

Parallel Solution Methods for Aerostructural Analysis and Design Optimization

Graeme J. Kennedy,*

University of Toronto Institute for Aerospace Studies, Toronto, ON, Canada

Joaquim R. R. A. Martins[†]

Department of Aerospace Engineering, University of Michigan, Ann Arbor, MI

This paper presents a comparison of methods for aerostructural analysis and optimization. The aerostructural analysis problem is solved in parallel using a panel method coupled to a finite-element solver. The coupled nonlinear aerostructural system is solved using a nonlinear block Gauss-Seidel, nonlinear block Jacobi, Newton-Krylov or approximate Newton-Krylov approach. The approximate Newton-Krylov method is shown to be an efficient and robust solution technique. An adjoint-based sensitivity method is developed that achieves a high-level of accuracy when compared to complex-step calculations. Three levels of parallelism are exploited within the present aerostructural optimization framework: optimization-level, system-level and discipline-level parallelism. The efficient and robust solution method and accurate gradient evaluation technique provide a powerful tool for aerostructural design optimization. Aerostructural induced drag minimization results are presented for a typical subsonic turbo-prop aircraft wing.

I. Introduction

Aerostructural design optimization using high-fidelity models is a computationally intensive multidisciplinary design optimization (MDO) problem.¹⁻⁴ Many authors have focused on using high-fidelity aerodynamic analysis in the aerostructural problem while using considerably smaller finite-element models.^{5,6} This is often justified since the primary area of interest in these studies is the aerodynamic performance of the aerostructural system. Developing techniques that work well when there is a computational imbalance between disciplines is important, since often the aerodynamic analysis requires many more degrees of freedom than the structural analysis — sometimes two or three orders of magnitude more. Several authors have even devised optimization techniques specifically designed to take advantage of this imbalance.⁷⁻⁹ However, it is also important to examine the consequences of larger and more sophisticated structural analyses and the implications this has on both the solution algorithms and the overall optimization problem. This is a significant problem as unconventional aircraft design and the increasing use of composites and advanced composite systems,¹⁰ may place additional demands on the structural analysis.

This paper presents aerostructural analysis and design sensitivity methods that are designed to be efficient when the aerodynamic and structural disciplines both require significant computational resources and time. This situation leads to different solution algorithms. These techniques are applied to the analysis and optimization of an aircraft wing. In the present work, the aerodynamic analysis uses a parallel panel code, TriPan, coupled to a parallel finite-element code called the Toolkit for the Analysis of Composite Structures (TACS). This finite-element code is specifically designed for the analysis of composite structures, including geometric nonlinearity and has analytic design variable sensitivity analysis capability. The inter-disciplinary coupling is handled by passing pressure and displacement values through a parametric surface representation of the outer mould line (OML) of the aircraft.

This paper is organized as follows. Section (II) reviews previous work performed in the field of aerostructural analysis and optimization, as well as load and displacement transfer schemes. Section (III) describes

*Ph.D. Candidate, AIAA Student Member

[†]Associate Professor, AIAA Senior Member

the relevant details of the aerostructural analysis components and optimization framework. The components include the aerodynamic and structural analysis, load and displacement transfer, and geometry parametrization. Sections (IV) and (V) outline the various algorithms used to solve the coupled, aerostructural system and the coupled adjoint equations respectively. Section (VI) presents results for an aerostructural optimization of a wing based on the Bombardier Q400.

II. Literature review

A. Aerostructural analysis and optimization

Several authors have previously developed methods for aerostructural analysis and design optimization. Reuther et al.⁵ and Martins¹¹ developed an aerostructural analysis and optimization framework that coupled a linear finite-element structural model to a finite-volume Euler CFD solver. They obtained a coupled solution using a pseudo-time marching scheme with periodic updates of the displaced shape. They used a structural model composed of either solid, three-dimensional elements, to represent a wind tunnel wing model, or shell and beam elements to represent a stiffened aircraft wing. Following the work of Brown,¹² they developed a systematic scheme to pass loads and displacements across the aircraft OML. Martins et al.² developed a sensitivity analysis of the aerostructural equations for both the adjoint and direct formulations, with a block Gauss–Seidel technique for solving the coupled adjoint system. Martins et al.¹ applied this framework to the optimization of a supersonic business jet.

Maute et al.³ performed an aerostructural analysis that coupled the Euler equations to a linear finite-element model. They employed a mesh movement strategy based on a spring analogy and a load and displacement transfer technique following the earlier work of Maman and Farhat¹³ and Farhat et al.¹⁴ They used a nonlinear block Gauss–Seidel method with relaxation for the solution of the coupled nonlinear equations. They present formulations of both the adjoint and direct methods for computing the sensitivities of the coupled aerostructural system. These linear equations are solved using a block Gauss–Seidel method that mirrors the method of solution for the coupled system. Two types of structural models are employed: an equivalent flat plate composed of either a single isotropic material or multiple composite layers and a full shell and stiffener model composed of shell and beam elements.

Maute and Allen¹⁵ develop an aerostructural optimization problem in which the internal structure of the wing box is parametrized using a single isotropic material with penalization (SIMP) approach to determine the topology of the optimal structure. The solution method is similar to those described previously.³

In order to improve the robustness and efficiency of the methods developed by Maute et al.,³ Barcelos et al.⁶ developed a class of Newton–Krylov–Schur methods for solving the coupled nonlinear fluid-structure-mesh movement problem. Their approach uses an approximate Newton’s method for the solution of the nonlinear coupled equations. At each iteration, a Schur complement approach is used to solve the coupled linear system that results from a linearization of the residual. They found that their technique is more robust and efficient than the original Gauss–Seidel method presented by Maute et al.³ More recently, Barcelos and Maute¹⁶ presented an aerostructural solution technique coupling the Navier–Stokes equations with a turbulence model to a linear structure and mesh movement strategy.

Aerostructural analysis techniques are a specialization of more general fluid-structure interaction solution methods. A complete review of this more general field of research is beyond the scope of this paper, however we discuss a few results that have contributed to our ideas.

Felippa et al.¹⁷ performed an extensive review of solution techniques for coupled nonlinear problems using partitioned solvers. Kim et al.¹⁸ developed a solution procedure for coupled multi-physics problems using a multi-level Newton’s method. They applied their approach to a coupled fluid-structure interaction problem, noting the importance of using accurate linearizations of the coupling terms.

In two papers, Biros and Ghattas^{19,20} presented a Lagrange–Newton–Krylov–Schur approach to the simultaneous solution of PDE-constrained optimization problems. They applied their approach to a design problem using the incompressible Navier–Stokes equations.

Heil et al.²¹ solved a time-dependent fluid-structure interaction problem by applying Newton’s method to a second-order backward difference discretization of the coupled system. The resulting equations were solved using a monolithic approach with both direct and iterative solvers. Their results demonstrate that a monolithic approach is competitive, even in cases of relatively weak coupling.

B. Load and displacement transfer schemes

One of the primary tasks in aerostructural analysis is to develop a scheme to couple the aerodynamic and structural disciplines into a single analysis. Many authors have devised load and displacement transfer techniques and a brief overview of some important contributions is covered below.

Maman and Farhat¹³ developed a method for the direct transfer of loads and displacements between fluid and structural meshes where the boundaries of the two domains may not be coincident everywhere. In their scheme, pressure from the fluid mesh is transferred to a projected point on the structural mesh, where a local normal is defined. For displacement transfer, structural displacements are projected back onto the fluid mesh. From the projected points, local interpolations are used to determine the values of the quantities of interest. In a second paper, Farhat et al.¹⁴ devised two methods for load and displacement transfer for transient problems. A method based upon consistent interpolation between coincident surfaces and a second, more general method based upon displacement and load transfer between discrete surfaces.

