Optim Eng
DOI 10.1007/s11081-008-9046-2

An asymmetric suboptimization approach to aerostructural optimization

Ian R. Chittick · Joaquim R.R.A. Martins

Received: 13 November 2007 / Accepted: 1 April 2008

Abstract An asymmetric suboptimization method for performing multidisciplinary design optimization is introduced. The objective of the proposed method is to improve the overall efficiency of aerostructural optimization, by simplifying the system-level problem, and thereby reducing the number of calls to a potentially costly aerodynamics solver. To guide a gradient-based optimization algorithm, an extension of the coupled sensitivity equations is developed to include post-optimality information from the structural suboptimization. The optimization of an aircraft wing is performed using linear aerodynamic and structural analyses, and a thorough performance comparison is made between the new approach and the conventional multidisciplinary feasible method. The asymmetric suboptimization method is found to be the more efficient approach when it adequately simplifies the system-level problem, or when there is a large enough discrepancy between disciplinary solution times.

Keywords Asymmetric suboptimization · Coupled post-optimality sensitivity analysis · Multidisciplinary design optimization

1 Introduction

Multidisciplinary design optimization (MDO) is an emerging field of engineering that uses numerical optimization methods to solve design problems involving more than one discipline. By considering the interactions between the multiple disciplines

I.R. Chittick is graduate student.
J.R.R.A. Martins is assistant Professor.

I.R. Chittick (✉) · J.R.R.A. Martins
Institute for Aerospace Studies, University of Toronto, Toronto, Canada
e-mail: ian.chittick@utoronto.ca

J.R.R.A. Martins
e-mail: martins@utias.utoronto.ca

during the optimization process, the optimum achieved using an MDO approach is superior to the one found by optimizing each discipline in sequence. MDO methods have been used in a wide variety of applications, but due to the complex and multidisciplinary nature of aircraft, interest in MDO has been particularly strong in the field of aerospace engineering.

1.1 Aerostructural optimization

One of the most common applications of MDO techniques is coupled aerodynamic and structural (aerostructural) optimization, because of the strong interactions between these two disciplines. Extensive research has been conducted in the field of MDO and its application to aerostructural design. The survey paper by Sobieszczanski-Sobieski and Haftka (1997) provides a comprehensive overview of the work accomplished in this area. Understandably, most of the early studies in aerostructural optimization centered on low-fidelity models, in order to demonstrate the benefits of MDO, and to advocate its use during the preliminary design stages. With the availability of ever faster computing platforms, application of these methods to high-fidelity aerostructural analysis and optimization has followed suit (Chattopadhyay and Pagaldipti 1995; Giunta 2000; Maute et al. 2001). However, due to the computational expenses incurred by more complex solvers, practical implementations have generally been limited to only a few design variables.

Different strategies have emerged to address the high computational costs associated with MDO. The use of variable-fidelity (Marduel et al. 2006) and variable-complexity (Thokala and Martins 2006) models has been examined to decrease computational expense, and the advent of the coupled adjoint method has drastically reduced the cost of performing high-fidelity sensitivity analysis with many design variables (Martins et al. 2005). The approach that we propose herein involves modifying the formulation of the MDO problem, in order to exploit the uneven load balance that is common in aerostructural analysis.

1.2 MDO architectures

The interdisciplinary coupling intrinsic to MDO tends to pose significant organizational and computational challenges, and there exist several different MDO architectures for dealing with this complexity. Work by Alexandrov and Kodiyalam (1998) serves as a good introduction to the various MDO methods, and the πMDO framework developed by Tedford and Martins (2006a, 2006b) presents a platform for standardized comparisons of the different formulations.

The conventional MDO architectures can be divided into two main classes: single-level formulations and multilevel formulations. Single-level formulations, such as multidisciplinary feasible (MDF) and simultaneous analysis and design (SAND), employ a single optimizer that is given control over the entire state of the system (Cramer et al. 1994; Tribes et al. 2005). These methods are often the easiest to implement, but can scale poorly with the number of design parameters and disciplines. Multilevel formulations, which include collaborative optimization (CO) (Kroo 1997; Braun and Kroo 1997) and bi-level system synthesis (BLISS) (Kodiyalam and

Sobieszczanski-Sobieski 2000; Sobieszczanski-Sobieski et al. 2003), divide the original problem into smaller subproblem optimizations. These approaches allow for a higher degree of discipline autonomy, but require a system-level optimizer to manage the interactions between the disciplines.

In this research, we examine a new hybrid architecture that shares characteristics with both the single-level and decompositional formulations. As with the MDF method, the new architecture enforces multidisciplinary feasibility after each evaluation of the coupled system. Similar to multilevel strategies, the new architecture involves a subspace optimization to help decrease the size of the system-level problem. The overall objective of the new approach is to decrease the cost of solving MDO problems that exhibit a large discrepancy between disciplinary solution times, as is often the case in high-fidelity aerostructural optimization.

1.3 Motivation

The trade-off between aerodynamic and structural efficiency in aircraft wings is governed by two main interactions. First, the structural weight affects the required lift, which in turn affects the lift-induced drag. Second, the aerodynamic loads affect the structural deformations, which in turn change the aerodynamic shape. Consequently, performing aerostructural analysis is an iterative process, and in order to obtain a converged state, several evaluations of the two disciplines are needed.

The aerodynamic discipline typically incorporates a computational fluid dynamics (CFD) analysis that solves a set of partial differential equations, which is an iterative process in itself. In contrast, the structural analysis usually consists of a linear finite-element solver with a symmetric stiffness matrix, requiring only a single linear solve operation to obtain the structural state. As a result, the computational cost of an aerodynamic solve is generally at least an order of magnitude greater than a structural solve.

The goal of the present work is to take advantage of this computational imbalance by solving a structural suboptimization problem within the aerostructural analysis module. The idea of combining an aerodynamic analysis with a structural optimization routine is not an entirely new concept. For example, Raveh et al. (2000) integrated maneuver load CFD computations with a structural optimization. However, this type of coupled analysis has not been explicitly formulated for an MDO application. The presence of the subspace optimization means that for each aerodynamic analysis one has to perform a structural optimization, which increases the computational cost of the coupled analysis. The advantage of the proposed method, however, is that it confines all of the structural design variables and constraints within the structural discipline. This simplifies the system-level problem, which should decrease the number of calls to the costly aerodynamic analysis. The overall objective is to achieve a more efficient approach to aerostructural optimization, without having to limit the interdisciplinary coupling. The interactions between the disciplines are modeled exactly, and the fidelity of the analyses is not compromised by any approximation technique, such as a response surface, which are sometimes used in hierarchical architectures (Sobieski and Kroo 2000; Kodiyalam and Sobieszczanski-Sobieski 2000).

2 Aerostructural model

Our work to date has focused on a lower-fidelity model, combining linear aerodynamic and structural analyses. Although the motivation for this research is rooted in the computational expense of nonlinear aerodynamic solvers, our simplified model has proven invaluable in testing the overall procedure and the multilevel convergence of the asymmetric suboptimization method. More importantly, it has allowed a sensitivity method for the new architecture to be developed and verified without being encumbered by long turn-around times in the analysis modules. The drawback of using a linear aerodynamic model is taken into account when discussing the results in Sect. 5.

