High-Fidelity Multidisciplinary Sensitivity Analysis and Design Optimization for Rotorcraft Applications

Li Wang[1] and Boris Diskin[2]
National Institute of Aerospace, Hampton, VA 23666

Robert T. Biedron[3] and Eric J. Nielsen[4]
NASA Langley Research Center, Hampton, VA 23666

Olivier A. Bauchau[5]
University of Maryland, MD 20742

A multidisciplinary sensitivity analysis of rotorcraft simulations involving tightly coupled high-fidelity computational fluid dynamics and comprehensive analysis solvers is presented and evaluated. A sensitivity-enabled fluid dynamics solver and a nonlinear flexible multibody dynamics solver are coupled to predict aerodynamic loads and structural responses of helicopter rotor blades. A discretely-consistent adjoint-based sensitivity analysis available in the fluid dynamics solver provides sensitivities arising from unsteady turbulent flows and unstructured dynamic overset meshes, while a complex-variable approach is used to compute structural sensitivities with respect to aerodynamic loads. The multidisciplinary sensitivity analysis is conducted through integrating the sensitivity components from each discipline of the coupled system. Accuracy of the coupled system is validated by conducting simulations for a benchmark rotorcraft model and comparing solutions with established analyses and experimental data. Sensitivities of lift computed by the multidisciplinary sensitivity analysis are verified by comparison with the sensitivities obtained by complex-variable simulations. Finally the multidisciplinary sensitivity analysis is applied to a constrained gradient-based design optimization for a HART-II rotorcraft configuration.

---

[1] Senior Research Engineer, AIAA Member, li.wang@nianet.org
[2] Research Fellow, AIAA Associate Fellow, bdiskin@nianet.org
[3] Research Scientist, robert.t.biedron@nasa.gov
[4] Research Scientist, AIAA Associate Fellow, eric.j.nielsen@nasa.gov
[5] Professor, AIAA Senior Member, obauchau@umd.edu

