Multi-point Time-dependent Aero-elastic Adjoint-based Aerodynamic Shape Optimization of Helicopter Rotors

Asitav Mishra
amishra3@uwyo.edu
Postdoctoral Research Associate

Dimitri Mavriplis
mavripl@uwyo.edu
Professor

Jay Sitaraman
jsitaram@uwyo.edu
Associate Professor

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

ABSTRACT

A formulation for sensitivity analysis of fully coupled time-dependent aeroelastic problems for multipoint design optimization is given in this paper. Rotor blade shape design for simultaneous optimization of hover and forward flight conditions is performed both for rigid blade and flexible blade configurations. Optimized blade shapes that reduce power requirements for both hover and forward flight while maintaining a trim condition are computed for the rigid blade problem. The flexible blade problem results in smaller improvements, due to incomplete convergence of the optimization problem, which may be due to more rapid error accumulation in the adjoint derived sensitivities as a result of incomplete convergence at each times step. Future work includes the inclusion of additional design points, the calculation of Pareto fronts, and the determination of optimal convergence tolerances for cost effective robust multipoint optimization.

INTRODUCTION

While the use of adjoint equations is now fairly well established in steady-state shape optimization (Refs. 1–5), only recently have inroads been made into extending them to unsteady flow problems. For high-fidelity rotorcraft optimization problems, unsteady adjoint methods are necessary particularly for forward flight conditions and problems with relative motion (rotor-fuselage). Time-dependent aerodynamic optimization for rotorcraft problems has been addressed in previous work by Nielsen et al. (Refs. 6–8) and by Mani et. al. (Ref. 9). The incorporation of aeroelastic effects has been pursued in previous work as detailed in references (Refs. 10, 11) for hover, and (Ref. 12) for forward flight.

Multipoint aerodynamic optimization is now common practice for steady-state fixed wing aerodynamic problems and has been demonstrated for steady aeroelastic problems (Ref. 13). For rotorcraft problems, multipoint aerodynamic optimization for rigid rotors has been demonstrated by Jones et al. (Ref. 8). The aim of current work is to demonstrate the potential for high-fidelity time-dependent adjoint-based multipoint optimization for flexible aeroelastic rotors. The overall goal is to enable realistic optimization for multiple operating conditions such as hover and forward flight. This work builds upon our previous work on time-dependent adjoint optimization for rigid rotors (Ref. 9), and flexible rotors, initially for hover (Ref. 11), and subsequently for forward flight with trim constraints (Ref. 12).

In the following sections, we first summarize the basic components of the coupled aeroelastic solver, and the coupled adjoint sensitivity approach used to drive the optimization procedure. We then describe the implementation of the multipoint optimization strategy. The results section includes a sample single point optimization problem, as well as a rigid and a flexible rotor multipoint optimization test cases. In a subsequent section we examine the accuracy and error growth of the adjoint-derived sensitivities and close with conclusions and a discussion of future work.

Presented at the AHS 71th Annual Forum, Virginia Beach, VA May 5–7, 2015.

ANALYSIS FORMULATION

Flow Solver 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. 9, 14, 15). The flow solver is based on the conservative form of the Navier-Stokes equations which may be written as:

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

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

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

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. 16) for all cases presented in this work. The time derivative term is discretized using a second-order accurate backward-difference formula (BDF2) scheme and the implicit residual is solved using Newton's method. 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. 17). For the remainder of this paper, the system of non-linear equations for the computational fluid dynamics (CFD) analysis problem will be represented by the generalized notation:

$$\mathbf{R}(\mathbf{U}, \mathbf{x}) = 0 \tag{3}$$

where the vector $\mathbf{U}$ denotes the flow values and Equation (3) denotes the simultaneous solution of all time steps.

Structural Dynamics Formulation

The rotorcraft blade structure is modeled by a non-linear bend-twist beam model, which is a suitable and widely utilized structural model for slender fixed and rotary wing aircraft structures within the context of an aeroelastic problem. The beam model has been developed previously and coupled to the NSU3D unstructured mesh Reynolds-averaged Navier-Stokes solver (Refs. 18,19). 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: $M\ddot{q} + C\dot{q} + K = F$, where $[M]$, $[C]$ and $[K]$ are mass, damping and stiffness matrices of the system of equations representing the beam. $\mathbf{F} = F(t)$ is the aerodynamic forcing vector and $\mathbf{q}$ represents the displacements along all degrees of freedom of the beam. 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: $[I]\dot{\mathbf{Q}} + [A]\mathbf{Q} = \mathbf{F}$, where $[I]$ is the identity matrix, $\mathbf{Q} = [\mathbf{q}, \dot{\mathbf{q}}]^T$, $\mathbf{F} = [0,\ [M]^{-1}F]^T$ 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:

$$\mathbf{J}(\mathbf{Q}, \mathbf{F}) = \mathbf{0} \tag{4}$$

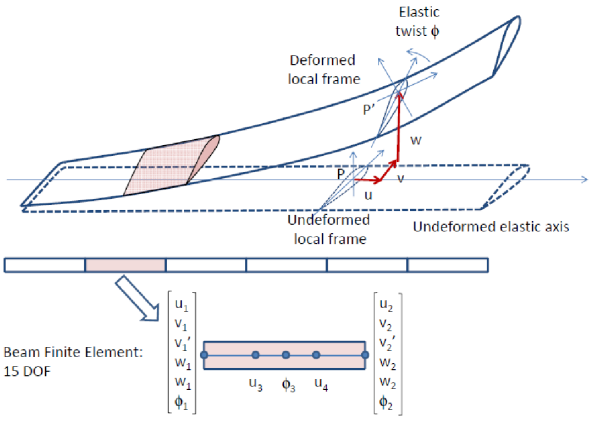

Fig. 1. Beam element with flap, lag, torsional and axial (total 15) degrees of freedom.

The beam model has been validated (Ref. 20) for the standard Hart-2 rotor case (Ref. 21) by comparing its natural frequency predictions with the predictions from other reliable computational structural dynamics (CSD) models, such as UMARC (Ref. 22) and DLR (Ref. 23).

Mesh deformation capability

In order to deform the CFD mesh in response to surface displacements generated by the solution of the CSD beam model, by cyclic pitching of the rotor blades, and due to blade design shape changes, a linear elastic analogy mesh deformation approach has been implemented. In this approach, the CFD 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. 18,24). 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:

$$\mathbf{G}(\mathbf{x}, \mathbf{D}) = \mathbf{0} \tag{5}$$

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

Prescribed Blade Motion

In order to perform a time-dependent simulation of a rotor in hover or forward flight, a prescribed motion capability is required. In the case of hover, a simple azimuthal solid body rotation of the entire rotor blade system and CFD mesh that contains the rotor is used. For forward flight, cyclic pitching motion must be superimposed on the azimuthal solid body rotation. For each blade, the time dependent pitch angle can be represented by a combination of mean pitch angle (θ_0) and several harmonic components of pitch angles (e.g. θ_{c_i} and

θ_{s_i} for i'th harmonic): $\theta = \theta_0 + \theta_{c_i} cos(i\psi) + \theta_{s_i} sin(i\psi)$, for harmonics of $i = [1, 2, \cdots]$. This introduces a new set of parameters, i.e. $\mathbf{D} = [\theta, \theta_{\mathbf{c_i}}, \theta_{\mathbf{s_i}}]$ that can be used to trim the rotor to achieve target thrust and moment values.

To incorporate rotor pitch actuation three additional equations need to be considered: i) a pitch actuation equation, represented by $\mathbf{S}^{\theta}(\mathbf{x}_{\mathbf{s}\theta}, \mathbf{x}) = \mathbf{x}_{\mathbf{s}\theta} - [\mathbf{T}^{\theta}(\theta(\mathbf{D}))]\mathbf{x} = \mathbf{0}$, and ii) a mesh motion equation to propagate surface displacements due to pitch to the interior mesh: $\mathbf{G}(\mathbf{x}_\theta, \mathbf{x}_{\mathbf{s}\theta}) = \mathbf{0}$. iii) an azimuthal mesh rotation equation:$\mathbf{S}^{\psi}(\mathbf{x}_\mathbf{p}, \mathbf{x}_\theta) = \mathbf{x}_\mathbf{p} - [\mathbf{T}^{\psi}]\mathbf{x}_\theta = \mathbf{0}$. Here $\mathbf{x}_{\mathbf{s}\theta}$ denotes the surface mesh coordinates after pitch and $\mathbf{x}_\theta$ denotes the interior mesh coordinates after mesh motion due to pitch, while $\mathbf{x}_\mathbf{p}$ denotes the interior mesh point coordinates after rotation. Here $[\mathbf{T}^{\theta}]$ and $[\mathbf{T}^{\psi}]$ are matrix representations for pitch actuation and rotation, respectively.