Brown¹² focused on the development of load and displacement transfer schemes where the structural and fluid model are non-conforming. Brown used displacement-interpolation functions that are either a continuous extension of the finite-element shape functions, or a rigid attachment to the nodal degrees of freedom. The load transfer from the fluid to structural model is constructed using the principle of virtual work. This technique is consistent — the sum of the forces on the fluid model is equal to the sum of the nodal forces on the structure, and is also conservative — the work done on the fluid model in moving through the displacements defined by the structure, is equal to the work performed on the structure.

In two recent papers, Allen et al.²² and Rendall and Allen²³ developed interpolation and mesh movement schemes that employ radial basis functions (RBFs). In these methods, the structural and aerodynamic models are embedded in an RBF volume. Structural displacements are transferred through the volume using RBF displacement interpolation. Load transfer is performed using the principle of virtual work. Consistency is achieved by ensuring that rigid body modes are preserved within the RBF interpolation.

C. The proposed framework

In this paper we depart from previous work by performing aerostructural optimization in a framework that exploits three levels of parallelism while using a multi-disciplinary feasible (MDF) approach.^{24,25} These levels of parallelism are summarized in Table (1). We use parallelism at the optimization-level, coordinating system-level process groups to analyze different flight conditions or load cases simultaneously. This level is embarrassingly parallel, requiring very little communication overhead. The second level of parallelism we exploit is system-level parallelism within the aerostructural analysis. We assume that due to memory or performance requirements, the aerodynamic and structural processes are split into non-overlapping process groups within the system-level process group. This second level involves coordination of the discipline-level analyses to perform an aerostructural solution or sensitivity calculation. The last level of parallelism we use is at the discipline-level. At this level, efficient single-disciplinary codes may be run on independent process groups. We use these three levels to perform efficient, parallel multidisciplinary analysis and optimization.

Table 1. Levels of parallelism and process groups within the present aerostructural optimization framework.

| Level | Process group | Parallelism |
|-------|--------------------|--|
| 1 | Optimization-level | Multiple flight conditions and load cases run concurrently |
| 2 | System-level | Interdisciplinary coupling and coordination |
| 3 | Discipline-level | Parallel discipline-level analysis |

III. Aerostructural analysis components

The following subsections outline the relevant details of the aerodynamic analysis, load and displacement transfer, structural analysis and geometric parametrization. These constitute the unassembled components of the aerostructural analysis.

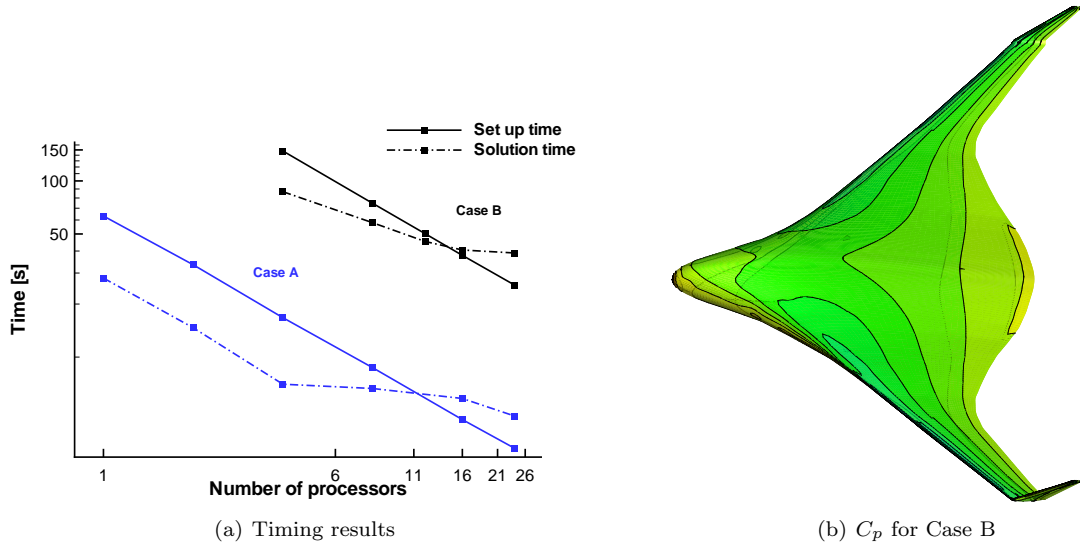


Figure 1. A timing study for TriPan using 1,2,4,8,12,16 and 24 processors for two cases. Case A has 8800 panels and Case B, a BWB, has 22638 panels.

A. Aerodynamic analysis

The aerodynamic analysis is performed using TriPan, an unstructured panel code for calculating the aerodynamic characteristics of inviscid, incompressible, external lifting flows. The code uses constant source and doublet singularity elements distributed over the surface of a body discretized with quadrilateral and triangular panels.^{26,27} Aerodynamic forces and moments are calculated using surface pressure integration. This approach has well known accuracy problems, especially when computing drag.²⁸ The discretized set of boundary conditions governing the doublet strengths are represented by the vector of aerodynamic residuals,

$$\mathcal{A}(w) = 0, \quad (1)$$

where w represents a vector of the doublet strengths. The linear system represented by Equation (1) is solved using the parallel, linear algebra routines in PETSc.^{29,30} A dense matrix storage format is used to store the aerodynamic influence coefficients for each panel. The matrix is split between processors such that the rows of the matrix are stored in contiguous segments. The ownership range of each segment is determined by performing a domain decomposition of the surface mesh. The dual graph of the mesh is constructed and then partitioned by the program METIS³¹ to determine processor assignment. For efficiency reasons, each processor has a local copy of the entire mesh. The surface mesh is relatively small, so copying the entire mesh is not too memory intensive, especially when compared with the dense aerodynamic influence coefficient matrix.

The solution time required for TriPan is plotted against the number of processors for two test cases in Figure 1(a). In this figure, Case A is a symmetric, rectangular wing with an aspect ratio of 8.0 and a NACA0012 profile, that is discretized with 8800 panels on the half-body. Case B is a blended wing body (BWB) discretized with 22628 panels. The surface pressure coefficient for the BWB, is shown in Figure 1(b). The matrix and residual assembly time scales ideally to within precision. After an initial communication of the state variables, each processor can independently determine all matrix and residual components that are locally owned without having to communicate any information.

The system of linear equations is solved using GMRES(30)³² with a block Jacobi ILU(0) preconditioner. We found this to be an effective method, requiring the least overall computational time, amongst a range of preconditioning options. The preconditioner is assembled by considering only those panels that are within a predetermined radius from the current panel centroid. The rationale for this construction is that the closest panels have the strongest effect on a given panel. The effect of a higher ILU preconditioner fill-level can be achieved by selecting a larger radius.

Figure (1(a)) shows the solution time required for Case A and Case B for between 1 and 24 processors. Note, that the solution time is often shorter than the time required for the matrix and residual assembly. The solution time for Case A scales well when fewer than 8 processors are used, but reaches some performance limit for larger number of processors. For Case B, the solution time decreases as more processors are added, but not at the ideal rate.

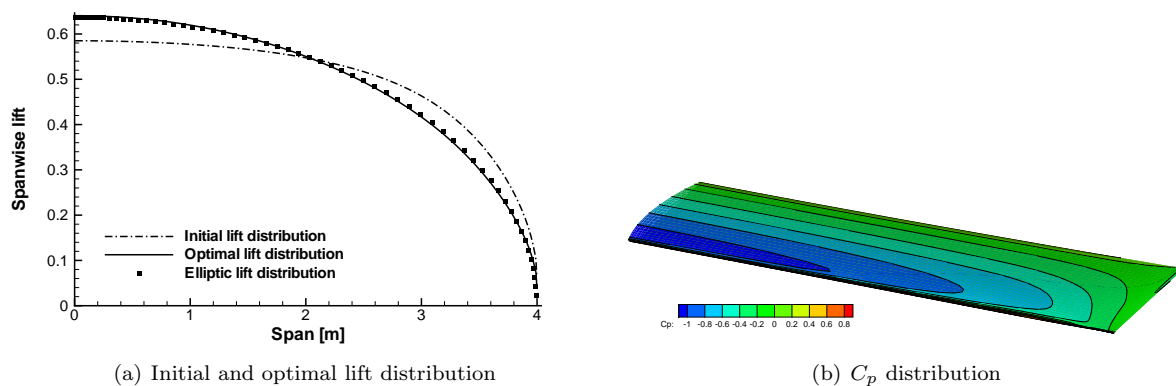


Figure 2. Initial and optimal lift distributions for a lift-constrained drag minimization test case run in TriPan. The figure on the right shows the optimal shape and C_p distribution.

