

Time-dependent adjoint-based aerodynamic shape optimization applied to helicopter rotors

Asitav Mishra ^{*} Karthik Mani [†] Dimitri Mavriplis [‡]
Jay Sitaraman [§]

Department of Mechanical Engineering, University of Wyoming, Laramie, WY 82071-3295.

ABSTRACT

A formulation for sensitivity analysis of fully coupled time-dependent aeroelastic problems is given in this paper. Both forward sensitivity and adjoint sensitivity formulations are derived that correspond to analogues of the non-linear aeroelastic analysis problem. Both sensitivity analysis formulations make use of the same iterative disciplinary solution techniques used for analysis, and make use of an analogous coupling strategy. The information passed between fluid and structural solvers is dimensionally equivalent in all cases, enabling the use of the same data structures for analysis, forward and adjoint problems. The fully coupled adjoint formulation is then used to perform rotor blade design optimization for a four bladed HART2 rotor in hover conditions started impulsively from rest. The effect of time step size and mesh resolution on optimization results is investigated. Results indicate that good optimization results can be obtained using time steps as large as 2 degrees, and that optimizations obtained on the coarse mesh level can be used to initialize fine mesh optimization problems in order to reduce overall computational effort. Future work will focus on aeroelastic optimizations for forward flight cases.

INTRODUCTION

In the recent past, the use of adjoint equations has become a popular approach for solving aerodynamic design optimization problems based on computational fluid dynamics (CFD) (Refs. 1–6). Adjoint equations are a very powerful tool in the sense that they allow the computation of sensitivity derivatives of an objective function to a set of given inputs at a cost which is essentially independent of the number of inputs. This is in contrast to the brute-force finite-difference method, where each input or design variable has to be perturbed individually to obtain a corresponding effect on the output. This is a tedious and costly process which is of little use when there are a large number of design variables or inputs. Another major shortcoming of the finite-difference method is that it suffers from step-size limitations which affect the accuracy of the computed gradients.

While the use of adjoint equations is now fairly well established in steady-state shape optimization, only recently have inroads been made into extending them to unsteady flow problems. Unsteady discrete adjoint-based shape optimization was initially demonstrated in the context of two-dimensional problems by Mani and Mavriplis (Ref. 7) and

also by Rumpfkeil and Zingg (Ref. 8). Preliminary demonstration of the method’s feasibility in three-dimensional problems was done by Mavriplis (Ref. 9). Full implementation in a general sense and application to large scale problems involving helicopter rotors was then carried out by Nielsen et. al. in the NASA FUN3D code (Refs. 10, 11).

Since engineering optimization is an inherently multidisciplinary endeavor, the next logical step involves extending adjoint methods to multidisciplinary simulations and using the obtained sensitivities for driving multidisciplinary optimizations. In the context of fixed and especially rotary wing aircraft, aeroelastic coupling effects can be very important and must be considered in the context of a successful optimization strategy.

The coupling of computational fluid dynamics (CFD) and computational structural dynamics (CSD) and the use of sensitivity analysis on such a system has been addressed in the past primarily from a steady-state standpoint (Refs. 12, 13). Until now, relatively little work has been done addressing unsteady aeroelastic optimization problems, mainly due to complexities in the linearization of coupled time-dependent systems. In previous work, we have derived the fully coupled adjoint problem for a two-dimensional aeroelastic airfoil problem and demonstrated the use of adjoint-derived sensitivities for performing time-dependent aeroelastic optimization including flutter suppression (Ref. 14). This formulation was subsequently extended to time-dependent three-dimensional aeroelastic problems in reference (Ref. 15). This work built upon a previously demonstrated time-dependent aerodynamic optimization capability that was applied to helicopter rotors in reference (Ref. 16) through the addition of a Hodges-Dowell

^{*}Postdoctoral Research Associate; amishra3@uwyo.edu

[†]Associate Research Scientist; kmani@uwyo.edu

[‡]Professor; mavripl@uwyo.edu

[§]Assistant Professor; jsitaram@uwyo.edu

Presented at the AHS 70th Annual Forum, Montreal, Quebec, CA May 20–22, 2014. Copyright © 2014 by the American Helicopter Society International, Inc. All rights reserved.

type beam finite-element model to simulate the rotor structure, and the development of the fully coupled discrete adjoint of the resulting aeroelastic system. This previous work emphasized the formulation, implementation, and verification of the adjoint sensitivity analysis approach for time dependent coupled aeroelastic problems. In the current work, we demonstrate the effectiveness of using this approach for performing aeroelastic optimization of a representative rotorcraft configuration.

Because high fidelity time-dependent optimization represents a computationally intensive approach, obtaining a suitable optimization result with manageable computational requirements is an important consideration. Therefore, a particular aspect of this work considers the suitability of optimization results obtained on relatively coarse meshes and using larger time steps, in order to reduce overall computational effort.

AERODYNAMIC ANALYSIS AND SENSITIVITY FORMULATION

Flow Solver Analysis Formulation

The base flow solver used in this work is the NSU3D unstructured mesh Reynolds-averaged Navier-Stokes solver. NSU3D has been widely validated for steady-state and time-dependent flows and contains a discrete tangent and adjoint sensitivity capability which has been demonstrated previously for optimization of steady-state and time-dependent flow problems. As such, only a concise description of these formulations will be given in this paper, with additional details available in previous references (Refs. 16–18). The flow solver is based on the conservative form of the Navier-Stokes equations which may be written as:

$$\frac{\partial \mathbf{U}(\mathbf{x}, t)}{\partial t} + \nabla \cdot \mathbf{F}(\mathbf{U}) = 0 \quad (1)$$

For moving mesh problems these are written in arbitrary Lagrangian-Eulerian (ALE) form as:

$$\frac{\partial V \mathbf{U}}{\partial t} + \int_{dB(t)} [\mathbf{F}(\mathbf{U}) - \dot{\mathbf{x}} \mathbf{U}] \cdot \mathbf{n} dB = 0 \quad (2)$$

Here V refers to the area of the control volume, $\dot{\mathbf{x}}$ is the vector of mesh face or edge velocities, and \mathbf{n} is the unit normal of the face or edge. The state vector \mathbf{U} consists of the conserved variables and the Cartesian flux vector $\mathbf{F} = (\mathbf{F}_x, \mathbf{F}_y, \mathbf{F}_z)$ contains both inviscid and viscous fluxes. The equations are closed with the perfect gas equation of state and the Spalart-Allmaras turbulent eddy viscosity model (Ref. 19) for all cases presented in this work.

The solver uses a vertex-centered median dual control volume formulation that is second-order accurate, where the inviscid flux integral S around a closed control volume is discretized as:

$$S = \int_{dB(t)} [\mathbf{F}(\mathbf{U}) - \dot{\mathbf{x}} \mathbf{U}] \cdot \mathbf{n} dB = \sum_{i=1}^{n_{edge}} \mathbf{F}_{e_i}^\perp (V_{e_i}, \mathbf{U}, \mathbf{n}_{e_i}) B_{e_i} \quad (3)$$

where B_e is the face area, V_e is the normal face velocity, \mathbf{n}_e is the unit normal of the face, and \mathbf{F}_e^\perp is the normal flux across the face. The normal flux across the face is computed using the second-order accurate matrix dissipation scheme (Ref. 20) as the sum of a central difference and an artificial dissipation term as shown below,

$$\mathbf{F}_e^\perp = \frac{1}{2} \{ \mathbf{F}_L^\perp(\mathbf{U}_L, V_e, \mathbf{n}_e) + \mathbf{F}_R^\perp(\mathbf{U}_R, V_e, \mathbf{n}_e) + \kappa^{(4)} [T] |\lambda| [T]^{-1} \{ (\nabla^2 \mathbf{U})_L - (\nabla^2 \mathbf{U})_R \} \} \quad (4)$$

where $\mathbf{U}_L, \mathbf{U}_R$ are the left and right state vectors and $(\nabla^2 \mathbf{U})_L, (\nabla^2 \mathbf{U})_R$ are the left and right undivided Laplacians computed for any element i as

$$(\nabla^2 \mathbf{U})_i = \sum_{k=1}^{neighbors} (\mathbf{U}_k - \mathbf{U}_i) \quad (5)$$

The time derivative term is discretized using a second-order accurate backward-difference formula (BDF2) scheme as:

$$\frac{\partial V \mathbf{U}}{\partial t} = \frac{\frac{3}{2} V^n \mathbf{U}^n - 2 V^{n-1} \mathbf{U}^{n-1} + \frac{1}{2} V^{n-2} \mathbf{U}^{n-2}}{\Delta t} \quad (6)$$

The index n is used to indicate the current time-level as the convention throughout the paper. The discretization of the BDF2 scheme shown in equation (6) is based on a uniform time-step size.

Denoting the spatially discretized terms at time level n by the operator $S^n(\mathbf{U}^n)$, the resulting system of non-linear equations to be solved for the analysis problem at each time step can be written as:

$$\mathbf{R}^n = \frac{\frac{3}{2} V^n \mathbf{U}^n - 2 V^{n-1} \mathbf{U}^{n-1} - \frac{1}{2} V^{n-2} \mathbf{U}^{n-2}}{\Delta t} + S^n(\mathbf{U}^n) = 0 \quad (7)$$

which in simplified form exhibiting the functional dependencies on \mathbf{U} and \mathbf{x} at different time levels is given as:

$$\mathbf{R}^n(\mathbf{U}^n, \mathbf{U}^{n-1}, \mathbf{U}^{n-2}, \mathbf{x}^n, \mathbf{x}^{n-1}, \mathbf{x}^{n-2}) = 0 \quad (8)$$

At each time step n , the implicit residual is linearized with respect to the unknown solution vector \mathbf{U}^n and solved for using Newton's method as:

$$\begin{aligned} \left[\frac{\partial \mathbf{R}^k}{\partial \mathbf{U}^k} \right] \delta \mathbf{U}^k &= -\mathbf{R}^k \\ \mathbf{U}^{k+1} &= \mathbf{U}^k + \delta \mathbf{U}^k \\ \delta \mathbf{U}^k &\rightarrow 0, \mathbf{U}^n = \mathbf{U}^k \end{aligned} \quad (9)$$

The Jacobian matrix is inverted iteratively using a line-implicit agglomeration multigrid scheme that can also be used as a preconditioner for a GMRES Krylov solver (Ref. 21).

Although the above equation denotes the solution at a single time level n , for the remainder of this paper we will use the generalized notation:

$$\mathbf{R}(\mathbf{U}, \mathbf{x}) = 0 \quad (10)$$

where the vector \mathbf{U} denotes the flow values at all time steps, and where each (block) row in this equation corresponds to the solution at a particular time step as given in equation (8). Equation (10) denotes the simultaneous solution of all time steps and is solved in practice by Newton's method using forward block substitution (i.e. forward integration in time) since each new time step depends on the previous two time levels.