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, and the prescribed blade motion. The system of equations to be solved at each time step can be written as:

$$\mathbf{S}^{\theta}(\mathbf{x}_{\mathbf{s}\theta}, \mathbf{x}, \mathbf{D}) = \mathbf{0} \tag{6}$$

$$\mathbf{G}(\mathbf{x}, \mathbf{x}_\mathbf{s}(\mathbf{Q})) = \mathbf{0} \tag{7}$$

$$\mathbf{S}^{\psi}(\mathbf{x}_\mathbf{p}, \mathbf{x}_\theta) = \mathbf{0} \tag{8}$$

$$\mathbf{R}(\mathbf{u}, \mathbf{x}) = \mathbf{0} \tag{9}$$

$$\mathbf{S}(\mathbf{F_B}, \mathbf{Q}, \mathbf{F}(\mathbf{x}, \mathbf{u})) = \mathbf{0} \tag{10}$$

$$\mathbf{J}(\mathbf{Q}, \mathbf{F_B}) = \mathbf{0} \tag{11}$$

$$\mathbf{S}'(\mathbf{x}_\mathbf{s}, \mathbf{Q}) = \mathbf{0} \tag{12}$$

where $\mathbf{S}$ and $\mathbf{S}'$ represent the residuals of the FSI equations, and 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, as well as on surface displacements arising from cyclic pitching and design shape changes. 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.

At the first coupling iteration, $\mathbf{x}_\mathbf{s}^\mathbf{c} = \mathbf{0}$ (superscript $\mathbf{c}$ denotes the coupling iteration index) 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 $\mathbf{u}$, 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 aero-structural problem at the given time step.

SENSITIVITY ANALYSIS FOR COUPLED AEROELASTIC PROBLEM

For gradient-based optimization, sensitivities of the objective with respect to the design parameters are required. Because the number of design parameters is generally very large, an adjoint procedure is used to compute these sensitivities in an efficient manner. While a forward linearization of the analysis problem scales linearly with the number of design variables, the cost of the adjoint calculation is virtually independent of the number of design variables. Both forward and adjoint sensitivity approaches have been implemented in this work. While the adjoint approach is used in all optimizations, the forward sensitivity method is used to verify the correctness of the adjoint implementation. The adjoint linearizations in each discipline are hand coded and verified using the complex step method (Ref. 25) as well as by comparison with the forward linearization.

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. We form and solve individual disciplinary adjoints in a coupled solution process that follows closely the coupled analysis solver. Analogous disciplinary solvers are reused for each disciplinary sensitivity problem, and the analysis coupling strategy is extended to the sensitivity formulation. Furthermore, the data transferred between disciplinary solvers should consist of vectors of the same dimension for the analysis, tangent and adjoint formulations. Each disciplinary solution procedure in the tangent or forward sensitivity formulation 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 formulation consists of solving a transposed system of equations of the forward sensitivity formulation. Once again, the solution of the various disciplinary adjoint equations 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 adjoint 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 adjoint variable $\Lambda_\mathbf{F}$ which is of the same dimension as the force inputs to the structural solver in the analysis problem.

For time dependent problems, the adjoint procedure results in a backward integration in time, starting at the final time step, and proceeding to the initial time step of the analysis run. This requires the solution to be written out to disk

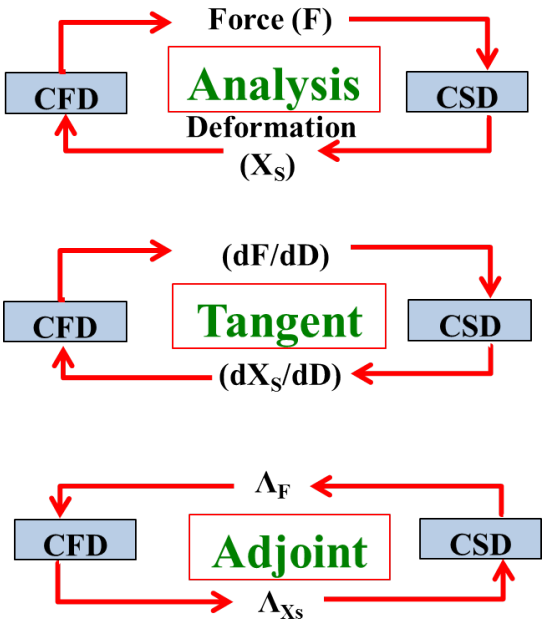

Fig. 2. Illustration of information flow for multidisciplinary analysis, forward sensitivity and adjoint formulations

at each time step in the analysis run, in order that it can be read back in by the adjoint solver as it proceeds backwards in time. Detailed description of the solution procedure for the coupled adjoint formulation can be found in the author's previous works (Refs. 10, 11) and a schematic of the whole adjoint formulation is depicted in Fig. 2.

OPTIMIZATION APPROACH

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 (3(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 (3(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.

In addition to the aforementioned design variables, inclusion of control parameters for pitching introduces a new set of

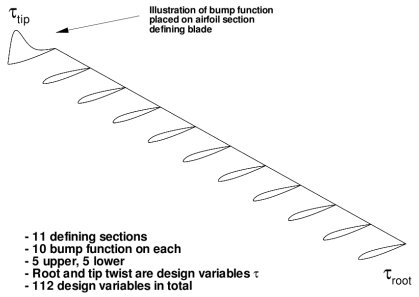

(a) Blade design parameters

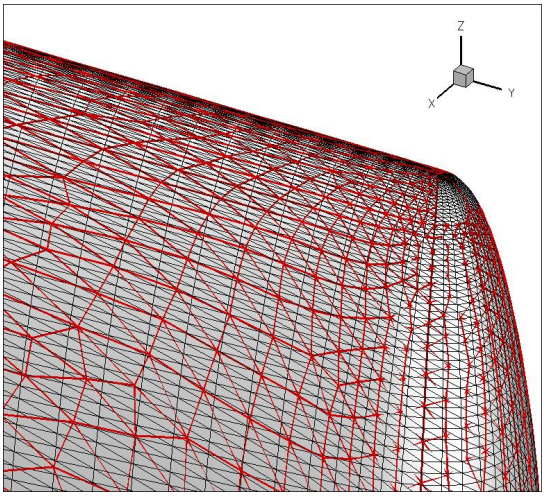

(b) Baseline structured blade mesh overlap with CFD mesh

Fig. 3. Illustration of (a) baseline blade with design parameters and (b) overlap in tip region between baseline blade structured mesh and CFD surface unstructured mesh.

(at least three) design parameters, i.e. $\mathbf{D} = [\theta_0, \theta_{c_i}, \theta_{s_i}]$. Design variable θ_0 is known as rotor collective pitch angle and the only pitch parameter used in hover. The parameters θ_{c_i} and θ_{s_i} are rotor cyclic pitch angles and are the additional pitch parameters required in forward flight to sustain lateral and longitudinal trims. This results in a total of $(112+1) = 113$ design variables in hover and $(112+3) = 115$ design variables in forward flight. Thus the blade transformation matrix denoted by $[\mathbf{T}^{\theta}(\theta)]$ is now linearized with respect to these additional design variables. These sensitivities, in addition to the above mentioned geometric sensitivities, are transferred onto the CFD interior mesh points from the CFD surface mesh points through the interpolation method described earlier.

Multipoint Unsteady Objective Function Formulation

A time-integrated objective function based on the time variation of the thrust (C_T), torque (C_Q) and moment coefficients (C_{M_X} and C_{M_Y}) is used for any general design point (such as hover or forward flight condition). The goal of the optimiza-

tion is to reduce the time-integrated torque coefficient while constraining the time-integrated thrust coefficient to the baseline rotor performance (both in hover and forward flight) as well as constraining the moment coefficients along roll (C_{M_X}) and pitch (C_{M_Y}) axes of the rotor to a trimmed value, i.e. zero moment values (only in forward flight or cruise). 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 global objective function is defined as:

$$L^g = L_{Shape} + L_{Trim} \quad (13)$$

$$L_{Shape} = \omega_1 \frac{1}{T} \sum_{n=1}^{n=N} \Delta t \left[\delta C_Q^n\right]^2 \quad (14)$$

$$L_{Trim} = \omega_2 \frac{1}{T} \sum_{n=1}^{n=N} \Delta t \left[\delta C_T^n\right]^2 + \omega_3 \left[\delta \bar{C}_{M_X}\right]^2 + \omega_4 \left[\delta \bar{C}_{M_Y}\right]^2 \quad (15)$$

$$\delta C_Q^n = (C_Q^n - C_{Q_{target}}^n) \quad (16)$$

$$\delta C_T^n = (C_T^n - C_{T\,target}^n) \text{ (hover)}$$

or

$$\delta C_T^n = (C_T^n - \bar{C}_{T\,target}) \text{ (cruise)} \quad (17)$$

$$\delta \bar{C}_{M_X} = \frac{1}{T} \sum_{n=1}^{n=N} \Delta t (C_{M_X}^n - C_{M_X\,target}^n) \quad (18)$$

$$\delta \bar{C}_{M_Y} = \frac{1}{T} \sum_{n=1}^{n=N} \Delta t (C_{M_Y}^n - C_{M_Y\,target}^n) \quad (19)$$