An adjoint sensitivity method is implemented within TriPan. Figure (2) shows the results from a lift-constrained drag-minimization problem. The wing has a span of 8.0 and a constant chord of 1.0 and is composed of a lofted stack of NACA2412 airfoils, initially without twist. The surface mesh consists of 4200 surface panels with 3600 wake panels. Four twist variables and angle of attack are allowed to vary. The lift constraint $C_L = 0.5$ is enforced. Figure (2(a)) shows the lift distribution of the optimal wing and the initial wing at an angle of attack that satisfies the lift constraint.

B. Load and displacement transfer

The present load and displacement transfer scheme follows the work of Brown.¹² In this approach, the displacements are extrapolated to the aerodynamic surface mesh through the use of rigid links. These links connect the nodes on the aerodynamic surface to the closest point on the structure. The nodal displacement on the aerodynamic surface may be written in terms of the displacements and, when appropriate, rotations defined by an element as follows,

$$\mathbf{u}_A = \mathbf{u}_S + \boldsymbol{\theta}_S \times \mathbf{r}, \quad (2)$$

where \mathbf{u}_A and \mathbf{u}_S are the displacements of the aerodynamic and structural points respectively, $\boldsymbol{\theta}_S$ are the rotations at the structural surface and \mathbf{r} is the vector of minimum distance connecting the points in the structure to the points on the aerodynamic surface. Equation (2) defines the displacements at the aerodynamic surface in terms of displacements in the structure. The method of virtual work can be used to determine the nodal forces and moments acting on the structure. The virtual work of the aerodynamic pressure forces is,

$$\delta W = \int_{S_A} p \hat{\mathbf{n}} \cdot \delta \mathbf{u}_A \, dS = \int_{S_A} p \hat{\mathbf{n}} \cdot \delta \mathbf{u}_S - p \hat{\mathbf{n}} \cdot (\mathbf{r} \times \delta \boldsymbol{\theta}_S) \, dS, \quad (3)$$

where p is the surface pressure and $\hat{\mathbf{n}}$ is the normal defined on the aerodynamic surface mesh. Note that the integration is performed on the aerodynamic surface. This expression may be used to determine a consistent and conservative load vector.

This technique is quite flexible in that any arbitrary aerodynamic and structural meshes can be combined. Even problems where the structure may lie outside the aerodynamic mesh, can be connected in this manner. However, there are two important issues that may arise with this scheme:

1. Adjacent points on the aerodynamic surface may cross, when the structure experiences large, rapidly varying rotations. No strain energy is associated with the deformation of the aerodynamic surface.

When the surface deforms, there is no stiffness counteracting the movement of adjacent surface points towards one another, other than the stiffness of the structure.

2. Large point moments can be produced when the structure lies far away from points on the aerodynamic surface. These large moments are required to maintain consistency and conservativeness. Usually, this indicates that the model is inadequate and the structure should move closer to the aerodynamic surface.

For many aerostructural problems where the meshes are close together over most of the domain, neither of these issues arise.

C. Structural analysis

The structural analysis is performed using the Toolkit for the Analysis of Composite Structures (TACS), a parallel, finite-element code developed by the first author, designed specifically for the analysis of stiffened, thin-walled, composite structures using either linear or geometrically nonlinear strain relationships. We typically use higher-order finite-elements as we have found that these provided better stress prediction capability. The residuals of the structural governing equations are written as,

$$\mathcal{S}(u) = \mathcal{S}_c(u) - \mathcal{F}(u), \quad (4)$$

where u is a vector of displacements and rotations, \mathcal{S}_c are the residuals due to conservative forces and internal strain energy and \mathcal{F} are the follower forces due to aerodynamic loads.

The Jacobian of the structural residuals involves two terms: the tangent stiffness matrix $\mathcal{K} = \partial\mathcal{S}_c/\partial u$ that is stored in memory and the derivative of the consistent force vector with respect to the structural displacements, that is computed using a matrix-free approach,

$$\frac{\partial\mathcal{S}}{\partial u} = \mathcal{K} - \frac{\partial\mathcal{F}}{\partial u}. \quad (5)$$

This second term has important implications for the parallel implementation. The stiffness matrix, is distributed over the group of structural processes. The derivative of the consistent force vector with respect to the displacements however, must be computed using operations that involve all aerostructural processes. This is due to the nature of the load-transfer scheme. As a result, calculations involving the Jacobian of the structural system require the synchronization of all processes and cannot be performed concurrently with aerodynamic calculations. It is often advantageous to discard the contribution of the follower forces during analysis. However, for sensitivity calculations this term cannot be ignored due to its contribution to the objective or constraint gradient.

While the matrices involved in structural problems are typically symmetric, the term $\partial\mathcal{F}/\partial u$ is non-symmetric due to the non-conservative nature of the aerodynamic forces. These non-symmetric matrices require different solution algorithms than those typically employed in a structural finite-element code. We use GMRES³² and the Krylov method GCROT^{33,34} to solve the non-symmetric, linear systems involving Equation (5).

Stress constraints are handled by applying a local failure constraint at each Gauss point in the finite-element model. These local failure constraints compute a load factor λ_k , required for that point to fail. The load factor implies that the current point will fail at λ_k times the current stress level. For a safe-life design, $\min\{t_k\} > F_S$ where F_S is a factor of safety. Instead of using the minimum value directly in the optimization, a Kreisselmier-Stienhauser (KS) constraint aggregation technique is applied to groups of these local constraints.^{1,35,36} We usually use multiple KS functions, aggregating amongst similar structural components rather than using a single KS function over all stress constraints in the entire structure. The KS function is computed as follows,

$$\lambda_{KS} = \min\{\lambda_k\} - \frac{1}{\rho} \ln \left[\sum_{i=1}^N \exp \{-\rho(\lambda_i - \min\{\lambda_k\})\} \right], \quad (6)$$

where ρ is a weighting parameter that controls the degree of approximation and λ_{KS} is the aggregated KS value. The advantage of this approach is that it reduces the number of constraints required in the optimization and it is a conservative approximation, in that λ_{KS} is a lower bound. Typically, values of ρ between 30 and 50 yield good results.³⁷

D. Geometry parametrization

Parametrization of the geometry is of fundamental importance to the aerostructural optimization problem. We employ a CAD-free approach to geometric parametrization to achieve a continuous and differentiable geometry and analysis problem. We use the implementation of Kenway and Martins³⁸ that is tightly integrated into the aerodynamic and structural solvers. This approach results in an efficient and accurate sensitivity analysis, something that is difficult to achieve with CAD. While we have worked with other geometric parametrization techniques, we use free-form deformation (FFD)³⁹ exclusively in this paper. In this approach, the points that define the geometry of a discipline are embedded in a B-spline volume. The B-spline volume control points may then be manipulated to modify the underlying embedded points. Depending on the order of the B-spline, these changes are continuous and may be differentiable.

IV. Aerostructural solution

The aerostructural system of equations is the concatenation of the aerodynamic and structural equations, represented by,

$$\mathcal{R} = \begin{bmatrix} \mathcal{A}(w, u, x) \\ \mathcal{S}(w, u, x) \end{bmatrix} = 0, \quad (7)$$

where \mathcal{A} and \mathcal{S} are the aerodynamic and structural residuals, w and u are the aerodynamic and structural state variables and x is a vector of design variables. Often, x will be omitted for brevity. Occasionally it will be convenient to combine the unknown state variables into a single vector, $q^T = [w^T, u^T]$.