The aerodynamic analysis employs an inviscid panel code to model the wing, which solves the system,

$$\boldsymbol{A}\boldsymbol{\Gamma} - \boldsymbol{v} = \boldsymbol{0}, \tag{1}$$

where $\boldsymbol{A}$ is the aerodynamic influence coefficients matrix, $\boldsymbol{\Gamma}$ is the vector of panel circulations, and $\boldsymbol{v}$ is the vector of panel boundary conditions, which is simply the local angle of attack of each panel. The aerodynamic discipline also enforces that the wing must produce the lift needed to maintain level flight, i.e.,

$$L - W = 0, \tag{2}$$

where L is the total wing lift and W is the total weight of the aircraft. This requirement can be satisfied by selecting the appropriate angle of attack (α).

The structural model consists of a single wing spar, which is modeled using frame finite-elements to represent a tube-shaped spar. The structural analysis is governed by the following equation,

$$\boldsymbol{K}\boldsymbol{u} - \boldsymbol{f} = \boldsymbol{0}, \tag{3}$$

where $\boldsymbol{K}$ is the stiffness matrix of the structure, $\boldsymbol{u}$ is the displacement vector and $\boldsymbol{f}$ is the vector of external forces. The number of aerodynamic panels and structural elements are variable parameters, and they dictate the fidelity of the discipline analyses. Figure 1 shows the wing discretized with 15 panels and elements, and the structural displacements resulting from a sample load distribution.

The simultaneous solution of (1) through (3) defines the state of the aerostructural system. The state variables of the coupled system are $\boldsymbol{\Gamma}$, α, and $\boldsymbol{u}$. The design variables that we use herein are the jig twist distribution of the wing ($\boldsymbol{\gamma}_{\text{jig}}$) and the wall thicknesses of the tube finite elements ($\boldsymbol{t}$). Note that the $\boldsymbol{A}$ matrix is based on fixed parameters, and the remaining variables that are not state variables have the following dependencies,

$$v = v(\boldsymbol{\gamma}_{\text{jig}}, \boldsymbol{u}, \alpha), \tag{4}$$
$$L = L(\boldsymbol{\Gamma}), \tag{5}$$
$$W = W(\boldsymbol{t}), \tag{6}$$
$$K = K(\boldsymbol{t}), \tag{7}$$
$$f = f(\boldsymbol{\Gamma}). \tag{8}$$

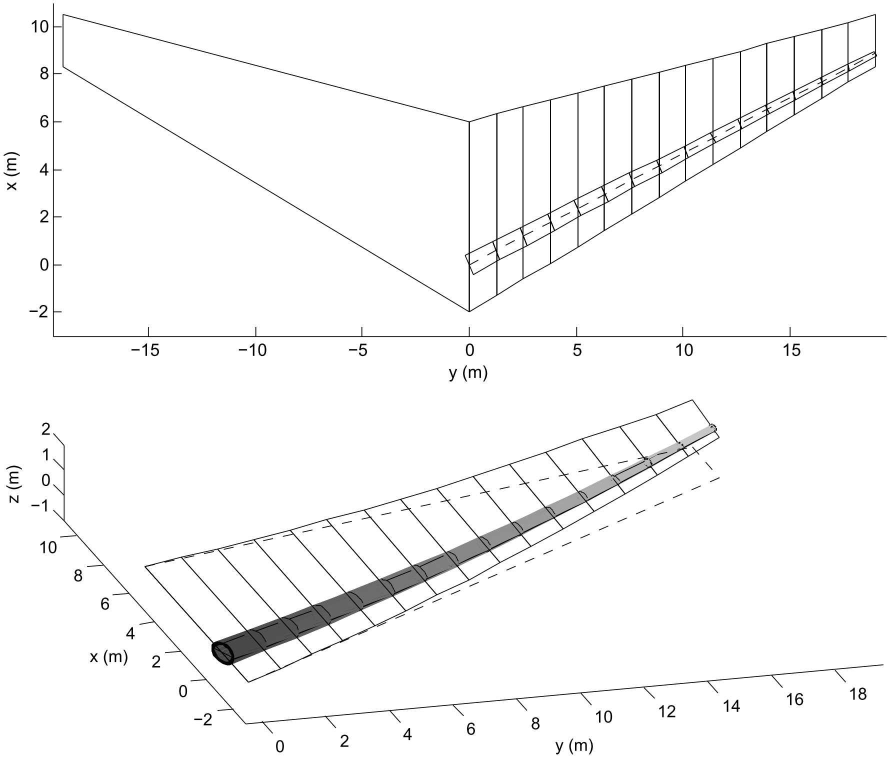

Fig. 1 Aerodynamic and structural discretization of the wing

The objective function of interest is the Breguet range equation. This expression represents the trade-off between the drag and the structural weight of the aircraft, and can be written as,

$$R(D, W) = \frac{V}{c}\frac{L(\boldsymbol{\Gamma})}{D(\boldsymbol{\Gamma})}\ln\frac{W(\boldsymbol{t})}{W(\boldsymbol{t}) - W_f}, \tag{9}$$

where V is the cruise velocity, c is the specific fuel consumption, L/D is the ratio of lift to drag, and $W/(W - W_f)$ is the ratio of initial and final cruise weights of the aircraft. The final weight is simply the structural weight without the fixed fuel weight, W_f. The optimization problem can be stated as follows,

$$\begin{aligned} \text{maximize:} \quad & R(D, W) \\ \text{w.r.t.:} \quad & \boldsymbol{\gamma}_{\text{jig}}, \boldsymbol{t} \\ \text{s.t.:} \quad & \boldsymbol{\sigma}_{\text{yield}} - \boldsymbol{\sigma}(\boldsymbol{u}) \geq \mathbf{0} \\ & \boldsymbol{t} - \boldsymbol{t}_{\text{min}} \geq \mathbf{0} \\ & \boldsymbol{r} - \boldsymbol{t} \geq \mathbf{0}, \end{aligned} \tag{10}$$

where the element stresses ($\boldsymbol{\sigma}$) cannot exceed by the yield stress of the material, and the structural design variables are constrained between the minimum gauge thickness ($\boldsymbol{t}_{\min}$) of the spar material and the radius of the beam elements ($\boldsymbol{r}$). In this research we focus on gradient-based optimization, and the optimizer used is SNOPT, which employs a sequential quadratic programming algorithm (Gill et al. 2005). More specifically we use pySNOPT, a version of the optimization code wrapped in Python (Alonso et al. 2004).