For the hover design point, $C_{T\,target}^n$ is a specified time dependent target thrust time history, and the moments terms are not included since they are expected to be zero. The collective pitch parameter (θ_0) helps sustain the trim target thrust state and the cyclic parameters are not used since the moment terms are absent. For the cruise design point, the mean target thrust coefficient value is specified to the Hart2 baseline mean trim value of $\bar{C}_{T\,target} = 4.4e-3$, and the target torque and moment values are set to zero. The weights ($\omega_i, i = [1,2,3,4]$) are included to equalize the difference in orders-of-magnitude between the thrust, torque and moment coefficients. The use of the pitch control parameters ($[\theta_0, \theta_{c_i}, \theta_{s_i}]$) as design variables and the use of the moment penalty terms ensure that the optimized rotor shape with optimized control parameters tends towards a final trimmed state when the rotor design cycles converge. This approach is computationally more efficient than the use of a hard constraint formulation, since this latter approach would require the computation of multiple adjoint problems at each design cycle, as opposed to the single adjoint required in the current formulation. However, the trim constraint may not be satisfied exactly in this approach.

RESULTS

Time Dependent Analysis Problem

The chosen multiple design points are for a four bladed HART2 rotor in two flight conditions: i) in hover with a 5 degree collective, with $M_{tip} = 0.638$ (rotational speed $\Omega = 1041$ RPM) and zero vertical climb velocity, and ii) a forward flight condition, with $M_{tip} = 0.638$ (rotational speed $\Omega = 1041$ RPM), $M_\infty = 0.095$ ($\mu = 0.15$), shaft tilt angle towards freestream, $\alpha_{shaft} = 5.4°$. For the forward flight design point, the pitch angle actuation is prescribed as: $\theta = \theta_0 + \theta_{1s} sin(\psi) + \theta_{1c} cos(\psi)$, with $\theta_0 = 5.0°, \theta_{1s} = -1.1°, \theta_{1c} = 2.0°$. In both the flight conditions, the rotor is impulsively started from rest, in an initially quiescent flow field, and rotated with the mesh as a solid body for two complete rotor revolutions. The objective for optimization is only computed over the second rotor revolution. This problem is solved both for a rigid blade model (using no structural model), as well as for a flexible blade model (using the beam structural model). Although the hover case is still highly transient over the first two rotor revolutions, previous work has shown that the improvements in performance obtained by optimizing over this short duration are sustained as the rotor approaches a steady hover state over a large number of revolutions (Ref. 11).

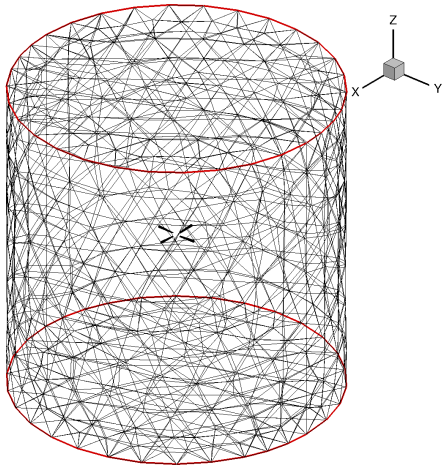

Fig. 4. Computational polyhedra mesh (2.32 million nodes) on HART2 rotor

The simulation makes use of a mixed element mesh made up of prisms, pyramids and tetrahedra consisting of approximately 2.32 million grid points, which is shown in Figure (4). The simulations are run for 10 and 5 rotor revolutions for hover and forward flight conditions, respectively, using a 2 degree time-step. For the rigid blade simulation in forward flight, the time-dependent mesh motion is determined by first pitching the blade about the blade axis followed by solving the mesh deformation equations and then 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.

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 moved ac-

cording to the deformations dictated by the new flexed blade coordinates determined from the structural beam code after the combined kinematics of pitch actuation followed by solid body rotation of the entire mesh have been performed. Thus, the flow now sees not only the pitched and rigidly rotated mesh (observed in rigid blade simulation in forward flight), but also the deformed mesh around the blades due to both pitching and structural deformations. This coupled fluid-structure interaction problem needs to be iterated until satisfactory convergence is achieved on flow, structure and mesh deformation problems within each time step.

The simulations were performed on the Yellowstone supercomputer at the NCAR-Wyoming Supercomputing Center (NWSC), with the analysis problem running in parallel on 1024 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 40 minutes of wall clock time per rotor revolution.

HART2 in Hover Figures 5 compare thrust and torque from the coupled solution with the rigid blade solution in hover (Ref. 11). The flexed blade results in prediction of more coning of the blade and slightly lower total thrust (Fig. 5(a)) as well as lower total torque magnitude (Fig. 5(b)).

HART2 in Forward Flight The effect of the CFD/CSD aeroelastic coupling in forward flight (Ref. 12) is clearly demonstrated in Figure 6, which compares the deformed blade shape and its corresponding C_p surface contours from the coupled simulation with that from the rigid blade simulation. From the figure it is noted that all four blades show different deformation characteristics due to corresponding different aerodynamic environment they experience in a forward flight. The blade attains the largest flap displacement at azimuth, $\psi = 180^\circ$ and the smallest flap displacement at azimuth, $\psi = 0^\circ$.

Figures (7) summarize performance comparison between flexible and rigid blades over 5 rotor revolutions. Besides showing consistent 4/rev harmonic content in the airloads, the figures also show that the rigid blades over predict thrust and power compared to flexible blades. This suggests that performance predictions might be erroneous if the aeroelastic nature of the rotor blades is not taken into account. Figure (7(c)) further compares the rotor longitudinal moments of the flexible rotor with the rigid one. It should be noted that neither the rigid nor the flexible rotor are in a moment trimmed state (zero mean moment) at the prescribed pitch parameters

Single Point Optimizations

Current work is built upon the previously developed optimization tool (Ref. 12) which was applied to the flexible HART2 rotor firstly in hover (Ref. 11) and more recently in forward flight (Ref. 12). The single point forward flight optimization

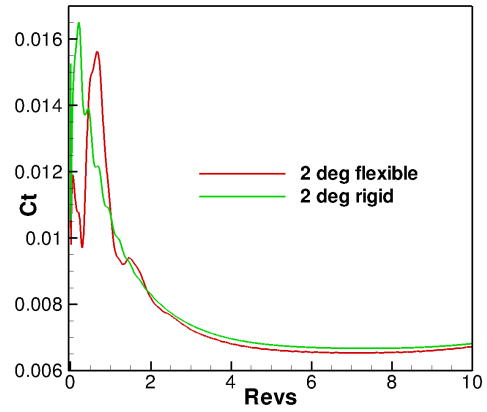

(a) Thrust coefficient vs time

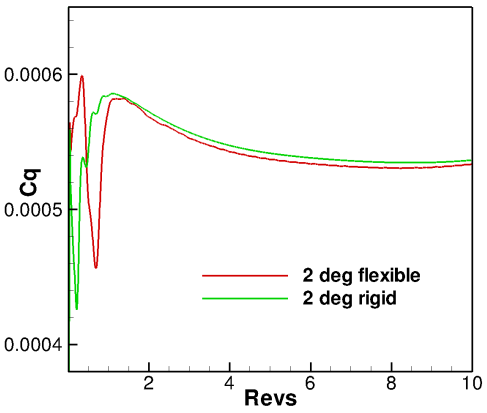

(b) Torque coefficient vs time

Fig. 5. HART2 rotor performance time history in hover: 2.32M, $\Delta t = 2^\circ$ (Ref. 11)

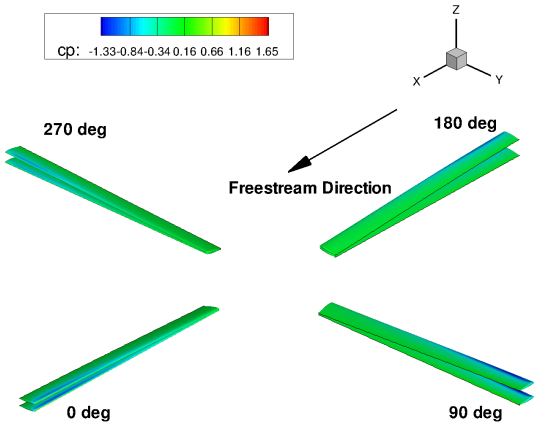

Fig. 6. Blade Deformation (C_p contours) on HART2 rotor in forward flight (Ref. 12)