During the solution procedure, a point is considered converged when the relative tolerance of both residuals is reduced below a specified tolerance, typically $\epsilon_r = 10^{-8}$, such that,

$$\begin{aligned} \|\mathcal{A}(w^{(n)}, u^{(n)})\|_2 &< \epsilon_r \|\mathcal{A}(w^{(0)}, u^{(0)})\|_2, \\ \|\mathcal{S}(w^{(n)}, u^{(n)})\|_2 &< \epsilon_r \|\mathcal{S}(w^{(0)}, u^{(0)})\|_2. \end{aligned} \quad (8)$$

The stopping criterion is applied to each discipline separately rather than the system of aerostructural residuals as a whole to avoid situations where the initial residual of one discipline is significantly greater than the initial residual of the other. The relative scaling of the two disciplines is very important but has not been addressed in this study. In addition, this condition is easier to apply when the disciplines are split across groups of processes.

Other authors have concentrated on aerostructural analysis techniques that are suitable for solving the coupled system when the aerodynamic and structural residuals are distributed across the same set of processors, or the structural residuals are on every process.^{1,16} In this paper, we focus on the situation where either memory or performance requirements dictate that the aerodynamic and structural solvers be split between groups of processors. This requires an additional level of parallelism reflected in the solution algorithm and its implementation.

We now present four methods we use to solve the aerostructural system: nonlinear block Gauss–Seidel, nonlinear block Jacobi, Newton–Krylov and approximate Newton–Krylov.

A. Nonlinear block Gauss–Seidel

Several authors have presented aerostructural solution algorithms that employ a nonlinear block Gauss–Seidel method.^{1,16} The nonlinear block Gauss–Seidel approach is to solve, in sequence, a linearization of each of the two disciplines to arrive at an update. This update is applied before the update for the next discipline is computed. We can write this procedure as follows,

$$\frac{\partial \mathcal{A}}{\partial w} \Delta w^{(n)} = -\mathcal{A}(w^{(n)}, u^{(n)}), \quad (9)$$

$$\mathcal{K}^{(n)} \Delta u^{(n)} = -\mathcal{S}(w^{(n+1)}, u^{(n)}), \quad (10)$$

where the updates $w^{(n+1)} = w^{(n)} + \alpha_A^{(n)} \Delta w^{(n)}$ and $u^{(n+1)} = u^{(n)} + \alpha_S^{(n)} \Delta u^{(n)}$ are applied immediately after each solution. The parameters $\alpha_A^{(n)}$ and $\alpha_S^{(n)}$ are adjustable damping factors for the aerodynamic and structural solutions. Note that the structural processes must be inactive while the aerodynamic update is computed and likewise, the aerodynamic processes must be inactive while the structural update is computed.

Solving Equations (9) and (10) to a tight tolerance typically does not lead to an efficient algorithm. As a result, we employ the following stopping criterion,

$$\|\mathcal{A}(w^{(n)}, u^{(n)}) + \frac{\partial \mathcal{A}}{\partial w} \Delta w^{(n)}\|_2 < \epsilon_{gs} \|\mathcal{A}(w^{(n)}, u^{(n)})\|_2, \quad (11)$$

with $\epsilon_{gs} = 10^{-3}$, with a similar approach for the structural solution.

The selection of the damping parameters is very important for robustness and efficiency of the algorithm. Since the solution strategy is employed within an optimization framework, analyses of slender, flexible structures subject to divergence should be expected in early design cycles. Adjustment of solution parameters must therefore take into account a wide range of possible analysis conditions. Here we choose $\alpha_A^{(n)} = 1$ and pick $\alpha_S^{(n)}$ dynamically based upon an Aitken acceleration scheme,⁶

$$\begin{aligned} \tilde{\alpha}_S &= \alpha_S^{(n-1)} \left(1 - \frac{(\Delta u^{(n)} - \Delta u^{(n-1)})^T \Delta u^{(n)}}{\|\Delta u^{(n)} - \Delta u^{(n-1)}\|_2^2} \right), \\ \alpha_S &= \max(\alpha_{low}, \min(1, \tilde{\alpha}_S)), \end{aligned} \quad (12)$$

where we choose $\alpha_{low} = 0.25$ and $\alpha_S^{(0)} = 0.5$. The effect of the second formula is to clip the update to the interval $(\alpha_{low}, 1)$.

B. Nonlinear block Jacobi

An alternative to the nonlinear block Gauss–Seidel approach is to solve for both updates simultaneously using current information. This leads to the nonlinear block Jacobi method, written as,

$$\frac{\partial \mathcal{A}}{\partial w} \Delta w^{(n)} = -\mathcal{A}(w^{(n)}, u^{(n)}), \quad (13)$$

$$\mathcal{K}^{(n)} \Delta u^{(n)} = -\mathcal{S}(w^{(n)}, u^{(n)}), \quad (14)$$

where the updates $w^{(n+1)} = w^{(n)} + \alpha_A^{(n)} \Delta w^{(n)}$ and $u^{(n+1)} = u^{(n)} + \alpha_S^{(n)} \Delta u^{(n)}$ are applied after both computations have completed. In this method the structural and aerodynamic groups perform computations concurrently, increasing the degree of parallelism within the algorithm. As in the nonlinear block Gauss–Seidel algorithm, Equations (13) and (14) are solve inexactly to a tolerance of $\epsilon_j = 10^{-3}$, according to Equations (11). The damping parameter $\alpha_A^{(n)} = 1$ is set for all iterations, while $\alpha_S^{(n)}$ is chosen according to Equation (12).

C. Newton–Krylov method

Newton’s method applied to Equation (7) results in the following linear system of equations for the update $\Delta q^{(n)}$,

$$\frac{\partial \mathcal{R}}{\partial q} \Delta q^{(n)} = -\mathcal{R}(q^{(n)}). \quad (15)$$

Newton’s method converges quadratically provided the starting point is sufficiently close to the solution and the Jacobian remains non-singular. In order to achieve convergence from points far from the solution, Newton’s method is often globalized with some strategy to ensure progress is made towards the solution until a suitable starting point is found. Typically these globalization strategies include pseudo-transient continuation or other forms of continuation.^{40,41} We have found that for this aerostructural problem, globalization is not necessary. For more difficult problems it may be required.

Solving Equation (15) inexactly for each update is typically more efficient than finding an accurate solution. We solve for the Newton update to a tolerance of $\epsilon_{nk} = 10^{-3}$,

$$\|\mathcal{R}(q^{(n)}) + \frac{\partial \mathcal{R}}{\partial q} \Delta q^{(n)}\|_2 < \epsilon_{nk} \|\mathcal{R}(q^{(n)})\|_2, \quad (16)$$

with the update $q^{(n+1)} = q^{(n)} + \Delta q^{(n)}$. In order to ensure that Equation (8) is satisfied, the following stopping criterion is used for the Newton–Krylov approach,

$$\|\mathcal{R}(q^{(n)})\|_2 < \epsilon_r \min(\|\mathcal{A}(q^{(0)})\|_2, \|\mathcal{S}(q^{(0)})\|_2). \quad (17)$$

Many authors have used Jacobian-free, Newton-Krylov methods where the matrix-vector products are calculated using a finite-difference calculation. These techniques have many well-known benefits.⁴² In this analysis however, the formation of the aerodynamic residuals is almost as costly as calculating the exact aerodynamic Jacobian, we therefore have employed exact computation of the block-diagonal components of the aerostructural Jacobian.