3 Architecture validation

The new asymmetric suboptimization approach to our aerostructural design problem is shown as Fig. 2. Throughout this work, the routine that converges a coupled set of disciplines will be referred to as the multidisciplinary analysis (MDA) module. Therefore, the MDA of the new architecture includes the aerodynamic analysis and the structural suboptimization. The two disciplines are linked within the MDA module through the exchange of coupling variables, which consist of the vector of external forces returned from the aerodynamic analysis, and the vector of displacements as determined by the structural suboptimization. A fixed-point iteration scheme is used to alternate between the two disciplines. Each call to the structural side results in a full gradient-based optimization, where, given a set of loads, the range is maximized with respect to the structural thicknesses, subject to the stress and thickness constraints. Since the aerodynamic parameters are held fixed during each suboptimization, maximizing the range is equivalent to minimizing the weight of the wing. Once the aerostructural system has converged, the suboptimization provides the value of the optimized structural weight (W^*) to the range calculation, while the aerodynamic solver provides the lift and drag values. As mentioned, the benefit of the new formulation is that it leaves the system-level with an unconstrained design space, with only the aerodynamic design variables to consider. The challenge, of course, is

Fig. 2 Asymmetric suboptimization method

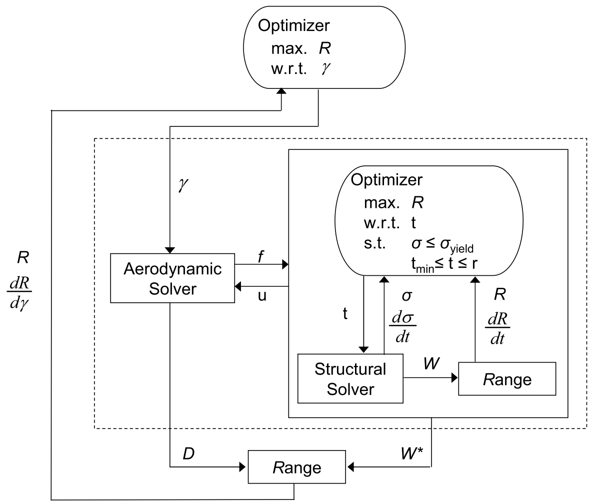

to be able accurately and efficiently calculate the system-level sensitivities, which is discussed in Sect. 4.

3.1 Reference formulations

The multidisciplinary feasible (MDF) method, shown as Fig. 3, is the traditional approach to solving MDO problems, and as such, it is well suited to providing a baseline solution to the design problem. In this monolithic method, a single optimizer is connected to the MDA, which again uses an iteration scheme to converge the aerodynamic and structural solvers. Without a suboptimization routine, all of the design variables and constraints involved in the problem need to be considered at the system-level. As a result, compared to the asymmetric suboptimization approach, the MDF optimizer faces a more challenging design space, thus requiring a larger set of sensitivity terms when determining an appropriate design step.

Sequential optimization, as the name suggests, involves optimizing the two disciplines in sequence. Figure 4 shows the flow of variables for the aerostructural system in a sequential configuration. The coupling variables are exchanged between the two separate optimization routines until a solution is reached. However, this approach does not converge to the true optimal solution. More specifically, the aerodynamic optimization is unaware of the effect that it has on the element stresses by varying the jig twist. The structural optimization is then limited to a design space dictated by the aerodynamic state, which always results in a convergence to an aerodynamic optimum as opposed to the true system optimum. Although not a valid MDO architecture, the sequential method is included in the following discussion to show how an incomplete consideration of the disciplinary interactions results in an inferior design. By including the structural optimization routine within the aerodynamic optimization loop, we will demonstrate that the asymmetric suboptimization method does not suffer the same shortcoming.

Fig. 3 MDF formulation

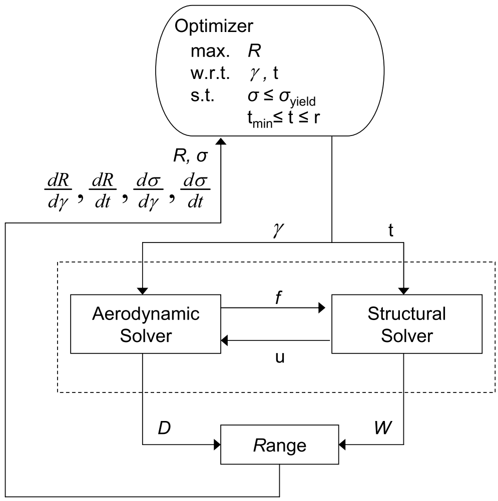

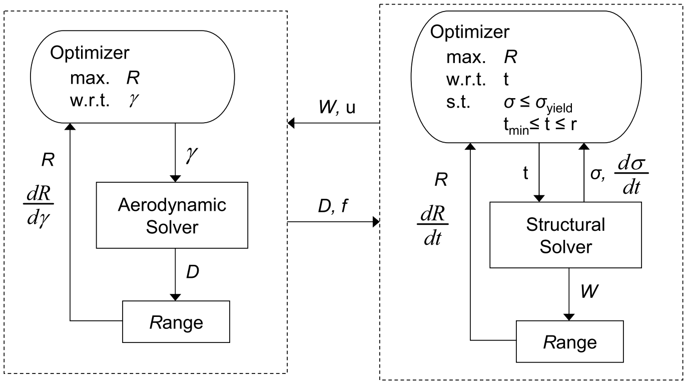

Fig. 4 Sequential method

Table 1 Range results for the three approaches

	MDF	AS	Sequential
L ($\times 10^6$ N)	1.16035	1.16035	1.20964
D ($\times 10^5$ N)	2.35625	2.35625	2.36796
W_i ($\times 10^6$ kg)	1.16035	1.16035	1.20965
W_f ($\times 10^5$ kg)	8.66352	8.66352	9.15648
R (km)	6690.32	6690.32	6613.89

3.2 Reference results

In order to show that our asymmetric suboptimization method is a valid alternative to the standard MDF architecture, our aerostructural model was optimized using both approaches. The results are shown in Fig. 5, along with the sequential results for comparison. These trials were run using 15 jig twist and thickness design variables.

The final designs achieved by the proposed method and the MDF method show very close agreement. The jig twist distributions of both methods result in a deflected wing that twists downwards in the same manner, which is expected for a swept configuration. Similarly, both methods exhibit an identical thickness distribution, which decreases from root to tip until the minimum gauge thickness is reached. In contrast, the sequential approach results in a slightly heavier configuration with a less pronounced wing deflection. The lift distribution graph indicates that the sequential approach does indeed result in a minimum drag design, since it agree very closely with the theoretical aerodynamic optimum, which is given by an elliptical curve. The asymmetric suboptimization and MDF methods, on the other hand, deviate from the elliptical distribution. Instead, the wing loading is shifted rootward, resulting in reduced wing bending moments. This allows for a lighter structure, and a greater range value, as summarized in Table 1. In other words, unlike the sequential method, the MDF and asymmetric suboptimization approaches choose to sacrifice aerodynamic performance in order to achieve greater overall system gains.