Mesh deformation capability

In order to deform the mesh for time-dependent problems a spring analogy and a linear elastic analogy mesh deformation approach have been implemented. The linear elasticity approach has proven to be much more robust and is used exclusively in this work. In this approach, the mesh is modeled as a linear elastic solid with a variable modulus of elasticity that can be prescribed either as inversely proportional to cell volume or to the distance of each cell from the nearest wall (Refs. 22, 23). The resulting equations are discretized and solved on the mesh in its original undeformed configuration in response to surface displacements using a line-implicit multigrid algorithm analogous to that used for the flow equations. The governing equations for mesh deformation can be written symbolically as:

$$\mathbf{G}(\mathbf{x}, \mathbf{D}) = \mathbf{0} \quad (11)$$

where \mathbf{x} denotes the interior mesh coordinates and \mathbf{D} denotes shape parameters that define the surface geometry.

AERODYNAMIC SENSITIVITY ANALYSIS FORMULATION

The basic sensitivity analysis implementation follows the strategy developed in references (Refs. 9, 17). Consider an arbitrary objective function L that is evaluated using the unsteady flow solution set \mathbf{U} and unsteady mesh solution set \mathbf{x} expressed as:

$$L = L(\mathbf{U}, \mathbf{x}) \quad (12)$$

Assuming that the state variables (i.e. \mathbf{U}, \mathbf{x}) are dependent on some input design parameters \mathbf{D} , the total sensitivity of the objective function L to the set of design inputs can be expressed as the inner product between the vector of state sensitivities to design inputs and the vector of objective sensitivities to the state variables as:

$$\frac{dL}{d\mathbf{D}} = \begin{bmatrix} \frac{\partial L}{\partial \mathbf{x}} & \frac{\partial L}{\partial \mathbf{U}} \end{bmatrix} \begin{bmatrix} \frac{\partial \mathbf{x}}{\partial \mathbf{D}} \\ \frac{\partial \mathbf{U}}{\partial \mathbf{D}} \end{bmatrix} \quad (13)$$

The non-linear flow residual operator and the linear elasticity mesh residual operator as described earlier provide the constraints which can be expressed in general form over the whole space and time domains as:

$$\mathbf{G}(\mathbf{x}, \mathbf{D}) = \mathbf{0} \quad (14)$$

$$\mathbf{R}(\mathbf{U}, \mathbf{x}) = \mathbf{0} \quad (15)$$

which when linearized with respect to the design inputs yields:

$$\begin{bmatrix} \frac{\partial \mathbf{G}}{\partial \mathbf{x}} & 0 \\ \frac{\partial \mathbf{R}}{\partial \mathbf{x}} & \frac{\partial \mathbf{R}}{\partial \mathbf{U}} \end{bmatrix} \begin{bmatrix} \frac{\partial \mathbf{x}}{\partial \mathbf{D}} \\ \frac{\partial \mathbf{U}}{\partial \mathbf{D}} \end{bmatrix} = \begin{bmatrix} -\frac{\partial \mathbf{G}}{\partial \mathbf{D}} \\ 0 \end{bmatrix} \quad (16)$$

These constitute the forward sensitivity or tangent sensitivity equations. In practice these equations may be solved by forward substitution following:

$$\frac{\partial \mathbf{G}}{\partial \mathbf{x}} \frac{\partial \mathbf{x}}{\partial \mathbf{D}} = -\frac{\partial \mathbf{G}}{\partial \mathbf{D}} \quad (17)$$

$$\frac{\partial \mathbf{R}}{\partial \mathbf{U}} \frac{\partial \mathbf{U}}{\partial \mathbf{D}} = -\frac{\partial \mathbf{R}}{\partial \mathbf{x}} \frac{\partial \mathbf{x}}{\partial \mathbf{D}} \quad (18)$$

$\frac{\partial \mathbf{G}}{\partial \mathbf{D}}$ represents the sensitivity of the surface mesh points to a change in the shape parameter \mathbf{D} which can be considered as a known input vector. Thus, the first equation may be solved to obtain the grid sensitivities $\frac{\partial \mathbf{x}}{\partial \mathbf{D}}$ which can then be used in the solution of the second equation to obtain the flow sensitivities $\frac{\partial \mathbf{U}}{\partial \mathbf{D}}$. In both cases, the required inversion of each disciplinary Jacobian matrix can be accomplished following the same solution procedure used for the corresponding analysis problem, i.e. following equation (9) for the flow equations. Since these equations represent the solution over the entire time domain, in practice this procedure is performed at each time step proceeding from the initial to final time step. The mesh and flow sensitivity vectors can then be substituted into equation (13) to obtain the complete sensitivity of the objective with respect to the design variable \mathbf{D} .

The forward sensitivity approach requires a new solution of equation (16) for each design parameter \mathbf{D} . On the other hand, the adjoint approach can obtain the sensitivities for any number of design inputs \mathbf{D} at a cost which is independent of the number of design variables. The adjoint problem can be obtained by pre-multiplying equation (16) by the inverse of the large coupled matrix and substituting the resulting expression for the sensitivities into equation (13) and defining adjoint variables as the solution of the system:

$$\begin{bmatrix} \frac{\partial \mathbf{G}^T}{\partial \mathbf{x}} & \frac{\partial \mathbf{R}^T}{\partial \mathbf{x}} \\ 0 & \frac{\partial \mathbf{R}^T}{\partial \mathbf{U}} \end{bmatrix} \begin{bmatrix} \Lambda_{\mathbf{x}} \\ \Lambda_{\mathbf{U}} \end{bmatrix} = \begin{bmatrix} \frac{\partial L^T}{\partial \mathbf{x}} \\ \frac{\partial L^T}{\partial \mathbf{U}} \end{bmatrix} \quad (19)$$

where $\Lambda_{\mathbf{U}}$ and $\Lambda_{\mathbf{x}}$ are the flow and mesh adjoint variables respectively. The final objective sensitivities can be obtained as:

$$\frac{dL^T}{d\mathbf{D}} = \begin{bmatrix} \frac{\partial \mathbf{G}^T}{\partial \mathbf{D}} & 0 \end{bmatrix} \begin{bmatrix} \Lambda_{\mathbf{x}} \\ \Lambda_{\mathbf{u}} \end{bmatrix} \quad (20)$$

The adjoint system can be solved using back-substitution




Fig. 1. 15 degrees of freedom beam element with flap, lag, torsional and axial degrees of freedom.

as:

$$\frac{\partial \mathbf{R}^T}{\partial \mathbf{U}} \Lambda_{\mathbf{U}} = \frac{\partial L^T}{\partial \mathbf{U}} \quad (21)$$

$$\frac{\partial \mathbf{G}^T}{\partial \mathbf{x}} \Lambda_{\mathbf{x}} = -\frac{\partial \mathbf{R}^T}{\partial \mathbf{x}} \Lambda_{\mathbf{U}} + \frac{\partial L^T}{\partial \mathbf{x}} \quad (22)$$

where once again the inversion of the transposed Jacobian matrices can be accomplished using the same iterative techniques as applied to the analysis and tangent sensitivity problems. Recalling that equation (19) applies over the entire time domain, the back-substitution procedure leads to a reverse integration in time, beginning with the last physical time step and proceeding to the initial time step. A more detailed description of the complete formulation is presented in (Ref. 14), where the procedure has been used to perform aerodynamic shape optimization for a rigid rotor.

BEAM MODEL: ANALYSIS AND ADJOINT OPTIMIZATION

A non-linear bend-twist beam model is a suitable and widely utilized structural model for slender fixed and rotary wing aircraft structures within the context of an aeroelastic problem. A bend-twist beam model, described below, has previously been developed and coupled to the NSU3D unstructured mesh Reynolds-averaged Navier-Stokes solver (Refs. 23, 24).

Beam Analysis Formulation

The non-linear governing equations of a slender beam are discretized using the Finite-element method (FEM) in space. Figure 1 shows a typical beam with 15 degrees of freedom for each element to accommodate bend wise, lag wise, axial and torsional displacements. The second order equation of motion for the beam can be expressed as:

$$[M]\ddot{\mathbf{q}} + [C]\dot{\mathbf{q}} + [K]\mathbf{q} = \mathbf{F} \quad (23)$$

Table 1. Comparison of Hart-II Natural Frequencies (Ref. 25)

Modes	Present Model	UMARC	DLR
Flap 1	1.104	1.112	1.125
Flap 2	2.802	2.843	2.835
Flap 3	5.010	5.189	5.168
Torsion 1	3.878	3.844	3.845

where $[M]$, $[C]$ and $[K]$ are mass, damping and stiffness matrices of the system of equations representing the beam. Vector $\mathbf{F} = \mathbf{F}(t)$ is the forcing vector. Vector \mathbf{q} represents the displacements along all degrees of freedom. This set of equations can be reduced to a first order system and solved using a second order backward difference formula (BDF2) time integration with standard Newton-type linearization and sub-iterations to efficiently invert the implicit system:

$$[I]\dot{\mathbf{Q}} + [A]\mathbf{Q} = \mathbf{F} \quad (24)$$

where $[I]$ is the identity matrix, $\mathbf{Q} = [\mathbf{q}, \dot{\mathbf{q}}]^T$, $\mathbf{F} = [0, [M]^{-1}\mathbf{F}]^T$ and $[A] = \begin{bmatrix} 0 & -[I] \\ [M]^{-1}[K] & [M]^{-1}[C] \end{bmatrix}$. The residual of the structural equations can be defined as: $\mathbf{J} = [I]\dot{\mathbf{Q}} + [A]\mathbf{Q} - \mathbf{F} = 0$, and can be expressed in a simplified form as:

$$\mathbf{J}(\mathbf{Q}, \mathbf{F}) = 0 \quad (25)$$

The beam model has been validated (Ref. 15) for the standard Hart-2 rotor case (Ref. 26) by comparing its natural frequency predictions with the predictions from other reliable CSD models. As shown in Table 1, the rotating natural frequencies compare well with those predicted by the UMARC (Ref. 27) and DLR (Ref. 28) structural codes, for the first 3 flap and first torsional frequency modes. Further, Fig. 2 shows that the frequency prediction over a range of operating rotor frequencies compare well with those predicted by the UMARC comprehensive code (Ref. 29).