Nomenclature

C_N = normal-force coefficient

C_{Mx} = rolling-moment coefficient

C_{My} = pitching-moment coefficient

C_P = power coefficient

C_Q = torque coefficient

C_T = thrust coefficient

$\mathbf{D}$ = design parameters

f = a scalar-valued function

$\mathbf{f}$ = CFD airloads

$\mathbf{g}$ = constraint functions

h = perturbation size

i = imaginary unit for complex numbers

k, n = time step indices

$\mathbf{K}$ = mesh elasticity matrix

$\overline{\mathbf{K}}$ = mesh elasticity matrix evaluated by input volume grid

L = Lagrangian functional

M = freestream Mach number

$\mathbf{n}$ = unit vector, $\mathbf{n} = (n_x, n_y, n_z)$

N = total number of time steps

N_b = number of rotor blades

P = rotor shaft power

Q = rotor shaft torque

$\overline{\mathbf{Q}}$ = initial flow solution

$\mathbf{Q}$ = CFD flow solution

r = radius coordinate

$\mathbf{r}$ = coordinate vector, $\mathbf{r} = (x, y, z)$

R = blade radius

$\mathbf{t}$ = vector for a translational motion

$\tilde{\mathbf{t}}_s$ = translation vector for surface nodes

T = rotor thrust

$\mathbf{T}$ = 4×4 transform matrix

$\mathbf{T}_r$ = transform matrix of averaged rigid motion

$\mathbf{T}_s$ = transform matrix of elastic surface motion

$\mathbf{u}$ = comprehensive analysis surface displacements

$\mathbf{W}$ = block matrix of rotation motion

$\mathbf{W}_r$ = rotation matrix of averaged rigid motion

x = independent variable

$\mathbf{X}$ = CFD grid coordinates

$\overline{\mathbf{X}}$ = input volume grid coordinates

$\hat{\mathbf{X}}$ = CFD reference grid coordinates

$\mathbf{X}_s$ = CFD surface grid coordinates

$\hat{\mathbf{X}}_s$ = CFD surface reference grid coordinates

β = flap angle

δ = lead-lag angle

Δt = time step size

κ = coefficient of UMUSCL scheme

$\mathbf{\Lambda}_F$ = load adjoint solution

$\mathbf{\Lambda}_G$ = grid adjoint solution

$\mathbf{\Lambda}_R$ = flow adjoint solution

Ω = angular speed of rotor

ϕ_0 = azimuth angle to move blade to reference frame

ψ = azimuth angle of blade

θ = pitch angle

θ_0 = collective pitch angle

θ_{1c} = lateral cyclic pitch angle

θ_{1s} = longitudinal cyclic pitch angle

I. Introduction

High-fidelity analysis methods for rotorcraft aeromechanics have become increasingly important to capture complex physics involving transonic flows, vortical wakes, reversed flows, blade-vortex interactions, and large amplitudes of elastic flapping, lagging and torsion motions. A rotorcraft analysis often requires many disciplines such as aerodynamics, aeroacoustics, structural dynamics and deformations, flight mechanics, and others to account for complex interactions of unsteady fluids with highly flexible rotor blades. Multidisciplinary rotorcraft comprehensive analysis (CA) tools [1–4] are often used to simulate rotorcraft aeromechanics. These tools are fast, but rely on low-fidelity aerodynamic models, such as lifting-line and vortex-wake models. The state of the art in high-fidelity rotorcraft analysis is represented by simulations that couple a CA model with a computational fluid dynamics (CFD) model [5–8].

Gradient-based optimization of rotorcraft simulations has been a focus of study in the past 20 years [9–11]. Most of the studies involving CA tools have been conducted by using finite-difference approximations for sensitivities. Finite-difference approaches are effective for computing sensitivity of many objective functions with respect to a few design parameters. The number of simulations required for such computations is proportional to the number of design parameters; hundreds of simulations may be needed for design and optimization of complex rotorcraft configurations. A finite-difference approach to sensitivity analysis may be acceptable for optimization of low-cost CA models, but it is not feasible for high-fidelity rotorcraft analysis, in which a single large-scale simulation involving a coupled CFD/CA system may require hundreds of thousands of CPU hours of high-performance computing. On the other hand, adjoint formulations developed in some state-of-the-art CFD solvers enable efficient computations of the sensitivity of an objective function to many design parameters [12–17]. The computational cost of an adjoint-based sensitivity analysis approximately equals the cost of a single simulation and does not increase with the number of design parameters. This property makes adjoint methods especially suitable for shape optimization where a few objective/constraint functions are used, but the number of design parameters is large. Discretely-consistent adjoint methods offer an additional advantage that the computed numerical sensitivities can be rigorously verified through comparisons with the exact sensitivities computed by a complex-variable finite-difference approach [18]. In contrast to the finite-difference methods employing real-valued perturbations, the complex-variable approach does not suffer from subtraction error cancellation and is capable of producing the true sensitiv-

ity of the discrete solution with respect to a chosen design parameter. Recently, adjoint-based sensitivity analyses of simplified rotorcraft simulations [15, 16, 19–21] have been demonstrated to guide design and optimization.

In the context of multidisciplinary simulations, the sensitivity analysis and the subsequent gradient-based optimization procedures must include sensitivities across all disciplines, thereby leading to a multidisciplinary sensitivity analysis. While discretely-consistent adjoint-based multidisciplinary sensitivity analysis is the ultimate goal, the implementation of a discretely-consistent formulation in a production-level code is a complex task that may require years of development effort. Moreover, development of a multidisciplinary adjoint formulation requires detailed expertise and unlimited access to the source code(s) of each discipline involved. Modern multidisciplinary applications often couple large legacy codes developed over many years by different groups and organizations. The codes often exchange data through specialized interfaces without exposing the internal treatment of the discipline solutions. In many cases, the source code(s) and/or expertise related to a particular discipline are simply not available. Even if the required expertise and source code(s) are available, the computational time may not be evenly distributed between disciplines. A useful multidisciplinary sensitivity analysis tool can couple efficient adjoint-based sensitivities computed for “heavier" disciplines such as CFD with sensitivities computed by “black-box” or “grey-box” methods for “lighter” disciplines such as CA. The “black-box” methods may involve real or complex-valued finite-difference approaches; the “grey-box” methods may involve automatic (algorithmic) differentiation methods [22–24]. The sensitivity analysis for each discipline can then be integrated within a multidisciplinary sensitivity analysis framework. A theoretical foundation for such integration has been presented in the literature [25]; and a practical implementation of a black-box formulation has been recently reported [26].