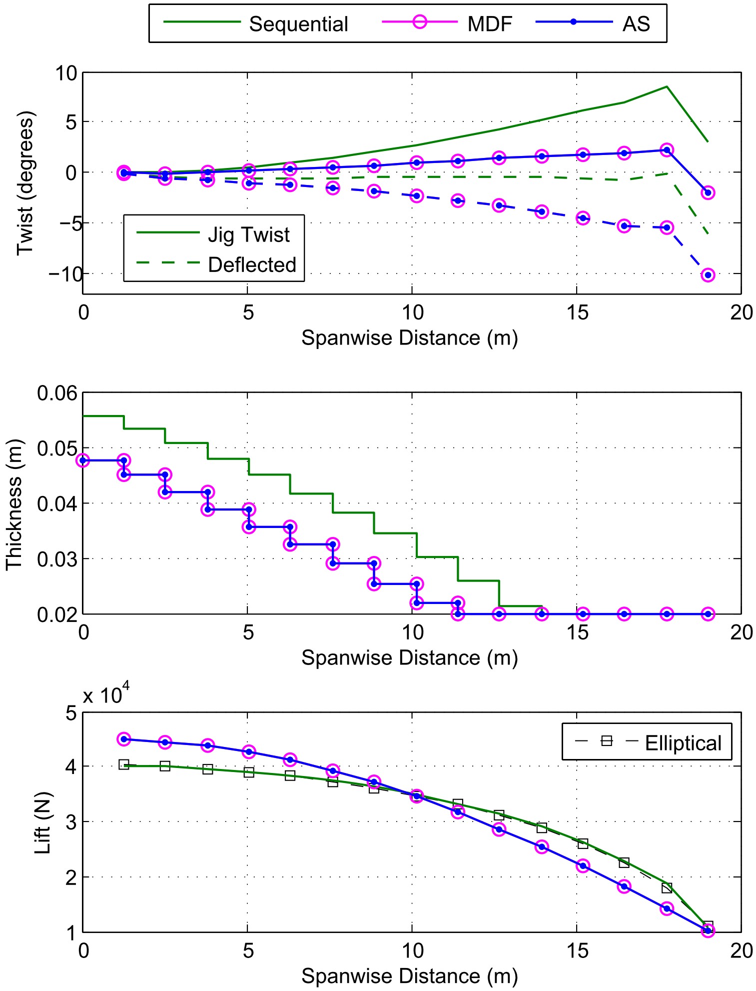

Fig. 5 Reference results for the three formulations. The distributions over the half-span of the wing are shown from root to tip

3.3 Reduced dimensionality

The fact that the asymmetric suboptimization method was able to avoid the suboptimal result of the sequential approach, and achieve the MDF solution, helps to validate the new approach. However, to form a better understanding of how the methods differ, the aerostructural model was simplified to a two design variable problem: a constant beam thickness value and a linear interpolation of the jig twist from a variable wing tip value to a fixed root twist of zero. This allowed the design space to be visualized using a contour plot, as shown in Fig. 6.

The contour lines indicate the range value for a given jig twist and thickness. These range values were obtained from a converged aerostructural solution, so to be

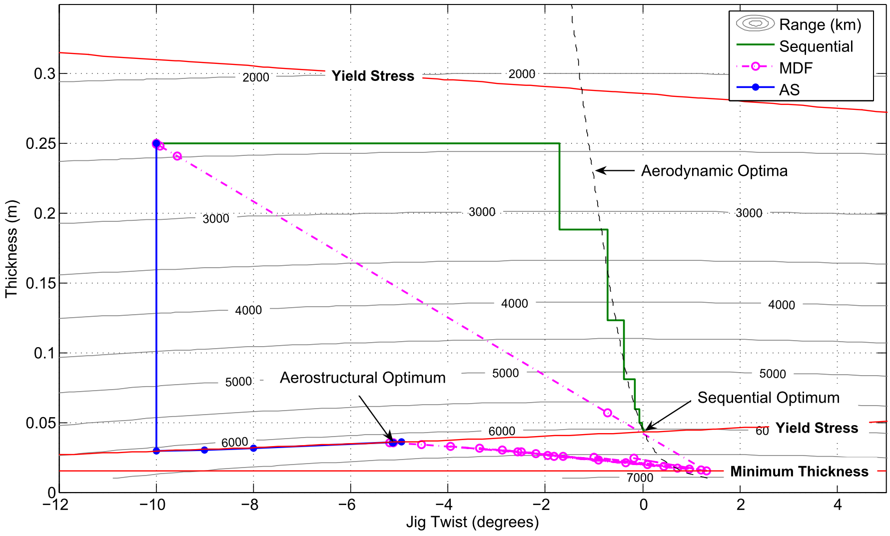

Fig. 6 Contour plot of the simplified aerostructural design space

more accurate, Fig. 6 displays the multidisciplinary feasible design space. Due to a decrease in the structural weight, the aircraft range increases in magnitude as the wing thickness is reduced. The maximum range for a given thickness value is shown with a dashed line, labeled as the aerodynamic optima. Also shown on the contour plot are the boundaries imposed by the stress constraints. The thinner tube elements result in more pronounced structural deflections, which translate into greater stress values. The stresses also slope upwards with higher thickness values, because the added weight increases the required lift and the consequent aerodynamic loading on the wing. Therefore, the feasible region of the design space falls between the top and bottom yield stress contours. For completeness, the minimum gauge thickness is also indicated near the bottom of the contour plot. Beginning from the same initial design point, the range was maximized using the sequential, MDF, and asymmetric suboptimization formulations, and the resulting convergence plots are shown on Fig. 6.

3.3.1 Sequential convergence

The sequential approach exhibits a staircase-style convergence. This is expected because each design variable is being considered in isolation, so the sequential method is only able to proceed along one axis at a time. Each aerodynamic optimization pushes the jig twist distribution towards the dashed line, while every subsequent structural optimization yields a reduced thickness value. Eventually, the method converges where the stress constraint intersects with an aerodynamic optimum. It should be noted that the sequential method actually operates in the individual discipline feasible design space, and does not represent a valid aerostructural state until it arrives at its solution. This explains why the aerodynamic optima and the stress constraints are not reached on Fig. 6 until the method has converged. The inherent flaw of the

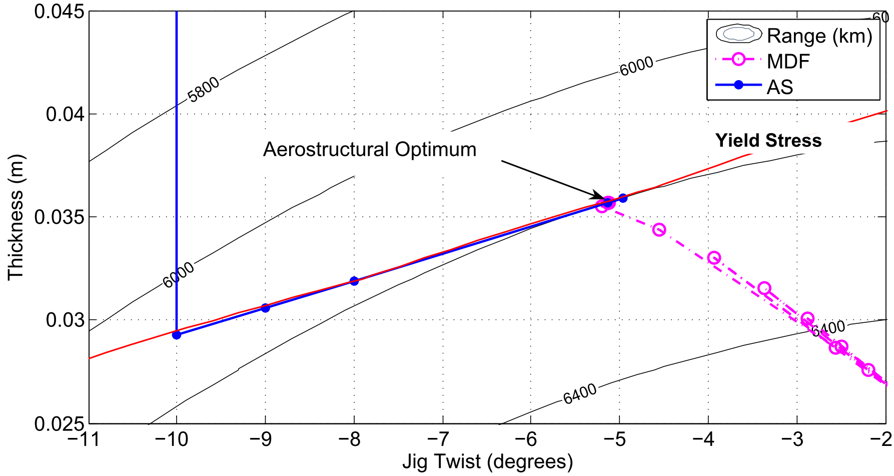

Fig. 7 Magnified view of the aerostructural solution

sequential formulation is an incomplete consideration of the disciplinary interactions. The aerodynamic optimization is unaware of how its state variables effect the structural constraint, and therefore has no reason to give up aerodynamic performance for structural improvement. As a result, the sequential method always converges to an aerodynamic optimum, and in the process, fails to see the broader picture afforded by the fully coupled design space.

3.3.2 MDF convergence

From the initial design point, the MDF method steps towards the unconstrained optimum and stops when the minimum thickness constraint is reached. Operating in the infeasible design space, the MDF method retreats back towards the breached stress constraint. Based on the sensitivity information computed at each intermediate design point, the optimizer is able to take a relatively direct route towards its final solution. Figure 7 shows a closer view of the MDF solution. From this plot it is evident that the MDF method arrives at the true aerostructural optimum, because at the final design point, the stress constraint is tangent to the range contour. In other words, no design step can be taken to improve the objective function without violating the yield stress. By continually considering the interactions between the aerodynamics and structures, the MDF formulation is able to achieve the design that maximizes the range of the coupled system.

