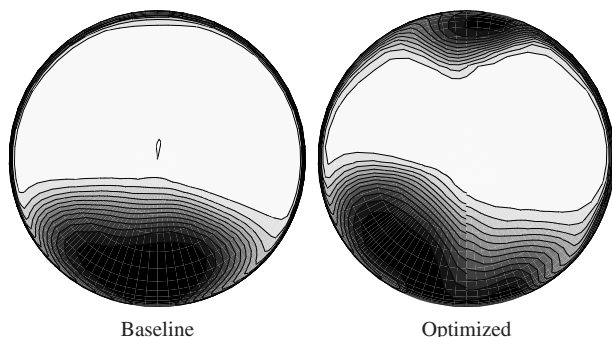


Fig. 16 Total pressure contours (dark is lower) on baseline and optimized outboard fan faces of BLI BWB.

preferable to the design, it exploited a flaw in the distortion descriptor to still satisfy the constraint.

Overall, the method proved to be successful in this problem by significantly improving the aircraft performance while addressing the many necessary constraints. Reducing airframe drag and increasing inlet performance seems to be the strong point of the method. However, for this case, the ability to significantly reduce distortion levels was found to be somewhat limited, simply due to the basic design of the inlet. To produce a large improvement in distortion levels, the overall layout and design of the inlet would have to be altered. However, the fine-tuning of any inlet could easily be done with this method in most cases. The results also demonstrated that the definition of distortion is very important to the optimization problem. An imperfect distortion descriptor could produce unwanted results such as those shown in Fig. 16. For a boundary-layer-ingesting inlet, stricter constraints on radial and circumferential distortion must be implemented to produce a desirable design.

Future Work

Although the inlet design method proved to be quite effective in its limited use, it was simply too slow and not robust enough to apply in a production mode. Many more design variables are desirable to determine a truly optimal design. Though the complex step method is quite robust in most cases, it is still computationally tedious and expensive to compute a complex solution for each desired gradient. Using an adjoint method, such as that used successfully in wing design by Jameson [18], can solve some of these problems. An adjoint solver could compute the gradients of an unlimited number of design variables for the cost of just a few solutions. Implementing the adjoint method would also allow for fine-grid optimizations and therefore more accurate, and perhaps more successful, results.

One other problem that must be addressed is the poor robustness of the method when an unsteady or separated flow feature is encountered. Because unsteady flow features are common when a thick wing boundary layer flow is ingested into a possibly choked inlet, sensitivities can sometimes be impossible to accurately compute, because the CFD code simply will not converge to an acceptable level. How this can be avoided is not obvious at this time, but the need to avoid this problem is clear. Active constraints to steer a design away from separated flow (which is undesirable anyway) may help alleviate this problem.

The BWB program has continued even up to the publication of this paper, and therefore the work presented here remains noteworthy. In fact, research on BLI inlets for the BWB is still ongoing at both NASA and The Boeing Company. Reference [19] gives some details of future planned research that includes the experimental investigation of a BWB with BLI inlets. So even though the work presented here was completed about a decade ago, the subject matter remains relevant.

Conclusions

A highly advanced methodology for designing inlets was successfully developed using multidisciplinary optimization tech-

niques. This inlet design method was applied to two blended wing body configurations: one with podded inlets and one with boundary-layer-ingesting inlets. The method was able to properly address the tight coupling between aerodynamic and propulsive efficiencies for this configuration. Although very coarse meshes and a limited number of design variables were used in the design problems, the method was able to significantly improve the overall performance (indicated by reduced fuel-burn rates) in all cases considered. On the other hand, fan-face distortion proved to be difficult to reduce to acceptable levels, due to the limited design space and the inherent design of the inlet itself. The method was able to produce these improved designs even though computational resources were severely limited.

When comparing the podded and boundary-layer-ingesting engine installations, the optimization results did show that the latter aircraft had a significant improvement in aircraft drag. However, this increase in aerodynamic efficiency came with a significant decrease in inlet performance, resulting in an overall higher fuel-burn rate. These initial results do not bode well for the boundary-layer-ingesting inlet, especially when the accompanying poor distortion levels are also taken into account. However, the limited design space selected for these optimization problems do not allow these results to be the final verdict on boundary layer ingestion as a design concept. Further investigation with this design tool and more extensive computational resources is necessary to truly ascertain the advantages and disadvantages of boundary-layer-ingesting inlets. This conclusion is reinforced by the fact that experimental research on a blended wing body configuration with boundary-layer-ingesting inlets is still planned at NASA in the decades to come.