Forward Sensitivity Formulation of Beam Model

The beam tangent (forward sensitivity) linearization is similar to the analysis problem. For a given function, L , its sensitivity with respect to a blade design parameter, D can be written as: $\frac{dL}{dD} = \frac{\partial L}{\partial D} + \frac{\partial L}{\partial \mathbf{Q}} \frac{\partial \mathbf{Q}}{\partial D}$. This requires solving for sensitivity of the beam state (\mathbf{Q}), which can be obtained by differentiating Eqn. (25) with respect to the design variable D and rearranging as:

$$\begin{bmatrix} \frac{\partial \mathbf{J}}{\partial \mathbf{Q}} \end{bmatrix} \frac{\partial \mathbf{Q}}{\partial D} = -\frac{\partial \mathbf{J}}{\partial \mathbf{F}} \frac{\partial \mathbf{F}}{\partial D} - \frac{\partial \mathbf{J}}{\partial D} \quad (26)$$

The last term on the right hand side is non zero for structural design parameters such as beam element stiffnesses, in which case the applied force does not change with the design parameter, making the first term on the right hand size zero. In the coupled aeroelastic case, using aerodynamic shape parameters that primarily affect the airloads on the structure, the first




Fig. 2. Fan plot comparing Beam model with UMARC (Ref. 29)

term on the right-hand side is non-zero while the second term vanishes. Solving for $\frac{\partial Q}{\partial D}$ in Eqn. (26), the forward sensitivity of the objective function $\frac{dL}{dD}$ can be obtained.

Adjoint Formulation of Beam Model

The adjoint formulation of the beam model can be derived by approaching the tangent formulation in the reverse (transpose) direction. Taking the transpose of the objective functional sensitivity yields:

$$\frac{dL^T}{dD} = \frac{\partial L^T}{\partial D} + \frac{\partial Q^T}{\partial D} \frac{\partial L^T}{\partial Q} \quad (27)$$

This requires solving for the transpose sensitivity of the beam state (Q). The solution of $\frac{\partial Q^T}{\partial D}$ can be derived from the transposed Eqn. (26):

$$\frac{\partial Q^T}{\partial D} = \left[-\frac{\partial J}{\partial D} - \frac{\partial J}{\partial F} \frac{\partial F}{\partial D} \right]^T \left[\frac{\partial J}{\partial Q} \right]^{-T} \quad (28)$$

Substituting the above into Eqn. (27):

$$\frac{dL^T}{dD} = \frac{\partial L^T}{\partial D} + \left[-\frac{\partial J^T}{\partial D} - \frac{\partial F^T}{\partial D} \frac{\partial J^T}{\partial F} \right] \left[\frac{\partial J}{\partial Q} \right]^{-T} \frac{\partial L^T}{\partial Q} \quad (29)$$

This requires solving for an adjoint vector Λ_Q defined as:

$$\left[\frac{\partial J}{\partial Q} \right]^T \Lambda_Q = \frac{\partial L^T}{\partial Q} \quad (30)$$

The above forms the adjoint formulation of the beam model. It is observed here again that the left hand side Jacobian term of the adjoint step is just the transpose of the Jacobian in the forward linearization. In this work, forward and adjoint formulations of the beam solver have been implemented and verified for both structural design parameters, and force-based

design parameters (as required for the coupled aeroelasticity problem). The implementation includes sensitivities for both static and dynamic beam motion sensitivities. The adjoint implementation has been verified for a tip loaded (harmonically varying tip force) unsteady beam. The adjoint derived sensitivities of the beam shape due to changes in design variables, such as element stiffnesses EI , GJ , and element masses, have been shown to compare well with sensitivities computed using complex variable methods and forward sensitivity methods to within machine precision (10^{-15}).

Having established the correctness and accuracy of the structural adjoint formulation, the potential for using this approach to drive time-dependent structural optimization problems has been explored using the large-scale bound constraint optimization tool (L-BFGS-B) (Ref. 30) as a precursor to their use in fully coupled aero-structural optimization problems.

FULLY COUPLED FLUID-STRUCTURE ANALYSIS FORMULATION

Fluid-structure interface (FSI)

In addition to the solution of the aerodynamic problem and the structural dynamics problem, the solution of the fully coupled time-dependent aeroelastic problem requires the exchange of aerodynamic loads from the CFD solver to the beam structure, which in turn returns surface displacements to the fluid flow solver. The governing equations for the FSI can be written in residual form as:

$$S(\mathbf{F}_b, \mathbf{Q}, \mathbf{F}(\mathbf{x}, \mathbf{u})) = \mathbf{F}_b - [T(\mathbf{Q})]\mathbf{F}(\mathbf{x}, \mathbf{u}) = 0 \quad (31)$$

$$S'(\mathbf{x}_s, \mathbf{Q}) = \mathbf{x}_s - [T(\mathbf{Q})]^T \mathbf{Q} = 0 \quad (32)$$

respectively for the forces transferred to the structural solver and displacements returned to the flow solver. In these equations, $[T]$ represents the transfer matrix which projects point-wise CFD surface forces $\mathbf{F}(\mathbf{x}, \mathbf{u})$ onto the individual beam elements resulting in the beam forces \mathbf{F}_b . The transpose of this matrix is used to obtain the CFD surface displacements \mathbf{x}_s from the beam degrees of freedom \mathbf{Q} . Also note that $[T]$ is a function of \mathbf{Q} since the transfer patterns change with the beam deflection.

General solution procedure

The aeroelastic problem consists of multiple coupled sets of equations namely, the mesh deformation equations, the flow equations (CFD), the beam model-based structural equations, and the fluid-structure interface transfer equations. The system of equations to be solved at each time step can be written as:

$$\mathbf{G}(\mathbf{x}, \mathbf{x}_s(\mathbf{Q})) = \mathbf{0} \quad (33)$$

$$\mathbf{R}(\mathbf{u}, \mathbf{x}) = \mathbf{0} \quad (34)$$

$$S(\mathbf{F}_B, \mathbf{Q}, \mathbf{F}(\mathbf{x}, \mathbf{u})) = \mathbf{0} \quad (35)$$

$$\mathbf{J}(\mathbf{Q}, \mathbf{F}_B) = \mathbf{0} \quad (36)$$

$$S'(\mathbf{x}_s, \mathbf{Q}) = \mathbf{0} \quad (37)$$

where \mathbf{S} and \mathbf{S}' represent the residuals of the FSI equations, and \mathbf{J} represents the residual of the structural analysis problem. Note that the mesh motion residual now depends also on any surface deflections \mathbf{x}_s introduced by the structural model.

Within each physical time step, solution of the fully coupled fluid structure problem consists of performing multiple coupling iterations on each discipline using the latest available values from the other disciplines. Thus the coupled iteration strategy proceeds as:

$$\left[\frac{\partial \mathbf{G}}{\partial \mathbf{x}} \right] \Delta \mathbf{x}^c = -\mathbf{G}(\mathbf{x}^{c-1}, \mathbf{x}_s^{c-1}) \quad (38)$$

$$\left[\frac{\partial \mathbf{R}}{\partial \mathbf{u}} \right] \Delta \mathbf{u}^c = -\mathbf{R}(\mathbf{u}^{c-1}, \mathbf{x}^c) \quad (39)$$

for the flow equations, where the superscript c denotes the coupling iteration index, and the variables are updated as $\mathbf{x}^c = \mathbf{x}^{c-1} + \Delta \mathbf{x}^c$ and $\mathbf{u}^c = \mathbf{u}^{c-1} + \Delta \mathbf{u}^c$. This is followed by the solution of the FSI and structural model as:

$$\mathbf{F}_b^c = [\mathbf{T}(\mathbf{Q}^c)] \mathbf{F}(\mathbf{x}^c, \mathbf{u}^c) \quad (40)$$

$$\left[\frac{\partial \mathbf{J}}{\partial \mathbf{Q}} \right] \Delta \mathbf{Q}^c = -\mathbf{J}(\mathbf{Q}^c, \mathbf{F}_b^c) \quad (41)$$

$$\mathbf{x}_s = [\mathbf{T}]^T \mathbf{Q}^c \quad (42)$$

In this implementation, subiterations are performed to converge the first two equations simultaneously for the final values of \mathbf{Q}^c , while the third equation corresponds to an explicit evaluation for the \mathbf{x}_s given the \mathbf{Q}^c values.

At the first coupling iteration, $\mathbf{x}_s^c = \mathbf{0}$ and solution of the mesh deformation equation is trivial, although non zero values of \mathbf{x}_s are produced at subsequent coupling iterations as the beam deflects in response to the aero loads. From a disciplinary point of view, the aerodynamic solver produces updated values of \mathbf{u} and \mathbf{x} , which are used to compute $\mathbf{F}(\mathbf{x}, \mathbf{u})$ pointwise surface forces. These surface forces are input to the FSI/structural model which returns surface displacements \mathbf{x}_s . These new surface displacements are then fed back into the mesh deformation equations and the entire procedure is repeated until convergence is obtained for the full coupled aerostructural problem at the given time step.

SENSITIVITY ANALYSIS FOR COUPLED AEROELASTIC PROBLEM

In the formulation of the sensitivity analysis for the coupled aeroelastic problem, it is desirable to mimic as closely as possible the solution strategies and data structures employed for the analysis problem. Thus, analogous disciplinary solvers can be reused for each disciplinary sensitivity problem, and the analysis coupling strategy can be extended to the sensitivity analysis formulation. Furthermore, the data transferred between disciplinary solvers should consist of vectors of the same dimension for the analysis, tangent and adjoint formulations. Starting with the forward sensitivity problem, the sen-

sitivity of an objective L can be written as:

$$\frac{dL}{d\mathbf{D}} = \left[\begin{array}{cc} \frac{\partial L}{\partial \mathbf{x}} & \frac{\partial L}{\partial \mathbf{u}} \end{array} \right] \left[\begin{array}{c} \frac{\partial \mathbf{x}}{\partial \mathbf{D}} \\ \frac{\partial \mathbf{u}}{\partial \mathbf{D}} \end{array} \right] \quad (43)$$

where the individual disciplinary sensitivities are given as the solution of the coupled system:

$$\left[\begin{array}{cccccc|cccc} \frac{\partial \mathbf{G}}{\partial \mathbf{x}} & 0 & 0 & 0 & 0 & \frac{\partial \mathbf{G}}{\partial \mathbf{x}_s} & \frac{\partial \mathbf{x}}{\partial \mathbf{D}} & \frac{\partial \mathbf{u}}{\partial \mathbf{D}} & \frac{\partial \mathbf{F}}{\partial \mathbf{D}} & \frac{\partial \mathbf{F}_b}{\partial \mathbf{D}} & \frac{\partial \mathbf{Q}}{\partial \mathbf{D}} & \frac{\partial \mathbf{x}_s}{\partial \mathbf{D}} \\ \frac{\partial \mathbf{R}}{\partial \mathbf{x}} & \frac{\partial \mathbf{R}}{\partial \mathbf{u}} & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ -\frac{\partial \mathbf{F}}{\partial \mathbf{x}} & -\frac{\partial \mathbf{F}}{\partial \mathbf{u}} & I & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & \frac{\partial \mathbf{S}}{\partial \mathbf{F}} & \frac{\partial \mathbf{S}}{\partial \mathbf{F}_b} & \frac{\partial \mathbf{S}}{\partial \mathbf{Q}} & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & \frac{\partial \mathbf{J}}{\partial \mathbf{F}_b} & \frac{\partial \mathbf{J}}{\partial \mathbf{Q}} & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & \frac{\partial \mathbf{S}'}{\partial \mathbf{Q}} & \frac{\partial \mathbf{S}'}{\partial \mathbf{x}_s} & 0 & 0 & 0 & 0 & 0 & 0 \end{array} \right] = \left[\begin{array}{c} -\frac{\partial \mathbf{G}}{\partial \mathbf{D}} \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{array} \right]$$