3.3.3 Asymmetric suboptimization convergence

Starting from the same initial design as the other two methods, the new architecture steps immediately to the yield stress constraint. The system-level optimizer then selects variable values that move the design along the constraint, converging quickly to the aerostructural optimum. A magnified view of the convergence is shown on Fig. 7, and it is apparent that the new architecture attains the same solution as the MDF method. The presence of the structural subspace optimizer ensures that the MDA module returns a design point that is not only multidisciplinary feasible, but also a structural optimum. One way of describing the convergence of the new architecture is that it travels along a series of structural optima until the aerostructural optimum is encountered. This differs fundamentally from the sequential approach, which converges towards the intersection of an aerodynamic optimum and a structural optimum.

Comparing the convergence of the new architecture against the MDF method, it is visibly apparent on the contour plot that the new architecture requires significantly fewer design steps. It needs to be remembered, however, that each design step requires several evaluations of the structural subspace optimization. Regardless, the new formulation succeeds in simplifying the system-level problem. From the point of view of the system-level optimizer, the structural design variables and constraints have moved behind the scenes, hidden within the analysis used to compute the range values.

4 Coupled post-optimality sensitivity analysis

Sensitivity analysis is a crucial consideration when performing gradient-based optimization, because the derivative calculations are often the most costly step within the optimization cycle. Finite-differencing is not a realistic sensitivity approach for most MDO formulations, since it requires repeated evaluations of the coupled analysis. This is especially true for the new architecture, where the computational cost of an MDA evaluation is particularly high, because of the suboptimization routine. Therefore, it was important to determine a more efficient means of obtaining the multidisciplinary sensitivities.

Due to the presence of the structural suboptimization, the new architecture seemed like a logical application of post-optimality sensitivity analysis. Standard post-optimality analysis allows for the change in the optimum solution with respect to a change in a previously fixed parameter to be attained, without having to perform a re-optimization (Braun et al. 1993). Unfortunately this method does not take into account the coupled nature of the aerostructural system at hand, and it became clear that an extension of the current theory was needed for this work.

The sensitivities of coupled systems can be computed using semi-analytical methods, such as the coupled direct sensitivity equations introduced by (Sobieszczanski-Sobieski 1990), or the coupled adjoint method (Martins et al. 2005). These methods allow for the system-level derivatives to be computed without having to re-solve the multidisciplinary analysis, which greatly reduces the cost and inaccuracy of finite-differencing performed on the entire system analysis. For the aerostructural system involving the coupled aerodynamic residuals ($\boldsymbol{\mathcal{A}}$) and structural residuals ($\boldsymbol{\mathcal{S}}$), as well as the aerodynamic state variables ($\boldsymbol{w}$) and the structural state variables ($\boldsymbol{u}$), the *direct sensitivity equation* can be written as,

$$\begin{bmatrix} \frac{\partial \boldsymbol{\mathcal{A}}}{\partial \boldsymbol{w}} & \frac{\partial \boldsymbol{\mathcal{A}}}{\partial \boldsymbol{u}} \\ \frac{\partial \boldsymbol{\mathcal{S}}}{\partial \boldsymbol{w}} & \frac{\partial \boldsymbol{\mathcal{S}}}{\partial \boldsymbol{u}} \end{bmatrix} \begin{bmatrix} \frac{\mathrm{d}\boldsymbol{w}}{\mathrm{d}\boldsymbol{x}} \\ \frac{\mathrm{d}\boldsymbol{u}}{\mathrm{d}\boldsymbol{x}} \end{bmatrix} = - \begin{bmatrix} \frac{\partial \boldsymbol{\mathcal{A}}}{\partial \boldsymbol{x}} \\ \frac{\partial \boldsymbol{\mathcal{S}}}{\partial \boldsymbol{x}} \end{bmatrix}, \tag{11}$$

where $\boldsymbol{x}$ are the global design variables of interest. The solution of this equation can then be substituted into the *total sensitivity equation*,

$$\frac{\mathrm{d}F}{\mathrm{d}\boldsymbol{x}} = \frac{\partial F}{\partial \boldsymbol{x}} + \frac{\partial F}{\partial \boldsymbol{w}}\frac{\mathrm{d}\boldsymbol{w}}{\mathrm{d}\boldsymbol{x}} + \frac{\partial F}{\partial \boldsymbol{u}}\frac{\mathrm{d}\boldsymbol{u}}{\mathrm{d}\boldsymbol{x}}, \tag{12}$$

to find the total derivative of the system-level objective function, F, with respect to the system-level design variables.

In this context, partial derivatives do not take into account the implicit dependence due to the solution of governing equations, while total derivatives do include this implicit dependence.

4.1 Aerodynamic residuals

For our aerostructural test case, the aerodynamic and structural residuals are easily identified. Since the aerodynamic analysis involves solving (1) and (2), the aerodynamic residuals are the system of equations that result from solving those two equations simultaneously, and can be written as,

$$\boldsymbol{\mathcal{A}} = \begin{bmatrix} \boldsymbol{A} & -\boldsymbol{e} \\ \boldsymbol{e}^T & 0 \end{bmatrix} \begin{bmatrix} \boldsymbol{\Gamma} \\ \alpha \end{bmatrix} - \begin{bmatrix} \boldsymbol{\gamma}_{\mathrm{jig}} + \boldsymbol{\gamma}_{\Delta} \\ nW/qb^2 \end{bmatrix} = \boldsymbol{0}. \tag{13}$$

The local incidence for each panel has been replaced by the individual contributions of jig twist, twist deflection and angle of attack, i.e., $\boldsymbol{v} = -(\boldsymbol{\gamma}_{\mathrm{jig}} + \boldsymbol{\gamma}_{\Delta} + \alpha\boldsymbol{e})$, where $\boldsymbol{e}$ is a vector of ones. The second row is a scalar equation that represents the lift constraint of (2), where n is the number of panels, q is the free stream dynamic pressure, and b is the wing span. The state variables for the aerodynamic residuals are $\boldsymbol{w}^T = [\boldsymbol{\Gamma}^T \quad \alpha]$, where α is the angle of attack of the aircraft in radians.

4.2 Structural residuals

The structural residuals are given simply by (3), i.e.,

$$\boldsymbol{\mathcal{S}} = \boldsymbol{K}\boldsymbol{u} - \boldsymbol{f} = \boldsymbol{0}. \tag{14}$$

The state variables for the structural residuals are the displacements, $\boldsymbol{u}$.

The coupled sensitivities of the preceding aerodynamic and structural residuals can be used in the MDF method to provide system-level gradients. However, for the asymmetric suboptimization architecture, this is not the case. As shown during the validation of the new architecture in Sect. 3, the presence of the subspace optimization changes the nature of the multidisciplinary analysis and significantly alters the design space of the system-level optimizer. Therefore, to be able to calculate accurate sensitivity information from the coupled system, it becomes necessary to take into account the added influence of the suboptimization routine.