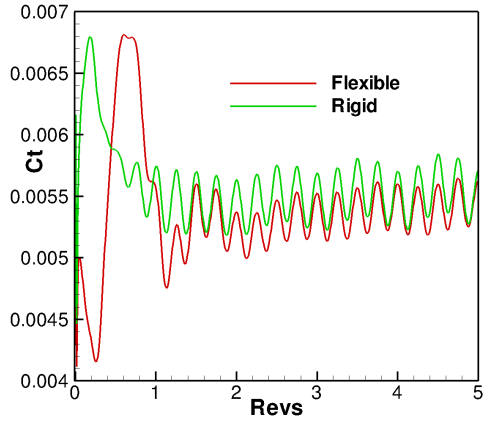

(a) C_T vs time

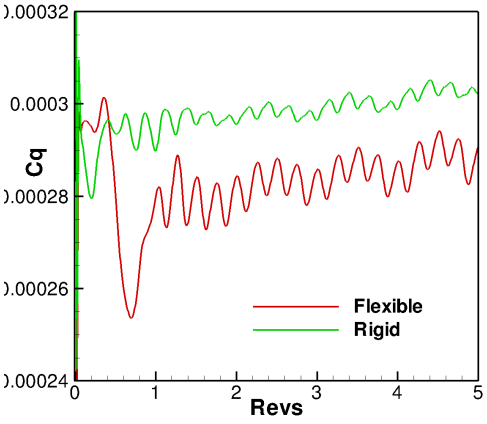

(b) C_Q vs time

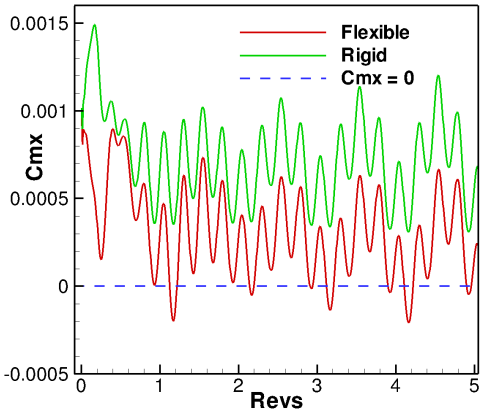

(c) Lateral moment vs time

Fig. 7. Performance and moment comparison of flexible blade with rigid blade

results are summarized herein. The simulation was run for two full rotor revolutions, starting impulsively from rest in quiescent flow. The objective was the time-integrated torque with thrust and moment penalty.

Figures 8 summarizes the optimized flexible blade shape as well as re-trimmed performance plots. Figure 8(a) shows the baseline trimmed flexible rotor is shape optimized to achieve a significant power reduction of approximately 5.0% with a thrust loss of approximately 2.6%. Some of this gain is lost upon re-trim. However, the final shape optimized and re-trimmed rotor achieves an overall 3.1% reduction in power compared to the baseline blade. Figure 8(c) shows one and half orders of magnitude gradient drop and consistent functional convergence even after 21 design iterations for the blade shape optimization stage. Figure 9 shows the shape optimized blade sections at 9 stations. The optimized blade shapes show the expected trend of a helicopter rotor in forward flight, i.e. thicker inboard and thinner outboard stations.

Multi-point Shape Optimization of Helicopter Rotor

The unsteady aeroelastic adjoint based design platform described so far has proven to be a reliable high-fidelity optimization tool to perform blade shape design optimization in hover (shown in previous works (Ref. 11)) as well as in forward flight conditions (Ref. 12). This section further explores how the formulation can be extended to perform multipoint blade shape design optimization of the HART2 rotor in simultaneous flight conditions of hover and forward flight (or cruise). A typical multipoint design cycle involves:

- Compute rotor performance at each design point
- Compute rotor sensitivities at each design point
- Build composite performance objective function
- Build composite performance objective sensitivities
- Perform optimization step
- Repeat cycle till convergence

Moreover, at any arbitrary design point, the performance optimization consists of three main stages: i) Trim ii) Shape/performance optimization and, iii) re-trim. The 'Trim' step involves trimming the rotor to a target rotor thrust and moment values (a specified target thrust time history in hover or a mean value of $C_T = 4.4e-3$ and zero longitudinal and lateral moments ($C_{M_Y}, C_{M_X} = 0$ in cruise). The objective function used in this step consists only of the thrust and moment terms in Eqn 13 ($\omega_1 = 0$) and the objective minimization is performed using only the pitch parameters (only collective in hover or all three, i.e. collective and two cyclics, in cruise) as design variables. In the second stage, blade shape optimization is performed by including the performance objective, i.e., C_Q term into the time dependent objective function to be minimized. However, even after optimization convergence in stage

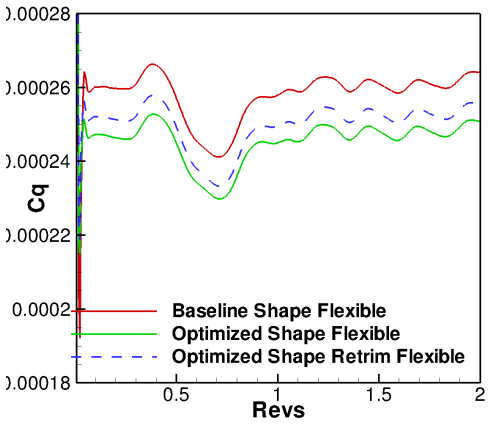

(a) Power time history

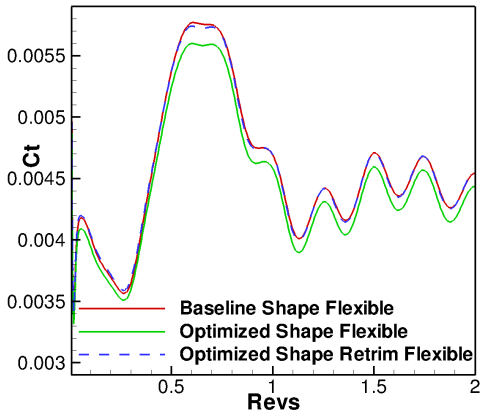

(b) Thrust time history

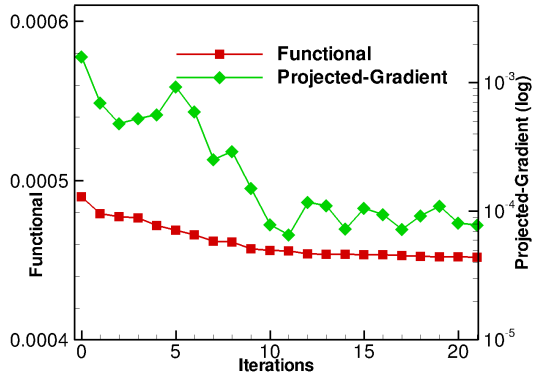

(c) Shape optimization convergence

Fig. 8. HART2 flexible blade shape optimization and re-trim

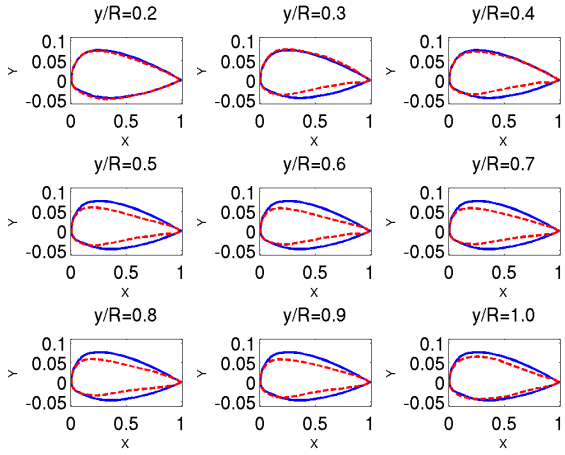

Fig. 9. HART2 flexible blade optimized blade sections; solid-baseline, dashed-optimized

two, the exact trim state is not maintained, since trim objectives are enforced only as weak penalty terms. Therefore, the last stage involves trimming the rotor back to the target thrust and moment values, using only the control pitch parameters (one in hover or three in cruise) as design inputs. This stage is otherwise referred to as the 're-trim' stage in this paper.

The optimization procedure used is the L-BFGS-B bounded reduced Hessian algorithm (Ref. 26). 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, a bound of $\pm 1.0^\circ$ of twist was set on the root and tip twist definitions, and a bound of $\pm 5.0^\circ$ of pitch angle was set on all the pitch parameters (collective and cyclics).

The multi-point design requires defining a combined objective function to account for all the design points, such as:

$$L_{mp} = \sum_i \omega_i L_i \tag{20}$$

where ω_i refers to the weight corresponding to a specific design point i. In the present work, for the two design points considered, namely, hover ($i = 1$) and forward flight ($i = 2$), $L_{mp} = \omega_1 L_{hover} + \omega_2 L_{cruise}$. The choice of weights for these two components has been found to be decisive in arriving at the optimal blade shape designs as optimization gains can be traded between the different design points. The different weights combination used are: $[\omega_1, \omega_2] = [0.5, 0.5]$, and $[0.25, 0.75]$.