The treatment of the off-diagonal blocks requires a more detailed discussion. The Jacobian, $\partial\mathcal{A}/\partial u$ is computed as the product of two terms,

$$\frac{\partial\mathcal{A}}{\partial u} = \frac{\partial\mathcal{A}}{\partial X_s} \frac{\partial X_s}{\partial u}, \quad (18)$$

where X_s are the aerodynamic nodal locations. The second term is calculated efficiently using a matrix-free approach. In this calculation information is transferred from the structural group of processes to the aerodynamic group of processes. The first term is a dense matrix that is fully populated due to the nature of the panel method. This matrix is also required for the adjoint equations. The formation of $\partial\mathcal{A}/\partial X_s$ is a time consuming operation, typically requiring 20 to 50 times the cost of forming the aerodynamic residuals. As a result, this term is not recalculated at every iteration. Instead, we recompute $\partial\mathcal{A}/\partial X_s$ when $\text{mod}(n, m) = 0$ for $m = 8$. This high value of m is chosen since we have found that the original matrix is often very effective, even in subsequent iterations, and that the matrix should be updated only if the aerostructural problem is difficult to solve. This modification renders this algorithm an *approximate* Newton method since the exact Jacobian is formed only every m th iteration. However, we retain this terminology to distinguish it from the variant described below where more approximations are used.

The other off-diagonal block $\partial\mathcal{S}/\partial w$, is formed through a series of matrix-free operations,

$$\frac{\partial\mathcal{S}}{\partial w} = \frac{\partial\mathcal{S}}{\partial f_A} \frac{\partial f_A}{\partial w}, \quad (19)$$

where f_A are the integrated forces over the aerodynamic surface. These integrated forces are transferred to the structural process group during the calculation.

At each iteration, the Newton update is determined using a preconditioned Krylov method. We use a preconditioner based upon a single application of block Jacobi. In this approach, each discipline applies its own block-preconditioner to its own set of equations. In general, the preconditioner could consist of a sub-Krylov method for each discipline. However, we have found that an effective approach is to limit the number of aerodynamic GMRES iterations to 5, without restart, and apply the structural preconditioner without any Krylov method. At the aerostructural level, Equation (15) is solved using right-preconditioned flexible GMRES (FGMRES)⁴³ with a maximum Krylov subspace size of 60. This set of preconditioning options may be changed and is set to approximately balance the time required for the aerodynamic and structural preconditioning operations for typical problem sizes.

D. Approximate Newton-Krylov method

We found that the Newton-Krylov method was very effective, but that the high-cost of forming the $\partial\mathcal{A}/\partial X_s$ term within the coupling derivative, made it significantly slower than the nonlinear block Gauss-Seidel approach. In order to achieve better performance, an approximate formulation of the term, $\partial\mathcal{A}/\partial X_s$ was implemented within the aerodynamic code. This approximation is constructed by computing the contributions from panels only within a given radius of one another, in a manner analogous to the technique used to form the approximate aerodynamic preconditioner. This reduces the computational cost of this term significantly. All remaining parameters associated with the Newton-Krylov method are also used for the approximate Newton-Krylov method.

E. Aerostructural analysis results

In order to test the efficiency of the aerostructural solution algorithms described in the previous section, we compare the solution time and residual histories for the aerostructural analysis of the N2A blended wing-body aircraft. The geometry is developed from the results produced by the Silent Aircraft Initiative.⁴⁴ The structural and aerodynamic meshes used for this result are shown in Figure (3). We use a structural model with 21360 degrees of freedom and an aerodynamic model with 4200 surface panels and 3600 wake panels.

The analyses were performed using different numbers of processors in the group of aerodynamic processes while using only one processor in the structural group. Figure (4) shows the parallel speed up for analyses




Figure 3. The structural and aerodynamic meshes used for the aerostructural analysis performance comparison.




Figure 4. Aerostructural solution times required for the nonlinear block Jacobi (J), Newton–Krylov (NK), nonlinear block Gauss–Seidel (GS) and approximate Newton–Krylov methods. The slope of the linear speed up is shown for comparison.




Figure 5. A plot of the residual history for the four different aerostructural analysis methods with 8 processors.

on 2, 4 and 8 processors compared with the ideal speed up. The approximate Newton–Krylov method is clearly the fastest, followed by the nonlinear block Gauss–Seidel, exact Newton–Krylov and block Jacobi approaches.

Figure (5) shows the residual histories for the cases run with 8 processes for all approaches. The long start up time required for the Newton–Krylov method is apparent. The approximate Newton–Krylov method requires more iterations but converges much faster than the Newton–Krylov method. The nonlinear block Gauss–Seidel method converges steadily, while the block Jacobi method converges through a series of oscillations. Figure (6) shows the residual histories for the Newton–Krylov and approximate Newton–Krylov method for 2, 4 and 8 processors.

V. Sensitivity analysis method

A. Theory

Efficient gradient-based optimization requires the accurate and efficient computation of the objective and constraint gradients. In the aerostructural optimization problem, there are typically fewer objective and constraint functions than there are design variables. Thus, an adjoint implementation of the sensitivity equations is appropriate. We have developed an aerostructural adjoint that is based entirely on analytic derivatives. This results in highly accurate results at low computational cost when compared with finite-difference of complex-step calculations.

The aerostructural adjoint equations can be written in the following form,

$$\frac{\partial \mathcal{R}^T}{\partial q} \psi = \frac{\partial f}{\partial q}, \quad (20)$$

where ψ is referred to as the adjoint vector and f is either an aerodynamic or structural function of interest. The total derivative is then determined using the additional computation,

$$\frac{df}{dx} = \frac{\partial f}{\partial x} - \psi^T \frac{\partial \mathcal{R}}{\partial x}. \quad (21)$$

We have implemented this approach for aerodynamic lift, drag and moments as well as the KS function (6), and the minimum load-factor, $f = \min\{\lambda_k\}$.




Figure 6. Residual histories for the approximate Newton–Krylov (ANK) and Newton–Krylov (NK) methods, for aerostructural analysis with 2,4 and 8 processors.

The first step is to compute the adjoint by solving the adjoint equations (20). We use either Gauss–Seidel or a Krylov method to solve these equations. The Gauss–Seidel iteration takes the following form,

$$\begin{aligned} \frac{\partial \mathcal{A}^T}{\partial w} \Delta \psi_A^{(n)} &= \frac{\partial f}{\partial w} - \frac{\partial \mathcal{A}^T}{\partial w} \psi_A^{(n)} - \frac{\partial \mathcal{S}^T}{\partial w} \psi_S^{(n)}, \\ \frac{\partial \mathcal{S}^T}{\partial u} \Delta \psi_S^{(n)} &= \frac{\partial f}{\partial u} - \frac{\partial \mathcal{S}^T}{\partial u} \psi_S^{(n+1)} - \frac{\partial \mathcal{A}^T}{\partial u} \psi_A^{(n)}, \end{aligned} \quad (22)$$

with the updates $\psi_A^{(n+1)} = \psi_A^{(n)} + \Delta \psi_A^{(n)}$ and $\psi_S^{(n+1)} = \psi_S^{(n)} + \Delta \psi_S^{(n)}$. The stopping criterion requires that the residual norm of both disciplines be reduced below a relative tolerance of $\epsilon_{rA} = 10^{-8}$ to ensure accurate gradient information.

The Krylov method solves the adjoint equation (20) using a fully-coupled approach. Matrix-vector products use the exact Jacobian-transpose of the coupled aerostructural system. One iteration of a transpose block Jacobi iteration is used as the preconditioner, with similar settings to those used in the Newton–Krylov solution method. The adjoint equation is again solved to a relative tolerance of $\epsilon_{rA} = 10^{-8}$.

Once the adjoint vector ψ has been determined, the total sensitivities must be computed using Equation (21). These require the partial derivative of the residuals with respect to the design variables. We have implemented the adjoint with geometric, structural and angle of attack as design variables. The geometric design variables tend to be the most difficult to compute requiring the following calculations,

$$\psi_A^T \frac{\partial \mathcal{A}}{\partial x} = \psi_A^T \frac{\partial \mathcal{A}}{\partial X_s} \frac{\partial X_s}{\partial x}, \quad (23)$$

$$\psi_S^T \frac{\partial \mathcal{S}}{\partial x} = \psi_S^T \frac{\partial \mathcal{S}_c}{\partial x} - \psi_S^T \frac{\partial \mathcal{F}}{\partial f_A} \frac{\partial f_A}{\partial X_s} \frac{\partial X_s}{\partial x}. \quad (24)$$

Computations must again take into consideration the distribution of the disciplines across multiple process groups. Equation (23) involves terms only on the aerodynamic processes, while Equation (24) involves terms on both structural and aerodynamic processes. As a result, this operation requires significantly more communication.