4.3 Structural optimization residuals

The structural optimization is performed in order to maximize the aircraft range by varying the internal wing thicknesses, for a given load distribution, subject to the stress and thickness constraints. Since the structural optimization is a constrained, gradient-based problem, the optimizer is working to satisfy the Karush–Kuhn–Tucker (KKT) conditions. The KKT conditions are necessarily satisfied at a structural optimum, but they do not entirely define the suboptimization, because an optimum must also be a structural solution. Therefore, combining the structural residuals (14) with the KKT conditions completes the picture, and allows us to fully describe the structural optimization residuals:

$$\mathcal{O} = \begin{cases} \boldsymbol{K}\boldsymbol{u} - \boldsymbol{f} & \equiv \mathcal{O}_{\mathcal{S}} \\ \frac{\mathrm{d}R}{\mathrm{d}\boldsymbol{t}} - \boldsymbol{\lambda}_\sigma^T \frac{\mathrm{d}\boldsymbol{\sigma}}{\mathrm{d}\boldsymbol{t}} + \boldsymbol{\lambda}_t^T \boldsymbol{I} = \boldsymbol{0} & \equiv \mathcal{O}_{\mathcal{L}} \\ \boldsymbol{\sigma}_{\text{yield}} - \boldsymbol{\sigma} - \boldsymbol{s}_{\boldsymbol{\sigma}}^2 = \boldsymbol{0} & \equiv \mathcal{O}_\sigma \\ \boldsymbol{t} - \boldsymbol{t}_{\min} - \boldsymbol{s}_{\boldsymbol{t}}^2 = \boldsymbol{0} & \equiv \mathcal{O}_t \\ \boldsymbol{s}_\sigma \boldsymbol{\lambda}_\sigma = \boldsymbol{0}, \quad \boldsymbol{s}_t \boldsymbol{\lambda}_t = \boldsymbol{0} & \equiv \mathcal{O}_{s\lambda}, \end{cases} \tag{15}$$

where $\boldsymbol{\lambda}^T = [\boldsymbol{\lambda}_\sigma^T \quad \boldsymbol{\lambda}_t^T]$ are the Lagrange multipliers for the stress and thickness constraints, respectively, and $\boldsymbol{s}^T = [\boldsymbol{s}_\sigma^T \quad \boldsymbol{s}_t^T]$ are the slack variables associated with those constraints. The radius constraints are omitted from the KKT equations because, although they are used to guide the optimizer away from non-physical solutions, they are not active at the optimum. The complete set of optimization state variables consists of four vectors, and thus $\boldsymbol{y}^T = [\boldsymbol{u}^T \quad \boldsymbol{t}^T \quad \boldsymbol{s}^T \quad \boldsymbol{\lambda}^T]$.

The multidisciplinary analysis of the asymmetric suboptimization method can be viewed as containing two separate disciplines: the aerodynamics and the structural optimization. Having identified the structural optimization residuals, $\mathcal{O}$, we are now able to present the corresponding coupled direct sensitivity equation as,

$$\begin{bmatrix} \frac{\partial \mathcal{A}}{\partial \boldsymbol{w}} & \frac{\partial \mathcal{A}}{\partial \boldsymbol{y}} \\ \frac{\partial \mathcal{O}}{\partial \boldsymbol{w}} & \frac{\partial \mathcal{O}}{\partial \boldsymbol{y}} \end{bmatrix} \begin{bmatrix} \frac{\mathrm{d}\boldsymbol{w}}{\mathrm{d}\boldsymbol{x}} \\ \frac{\mathrm{d}\boldsymbol{y}}{\mathrm{d}\boldsymbol{x}} \end{bmatrix} = - \begin{bmatrix} \frac{\partial \mathcal{A}}{\partial \boldsymbol{x}} \\ \frac{\partial \mathcal{O}}{\partial \boldsymbol{x}} \end{bmatrix}, \tag{16}$$

where the total sensitivity equation is,

$$\frac{\mathrm{d}F}{\mathrm{d}\boldsymbol{x}} = \frac{\partial F}{\partial \boldsymbol{x}} + \frac{\partial F}{\partial \boldsymbol{w}} \frac{\mathrm{d}\boldsymbol{w}}{\mathrm{d}\boldsymbol{x}} + \frac{\partial F}{\partial \boldsymbol{y}} \frac{\mathrm{d}\boldsymbol{y}}{\mathrm{d}\boldsymbol{x}}. \tag{17}$$

Recall that the system-level design variables ($\boldsymbol{x}$) for the asymmetric suboptimization method are the jig twists of each panel, whereas for the MDF approach, they include the structural thicknesses as well.

We refer to the preceding approach as the coupled post-optimality sensitivity (CPOS) method, because it relies on post-optimality sensitivity information from the converged suboptimization. Equations (16) and (17) constitute the direct formulation of the CPOS method. Another valid approach is the adjoint formulation of the CPOS equations, which offers the advantage of having a computational cost that is essentially independent of the number of system-level design variables. However, the

Table 2 Sensitivity analysis comparison

$dF/d\mathbf{x}$	Finite-difference	CPOS direct
$dR/d\gamma_2$	−2141.0892	−2141.0912
$dR/d\gamma_3$	2643.3319	2643.3328
$dR/d\gamma_4$	6742.0192	6742.0344
$dR/d\gamma_5$	−1007.6901	−1007.6931
Elapsed time (s)	190.41	6.13

results presented in this paper were obtained using the direct method. The adjoint version of the CPOS method will be investigated in a later paper. For a complete derivation of the individual partial derivative terms in (16), please refer to our previous work (Chittick and Martins 2007).

A disadvantage of the CPOS method is that it requires higher-order sensitivities. This occurs due to the first KKT condition, denoted as $\mathcal{O}_{\mathcal{L}}$ in (15). The $d\boldsymbol{\sigma}/d\boldsymbol{t}$ term in the suboptimization residuals is explicitly dependent on both $\boldsymbol{\Gamma}$ and $\boldsymbol{t}$, which leads to second-order derivatives when evaluating the partial terms of (16). Although costly to compute, these derivatives do not disqualify the CPOS approach. As we will show in the following section, the CPOS method offers significant computational savings over the finite-difference approach.

4.4 CPOS validation

The CPOS equations that we developed for the asymmetric suboptimization method are verified against the finite-difference approach in Table 2. The results correspond to an analysis with five jig twist and five thickness design variables, as well as five panels and five elements. The jig twist of the first panel was held fixed. The partial derivatives in the CPOS equations are calculated using the complex-step derivative approximation (Squire and Trapp 1998; Martins et al. 2003), which provides numerically exact values.

As shown in Table 2, the CPOS results agree with finite-differences to at least five digits. In all likelihood, the CPOS results are more accurate due to subtractive cancellation errors in the finite-difference estimates. This means that the CPOS method accurately models the response of our tightly coupled system, without having to reconverge the analysis or even once re-optimize the subspace problem. As a result, the computational time of the CPOS method was only 3% of the time required by finite-differences.