Optimizing on multiple design flight conditions requires solving for an equivalent number of adjoint problems and hence the computational cost increases proportionally with the number of flight conditions. However, since the different design points are independent of each other, they can run in parallel. The present work implements this strategy to run multiple design points in parallel to ensure optimal

use of available computational resources for a limited wall clock time (e.g. 12 hours on Yellowstone supercomputer at the NCAR-Wyoming Supercomputing Center (NWSC)). For example, the serial multi-point simulation on 1024 cores takes 190 minutes per one design cycle, whereas the parallel multi-point simulation on 2048 cores (1024 cores for each design point) takes 100 minutes per one design cycle, which is approximately twice as fast as the serial run.

Currently, both rigid and flexible rotor multipoint design problems which combine both hover and forward flight conditions have been performed. For the hover flight design point, the baseline HART2 rotor has a 5 degree collective and no cyclics. The baseline HART2 rotor for the cruise design point has the pitch control parameters corresponding to the trimmed rotor state for a target mean thrust of $C_T = 4.4e-3$ and zero moments. These baseline pitch design parameters for the cruise design point for the rigid blade are: $\theta_0 = 4.23^\circ, \theta_{c_1} = 0.9^\circ, \theta_{s_1} = -1.77^\circ$. The hover design point makes use of 112 shape design parameters and only one pitch design parameter, i.e. the collective pitch, and hence a total of 113 design parameters. On the other hand, the cruise design point makes use of 112 shape design parameters and all the three pitch design parameters and hence a total of 115 design parameters. However, the collective blade pitch parameter for hover and for cruise are treated as separate design variables for the optimization, whereas the shape design parameters are common to both hover and cruise design points. Thus in total, the optimization operates on 116 design variables (i.e. 112 shape parameters, 3 pitch parameters for cruise, and 1 pitch parameter for hover). Since, the final objective of each design point is a function of both the shape and trim components ($L_i = L_{Shape_i} + L_{Trim_i}$) as described in Eqn. 13, including the pitch design parameters enables the shape optimized rotor to approximate the target trimmed state at each of the design points (hover and cruise).

Figures 10(a) and 10(b) show that the final blade shape optimized using weight combination $[\omega_1, \omega_2] = [0.5, 0.5]$ results in a 4.6% power reduction with a thrust penalty of 2.3% in hover, and a 3.9% power reduction with a thrust penalty of 3.3% in forward flight condition. After re-trim the optimized blade shape achieves power gains of 1.43% in hover and .38% in cruise. Figure 10(c) shows the design convergence for the shape optimization. The gradient has dropped by more than 1 order of magnitude and correspondingly, the combined objective has also converged significantly. Finally, Figure 11 compares the optimized blade shape with the baseline one at 9 different stations. Figures 12(a) and 12(b) show similar results optimized using weight combination $[\omega_1, \omega_2] = [0.25, 0.75]$. After re-trim the optimized blade shape achieves power gains of 1.5% in hover and .75% in cruise. Adding more weight to the cruise design point has resulted in more favorable power gain in forward flight. Figure 12(c) shows the design convergence for the shape optimization. The gradient has dropped by more than 1 order of magnitude and the objective has also converged significantly. Finally, Figure 13 compares the optimized blade shape with the baseline one at 9 different stations. The trend is similar to that observed for the first multi-point

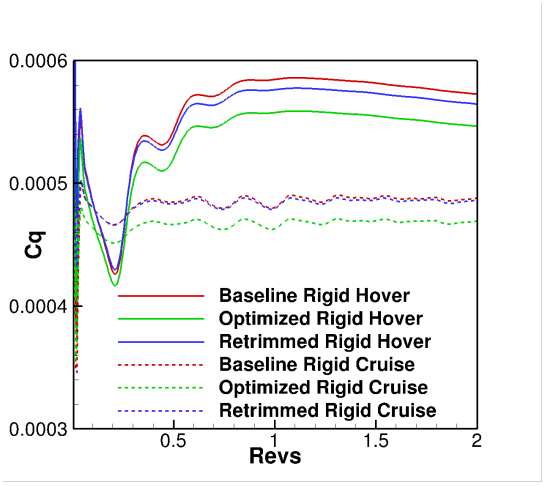

(a) Power time history

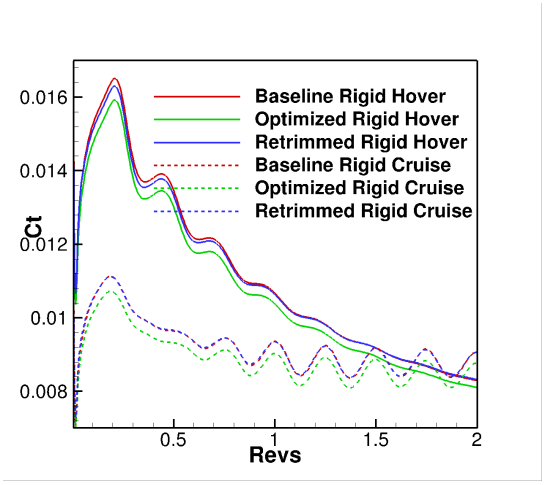

(b) Thrust time history

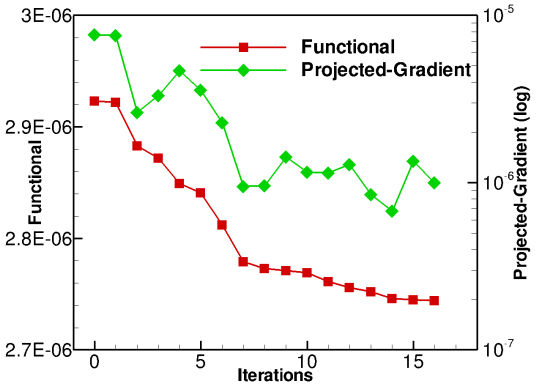

(c) Shape optimization convergence

Fig. 10. HART2 multi-point ($[\omega_1, \omega_2] = [0.5, 0.5]$) rigid blade shape optimization

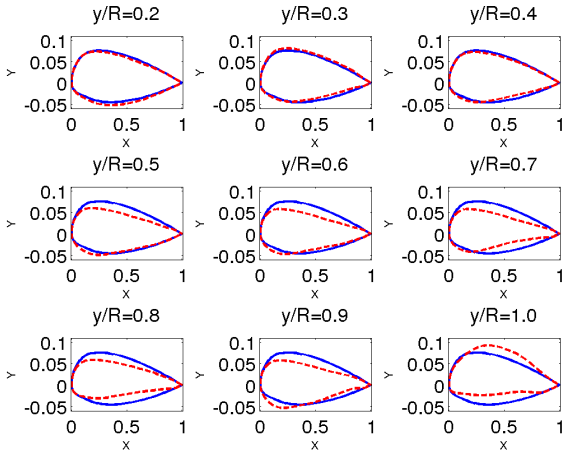

Fig. 11. HART2 multi-point ($[\omega_1,\omega_2] = [0.5,0.5]$) rigid blade optimized blade sections; solid-baseline, dashed-optimized

weight combination.

The multi-point coupled design optimization is further performed on flexible blades. Figures 14(a) and 14(b) show that the blade shape optimized using weight combination $[\omega_1,\omega_2] = [0.5,0.5]$ results in a 1.01% power reduction with a thrust penalty of 0.8% in hover, and a 0.92% power reduction with a thrust penalty of 1.1% in forward flight condition. Figure 14(c) shows that the objective is still dropping, but not significantly. However, as Fig. 15 shows, the optimized blade section at the $y/R = 0.9$ station has already evolved into a shape that reduces drag compared to the baseline section. Similar results are shown for the other multi-point weight combination $[\omega_1,\omega_2] = [0.25,0.75]$. Figures 16(a) and 16(b) show that the optimized shape results in a 1.2% power reduction with a thrust penalty of 1.05% in hover, and a 1.1% power reduction with a thrust penalty of 1.3% in forward flight condition. Definitely, adding more weight to the cruise design point has again resulted in more favorable power gain in forward flight as noted for the CFD only (rigid blade) optimization. However, it should be noted that in the cases shown here, the retrim phase was not performed, since the flexible multipoint optimizations tended to stall after a small number of design cycles and the optimization problem is not suitably converged.