Acknowledgments

The author would like to acknowledge McDonnell Douglas Corporation (now The Boeing Company) for providing financial support and Bob Liebeck specifically for his help in identifying the problem that motivated this research. Ilan Kroo of Stanford University provided a great deal of guidance in this research as well. NASA Ames Research Center provided most of the computer time on their Silicon Graphics Origin 2000 supercomputers. Bob Biedron of NASA Langley Research Center also contributed to applying the complex variable method with the CFL3D code.

References

- [1] Smith, A. M. O., and Roberts, H. E., "The Jet Airplane Utilizing Boundary Layer Air for Propulsion," *Journal of the Aeronautical Sciences*, Vol. 14, No. 2, Feb. 1947, pp. 97–109.
- [2] Lynch, F. T., "A Theoretical Investigation of the Effect of Ingesting Boundary Layer Air on Turbofan Engine Fuel Consumption," Douglas Aircraft Co., Rept. SM 23981, Long Beach, CA, May 1960.
- [3] Douglass, W. M., "Propulsive Efficiency with Boundary Layer Ingestion," McDonnell Douglas Corp., Rept. MDC J0860, Long Beach, CA, Aug. 1970.
- [4] Smith, L. H., "Wake Ingestion Propulsion Benefit," *Journal of Propulsion and Power*, Vol. 9, No. 1, 1993, pp. 74–82. doi:10.2514/3.11487
- [5] Liebeck, R. H., Page, M. A., and Rawdon, B. K., "Blended-Wing-Body Subsonic Commercial Transport," AIAA Paper 98-0438, Jan. 1998.
- [6] Rodriguez, D. L., "A Multidisciplinary Optimization Method for Designing Boundary Layer Ingesting Inlets," Ph.D. Thesis, Stanford Univ., Stanford, CA, 2001.
- [7] Rumsey, C., Sanetrik, M., Biedron, R., Melson, N., and Parlette, E., "Efficiency and Accuracy of Time-Accurate Turbulent Navier–Stokes Computations," *Computers and Fluids*, Vol. 25, No. 2, 1996, pp. 217–236. doi:10.1016/0045-7930(95)00043-7
- [8] Spalart, P., and Allmaras, S., "A One-Equation Turbulence Model for Aerodynamic Flows," AIAA Paper 92-0439, Jan. 1992.
- [9] Rodriguez, D. L., "Validation of CFD Methods on a Boundary Layer Ingesting Inlet," AIAA Paper 98-0925, Jan. 1998.
- [10] Klann, J., and Snyder, C., "NEPP Programmer's Manual," NASA TM 106575, 1994.
- [11] Gill, P. E., Murray, W., Saunders, M. A., and Wright, M. H., "User's Guide for NPSOL 5.0: A Fortran Package for Nonlinear

- Programming,” TR SOL 94, Stanford Univ., Stanford, CA, 1994.
- [12] Livesey, J. L., “Flow Property Averaging Methods for Compressible Internal Flows,” AIAA Paper 82-0135, Jan. 1982.
 - [13] Lyness, J. N., “Numerical Algorithms Based on the Theory of Complex Variables,” *Proceedings of the ACM 22nd National Conference*, Thomas Book, Washington, D.C., 1967, pp. 124–134.
 - [14] Lyness, J. N., and Moler, C. B., “Numerical Differentiation of Analytics Functions,” *SIAM Journal on Numerical Analysis*, Vol. 4, 1967, pp. 202–210.
doi:10.1137/0704019
 - [15] Sherman, L. L., Taylor, A. C., Green, L. L., Newman, P. A., Hou, G., and Korivi, V. M., “First- and Second-Order Aerodynamic Sensitivity Derivatives via Automatic Differentiation with Incremental Iterative Methods,” AIAA Paper 94-4262, Sept. 1994.
 - [16] Rodriguez, D. L., “A Multidisciplinary Optimization Method for Designing Inlets Using Complex Variables,” AIAA Paper 2000-4875, Sept. 2000.
 - [17] Anabtawi, A., “Experimental Investigation of Boundary Layer Ingestion into Diffusing Inlets,” Ph.D. Thesis, Univ. of Southern California, Los Angeles, 1999.
 - [18] Jameson, A., “Aerodynamic Design Via Control Theory,” NASA CR-181749, 1988.
 - [19] Warwick, G., “Shaping the Future,” *Aviation Week and Space Technology*, Feb. 2009, pp. 50–53.