The first and second equations correspond to equations for the mesh and flow variable sensitivities, as previously described for the aerodynamic solver, and the third equation corresponds to the construction of the surface force sensitivities given these two previous sensitivities. The fourth equation denotes the sensitivity of the FSI transfer from the fluid to the structural solver, while the fifth equation corresponds to the sensitivity of the structural solver. Finally, the last equation corresponds to the sensitivity of the FSI transfer from the structural solver back to the flow solver. This coupled system of sensitivities can be solved analogously to the coupled analysis problem as:

$$\left[\frac{\partial \mathbf{G}}{\partial \mathbf{x}} \right] \frac{\partial \mathbf{x}^c}{\partial \mathbf{D}} = -\frac{\partial \mathbf{G}}{\partial \mathbf{x}_s} \frac{\partial \mathbf{x}^{c-1}}{\partial \mathbf{D}} - \frac{\partial \mathbf{G}}{\partial \mathbf{D}} \quad (44)$$

$$\left[\frac{\partial \mathbf{R}}{\partial \mathbf{u}} \right] \frac{\partial \mathbf{u}^c}{\partial \mathbf{D}} = -\frac{\partial \mathbf{R}}{\partial \mathbf{x}} \frac{\partial \mathbf{x}^c}{\partial \mathbf{D}} \quad (45)$$

followed by the explicit evaluation of the surface force sensitivities as:

$$\frac{\partial \mathbf{F}^c}{\partial \mathbf{D}} = \frac{\partial \mathbf{F}}{\partial \mathbf{x}} \frac{\partial \mathbf{x}^c}{\partial \mathbf{D}} + \frac{\partial \mathbf{F}}{\partial \mathbf{u}} \frac{\partial \mathbf{u}^c}{\partial \mathbf{D}} \quad (46)$$

These sensitivity evaluations are all implemented within the flow solver. They are followed by the solution of the remaining components of the system as:

$$\left[\begin{array}{cc} \frac{\partial \mathbf{S}}{\partial \mathbf{F}_b} & \frac{\partial \mathbf{S}}{\partial \mathbf{Q}} \end{array} \right] \left[\begin{array}{c} \frac{\partial \mathbf{F}_b^c}{\partial \mathbf{D}} \\ \frac{\partial \mathbf{Q}^c}{\partial \mathbf{D}} \end{array} \right] = \left[\begin{array}{c} -\frac{\partial \mathbf{S}}{\partial \mathbf{F}} \frac{\partial \mathbf{F}^c}{\partial \mathbf{D}} \\ 0 \end{array} \right] \quad (47)$$

These two equations are solved simultaneously to obtain the structural sensitivities $\frac{\partial \mathbf{Q}}{\partial \mathbf{D}}$ which are then used to evaluate the surface mesh sensitivities explicitly as:

$$\frac{\partial \mathbf{x}_s}{\partial \mathbf{D}} = -\frac{\partial \mathbf{S}'}{\partial \mathbf{Q}} \frac{\partial \mathbf{Q}}{\partial \mathbf{D}} \quad (48)$$

where the fact that $\frac{\partial \mathbf{S}'}{\partial \mathbf{x}_s} = [\mathbf{I}]$ (identity matrix) has been used. These new surface mesh sensitivities are then fed back into the first equation in the system initiating the next coupling iteration. As can be seen, each disciplinary solution procedure requires the inversion of the same Jacobian matrix as the corresponding analysis problem, which is done using the same iterative solver. Furthermore, the fluid-structure coupling requires the transfer of the force sensitivities $\frac{\partial \mathbf{F}}{\partial \mathbf{D}}$ from the flow to the structural solver, and the surface mesh sensitivities $\frac{\partial \mathbf{x}_s}{\partial \mathbf{D}}$ from the structural solver back to the fluid solver, which are of the same dimension as the force and surface displacements transferred in the analysis problem, respectively.

The corresponding adjoint problem can be obtained by pre-multiplying equation (44) by the inverse of the large coupling matrix and substituting this into equation (43), transposing the entire system, and defining adjoint variables as solutions to the following coupled system:

$$\begin{bmatrix} \frac{\partial \mathbf{G}^T}{\partial \mathbf{x}} & \frac{\partial \mathbf{R}^T}{\partial \mathbf{x}} & -\frac{\partial \mathbf{F}^T}{\partial \mathbf{x}} & 0 & 0 & 0 \\ 0 & \frac{\partial \mathbf{R}^T}{\partial \mathbf{u}} & -\frac{\partial \mathbf{F}^T}{\partial \mathbf{u}} & 0 & 0 & 0 \\ 0 & 0 & \mathbf{I} & \frac{\partial \mathbf{S}^T}{\partial \mathbf{F}} & 0 & 0 \\ 0 & 0 & 0 & \frac{\partial \mathbf{S}^T}{\partial \mathbf{F}_b} & \frac{\partial \mathbf{J}^T}{\partial \mathbf{F}_b} & 0 \\ 0 & 0 & 0 & \frac{\partial \mathbf{S}^T}{\partial \mathbf{Q}} & \frac{\partial \mathbf{J}^T}{\partial \mathbf{Q}} & \frac{\partial \mathbf{S}'^T}{\partial \mathbf{Q}} \\ \frac{\partial \mathbf{G}^T}{\partial \mathbf{x}_s} & 0 & 0 & 0 & 0 & \frac{\partial \mathbf{S}'^T}{\partial \mathbf{x}_s} \end{bmatrix} \times \begin{bmatrix} \Lambda_{\mathbf{x}} \\ \Lambda_{\mathbf{u}} \\ \Lambda_{\mathbf{F}} \\ \Lambda_{\mathbf{F}_b} \\ \Lambda_{\mathbf{Q}} \\ \Lambda_{\mathbf{x}_s} \end{bmatrix} = \begin{bmatrix} \frac{\partial \mathbf{L}^T}{\partial \mathbf{x}} \\ \frac{\partial \mathbf{L}^T}{\partial \mathbf{u}} \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix} \quad (49)$$

This system can be solved starting with the last equation and

proceeding to the first equation as:

$$\Lambda_{\mathbf{x}_s}^c = -\frac{\partial \mathbf{G}^T}{\partial \mathbf{x}_s} \Lambda_{\mathbf{x}}^{c-1} \quad (50)$$

followed by the solution of the structural adjoints

$$\begin{bmatrix} \frac{\partial \mathbf{S}^T}{\partial \mathbf{F}_b} & \frac{\partial \mathbf{J}^T}{\partial \mathbf{F}_b} \\ \frac{\partial \mathbf{S}^T}{\partial \mathbf{Q}} & \frac{\partial \mathbf{J}^T}{\partial \mathbf{Q}} \end{bmatrix} \begin{bmatrix} \Lambda_{\mathbf{F}_b}^c \\ \Lambda_{\mathbf{Q}}^c \end{bmatrix} = \begin{bmatrix} 0 \\ -\frac{\partial \mathbf{S}'^T}{\partial \mathbf{Q}} \Lambda_{\mathbf{x}_s}^c \end{bmatrix} \quad (51)$$

followed by the explicit construction of the pointwise CFD surface force adjoint:

$$\Lambda_{\mathbf{F}}^c = -\frac{\partial \mathbf{S}^T}{\partial \mathbf{F}} \Lambda_{\mathbf{F}_b}^c \quad (52)$$

and ending with the solution of the mesh and flow adjoints as:

$$\begin{bmatrix} \frac{\partial \mathbf{G}^T}{\partial \mathbf{x}} & \frac{\partial \mathbf{R}^T}{\partial \mathbf{x}} \\ 0 & \frac{\partial \mathbf{R}^T}{\partial \mathbf{u}} \end{bmatrix} \begin{bmatrix} \Lambda_{\mathbf{x}}^c \\ \Lambda_{\mathbf{u}}^c \end{bmatrix} = \begin{bmatrix} \frac{\partial \mathbf{L}^T}{\partial \mathbf{x}} + \frac{\partial \mathbf{F}^T}{\partial \mathbf{x}} \Lambda_{\mathbf{F}}^c \\ \frac{\partial \mathbf{L}^T}{\partial \mathbf{u}} + \frac{\partial \mathbf{F}^T}{\partial \mathbf{u}} \Lambda_{\mathbf{F}}^c \end{bmatrix} \quad (53)$$

Once again, the solution of the various disciplinary adjoints requires the inversion of the corresponding disciplinary Jacobians (transposed in this case) which can be accomplished using the same iterative solvers as for the analysis and forward sensitivity problems. Additionally, the input to the structural adjoint problem consists of the variable $\Lambda_{\mathbf{x}_s}$, which is the same dimension as the surface displacements output from the structural analysis solver, while the output of the structural adjoint solver consists of the variable $\Lambda_{\mathbf{F}}$ which is of the same dimension as the force inputs to the structural solver in the analysis problem.

Verification of Coupled Aeroelastic Sensitivity

The forward and adjoint sensitivities for the coupled aeroelastic problem are verified using the complex step method. Any function $f(x)$ operating on a real variable x can be utilized to compute the derivative $f'(x)$ by redefining the input variable x and all intermediate variables used in the discrete evaluation of $f(x)$ as complex variables. For a complex input, the function when redefined as described produces a complex output. The derivative of the real function $f(x)$ can be computed by expanding the complex operator $f(x+ih)$ as:

$$f(x+ih) = f(x) + ihf'(x) + \dots \quad (54)$$

from which the derivative $f'(x)$ can be easily determined as:

$$f'(x) = \frac{\text{Im}[f(x+ih)]}{h} \quad (55)$$

As in the case of finite-differencing, the complex step-based differentiation also requires a step size. However, unlike finite-differencing the complex step method is insensitive to small step sizes since no differencing is required. In theory it is possible to verify forward and adjoint-based gradients using the complex step method to machine precision. With this in mind, a complex version of the complete coupled aerostructural analysis code has been constructed through scripting of the original source code to redefine variables from real to complex types and to overload a small number of functions for use with complex variables.


RESULTS

Time Dependent Analysis Problem


The chosen test case consists of the four bladed HART2 rotor with a 5 degree collective in a hover condition. The rotor is impulsively started from rest, in an initially quiescent flow field, and rotated with the mesh as a solid body for a fixed number of revolutions. This problem is solved both for a rigid blade model (using no structural model), as well as for a flexible blade mode (using the beam structural model). For the latter, the flow is solved in tight coupling mode with the beam solver. A series of progressively finer time step simulations are performed, and the simulation is carried out on a coarse and a fine mesh, in order to assess the level of temporal and spatial discretization errors.