Accuracy of Sensitivities

One of the possible causes of incomplete optimization convergence is related to the accuracy of the computed sensitivities. In previous work, we have verified that the computed sensitivities are indeed accurate to machine precision by comparing the adjoint computed sensitivities with those obtained using the forward or tangent sensitivity approach, and by comparing with the complex step method (Ref. 25). Although these tests confirm the correctness of the adjoint implementation, they require that all equations be converged to machine precision at each time step, and do not provide any indication of

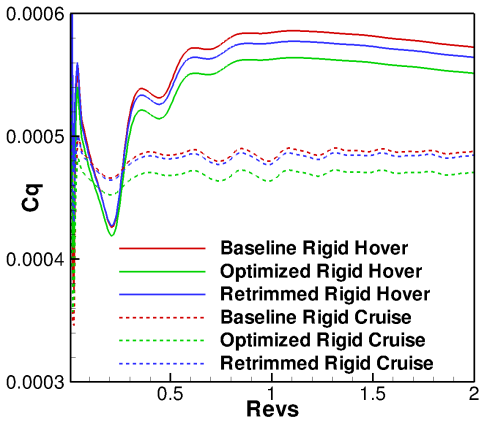

(a) Power time history

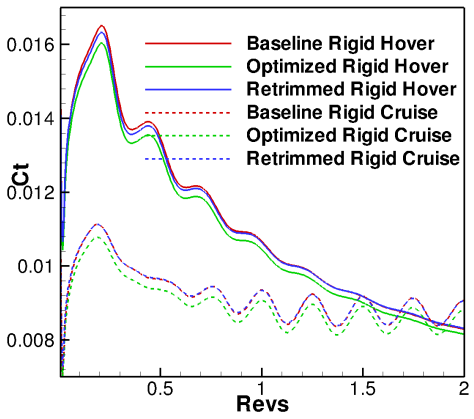

(b) Thrust time history

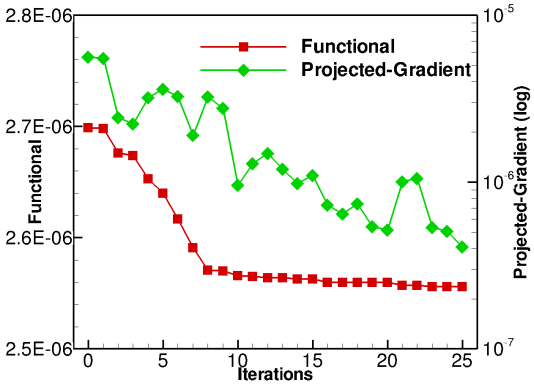

(c) Shape optimization convergence

Fig. 12. HART2 multi-point ($[\omega_1,\omega_2] = [0.25,0.75]$) rigid blade shape optimization

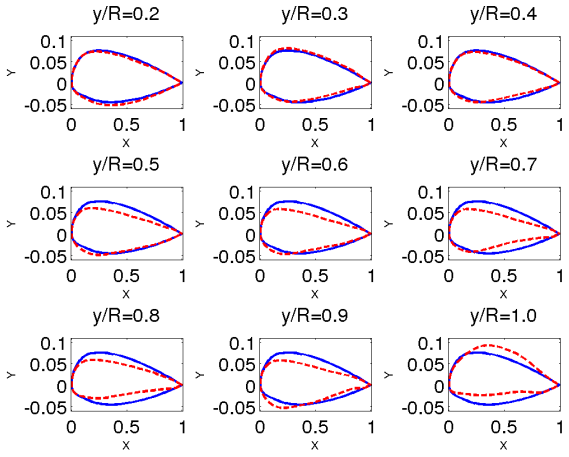

Fig. 13. HART2 multi-point ($[\omega_1, \omega_2] = [0.25, 0.75]$) rigid blade optimized blade sections; solid-baseline, dashed-optimized

round-off error growth as a function of time, or of inaccuracies due to incomplete convergence of the adjoint equations at each time step.

In previous work we have investigated the accuracy of the sensitivities over up to 5 time steps (Ref. 12). This verification is extended in Table (1), where the sensitivity of the single point forward flight objective L^g with respect to the collective pitch parameter obtained from the complex step method (C), the tangent method (T) and the adjoint method (A), are compared up to 180 time steps, which represents a full rotor revolution (using 2^o time steps). Both the rigid and the fully coupled aeroelastic problems are converged to machine zero at each time step in order to avoid contaminating the sensitivity values with errors due to incomplete convergence. As can be seen from the table, the sensitivities agree to at least 11 significant digits for both the rigid and flexible aeroelastic rotor cases, and the error does not appear to grow with the number of time steps.

Next, we examine the effect of partial convergence on the accuracy of the sensitivities. Considering the computational domain with 2.32 million points shown earlier for the forward flight condition, typical flow convergence for both rigid and flexible cases for a physical time step is shown in Figures 18(a) and 18(b). Both rigid as well as flexible blade simulations are converged by two orders of magnitude at each time step with the rigid blade requiring 60 non-linear iterations and the flexible blade requiring 6 coupling with 10 non-linear iterations per cycle. Corresponding to this order of convergence the accuracy of the tangent and adjoint sensitivities compared to the complex sensitivities (considered to be exact) are shown in Fig. 19(a). As can be noted, the errors accumulate over time within one rotor revolution and more interestingly, the adjoint error can accumulate to as large as 50% of the actual

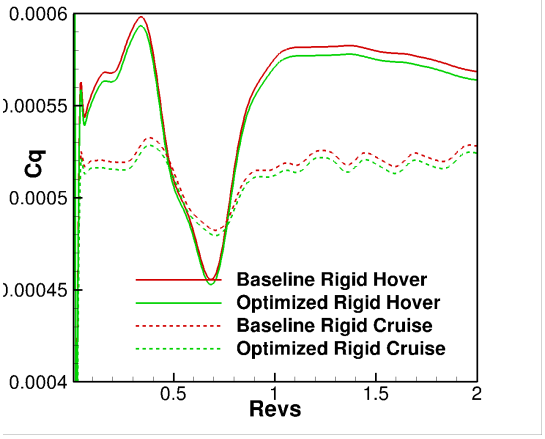

(a) Power time history

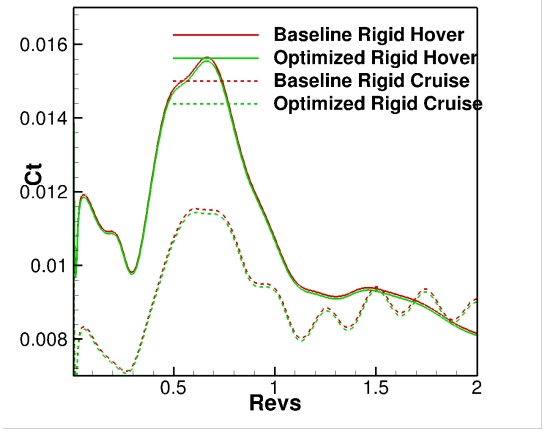

(b) Thrust time history

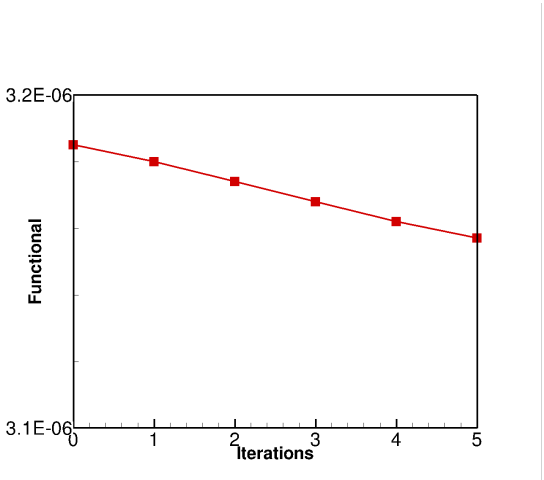

(c) Shape optimization convergence

Fig. 14. HART2 multi-point ($[\omega_1, \omega_2] = [0.5, 0.5]$) flex blade shape optimization

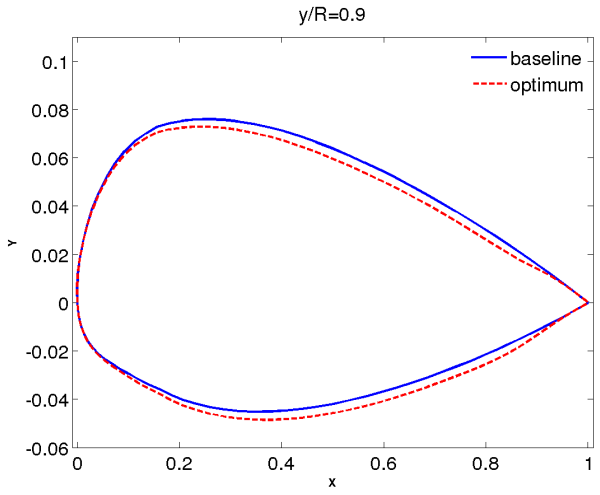

Fig. 15. HART2 multi-point ($[\omega_1, \omega_2] = [0.5, 0.5]$) flex blade optimized blade at $y/R = 0.9$; solid-baseline, dashed-optimized

(complex) sensitivity values (see Fig. 19(b)). In this case, the adjoint error accumulates faster than the tangent error as the simulation advances in time. Figure 19(c) shows that the sensitivity error accumulation is faster in the coupled CFD/CSD (flexible blade) than the CFD only (rigid blade) simulation. The coupled adjoint sensitivity error at the end of one rotor revolution is more than an order of magnitude larger than the corresponding error in the rigid CFD-only simulation. These results are somewhat surprising since the convergence of the analysis problem is very similar for both the rigid and flexible blade problems, as seen in Figures 18(a) and 18(b).