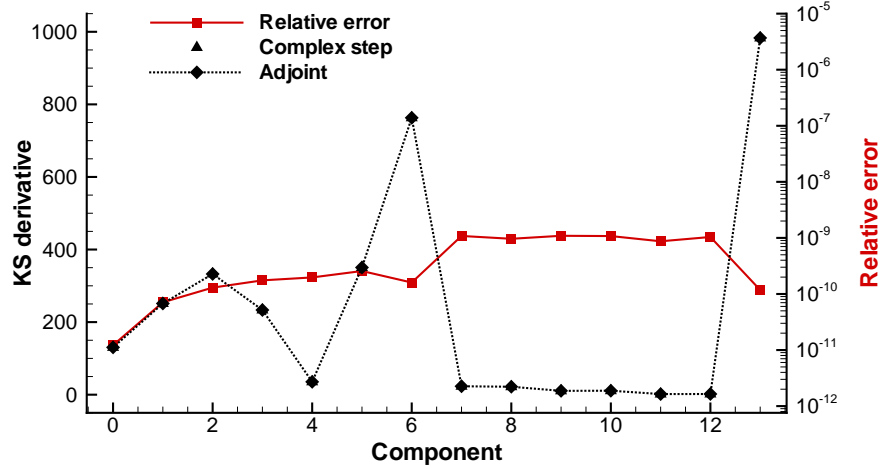


Figure 7. Aerostructural sensitivity study comparison with the complex step method

B. Sensitivity accuracy study

We verify the accuracy of the derivatives for aerodynamic and structural functions of interest using the complex-step method.⁴⁵ In this approach, complex arithmetic is used throughout the entire code and the total derivative is calculated as follows,

$$\frac{df}{dx_i} = \frac{\text{Im}(f(x + he_i))}{h} + O(h^2), \quad (25)$$

where h is a step size and e_i is the i th Cartesian basis vector. The advantage of this formula (25) is that it does not suffer from subtractive cancellation. As a result, very small step sizes may be used, yielding gradients accurate to machine precision. Typically we choose $h = 10^{-30}$.




Figure 8. Aerostructural sensitivity study comparison with the complex step method

We compare the complex-step calculations to the adjoint implementation for a small aerostructural problem with 566 surface panels and 1956 structural degrees of freedom. A small case is chosen to allow rapid testing of all aerodynamic and structural functions. The first five design variables are aerodynamic twist

variables, the next eight are structural thicknesses and final design variable is the angle of attack. Figure (7) shows results for the lift calculation, normalized with the dynamic pressure while Figure (8) shows the results for the KS function. The results demonstrate the high-degree of accuracy of all components. They also show the large difference in magnitudes of the gradient components both within the gradient of a single function as well as between aerodynamic and structural functions. These large differences are a common characteristic of aerostructural optimization.

VI. Results

We apply our aerostructural optimization framework to lift-constrained, drag-minimization of a wing based loosely on planform and weight data from the Bombardier Q400. This data is partially obtained from the manufacturer’s specifications.⁴⁶ The basic geometric parameters and weight build-up are listed in Table (2). We choose a design altitude of 6705m (22,000ft), with a cruise Mach number of $M = 0.58$. The initial wing is composed of a lofted stack of NACA2412 airfoils, in two wing segments. The first segment is without taper from the root to 40% of the semi-span. The second segment follows from 40% of the semi-span to the tip with a taper ratio of 0.54. The initial wing has no twist. The aerodynamic mesh consists of 80 spanwise panels and 70 panels around the section. The wake consists of 60 downstream panels and is projected aft in the free-stream direction. This results in a total of 5600 surface panels and 4800 wake panels. Figure (11(b)) shows the aerodynamic solution at the initial analysis point of the aerostructural optimization problem.

Table 2. Geometry and weight data for the wing optimization

| Parameter | Value | Parameter | Value |
|------------------------|-------|-----------|-------|
| Span [m] | 28.40 | MTOW [kg] | 29275 |
| Area [m ²] | 68.44 | OEW [kg] | 17185 |
| Aspect ratio | 11.78 | Fuel [kg] | 5482 |

We use a structural model of the wing that includes the spars, ribs and skin. The internal wing-structure consists of two spars, the first at 15% of the local chord and the second at 50% of the local chord, and 20 chordwise ribs that extend from 15% to 75% of the chord. These structural components are modeled with 912 fourth-order shell elements with 46290 structural degrees of freedom. Figure (11) shows an exploded of the structural model. We use a material model based on 7075 Aluminum with a Poisson’s ratio $\nu = 0.33$ and Young’s modulus $E = 70\text{GPa}$. We use a design yield stress of $\sigma_{YS} = 320\text{MPa}$.

We calculate the weight of the aircraft based on the weight of the finite-element model plus a fixed-weight. We choose this fixed weight to be MTOW discounted by the approximate wing weight, that we estimate to be 8% of MTOW or 14% of OEW. The fixed-weight coefficient is $C_{\mu_{\text{fixed}}} = 0.3833$, where $C_{\mu_{\text{fixed}}} = m_{\text{fixed}}g/q_{\infty}S_{ref}$, and g is the acceleration due to gravity. We use two analysis points: an on-design flight condition where we enforce $C_L = C_{\mu}$ and a maneuver condition where we enforce $C_L = 2.5C_{\mu}$. These two analyses are run concurrently and in parallel. We use three KS constraint functions (6) at the maneuver point: one each for the front and rear spar and one combined for the top and bottom skins with a factor of safety of $F_S = 2.0$.

Table 3. Optimization results initial and final values of C_L , C_D , and C_{μ} . Note that the lift-constraint is not satisfied at the initial point.

| | Aerostructural | | Aero only |
|-----------|----------------|---------------|---------------|
| | Initial value | Optimal value | Optimal value |
| C_L | 0.6945 | 0.4331 | 0.4331 |
| C_D | 0.01115 | 0.004796 | 0.004752 |
| α | 5.0° | 2.75° | 2.72° |
| C_{μ} | 0.4722 | 0.4331 | N/A |

We use a total of 86 design variables: an angle of attack variable for each flight condition, seven spanwise twist variables and 76 structural thicknesses: 19 each for the thicknesses of each segment of the front and

trailing edge spar and an additional 38 skin thicknesses. We link the skin thickness design variables on the top and bottom at the same spanwise segment. We constrain the thickness of the spars $t \geq 5\text{mm}$ and the thickness of the skins $t \geq 2.5\text{mm}$. We do not include the leading or trailing edge in the structural design problem as these portions of the structure are not accurately represented in the existing finite-element model. We also exclude the chordwise ribs from design problem.

The results of the optimization are summarized in Table (3), along with the results from an aerodynamic-only optimization problem of minimizing drag with a lift-constraint imposed to match the optimal aerostructural lift. The aerodynamic-only optimization and the aerostructural optimization results match closely.

The optimal structural thickness distribution is shown in Figure (9), with the structural mesh shown. The optimizer has made the thickness of the skins and spars increase towards the root with the skin consistently thicker than the spars. This reflects the greater bending-resistance produced by increasing the skin thickness rather than the spar thickness.




Figure 9. The optimal thickness distribution for the wing case. Thicknesses are measured in millimeters.

A summary of the optimal operating condition and maneuver condition results are shown in Figures (11) and (12) respectively. These show the von Mises stress distributions, Figures (11(a)) and (12(a)), the surface C_p distributions, Figures (11(b)) and (12(b)), and the lift and twist distributions with reference elliptic distributions, Figures (11(c)) and (12(c)).