The baseline simulation (coarse mesh) makes use of a mixed element mesh made up of prisms, pyramids and tetrahedra consisting of approximately 2.32 million grid points and is shown in Figures (3(a)), (3(b)) and (3(c)), where the rigid blade simulation is compared with the coupled CFD/CSD simulation. The simulations are run for 10 rotor revolutions using a 2 degree time-step for 1,800 time-steps starting from freestream initialization. For the rigid blade simulation, the time-dependent mesh motion is determined by rotating the entire mesh as a solid body at each time step. The unsteady Reynolds-averaged Navier-Stokes equations are solved at each time step in ALE form, using the Spalart-Allmaras turbulence model. Figures (4(a)) and (4(b)) shows a snapshot of the pressure coefficient contours on the rotor at the end of a single revolution for both the rigid and flexible blades.


The coupled CFD/CSD simulation is run in a similar manner. However, the flow solution (CFD) is coupled with the beam solver (CSD) at every time step by appropriately exchanging, a) airloads information from the flow domain to the beam and b) blade deformation information from the beam to the flow domain, at the fluid-structure interface (i.e. blade surface). In this coupled simulation, the mesh is first moved according the deformations dictated by the new flexed blade coordinates determined from the structural beam code before the solid body rotation of the entire mesh is performed. Thus, the flow now sees not only the rigidly rotated mesh (observed in rigid blade simulation), but also the deformed mesh around the blades. This coupled fluid-structure interaction problem needs to be iterated until satisfactory convergence is achieved



(a) Computational Domain

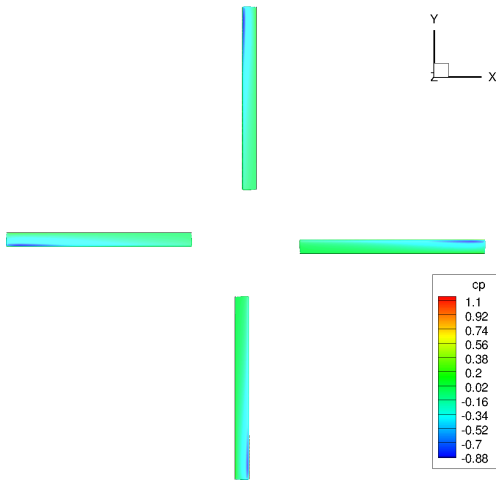


(b) Planform view




(c) Zoomed in view

8 **Fig. 3. HART2 rotor mesh consisting of 2.32 million points used in the optimization example.**



(a) Planform view



(b) Zoomed in view

Fig. 4. C_p contours for the baseline HART2 rotor, with flexible and rigid blades, in hover after one revolution. The mesh consists of 2.32 million vertices.

on both flow and mesh deformation within each time step. This kind of CFD/CSD coupling done within every time step is known as *tight coupling*.

The simulations were performed on the Yellowstone supercomputer at the NCAR-Wyoming Supercomputing Center (NWSC), with the analysis problem running in parallel on 512 cores. Each time step used 6 coupling iterations, and each coupling iteration used 10 non-linear flow iterations with each non-linear iteration consisting of a three-level line-implicit multigrid cycle. The typical simulation at this level of resolution requires approximately one hour of wall clock time per rotor revolution.

Figures 5 through 8 summarize the overall convergence of the rigid and aero-elastic coupling analysis formulations. Figure 5 shows the typical flow and turbulence residual convergence within a single time step for the rigid rotor case (no structural model), while Figure 6 depicts convergence of the flow and turbulence residuals at the same time step for the coupled aeroelastic case. In this case, the jumps in residual




Fig. 5. Flow and turbulence residual convergence at a given time step for rigid analysis




Fig. 6. Flow and turbulence residual convergence at a given time step for coupled aeroelastic analysis




Fig. 7. Mesh deformation residual convergence at a given time step for coupled aeroelastic analysis

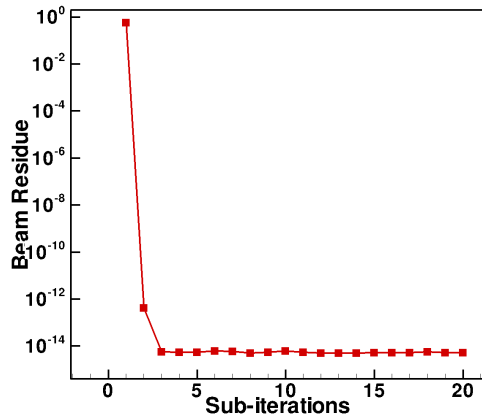

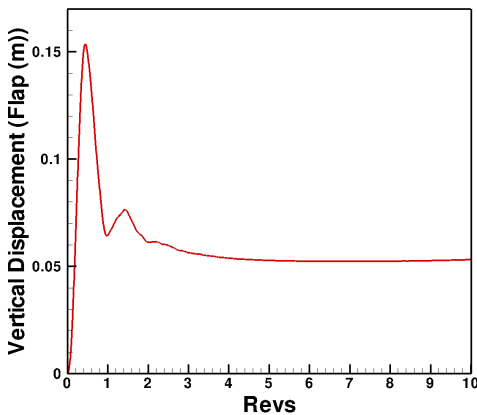


Fig. 8. Residual convergence of beam and overall FSI in one coupling iteration



(a) Blade Deformation



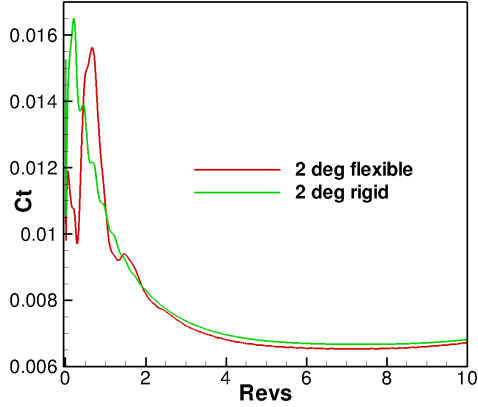
(b) Blade tip vs time

Fig. 9. HART2 blade deformation

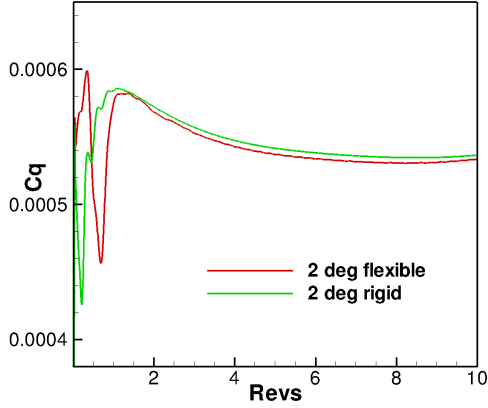
values at the start of new coupling iterations are clearly visible, although these jumps become smaller as the coupling procedure converges, and the overall residual histories closely follow those of the rigid rotor case after the first few coupling iterations. (The drop in density residual at each new coupling iteration is due to the imposition of a small start-up CFL value at each new coupling iteration.) Figure 7 depicts the convergence of the mesh deformation residual for the same time step, also showing jumps in the residual at the start of each new coupling iteration. Solution of the mesh deformation equations terminates when the residuals reach a prescribed tolerance of $1.e - 08$, thus the variable number of iterations per coupling cycle. Most notable is the fact that the initial mesh deformation residual decreases at each new coupling iteration, providing a measure of the convergence of the entire coupling procedure. Figure 8 illustrates the convergence of the coupled beam/FSI residual (i.e. equations (35) and (36)), showing rapid convergence to machine zero in a small number of iterations within a single CFD/CSD coupling iteration. The corresponding beam residual drop is observed to be of 15 orders of magnitude, as shown in Fig. 8. The effect of the CFD/CSD aeroelastic coupling is clearly demonstrated in Figure 9, which compares the deformed blade shape from the coupled simulation with that from the rigid blade simulation. As Figure 9(b) shows, the blade initially deforms onto larger tip flap values ($\sim 16cm$) before settling into a lower value of $\sim 6cm$. Further Figure 10 compares thrust, torque and Figure of Merit (FM) values from the coupled solution with the rigid blade solution. The flexed blade results in prediction of more coning of the blade and slightly lower total thrust (Fig. 10(a)) as well as lower total torque magnitude (Fig. 10(b)). Therefore, it is evident that the rigid body model might lead to over prediction of required rotor torque. It is noteworthy that relatively large differences are observed between the rigid and flexible rotor simulations in the first rotor revolution, even though these diminish as the rotor approaches the steady hover condition. Therefore, it is of particular interest to test the coupled aeroelastic optimization capability in this region of the simulation.

In order to study the effect of time-step size and grid resolution, analogous simulations were performed using smaller time steps and a finer grid. In total, three simulations using the 2.3 million point grid described above were performed using time step sizes of $\Delta t = 2.0^\circ$ and 1.0° and 0.5° . The number of subiterations per time step was held constant (6 coupling iterations with 10 flow subiterations per coupling iteration) for all three runs, thus providing slightly lower converged residual levels for the smaller time step runs.

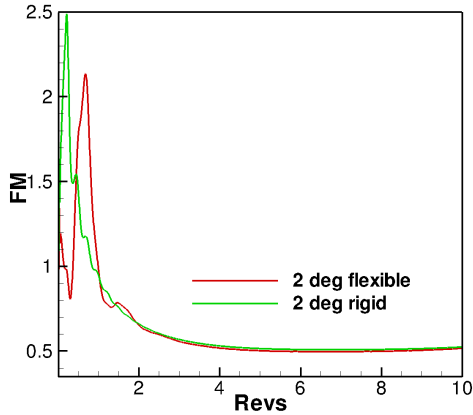
In order to assess the effect of grid resolution, a simulation using a 2.0° time step on a finer mesh with 11.4 million grid nodes with mixed elements was also performed. For the fine mesh, the number of subiterations per coupling iteration was raised to 25, in order to obtain the same relative decrease in residuals at each time step compared to the equivalent coarse grid case. Figure 11 compares computed thrust and torque histories from all these cases for the first revolution of the rotor, where transient effects are important. The



(a) Thrust coefficient vs time



(b) Torque coefficient vs time



(c) Figure of Merit vs time

Fig. 10. HART2 rotor performance time history: 2.32M, $\Delta t = 2^\circ$

airload time histories obtained using the three different time step sizes on the coarse grid are in relatively close agreement, with the difference between the 1.0° and 0.5° time step cases being smaller than that between the two larger time step cases, which is indicative of temporal error convergence. More appreciable differences are observed between the coarse and fine mesh runs, suggesting that the 2.3 million point mesh may be too coarse for accurate prediction of airloads for this case. However, one of the objectives of this work is to investigate the effectiveness of performing optimizations using less costly coarser meshes and larger time steps.