5 Results and discussion

The CPOS method was applied in a performance study comparing the asymmetric suboptimization method against the traditional MDF approach. For both architectures, a fixed-point iteration scheme was used to converge the MDA, and the optimization tolerance was set at 10^{-6}. The corresponding coupled sensitivity formulations were used to provide the system-level gradients. The MDO trials were

Table 3 Trials of variable dimensionality comparing the MDF method and the asymmetric suboptimization architecture; times are measured in seconds

Trial	N_γ	N_t	N_{panel}	N_{elem}	Method	MDA Evaluations	Avg. MDA Iter per Eval.	Elapsed Time	Structural Time	Aerodynamic Time	MDA Sensitivities Time	Required T_A/T_S
1	5	5	30	30	MDF	42	19.31	1795.08	217.34	53.61	1521.25	**4.12**
					AS	13	11.46	2482.53	1198.49	9.92	1273.63	
2	5	**10**	30	30	MDF	44	19.57	1946.54	230.59	56.59	1655.93	**9.57**
					AS	12	11.50	3751.31	1771.32	9.05	1970.11	
3	5	**15**	30	30	MDF	46	19.26	2108.94	237.71	58.76	1808.37	**19.30**
					AS	13	11.38	5881.22	2727.80	9.78	3142.72	
4	**10**	5	30	30	MDF	55	19.27	2377.41	284.14	69.99	2019.12	**3.87**
					AS	17	12.47	3200.29	1483.58	13.97	1701.53	
5	**15**	5	30	30	MDF	83	18.94	3623.88	421.59	103.96	3084.45	**2.22**
					AS	23	12.22	4307.06	1974.65	18.47	2312.28	
6	5	5	30	**60**	MDF	47	19.45	5143.85	504.19	60.67	4575.06	**0.09**
					AS	12	10.83	4769.04	1916.69	8.57	2842.93	
7	5	5	30	**90**	MDF	49	20.00	10068.43	839.88	64.89	9158.84	**0.04**
					AS	12	10.08	7961.15	2940.91	8.01	5012.18	
8	5	5	**60**	30	MDF	41	19.29	4386.45	230.92	193.33	3957.21	**0.44**
					AS	12	11.83	3616.58	1044.63	35.06	2534.70	
9	5	5	**90**	30	MDF	38	18.74	8522.87	219.92	386.73	7909.54	**0.30**
					AS	11	10.73	5053.07	872.88	64.25	4112.10	

performed based on four separate parameters: the number of aerodynamic design variables (N_γ), the number of structural design variables (N_t), the number of aerodynamic panels (N_{panel}), and the number of structural elements (N_{elem}). These four parameters were varied independently in order to determine the effect that each one had on the architectures being compared. The computations were run on a 1.5 GHz Itanium 2 processor in an SGI Altix with 32 GB RAM, and the results are presented in Table 3.

5.1 Reference trial

Trial 1, with five aerodynamic and five structural design variables, 30 panels and 30 elements, was selected as the reference case. Table 3 shows that for this trial, the MDF architecture required over three times as many system-level iterations as the new architecture. As discussed during our preliminary study, by redistributing the design variables and constraints, the asymmetric suboptimization method simplifies the system-level problem, thus needing fewer MDA evaluations to find the optimum. Trial 1 also shows the new architecture requiring on average eight fewer iterations than the MDF method to converge the MDA. This occurs because the presence of the structural optimization in the MDA module makes the wing stiffer than it would be otherwise (because the stress constraints need to be satisfied). The stiffer wing limits the structural displacements, resulting in better convergence of the aerostructural cycle.

The timings of the first trial, listed in Table 3, do not show the true potential of the new architecture. For the reference trial, the new architecture takes almost 700 seconds longer overall than the MDF method. This is largely due to the total time spent performing structural analysis and optimization, which is the consequence of having an optimizer in the MDA module. In addition, the time required to calculate the system-level sensitivities after each MDA solution is greater for the new architecture, due to the aforementioned second-order derivatives. However, with fewer MDA iterations, the asymmetric suboptimization succeeds in reducing the amount of time performing aerodynamics analysis, accomplishing its intended purpose. Unfortunately, due to the linear aerodynamics model being used in our analysis, this does not translate into an advantage. Therefore, we now derive the performance metric used to evaluate the new approach: the ratio of the aerodynamic solver time to structural solver time required to make the asymmetric suboptimization and MDF cost the same overall.

The total time required by MDF can be broken down as follows,

$$T_{\text{total}} = N(T_{\text{A}} + T_{\text{S}}) + T_{\text{O}}, \tag{18}$$

where N is the total number of aerostructural iterations, T_{A} is the time for an aerodynamic analysis, T_{S} is the time for a structural analysis, and T_{O} represents the overhead, which is the rest of the time and is composed primarily of the time spent computing the MDA sensitivities. Similarly, the total time of the new method can be broken down as,

$$T'_{\text{total}} = N'(T_{\text{A}} + T_{\text{SO}}) + T'_{\text{O}}, \tag{19}$$

where T_{SO} is the time required for a single structural suboptimization. Equating the two total times and solving for the ratio yields,

$$\frac{T_{\text{A}}}{T_{\text{S}}} = \frac{1}{N' - N}\left[\frac{T_{\text{O}} - T'_{\text{O}}}{T_{\text{S}}} + N - N'\frac{T_{\text{SO}}}{T_{\text{S}}}\right]. \tag{20}$$

This equation expresses the ratio of computational cost between the aerodynamic and structural solvers that is needed for the new method to break even with the MDF approach. A ratio greater than the one calculated using this formula indicates a situation where the new method is the more efficient approach.

For Trial 1, the computed ratio is 4.12, and is shown as the last column in Table 3. Since higher-fidelity scenarios involve an aerodynamic solution time at least an order of magnitude greater than the structural solution, *a time ratio less than ten is a favorable result for the new architecture*. Certain high-fidelity frameworks might even encounter aerodynamic-to-structural time ratios of up to 10^3, depending on the particular solvers being used.

5.2 Increasing the number of structural design variables

Trials 2 and 3 in Table 3 show the effect of increasing the number of structural design variables. For the MDF method, the added design variables result in a few additional optimization iterations, which translates into proportional increases in the structural, aerodynamic, and sensitivity analysis times. For the asymmetric suboptimization method, the additional structural variables do not change the system-level problem.

The extra structural variables complicate the subspace optimization of the new architecture, causing longer structural times. The constraint gradients are the main cause for this delay, because they are computed using the complex-step method, which scales with the number of design variables. Increasing N_t also slows down the CPOS calculations. The thickness variables act as state variables of the structural optimization residuals, so increasing N_t increases the size of the sensitivity matrices. Consequently, increasing N_t exhibits an unfavorable trend for the asymmetric suboptimization method, and the required aerodynamic-to-structural time ratio increases significantly.

5.3 Increasing the number of aerodynamic design variables

Trials 4 and 5 in Table 3 show the effect of increasing the number of aerodynamic design variables. Increasing the number of jig twists has a considerable effect on the number of MDF iterations required. The number of system-level iterations of the new architecture increases as well, but not as drastically, due to the absence of constraints.

For both architectures the structural, aerodynamic, and MDA sensitivity times increase relative to the reference trial. However, this increase is directly proportional to the increase in MDA evaluations, and as a result the trend favors the asymmetric suboptimization approach. The required aerodynamic-to-structural time ratio decreases, and a further increase in N_γ would result in the new formulation being more computationally efficient than the MDF approach.

5.4 Increasing the structural model fidelity

Trials 6 and 7 demonstrate what happens with a larger number of structural elements. This increases both the fidelity of the structural analysis and the number of stress constraints. The MDF method experiences a slight increase in MDA evaluations due to the extra constraints, while the system-level problem of the new architecture remains unchanged.