At the operating point the optimal and equivalent elliptic lift distributions are very close, Figure (11(c)). The distributions deviate slightly at the root, where the aerostructural optimum is slightly more loaded, and near the tip where the optimum is slightly less loaded. The optimal lift distribution reflects the relative importance of reducing the drag directly by modifying the deflected twist distribution or by reducing drag by shifting the lift-distribution inboard to reduce stresses and as a result structural weight. The relative importance of weight or drag reduction depends on many factors including the structural model used, the relative weights of the different aircraft components and the geometric variables given to the optimizer. This reflects the many multidisciplinary decisions involved in aircraft design.

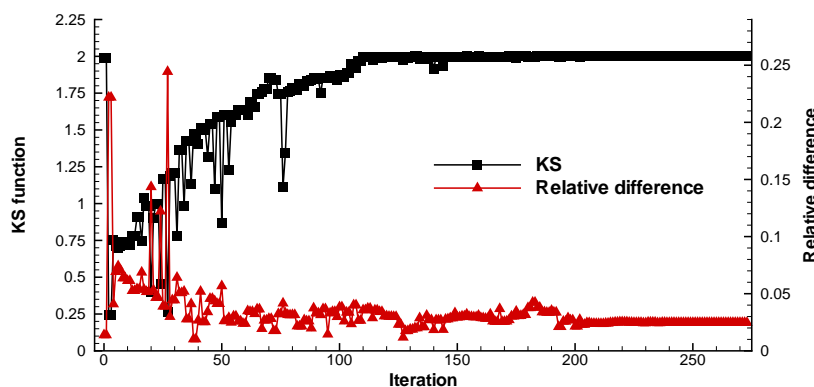


Figure 10. The KS function for the upper and lower skin and the relative difference between the KS value and the minimum failure load factor, plotted against the iteration count for the optimization. Note that these include line search iterations.

One important aspect of the optimization is ensuring that the structural stress constraints are satisfied. Figure (10) shows a comparison of the KS constraint function for the skins and the relative difference between

the KS constraint and the minimum value of the load factor. The relative error is defined as,

$$\text{Relative error} = \frac{\lambda_{KS} - \min\{\lambda_k\}}{\min\{\lambda_k\}}. \quad (26)$$

The relative error is always positive and therefore the KS function provides a conservative estimate of the minimum failure load factor. Furthermore, the relative error is about 2.5% at the final point. This error could be altered by modifying the parameter ρ to achieve a tighter bound.

VII. Conclusion

In this paper we have presented an aerostructural analysis and optimization framework that uses three-levels of parallelism to achieve efficient design optimization. Four aerostructural solution methods were compared and the approximate Newton–Krylov approach was demonstrated to be an efficient and robust solution technique. We also presented an adjoint sensitivity evaluation method that demonstrated a high level of accuracy. These tools can be used for efficient gradient-based aerostructural design optimization. We demonstrated our approach on the aerostructural design optimization of a typical subsonic, turboprop aircraft wing.


Acknowledgments

The authors gratefully acknowledge the financial assistance of the Natural Sciences and Engineering Research Council (NSERC). Computations were performed on the General Purpose Cluster supercomputer at the SciNet HPC Consortium. SciNet is funded by: the Canada Foundation for Innovation under the auspices of Compute Canada; the Government of Ontario; Ontario Research Fund Research Excellence; and the University of Toronto.


References

- ¹Martins, J. R. R. A., Alonso, J. J., and Reuther, J. J., “High-fidelity aerostructural design optimization of a supersonic business jet,” *Journal of Aircraft*, Vol. 41, No. 3, 2004, pp. 523–530.
- ²Martins, J. R. R. A., Alonso, J. J., and Reuther, J. J., “A coupled–adjoint sensitivity analysis method for high–fidelity aero–structural design,” *Optimization and Engineering*, Vol. 6, 2005, pp. 33–62.
- ³Maute, K., Nikbay, M., and Farhat, C., “Coupled analytical sensitivity analysis and optimization of three-dimensional nonlinear aeroelastic systems,” *AIAA Journal*, Vol. 39, No. 11, 2001, pp. 2051–2061.
- ⁴Alonso, J. J., LeGresley, P., van der Weide, E., Martins, J. R. R. A., and Reuther, J. J., “pyMDO: A framework for high-fidelity multi-disciplinary optimization,” *AIAA Paper* 2004-4480, Aug. 2004.
- ⁵Reuther, J. J., Alonso, J. J., Martins, J. R. R. A., and Smith, S. C., “A coupled aero-structural optimization method for complete aircraft configurations,” AIAA paper 99-0187, 1999.
- ⁶Barcelos, M., Bavestrello, H., and Maute, K., “A Schur-Newton-Krylov solver for steady-state aeroelastic analysis and design sensitivity analysis,” *Computer Methods in Applied Mechanics and Engineering*, Vol. 195, No. 17-18, 2006, pp. 2050 – 2069.
- ⁷Chittick, I. R. and Martins, J. R. R. A., “Aero-structural optimization using adjoint coupled post-optimality sensitivities,” *Structures and Multidisciplinary Optimization*, 2007.
- ⁸Chittick, I. R. and Martins, J. R. R. A., “An asymmetric suboptimization approach to aerostructural optimization,” *Optimization and Engineering*, Vol. 10, No. 1, March 2009, pp. 133–152.
- ⁹Kennedy, G. J., Martins, J. R. R. A., and Hansen, J. S., “Aerostructural optimization of aircraft structures using asymmetric subspace optimization,” *12th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference, Victoria, BC*, AIAA, 2008.
- ¹⁰Li, V. and Velicki, A., “Advanced PRSEUS structural concept design and optimization,” *12th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference, Victoria BC*, No. 2008-5840, 2008.
- ¹¹Martins, J. R. R. A., *A coupled-adjoint method for high-fidelity aero-structural optimization*, Ph.D. thesis, Stanford University, Stanford, CA 94305, Oct. 2002.
- ¹²Brown, S., “Displacement extrapolation for CFD+CSM aeroelastic analysis,” AIAA Paper 97-1090, 1997.
- ¹³Maman, N. and Farhat, C., “Matching fluid and structure meshes for aeroelastic computations: a parallel approach,” *Computers and Structures*, Vol. 54, No. 4, 1995, pp. 779–785.
- ¹⁴Farhat, C., Lesoinne, M., and LeTallec, P., “Load and motion transfer algorithms for fluid/structure interaction problems with non-matching discrete interfaces: Momentum and energy conservation, optimal discretization and application to aeroelasticity,” *Computer methods in applied mechanics and engineering*, 1996, pp. 95–115.
- ¹⁵Maute, K. and Allen, M., “Conceptual design of aeroelastic structures by topology optimization,” *Structural and multidisciplinary optimization*, Vol. 27, 2004, pp. 27–42.