Time Dependent Optimization Problem

Geometry Parameterization In order to obtain sensitivities with respect to a set of shape parameters that are well suited for design optimization purposes, a baseline blade is constructed by stacking 11 airfoil section along the span. Each airfoil contains 10 Hicks-Henne bump functions, 5 on the upper surface, and 5 on the lower surface, that can be used to modify the airfoil shape. Additionally, the twist values of the blade at the root and tip airfoil sections are also used as design variables resulting in a total of 112 design variables. Figure (12(a)) provides an illustration of the baseline blade design setup. A high density structured mesh is generated about this blade geometry, which is then rotated and translated to match each individual blade in the CFD mesh, as shown in Figure (12(b)). Interpolation patterns between each unstructured mesh surface point and the baseline structured mesh are determined in a preprocessing phase. These interpolation patterns are then used to interpolate shape changes from the baseline blade to all four blades in the CFD mesh (as determined by changes in the design variables) and to transfer sensitivities from the surface CFD mesh points to the design variables using the chain rule of differentiation.


Unsteady Objective Function Formulation A time-integrated objective function based on the time variation of the thrust (C_T) and torque (C_Q) coefficients is used for this test case. The goal of the optimization is to reduce the time-integrated torque coefficient while constraining the time-integrated thrust coefficient to the baseline rotor performance. The objective function is based on the summation of the differences between a target and a computed objective value at each time level n . Mathematically the local objective function at each time-step in the integration range is defined as:

$$L^n = (\delta C_T^n)^2 + 10(\delta C_Q^n)^2 \quad (56)$$


$$\delta C_T^n = (C_T^n - C_{T_{target}}^n) \quad (57)$$

$$\delta C_Q^n = (C_Q^n - C_{Q_{target}}^n) \quad (58)$$


where the target thrust coefficient values at each time-step in the integration range are set from the baseline HART2 rotor values and the target torque values are set to zero. The weight of 10 on the torque coefficient is necessary to equalize the difference in orders-of-magnitude between the thrust and torque



(a) Thrust coefficient vs time




(b) Torque coefficient vs time




(c) Figure of Merit vs time

Fig. 11. HART2 rotor performance time history all cases



(a) Blade design parameters



(b) Baseline structured blade mesh overlap with CFD mesh

Fig. 12. Illustration of (a) baseline blade with design parameters and (b) overlap in tip region between baseline blade structured mesh and CFD surface unstructured mesh.

coefficients. The global or time-integrated objective is then constructed using equal unit weights at each time-step as:

$$L^g = \frac{1}{T} \sum_{n=1}^{n=N} \Delta t L^n \quad (59)$$

This so-called penalty formulation for constraining the thrust produced by the rotor is less rigorous than imposing a hard constraint on thrust that must be maintained at each design cycle, since small decreases in thrust can be traded off for decreases in torque using the penalty formulation. However, this approach is computationally more efficient than the use of a hard constraint formulation, since this latter approach would require the computation of two separate adjoint problems at each design cycle, as opposed to the single adjoint required in the current formulation.

Fully Coupled CFD/CSD Adjoint: Unsteady Sensitivity Verification The fully coupled CFD/CSD adjoint formulation was verified by first comparing the tangent sensitivities

Table 2. Uncoupled forward linearization verification

n^a	Method	Uncoupled (Rigid)
1	Complex	4.0965521828586 29 E-006
	Tangent	4.0965521828586 30 E-006
2	Complex	6.149890891151 360 E-006
	Tangent	6.149890891151 350 E-006
3	Complex	7.459751648971 305 E-006
	Tangent	7.459751648971 288 E-006
4	Complex	8.387775876237 277 E-006
	Tangent	8.387775876237 260 E-006
5	Complex	9.1689005477400 15 E-006
	Tangent	9.1689005477400 23 E-006

^a n =number of time step

with those obtained from complex step method. Subsequently, the adjoint formulation was verified by comparing its sensitivities with those obtained from the tangent as well as the complex step method. The tangent formulation was verified for perturbations on one geometric design parameter, namely, blade twist at the tip, for both uncoupled (rigid blade) as well as coupled simulations. As described earlier, a complex perturbation of size 1×10^{-100} is introduced on twist at the beginning of the analysis run. The derivatives of the functional (L^n) are evaluated and compared with those obtained using the forward linearization procedure at every time instance for up to 5 time steps. In both cases, for the complex step method, as well as for the forward linearization approach, the fully coupled aeroelastic problem is converged to machine zero at each time step in order to avoid contaminating the sensitivity values with errors due to incomplete convergence.

Tables (2) and (3), respectively, compare rigid aerodynamic only derivatives and coupled aeroelastic derivatives ($\frac{\partial L^n}{\partial D}$) obtained from the complex analysis run with those from the forward linearization run for the first 5 time steps. The uncoupled tangent verification serves as a sanity check of the new forward sensitivity formulation when the structural code is switched off. As can be seen from the tables, the forward tangent sensitivities and the complex step sensitivities agree to 12 significant digits for both the rigid and flexible aeroelastic rotor cases.

Having validated the forward sensitivity with that from the complex step method, the adjoint formulation was validated with these two formulations. Table 3 compares the three formulations for 5 unsteady time steps as before. The sensitivity values from the adjoint formulation matches to twelve significant digits with the other two formulations. This provides the confidence in the present methodology to be used as an effective aeroelastic optimization tool.

Table 3. Coupled adjoint linearization verification

n	Method(n)	Coupled (Aeroelastic)
1	Complex	3.690007037237 534 E-006
	Tangent	3.690007037237 471 E-006
	Adjoint	3.690007037237 598 E-006
2	Complex	5.150483530831 191 E-006
	Tangent	5.150483530831 145 E-006
	Adjoint	5.150483530831 289 E-006
3	Complex	5.828069793498 591 E-006
	Tangent	5.828069793498 538 E-006
	Adjoint	5.828069793498 741 E-006
4	Complex	6.05621108634 4925 E-006
	Tangent	6.05621108634 4902 E-006
	Adjoint	6.05621108634 5518 E-006
5	Complex	6.026286742020 757 E-006
	Tangent	6.026286742020 644 E-006
	Adjoint	6.026286742020 636 E-006

Helicopter Blade Optimization The optimization framework is applied to the flexible HART2 rotor using the same time-dependent test case as described in the analysis section. The simulation is run for one full rotor revolution, starting impulsively from rest in quiescent flow. The objective consists of the time-integrated torque with thrust penalty, as described previously. However, the objective to be minimized is only integrated over the last 60 degrees of revolution, in order to avoid the optimization process from focusing on start-up transients. Limiting the optimization problem to a single rotor revolution reduces overall simulation cost. However, it is clear that the rotor has not achieved its steady state hover condition within this simulation time frame, and thus the current approach may not be ideal for optimizing hover conditions. However, setting up the optimization problem in this manner allows the assessment of the ability of the optimization framework to handle problems with large transients in both aerodynamics and structural deflections.

The optimization problem is solved on the coarse grid (2.3 million points) using the three different time step sizes (2, 1, and 0.5 degrees) as well as on the fine grid (11 million points) using the 2 degree time step. One of the objectives of this study is to assess sensitivity of the optimization process to time step sizes and mesh resolution, and to determine the effectiveness of optimizations performed using lower cost coarser resolution simulations.

The optimization procedure used is the L-BFGS-B bounded reduced Hessian algorithm (Ref. 30). Each request by the optimization driver for a function and gradient value results in a single forward time-integration of the analysis solver and a single backward integration in time of the adjoint solver. A bound of $\pm 5\%$ chord for each defining airfoil section was set on the Hicks-Henne bump functions, and a bound of $\pm 1.0^\circ$ of twist was set on the root and tip twist definitions.

The coarse grid optimizations were performed on the Mount Moran computer cluster at the Advanced Research

Table 4. Comparison of simulation parameters for 2.3M grid case

Δt (deg)	t_f^a (hr:min)	Iter ^b	nf ^c
2.0	1:50	72	78
1.0	3:40	38	44
0.5	7:20	22	23

^a t_f =time per function/gradient call

^bIter=maximum number of design iterations

^cnf=maximum number of function calls

Computing Center (ARCC) at the University of Wyoming with the simulations (analysis/adjoint) running in parallel on 512 cores. Each time step in the analysis problem employed 6 coupling cycles. Each coupling cycle used 10 nonlinear iterations with each nonlinear iteration requiring 3 linear multigrid cycles per Krylov vector.

Table 4 summarizes various parameters from the simulations performed for the three time step sizes considered for the coarse grid run. As observed from the table, reducing the time step size by a factor of 2 requires twice as long to complete every function/gradient call. This results in fewer design iterations being completed within a given amount of wall clock time.

Figure 13 shows the residual convergence for a typical unsteady adjoint time step for $\Delta t = 2^\circ$. The figure shows the residual drops by 5 orders of magnitude over 6 coupling cycles. Figures 14(a) and 14(b) show the functional and gradient convergence over design cycles for all the time step sizes. The figures show that the functional values have achieved significant convergence in all the three cases and the values are not significantly different from each other. The gradient values for the largest time step ($\Delta t = 2^\circ$), having run for largest number of design iterations (72), has converged beyond two orders of magnitude. The gradient values for the other two smaller time steps have dropped by around one and half orders of magnitude. Overall, there is good convergence for all the time steps considered.

Figures 15 compare the optimized Hart2 rotor airload time histories with those from the baseline rotor for $\Delta t = 2^\circ$. The optimized rotor results in a thrust penalty of 1.8% (Fig. 15(a)) and an improved torque reduction by 4.9% (Fig. 15(b)). The overall performance gain in terms of Figure of Merit is 2.2% (Fig. 15(c)). The figures also show the range in the time domain over which the objective function is being computed and optimized. This range corresponds to the last 60° of rotor revolution. All the optimization cases discussed in this work consistently use this same range for computing their respective objective functions.