As N_{elem} increases, the number of structural degrees-of-freedom multiplies. These additional states affect the structural suboptimization, increasing the time required to solve the structural residuals and to compute the constraints and their sensitivities. However, the increase in time needed to evaluate the MDA sensitivities is the dominating factor. Since the MDF method takes roughly four times as many MDA evaluations as the new architecture, the total time dedicated to calculating the MDF sensitivities is much greater. As a result, the asymmetric suboptimization method becomes the more computational efficient approach for both trials. The aerodynamic-to-structural time ratio of less than one indicates that even with an aerodynamic solver that is less costly than the structures, the asymmetric suboptimization method can still be the superior approach.

5.5 Increasing the aerodynamic model fidelity

Finally, Trials 8 and 9 reveal the effect of an increased number of panels, and consequently, an increased number of aerodynamic states. This leads to a more costly aerodynamic analysis, and Trial 9 is the first MDF trial where the aerodynamic analysis is more expensive to compute than the structural analysis.

The larger aerodynamic state vector also increases the MDA sensitivity times of both architectures. Similar to the previous case, the total MDF sensitivity time increases more rapidly due to the higher number of MDA evaluations. Thus the required aerodynamic-to-structural time ratio decreases with increasing panels, and once again, the asymmetric suboptimization method proves to be the faster approach for both trials.

6 Conclusions

We presented an asymmetric suboptimization approach to aerostructural optimization, motivated by the unequal load balance between the disciplinary solvers. We demonstrated that the proposed method converges along a series of structural optima, and in the process, is able to achieve the true aerostructural optimum. Compared to the traditional MDF method, the new approach possesses a simplified system-level problem, resulting in fewer calls to the MDA module, fewer total iterations, and fewer evaluations of the aerodynamic analysis.

We introduced the coupled post-optimality sensitivity method, which expands the standard coupled sensitivity equations to include the structural optimization residuals. The resulting method allows the coupled sensitivities to be computed without having to re-converge the analysis, making the asymmetric suboptimization method a viable approach.

A performance comparison was made between the asymmetric suboptimization method and the MDF method. The study showed that the new method exhibits an unfavorable trend when the number of structural design variables are increased, and a favorable trend when either the number of aerodynamic design variables or the fidelity of the analyses are increased. This indicates that the asymmetric suboptimization method becomes the more efficient approach either when it adequately reduces the number of MDA evaluations, or when there is a large enough discrepancy between disciplinary solution times.

The results confirm the authors' prediction that the asymmetric suboptimization method becomes more advantageous as the computational expense of the aerodynamic analysis is increased. Future work will incorporate a more complex aerodynamics solver, coupled with a composite wing-box structure involving both stress and buckling analysis. The eventual goal is a high-fidelity implementation of the architecture and sensitivity method.

References

Alexandrov NM, Kodiyalam S (1998) Initial results of an MDO evaluation survey. AIAA paper 98-4884
Alonso JJ, LeGresley P, van der Weide E, Martins JRRA, Reuther JJ (2004) pyMDO: a framework for high-fidelity multi-disciplinary optimization. In: Proceedings of the 10th AIAA/ISSMO multidisciplinary analysis and optimization conference, Albany, NY, AIAA 2004-4480, 2004
Braun RD, Kroo IM (1997) Development and application of the collaborative optimization architecture in a multidisciplinary design environment. In: Alexandrov N, Hussaini MY (eds) Multidisciplinary design optimization: state of the art. SIAM, Philadelphia, pp 98–116
Braun RD, Kroo IM, Gage PJ (1993) Post-optimality analysis in aerospace vehicle design. In: Proceedings of the AIAA aircraft design, systems and operations meeting, Monterey, CA, AIAA 93-3932, 1993
Chattopadhyay A, Pagaldipti N (1995) A multidisciplinary optimization using semi-analytical sensitivity analysis procedure and multilevel decomposition. Comput Math Appl 29(7):55–66
Chittick IR, Martins JRRA (2007) A new subspace optimization method for aero-structural design. In: Proceedings of the 3rd AIAA multidisciplinary design optimization specialist conference, Waikiki, HI, AIAA 2007-1867, 2007
Cramer EJ, Dennis JE, Frank PD, Lewis RM, Shubin GR (1994) Problem formulation for multidisciplinary optimization. SIAM J Optim 4(4):754–776
Gill PE, Murray W, Saunders MA (2005) User's guide for SNOPT version 7, a Fortran package for large-scale nonlinear programming. Systems Optimization Laboratory, Stanford University, Stanford, CA 94305, 2005
Giunta AA (2000) A novel sensitivity analysis method for high fidelity multidisciplinary optimization of aero-structural systems. AIAA Paper 2000-0683
Kodiyalam S, Sobieszczanski-Sobieski J (2000) Bilevel integrated system synthesis with response surfaces. AIAA J 38(8):1479–1485
Kroo IM (1997) MDO for large-scale design. In: Alexandrov N, Hussaini MY (eds) Multidisciplinary design optimization: state-of-the-art. SIAM, Philadelphia, pp 22–44
Marduel X, Tribes C, Trépanier J-Y (2006) Variable-fidelity optimization: efficiency and robustness. Optim Eng 7:479–500
Martins JRRA, Sturdza P, Alonso JJ (2003) The complex-step derivative approximation. ACM Trans Math Softw 29(3):245–262
Martins JRRA, Alonso JJ, Reuther JJ (2005) A coupled-adjoint sensitivity analysis method for high-fidelity aero-structural design. Optim Eng 6(1):33–62
Maute K, Nikbay M, Farhat C (2001) Coupled analytical sensitivity analysis and optimization of three-dimensional nonlinear aeroelastic systems. AIAA J 39(11):2051–2061
Raveh DE, Levy Y, Karpel M (2000) Structural optimization using computational aerodynamics. AIAA J 38(10):1974–1982
Sobieski IP, Kroo IM (2000) Collaborative optimization using response surface estimation. AIAA J 38(10):1931–1938
Sobieszczanski-Sobieski J (1990) Sensitivity of complex, internally coupled systems. AIAA J 28(1):153–160
Sobieszczanski-Sobieski J, Haftka RT (1997) Multidisciplinary aerospace design optimization: survey of recent developments. Struct Optim 14(1):1–23
Sobieszczanski-Sobieski J, Altus TD, Phillips M, Sandusky R (2003) Bilevel integrated system synthesis for concurrent and distributed processing. AIAA J 41(10):1996–2003
Squire W, Trapp G (1998) Using complex variables to estimate derivatives of real functions. SIAM Rev 40(1):110–112
Tedford NP, Martins JRRA (2006a) Comparison of MDO architectures within a universal framework. In: Proceedings of the 2nd AIAA multidisciplinary design optimization specialist conference, Newport, RI, AIAA 2006-1617, 2006
Tedford NP, Martins JRRA (2006b), On the common structure of MDO problems: a comparison of architectures. In: Proceedings of the 11th AIAA/ISSMO multidisciplinary analysis and optimization conference, Portsmouth, VA, AIAA 2006-7080, 2006
Thokala P, Martins JRRA (2006) Variable complexity methods applied to airfoil design. Eng Optim 39(3):271–286
Tribes C, Dube JF, Trépanier JY (2005) Decomposition of multidisciplinary optimization problems: formulations and application to a simplified wing design. Eng Optim 37(8):775–796