Fig. 20(a) depicts the accuracy of the sensitivities computed on a coarser mesh (107,000 mesh points) using the same level of convergence (2 orders of magnitude). As can be seen, the error grows faster for the finer mesh (Fig. 19(a)) than the coarse mesh for a fixed convergence tolerance (2 orders). However, if the convergence tolerance is made tighter, for example, by converging the solution by three orders of magnitude, the error growth in the sensitivities reduces significantly as can be seen from Figures 20(b) and 20(c).

These results indicate that, although the analysis problem displays similar convergence behavior for the rigid and flexible rotor cases, the sensitivities for the flexible rotor problem are more susceptible to error growth due to incomplete convergence, and tighter convergence tolerances may be required to ensure successful optimization outcomes. Furthermore, error growth may be more prevalent on finer meshes even for similar convergence tolerances, implying that further work on quantifying these effects is required.

CONCLUSIONS AND FUTURE WORK

In this work, a previously developed discrete adjoint formulation for time-dependent tightly coupled aeroelastic three-dimensional problems of flexible rotors has been extended to perform multi-point design optimization, both for rigid as well

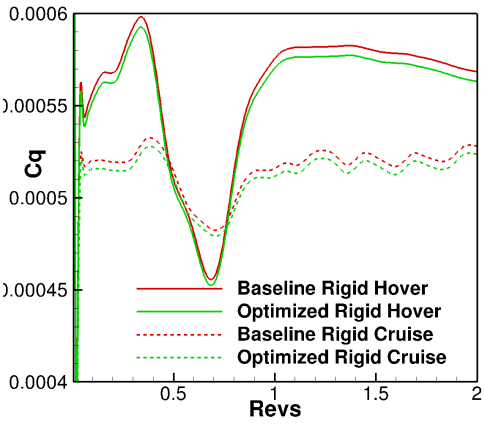

(a) Power time history

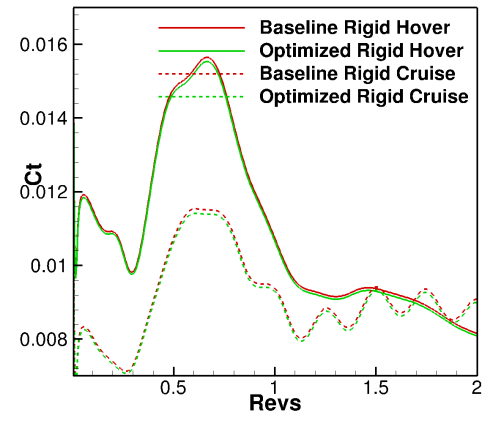

(b) Thrust time history

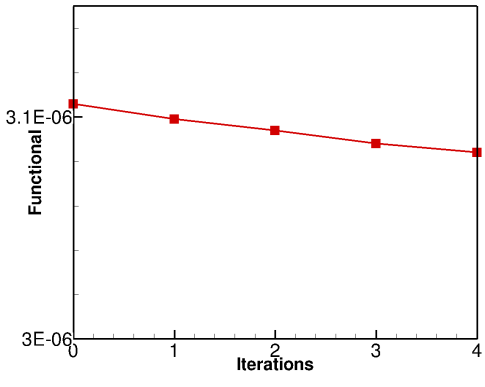

(c) Shape optimization convergence

Fig. 16. HART2 multi-point ($[\omega_1, \omega_2] = [0.25, 0.75]$) flex blade shape optimization

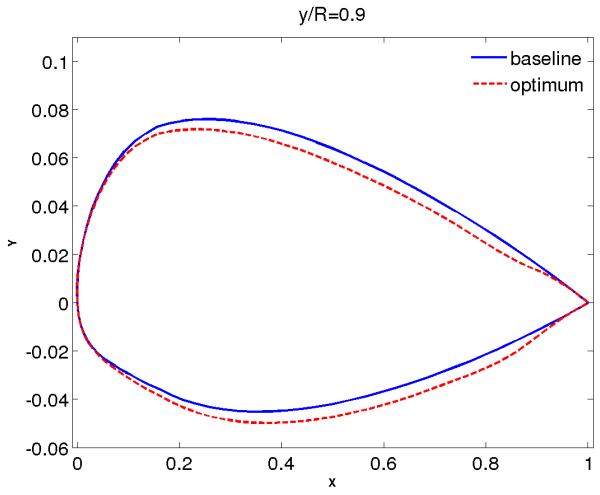

Fig. 17. HART2 multi-point ($[\omega_1, \omega_2] = [0.25, 0.75]$) flex blade optimized blade at $y/R = 0.9$; solid-baseline, dashed-optimized

as flexible blades. The current formulation is built upon a previously developed adjoint formulation for rotor problems in hover as well as forward flight conditions, by implementing the capability to solve for the multiple design points in parallel. The optimization minimizes a composite objective function constructed using a penalty formulation with several weight combinations for the multiple design points. The flexible multipoint design optimization is found to be more stiff than rigid and/or single point design optimization. A possible cause may be due to more rapid growth in the adjoint sensitivity errors due to incomplete convergence at each time step.

In the future, the empirically chosen design point weights will require further investigation for more effective multipoint design. Further work is required to find an optimal convergence tolerance to balance computational cost and sensitivity error accumulation. Future work will also consider a range of design point weights in the interest of computing Pareto fronts between complementary design points.

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

[1]Jameson, 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.

[2]Jameson, 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.

Table 1. Adjoint sensitivity verification for Forward Flight

n	Uncoupled (Rigid)	Coupled (Aero-elastic)
1 (C)	8.135662924562**930**E-005	7.569817143673**123**E-005
(T)	8.135662924562**923**E-005	7.569817143673**061**E-005
(A)	8.135662924562**884**E-005	7.569817143672**761**E-005
2 (C)	7.303561433847**516**E-005	6.040142774935**852**E-005
(T)	7.303561433847**498**E-005	6.040142774935**835**E-005
(A)	7.303561433847**042**E-005	6.040142774935**570**E-005
3 (C)	1.579075715783**277**E-005	-4.959909870786**381**E-006
(T)	1.579075715783**178**E-005	-4.959909870787**765**E-006
(A)	1.579075715780**985**E-005	-4.959909870785**228**E-006
4 (C)	-3.279789405316**752**E-005	-6.193896187819**846**E-005
(T)	-3.279789405316**760**E-005	-6.193896187820**087**E-005
(A)	-3.279789405324**491**E-005	-6.193896187820**419**E-005
5 (C)	-7.466264294860**007**E-005	-1.142069116982**308**E-004
(T)	-7.466264294859**811**E-005	-1.142069116982**308**E-004
(A)	-7.466264294878**836**E-005	-1.142069116982**432**E-004
180 (C)	-2.606225487510**254**E-003	-5.176189427439**016**E-003
(T)	-2.606225487510**356**E-003	-5.176189427439**005**E-003
(A)	-2.606225487528**119**E-003	-5.176189427434**507**E-003

[3]Nielsen, E. and Anderson, W., "Recent Improvements in Aerodynamic Design Optimization on Unstructured Meshes," *AIAA Journal*, Vol. 40, (6), June 2002, pp. 1155–1163.

[4]Jameson, A. and Martinelli, L., "Aerodynamic shape optimization techniques based on control theory," *Computational Mathematics Driven by Industrial Problems*, edited by R. Burkard, A. Jameson, G. Strang, P. Deuflhard, J.-L. Lions, V. Capasso, J. Periaux, and H. Engl, Vol. 1739, Lecture Notes in Mathematics, Springer Berlin Heidelberg, 2000, pp. 151–221.
doi: 10.1007/BFb0103920

[5]Giles, M., Duta, M., Muller, J., and Pierce, N. A., "Algorithm Developments for Discrete Adjoint Methods," *AIAA Journal*, Vol. 41, (2), February 2003, pp. 198–205.

[6]Nielsen, 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.

[7]Nielsen, 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.

[8]Jones, W. T., Nielsen, E., Lee-Rausch, E. M., and Acree, C. W., "Multi-point Adjoint-Based Design of Tilt-Rotors in a Noninertial Reference Frame," AIAA Paper 2014 – 0290, January 2014.

[9]Mani, 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.