Figures 16 and 17 further compare the airloads from the $\Delta t = 1.0^\circ$ and 0.5° cases, respectively. The results are consistent with what was observed from the run with 2° time step. Each of these optimization results in reduced final torque coefficient values (5.3% and 5.16%, respectively) along with

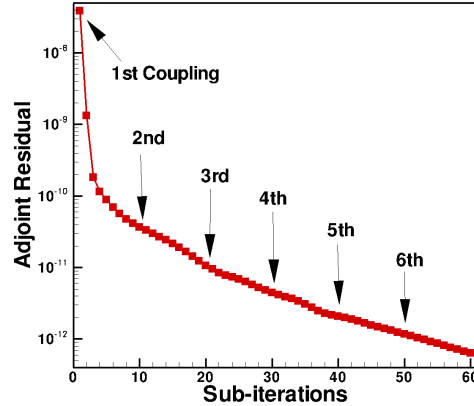




Fig. 13. Residual convergence in a typical adjoint time step




(a) Functional Convergence

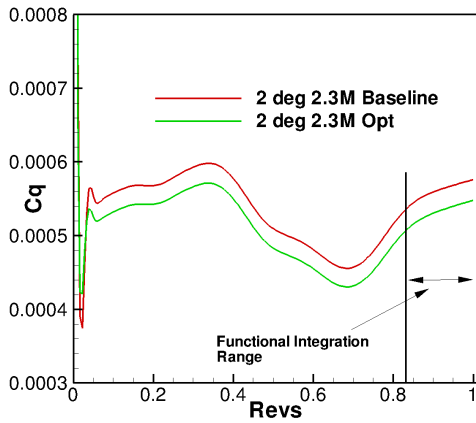


(b) Gradient Convergence


Fig. 14. Convergence in design optimization; Grid size: 2.32M



(a) Thrust coefficient vs time



(b) Torque coefficient vs time



(c) Figure of Merit vs time

Fig. 15. Performance time history on optimized and baseline HART2 rotor; 2.32 million grid; $\Delta t = 2.0^\circ$

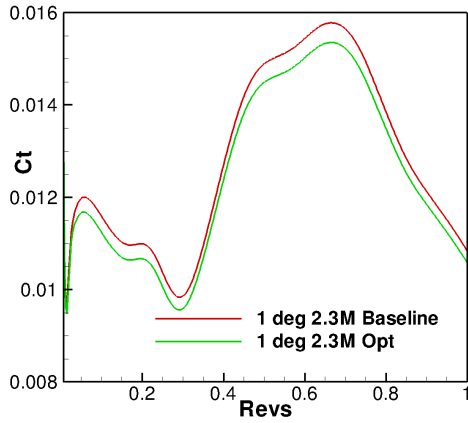
smaller thrust penalties (2.3% and 2.4%, respectively). The overall Figure of Merit improvement (2.2% and 1.6%, respectively) is also consistent with the results from the optimization case with 2° time step.

Finally, Figures 18 compare the optimized rotor blade sections with the baseline sections for the three different optimization cases using different time step sizes. The design process seeks to reduce the thickness of the outboard blade sections while producing smaller changes on the inboard sections. The blade geometry modifications are consistent with an aerodynamically designed efficient rotor blade geometry, where the inboard stations are thicker to provide larger Cl_{max} values near low dynamic pressure root sections and the outboard stations are thinner so as to reduce compressibility effects at the tip. The differences in the final design shapes obtained using the three different time steps are small, and all designs display similar trends. At most sections, the 1 degree and 0.5 degree time step designs are closest to each other, although this is not observed at the inner most station. Some of the differences in the design state are likely due to differences in convergence of the design process, as well as to differences in the computed objective and sensitivities due to the change in time step value.


However, the blade design optimization studies performed on this grid reveal that the overall design process is relatively insensitive to the time step size (over the range considered). Therefore, in order to reduce overall computational effort, design optimization using a relatively large time step of 2° appears to be feasible. Furthermore, the performance gains realized through the optimization process based on a single rotor revolution appear to result in equivalent gains in the long-time integration of the rotor in hover. This is illustrated in Figure 19, where the performance improvements of the optimized geometry are maintained for the rotor after 10 revolutions in hover condition.

An improved Hart2 rotor blade design optimization is performed using the finer mesh of 11.4 million grid points with a time step of $\Delta t = 2^\circ$. Use of the larger time step on the fine grid is motivated by the results obtained on the coarser grid, as well as by the overall desire to reduce computational cost. Two types of design iterations are performed: i) one that starts with baseline blade design (referred to in figures as D_{base}) and ii) another one that starts with the converged blade design from the coarse mesh (2.32 million) obtained using $\Delta t = 2^\circ$. The optimizations were performed on the NCAR-Wyoming supercomputer with the simulations (analysis/adjoint) running in parallel on 2048 cores. Each time step in the analysis problem employed 6 coupling cycles. Each coupling cycle used 25 nonlinear iterations with each nonlinear iteration requiring 3 linear multigrid cycles per Krylov vector. A typical design cycle (function/gradient computation) required approximately 5 hours total run time.

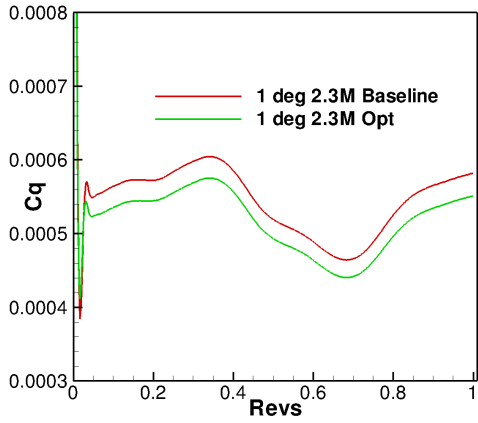
Figures 20(a) and 20(b) compare the functional and gradient convergence for the two fine mesh optimization cases described earlier. Both the functional and gradient for the restarted case start out at lower values than those of the base-



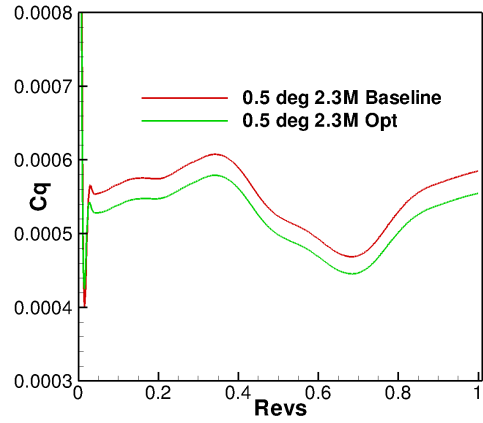
(a) Thrust coefficient vs time




(a) Thrust coefficient vs time




(b) Torque coefficient vs time



(b) Torque coefficient vs time



(c) Figure of Merit vs time



(c) Figure of Merit vs time

Fig. 16. Performance time history on optimized and baseline HART2 rotor; 2.32 million grid; $\Delta t = 1.0^\circ$

Fig. 17. Performance time history on optimized and baseline HART2 rotor; 2.32 million grid; $\Delta t = 0.5^\circ$

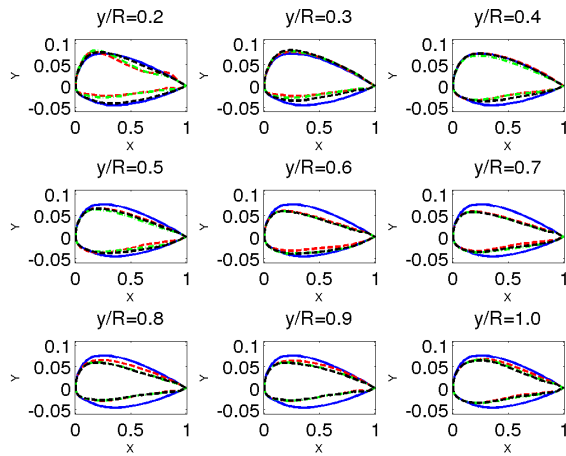



Fig. 18. Optimized and baseline Hart2 rotor blade sections ($y/R = [0.2 : 1.0]$); 2.32M grid; **blue solid - baseline, red dashes - optimization using $\Delta t = 2.0^\circ$, green dash-dots - optimization using $\Delta t = 1.0^\circ$, black dashes - optimization using $\Delta t = 0.5^\circ$,**

line optimization case (non-restarted), indicating that the optimization performed on the coarse mesh produces better performance when analyzed on the finer mesh, as expected. The design produced on the coarse mesh is further optimized on the fine mesh during the restarted optimization run. Alternatively, the baseline fine mesh design case starts with higher objective and gradient values, but converges towards the design produced by the restarted case over 4 design cycles.


These results suggest that computational savings can be realized in the design process by performing initial optimization on a coarser mesh, possibly using a larger time step for additional savings, prior to refining the design at higher resolution.

Figures 21 compare the optimized Hart2 rotor load time history for the optimization runs from the two fine mesh optimization cases (baseline and restarted) along with the performance of the design produced on the coarser grid, which constitutes the starting geometry for the restarted design case. Specifically, the coarse mesh optimization, restarted fine mesh optimization and baseline fine mesh optimization result in 4.8%, 7.6% and 2.0% reductions in torque with 1.8%, 4.4% and 1.4% reductions in thrust, respectively. As can be seen, the bulk of the improvement in power reduction is observed in the performance of the design obtained exclusively on the coarse mesh, with small additional improvements observed for both the restarted and baseline optimization cases.


Figures 22 compare the optimized rotor blade shapes at the various sections for the coarse grid optimization, and both fine grid optimizations (baseline and restarted), showing relatively small differences for these three optimization cases. However, at certain stations the deformations generated by the fine mesh baseline optimization problem are considerably smaller than those in the other optimization runs, possibly due to the small number of design iterations used in this case.



(a) Thrust coefficient vs time

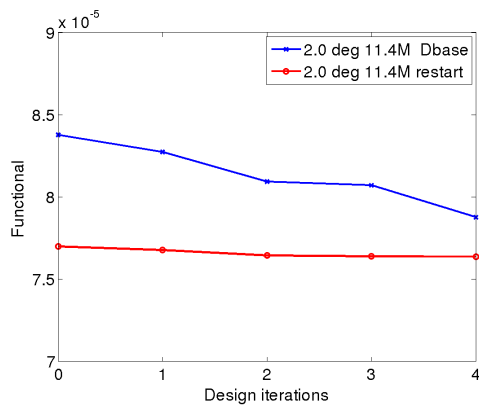


(b) Torque coefficient vs time




(c) Figure of Merit vs time

Fig. 19. Performance time history on optimized and baseline HART2 rotor over 10 revs; 2.32 million grid; $\Delta t = 2.0^\circ$



(a) Functional Convergence



(b) Gradient Convergence

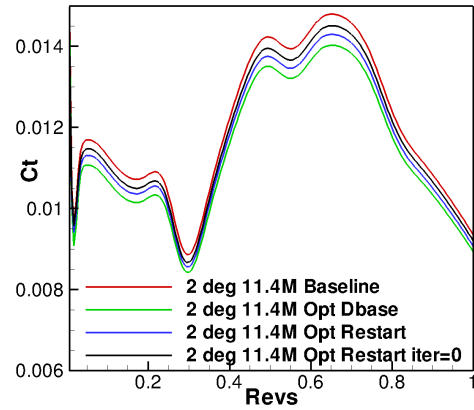
Fig. 20. Convergence in design optimization; grid size = 11.4 million

CONCLUSIONS AND FUTURE WORKS


In this work, a discrete adjoint formulation for time-dependent tightly coupled aeroelastic three-dimensional problems has been developed and demonstrated. The formulation is designed to reuse as much as possible the original coupled aeroelastic data-structures and solution strategies used for the analysis problem, thus simplifying implementation and verification. The adjoint sensitivities were used to perform a shape optimization problem for a flexible time-dependent rotor problem. Optimizations were performed on a coarse mesh using three different time step sizes, demonstrating relatively low sensitivity of the final produced design to the time step size. Further optimization using a finer mesh illustrated the beneficial effects of performing initial designs using lower mesh resolution in order to reduce overall computational effort.