- ¹⁶Barcelos, M. and Maute, K., “Aeroelastic design optimization for laminar and turbulent flows,” *Computer Methods in Applied Mechanics and Engineering*, Vol. 197, No. 19-20, 2008, pp. 1813 – 1832.
- ¹⁷Felippa, C. A., Park, K. C., and Farhat, C., “Partitioned analysis of coupled mechanical systems,” *Computer Methods in Applied Mechanics and Engineering*, Vol. 190, No. 24-25, 2001, pp. 3247 – 3270.
- ¹⁸Kim, J., Aluru, N., and Tortorelli, D., “Improved multi-level Newton solvers for fully coupled multi-physics problems,” *International Journal for Numerical Methods in Engineering*, Vol. 58, 2003, pp. 463–480.
- ¹⁹Biros, G. and Ghattas, O., “Parallel Lagrange–Newton–Krylov–Schur methods for PDE-constrained optimization. Part I: The Krylov–Schur solver,” *SIAM Journal on Scientific Computing*, Vol. 27, No. 2, 2005, pp. 687–713.
- ²⁰Biros, G. and Ghattas, O., “Parallel Lagrange–Newton–Krylov–Schur methods for PDE-constrained optimization. Part II: The Lagrange–Newton Solver and its application to optimal control of steady viscous flows,” *SIAM Journal on Scientific Computing*, Vol. 27, No. 2, 2005, pp. 714–739.
- ²¹Heil, M., Hazel, A., and Boyle, J., “Solvers for large-displacement fluidstructure interaction problems: segregated versus monolithic approaches,” *Computational Mechanics*, Vol. 43, 2008, pp. 91–101, 10.1007/s00466-008-0270-6.
- ²²Allen, C. and Rendall, T., “Unified approach to CFD-CSD interpolation and mesh motion using radial basis functions,” *25th AIAA Applied Aerodynamics Conference*, No. 2007-3804, 2007.
- ²³Rendall, T. and Allen, C., “Unified fluid–structure interpolation and mesh motion using radial basis functions,” *International journal for numerical methods in engineering*, Vol. 74, 2008, pp. 1519–1559.
- ²⁴Sobieszcanski-Sobieski, J. and Haftka, R., “Multidisciplinary aerospace design optimization: survey of recent developments,” *Structural Optimization*, Vol. 14, No. 1, 1997, pp. 1–23.
- ²⁵Cramer, E., Dennis, J., Frank, P., Lewis, R., and Shubin, G., “Problem formulation for multidisciplinary optimization,” *SIAM Journal on Optimization*, Vol. 4, No. 4, 1994, pp. 754–776.
- ²⁶Katz, J. and Plotkin, A., *Low-Speed Aerodynamics*, McGraw–Hill Inc., 1991.
- ²⁷J.L.Hess and Smith, A., “Calculation of potential flow about arbitrary bodies,” *Progress in Aerospace Sciences*, Vol. 8, 1967, pp. 1–138.
- ²⁸Smith, S., “A computational and experimental study of nonlinear aspects of induced drag,” Tech. rep.
- ²⁹Balay, S., Buschelman, K., Eijkhout, V., Gropp, W. D., Kaushik, D., Knepley, M. G., McInnes, L. C., Smith, B. F., and Zhang, H., “PETSc Users Manual,” Tech. Rep. ANL-95/11 - Revision 2.1.5, Argonne National Laboratory, 2004.
- ³⁰Balay, S., Gropp, W. D., McInnes, L. C., and Smith, B. F., “Efficient Management of Parallelism in Object Oriented Numerical Software Libraries,” *Modern Software Tools in Scientific Computing*, edited by E. Arge, A. M. Bruaset, and H. P. Langtangen, Birkhäuser Press, 1997, pp. 163–202.
- ³¹Karypis, G. and Kumar, V., “A fast and high quality multilevel scheme for partitioning irregular graphs,” *SIAM Journal on Scientific Computing*, Vol. 20, No. 1, 1998, pp. 359–392.
- ³²Saad, Y. and Schultz, M. H., “GMRES: A generalized minimal residual algorithm for solving nonsymmetric linear systems,” *SIAM Journal on Scientific and Statistical Computing*, Vol. 7, No. 3, 1986, pp. 856–869.
- ³³de Strurler, E., “Nested Krylov methods based on GCR,” *Journal of Computational and Applied Mathematics*, Vol. 67, 1996, pp. 15–41.
- ³⁴de Strurler, E., “Strategies for optimal Krylov subspace methods,” *SIAM Journal on Numerical Analysis*, Vol. 36, No. 3, pp. 864–889.
- ³⁵Wrenn, G., “An indirect method for numerical optimization using the Kreisselmeier-Steinhauser function,” NASA Technical Report CR-4220, 1989.
- ³⁶Raspanti, C., Bandoni, J., and Biegler, L., “New strategies for flexibility analysis and design under uncertainty,” *Computers and Chemical Engineering*, Vol. 24, 2000, pp. 2193–2209.
- ³⁷Poon, N. M. K. and Martins, J. R. R. A., “An adaptive approach to constraint aggregation using adjoint sensitivity analysis,” *Structures and Multidisciplinary Optimization*, Vol. 30, No. 1, 2007, pp. 61–73.
- ³⁸Kenway, G. K., Kennedy, G. J., Mader, C. A., Yan, B., and Martins, J. R., “A CAD-free approach to high-fidelity aerostructural optimization,” *13th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference, Forth Worth, TX, AIAA*, 2010.
- ³⁹Sederberg, T. W. and Parry, S. R., “Free-form deformation of solid geometric models,” *SIGGRAPH Comput. Graph.*, Vol. 20, No. 4, 1986, pp. 151–160.
- ⁴⁰Kelley, C. T., David, and Keyes, D. E., “Convergence analysis of pseudo-transient continuation,” *SIAM J. Num. Anal.*, Vol. 35, 1997, pp. 508–523.
- ⁴¹Allgower, E. L. and Georg, K., “Continuation and path following,” *Acta Numerica*, 1992, pp. 1–64.
- ⁴²Knoll, D. and Keyes, D., “Jacobian-free Newton-Krylov methods: a survey of approaches and applications,” *Journal of Computational Physics*, Vol. 193, No. 2, 2004, pp. 357–397.
- ⁴³Saad, Y., “A flexible inner-outer preconditioned GMRES algorithm,” *SIAM Journal on Scientific Computing*, Vol. 14, No. 2, 1993, pp. 461–469.
- ⁴⁴J.I.Hileman, Spakovszky, Z., Drela, M., Sargent, M., and Jones, A., “Airframe design for silent fuel-efficient aircraft,” *Journal of Aircraft*, Vol. 47, No. 3, 2010, pp. 956–969.
- ⁴⁵Martins, J. R. R. A., Sturdza, P., and Alonso, J. J., “The complex-step derivative approximation,” *ACM Transactions on Mathematical Software*, Vol. 29, No. 3, Sept. 2003, pp. 245–262.
- ⁴⁶Bombardier, “Q400 - Specifications,” <http://www2.bombardier.com/q400/en/specifications.jsp>, August 2010.



(a) von Mises stresses at the operating condition




(b) C_p surface contours at the operating condition




(c) Lift and twist distribution at the operating condition

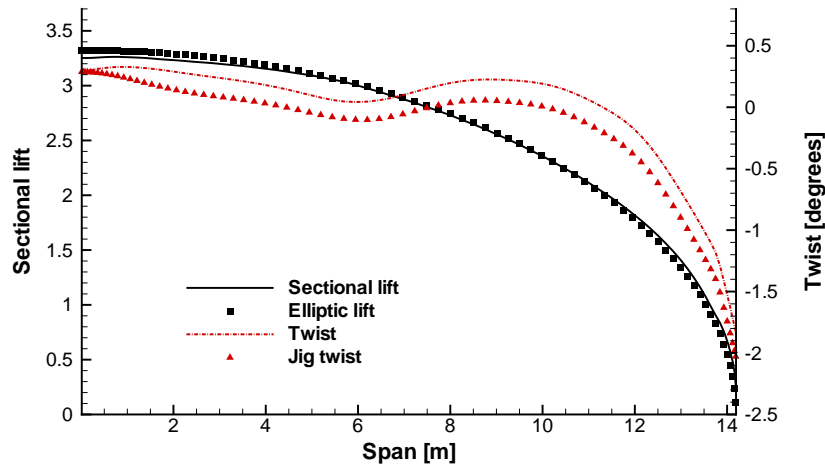
Figure 11. The operating condition for the optimal wing with $C_L = 0.4331$. The top figure shows the distribution of the von Mises stresses. The structure is shown in an exploded view so that the rib and spar structure are visible. The middle figure shows the surface C_p distribution. The bottom figure shows the lift-distribution, a reference elliptic distribution and the jig and flying twist.



(a) von Mises stresses at the maneuver condition



(b) C_p surface contours at the maneuver condition



(c) Lift and twist distribution at the maneuver condition

Figure 12. The maneuver condition for the optimal wing with $C_L = 1.083$. The top figure shows the distribution of the von Mises stresses. The structure is shown in an exploded view so that the rib and spar structure are visible. The middle figure shows the surface C_p distribution. The bottom figure shows the lift-distribution, a reference elliptic distribution and the jig and flying twist.