[10]Mishra, A., Mani, K., Mavriplis, D. J., and Sitaraman, J., "Time-dependent Adjoint-based Optimization for Coupled

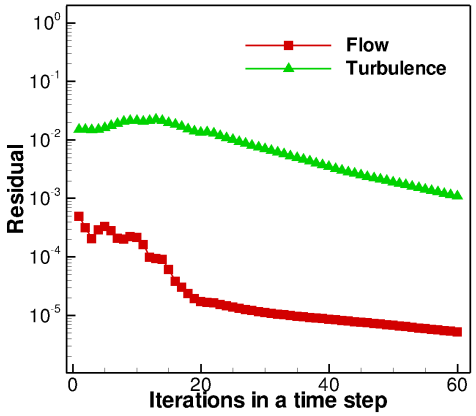

(a) Rigid Blade

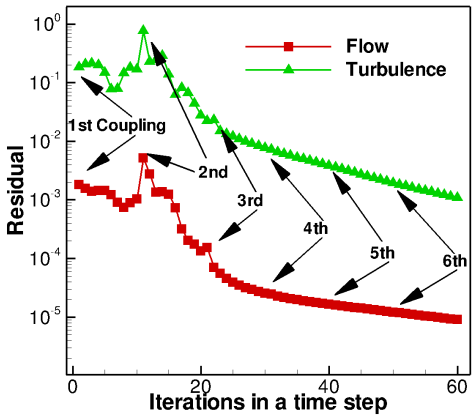

(b) Flexible Blade

Fig. 18. Flow and turbulence residual convergence at a given time step for rigid and flexible blade analysis

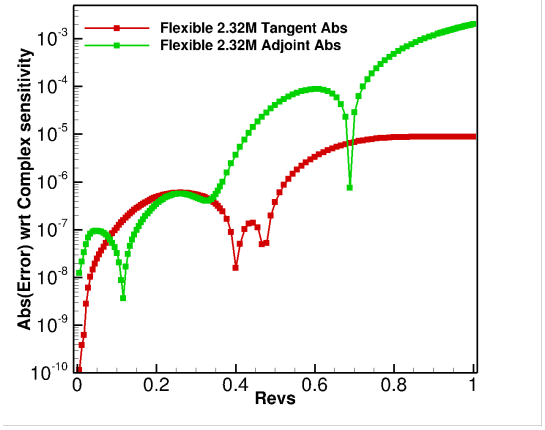

(a) Error CFD/CSD

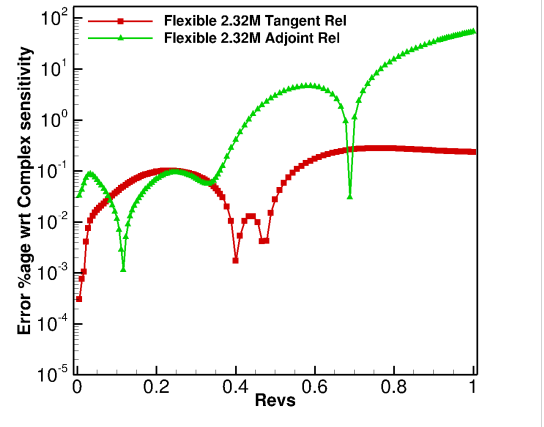

(b) Relative Error (%-age) - CFD/CSD

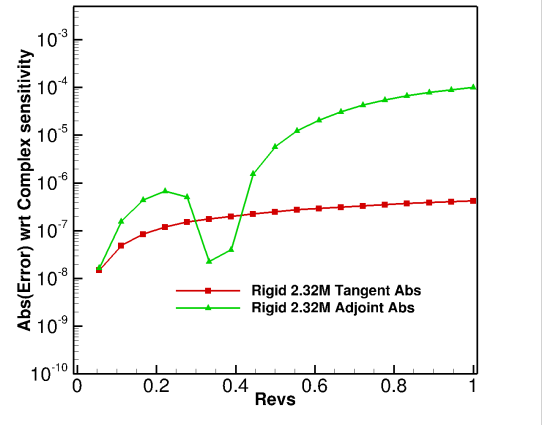

(c) Error CFD only

Fig. 19. Accuracy of sensitivity in 2.32 million nodes grid

Aeroelastic Problems," 21st AIAA CFD Conference, San Diego, CA, AIAA Paper 2013–2906, June 24–27 2013.

[11]Mishra, A., Mani, K., Mavriplis, D. J., and Sitaraman, J., "Time-dependent Adjoint-based Aerodynamic Shape Optimization Applied to Helicopter Rotors," 70th American Helicopter Society Annual Forum, Montreal, QC, CA, May 20–22 2014.

[12]Mishra, A., Mavriplis, D. J., and Sitaraman, J., "Time-dependent Aero-elastic Adjoint-based Aerodynamic Shape Optimization of Helicopter Rotors in Forward Flight," 56st AIAA Aerospace Sciences Meeting, Kissimmee, FL, AIAA Paper 2015-1130, January 5–9 2015.

[13]Kenway, G. K. W. and Martins, J. R. R. A., "Multipoint High-Fidelity Aerostructural Optimization of a Transport Aircraft Configuration," *Journal of Aircraft*, Vol. 51, (1), doi: 10.2514/1.C032150, 2014, pp. 144–160.

[14]Mavriplis, D. J., "Solution of the Unsteady Discrete Adjoint for Three-Dimensional Problems on Dynamically De-

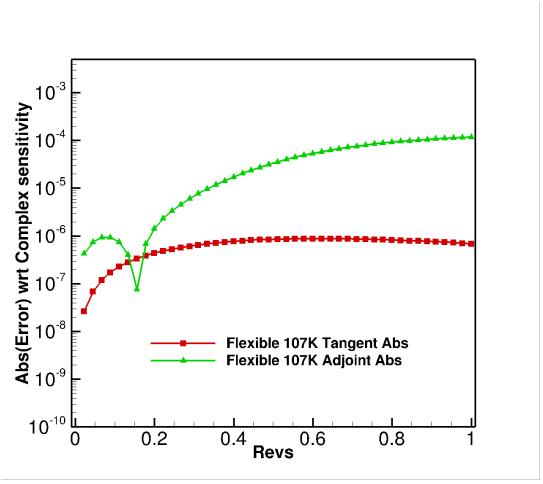

(a) Error CFD/CSD

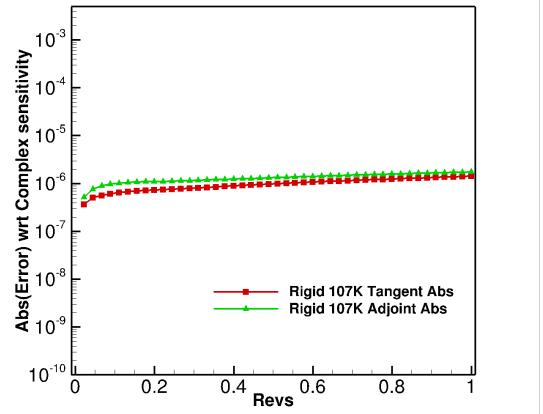

(b) Error CFD only

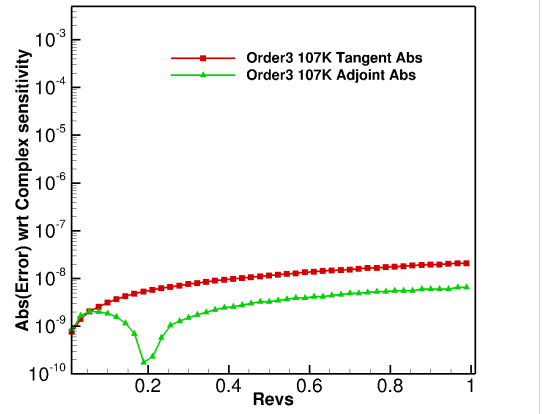

(c) Error CFD only (3 orders convergence)

Fig. 20. Accuracy of sensitivity in 0.107 million nodes grid

forming Unstructured Meshes," Proceedings of the 46th Aerospace Sciences Meeting and Exhibit, Reno NV, AIAA Paper 2008–0727, 2008.

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

[16]Spalart, P. R. and Allmaras, S. R., "A One-equation Turbulence Model for Aerodynamic Flows," *La Recherche Aérospatiale*, Vol. 1, 1994, pp. 5–21.

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

[18]Mavriplis, 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.

[19]Schuster, 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.

[20]Mishra, 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.

[21]Yu, 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.

[22]Chopra, I. and Bir, G., "University of Maryland Advanced Rotor Code: UMARC," American Helicopter Society Aeromechanics Specialists Conference, January 1994.

[23]Heckmann, 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.

[24]Yang, 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.

[25]Lyness, J. and Moler, C., "Numerical Differentiation of Analytic Functions," *SIAM Journal of Numerical Analysis*, Vol. 4, 1967, pp. 202–210.

[26]Zhu, 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.