Further work is required to determine the most effective strategies for time-dependent optimization. Performing initial optimizations at lower resolution holds tremendous potential for reducing overall computational effort in the context of time-dependent problems, since a reduction of resolution by a factor of two in all space dimensions and in time corresponds to a reduction of 16 in computational cost.


Although the current optimization example was based on



(a) Thrust coefficient vs time



(b) Torque coefficient vs time



(c) Figure of Merit vs time

Fig. 21. Performance time history on optimized and baseline HART2 rotor; grid size = 11.4 million. Opt Dbase denotes optimization starting from baseline geometry, Opt Restart denotes optimization starting with coarse mesh design

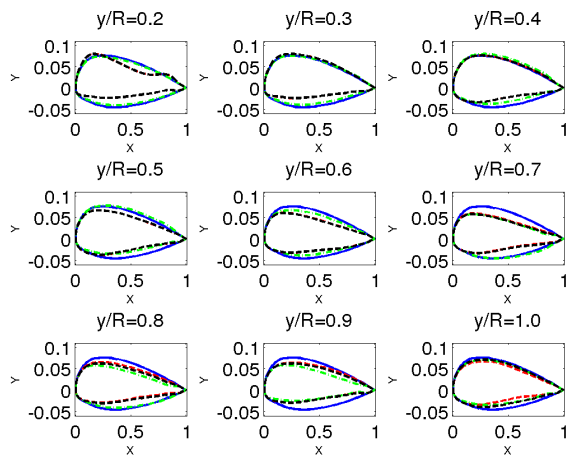


Fig. 22. Optimized and baseline Hart2 rotor blade sections ($y/R = [0.2 : 1.0]$); fine versus coarse blue solid - baseline, red dashes - optimization obtained on coarse 2.3M mesh, green dash-dots - optimization starting from baseline on 11.4 M Dbase, black dashes - optimization starting with coarse mesh design on 11.4 M Restart,

a rotor in hover conditions, only one rotor revolution has been simulated and the objective to be optimized consisted of time-integrated airloads over the final portion of this first revolution. One of the motivations for this approach was to demonstrate the potential of the current methodology to handle time-dependent problems with strong aeroelastic transients, such as those present in the first revolution of the impulsively started rotor. The realized reduction in the objective was shown to translate into improved hover performance as observed after a larger number of rotor revolutions. However, a more effective approach for pure hover performance optimization would be to formulate the problem in a non-inertial rotating reference frame. In future work, the current approach will be applied to the optimization of a flexible rotor at forward flight conditions, including cyclic pitching with trim constraints.

ACKNOWLEDGEMENTS

This work was partly funded by the Alfred Gessow Rotorcraft Center of Excellence through a subcontract with the University of Maryland. Computer resources were provided by the University of Wyoming Advanced Research Computing Center and by the NCAR-Wyoming Supercomputer Center.

REFERENCES

¹Jameson, A., “Aerodynamic Shape Optimization using the Adjoint Method,” VKI Lecture Series on Aerodynamic Drag Prediction and Reduction, von Karman Institute of Fluid Dynamics, Rhode St Genese, Belgium, 2003.

²Jameson, A. and Vassberg, J., “Computational Fluid Dynamics for Aerodynamic Design: Its Current and Future Impact,” Proceedings of the 39th Aerospace Sciences Meeting and Exhibit, Reno NV, AIAA Paper 2001–0538, 2001.

³Nadarajah, S. and Jameson, A., “A Comparison of the Continuous and Discrete Adjoint Approach to Automatic Aerodynamic Optimization,” Proceedings of the 38th Aerospace Sciences Meeting and Exhibit, Reno NV, AIAA Paper 2000–0667, 2000.

⁴Nielsen, E. and Anderson, W., “Recent Improvements in Aerodynamic Design Optimization on Unstructured Meshes,” *AIAA Journal*, Vol. 40-6, June 2002, pp. 1155–1163.

⁵Jameson, A., Alonso, J., Reuther, J., Martinelli, L., and Vassberg, J., “Aerodynamic shape optimization techniques based on control theory,” AIAA Paper 98–2538, 1998.

⁶Giles, M., Duta, M., and Muller, J., “Adjoint Code Developments Using Exact Discrete Approach,” AIAA Paper 2001–2596, 2001.

⁷Mani, K. and Mavriplis, D. J., “Unsteady Discrete Adjoint Formulation for Two-Dimensional Flow Problems with Deforming Meshes,” *AIAA Journal*, Vol. 46-6, June 2008, pp. 1351–1364.

⁸Rumpfkeil, M. and Zingg, D., “A General Framework for the Optimal Control of Unsteady Flows with Applications,” 45th AIAA Aerospace Sciences Meeting and Exhibit, Reno, NV, January, AIAA Paper 2007–1128, 2007.

⁹Mavriplis, D. J., “Solution of the Unsteady Discrete Adjoint for Three-Dimensional Problems on Dynamically Deforming Unstructured Meshes,” Proceedings of the 46th Aerospace Sciences Meeting and Exhibit, Reno NV, AIAA Paper 2008–0727, 2008.

¹⁰Nielsen, E. J., Diskin, B., and Yamaleev, N., “Discrete Adjoint-Based Design Optimization of Unsteady Turbulent Flows on Dynamic Unstructured Grids,” *AIAA Journal*, Vol. 48-6, June 2010, pp. 1195–1206.

¹¹Nielsen, E. J., Lee-Rausch, E. M., and Jones, W. T., “Adjoint-Based Design of Rotors in a Noninertial Reference Frame,” *Journal of Aircraft*, Vol. 47-2, March-April 2010, pp. 638–646.

¹²Martins, J. R. R. A. and Lambe, A. B., “Multidisciplinary Design Optimization: A Survey of Architectures,” *AIAA Journal*, Vol. 51, 2013, pp. 2049–2075. doi: 10.2514/1.J051895

¹³Kenway, G. K. W., Kennedy, G. J., and Martins, J. R. R. A., “Scalable parallel approach for high-fidelity steady-state aeroelastic analysis and adjoint derivative computations,” *AIAA Journal*, (In press), 2013.

¹⁴Mani, K. and Mavriplis, D. J., “Adjoint based sensitivity formulation for fully coupled unsteady aeroelasticity problems,” *AIAA Journal*, Vol. 47, (8), August 2009, pp. 1902–1915.

¹⁵Mishra, A., Mani, K., Mavriplis, D. J., and Sitaraman, J., “Helicopter Rotor Design using Adjoint-based Optimization in a Coupled CFD-CSD Framework,” 69th American

Helicopter Society Annual Forum, Phoenix, AZ, May 21–23 2013.

¹⁶Mani, K. and Mavriplis, D. J., “Geometry Optimization in Three-Dimensional Unsteady Flow Problems using the Discrete Adjoint,” 51st AIAA Aerospace Sciences Meeting, Grapevine, TX, AIAA Paper 2013-0662, January 2013.

¹⁷Mavriplis, D. J., “Discrete Adjoint-Based Approach for Optimization Problems on Three-Dimensional Unstructured Meshes,” *AIAA Journal*, Vol. 45-4, April 2007, pp. 741–750.

¹⁸Mavriplis, D. J., “Solution of the Unsteady Discrete Adjoint for Three-Dimensional Problems on Dynamically Deforming Unstructured Meshes,” Proceedings of the 46th AIAA Aerospace Sciences Meeting, Reno, NV, AIAA Paper 2008-0727, 2008.

¹⁹Spalart, P. R. and Allmaras, S. R., “A One-equation Turbulence Model for Aerodynamic Flows,” *La Recherche Aéronautique*, Vol. 1, 1994, pp. 5–21.

²⁰Mavriplis, D. J., “Unstructured-Mesh Discretizations and Solvers for Computational Aerodynamics,” *AIAA Journal*, Vol. 46-6, June 2008, pp. 1281–1298.

²¹Mavriplis, D. J., “Multigrid Strategies for Viscous Flow Solvers on Anisotropic Unstructured Meshes,” *Journal of Computational Physics*, Vol. 145, (1), September 1998, pp. 141–165.

²²Yang, Z. and Mavriplis, D. J., “A Mesh Deformation Strategy Optimized by the Adjoint Method on Unstructured Meshes,” *AIAA Journal*, Vol. 45, (12), 2007, pp. 2885–2896.

²³Mavriplis, D. J., Yang, Z., and Long, M., “Results using NSU3D for the first Aeroelastic Prediction Workshop,” Proceedings of the 51st Aerospace Sciences Meeting and Exhibit, Grapevine TX, AIAA Paper 2013-0786, 2013.

²⁴Schuster, D. M., Chwalowski, P., Heeg, J., and Wieseman, C. D., “Summary of Data and Findings from the First Aeroelastic Prediction Workshop,” Seventh International Conference on Computational Fluid Dynamics (ICCFD7), July 9-13 2012.

²⁵Kumar, A. A., Vishwamurthy, S. R., and Ganguli, R., “Correlation of helicopter rotor aeroelastic response with HART-II wind tunnel test data,” *Aircraft Engineering and Aerospace Technology: An International Journal*, Vol. 82, (4), 2010, pp. 237–248.

²⁶Yu, Y. H., Tung, C., van der Wall, B., Pausder, H.-J., Burley, C., Brooks, T., Beaumier, P., Delrieux, Y., Mercker, E., and Pengel, K., “The HART-II Test: Rotor Wakes and Aeroacoustics with Higher-Harmonic Pitch Control (HHC) Inputs - The Joint German/French/Dutch/US Project-,” 58th American Helicopter Society Annual Forum, Montreal, Canada, June 11–13 2002.

²⁷Chopra, I. and Bir, G., “University of Maryland Advanced Rotor Code: UMARC,” American Helicopter Society Aeromechanics Specialists Conference, January 1994.

²⁸Heckmann, A. Otter, M., Dietz, S., and Lopez, J. D., “The DLR FlexibleBodies Library to Model Large Motions of Beams and of flexible bodies exported from finite element programs,” Modelica, September 4–5 2006.

²⁹Smith, M. J., Lim, J. W., van der Wall, B. G., Baeder, J. D., Bierdrion, R. T., Boyd Jr., D. D., Jayaraman, B., Junk, S. N., and Min, B.-Y., “Adjoint-based Unsteady Airfoil Design Optimization with application to Dynamic Stall,” 68th American Helicopter Society Annual Forum, Fort Worth, TX, May 1–3 2012.

³⁰Zhu, C., Byrd, R. H., Lu, P., and Nocedal, J., “L-BFGS-B - FORTRAN Subroutines for Large-scale Bound Constrained Optimization,” Technical report, Department of Electrical Engineering and Computer Science, December 31 1994.