The goal of this paper is to present and evaluate a multidisciplinary sensitivity analysis approach for high-fidelity CFD/CA simulations based on a tight coupling methodology. An unstructured-grid, highly-scalable CFD solver, FUN3D [7, 27, 28], and a nonlinear flexible multibody dynamics CA solver, DYMORE [1], are coupled to predict the rotor airloads and structural responses of helicopter rotor blades. For overset meshes, the domain-connectivity information is established via the DiRTlib and SUGGAR++ software libraries [29, 30]. A discretely-consistent adjoint-based sensitivity analysis available in FUN3D provides sensitivities arising from unsteady turbulent flows and unstructured dynamic overset meshes, while a complex-

variable finite-difference approach is used to compute sensitivities of surface displacements with respect to aerodynamic loads. The computations of the sensitivity of structural responses with respect to different air-load components are "embarrassingly parallel" and can be simultaneously conducted by different processors. In contrast to the analysis procedure, in which DYMORE operates in the serial execution mode, the sensitivity analysis of DYMORE is distributed to all the processors used in the CFD analysis. The multidisciplinary sensitivity analysis integrates all sensitivity components from the coupled systems.

The paper is organized as follows. The technical approach is covered in Section II and followed by numerical results in Section III. Section IV concludes with a summary of the work and future plans.

II. Technical Approach

In this section, an overview of FUN3D and DYMORE capabilities is provided. Tight and loose coupling approaches for multidisciplinary analysis are introduced. The complex-variable approach to sensitivity analysis is briefly described. Details of mathematical formulations for the tight-coupling analysis and for the multidisciplinary sensitivity analysis are given. Correctness of the sensitivity analysis implementation is verified, and the design optimization procedure is described.

A. Computational Fluid Dynamics and Comprehensive Analysis Solvers

Solutions of the Reynolds-averaged Navier-Stokes (RANS) equations are computed with the FUN3D flow solver [7, 27, 28]. A standard one-equation Spalart-Allmaras turbulence model [31] is used for the current simulations. FUN3D is a finite-volume, node-centered, unstructured-grid RANS solver, which is widely used for high-fidelity analysis and adjoint-based design of complex turbulent flows [15, 17, 32]. FUN3D solves the governing flow equations on mixed-element grids; the elements can be tetrahedra, pyramids, prisms, and hexahedra. At median-dual control-volume faces, the inviscid fluxes are computed by using an approximate Riemann solver. In the current study, Roe's flux difference splitting [33] is used. For second-order accuracy, solutions at dual faces are obtained by a UMUSCL (Unstructured Monotonic Upstream-centered Scheme for Conservation Laws) scheme [34, 35] setting $\kappa = 0.5$ for the flow equations and $\kappa = 0$ for the turbulence model. The viscous fluxes are discretized with a finite-volume formulation, in which the Green-Gauss theorem is used to compute gradients on the dual faces for tetrahedral meshes; this is equivalent to a Galerkin type approximation. For nontetrahedral meshes, the Green-Gauss (cell-based) gradi-

ents are combined with edge-based gradients to improve the h-ellipticity of viscous operators. The diffusion term in the turbulence model is handled in the same fashion as the viscous terms in the flow equations. The vorticity-based source term for the turbulence model is computed by using velocity gradients evaluated by the Green-Gauss method on dual volumes. To advance the equations in time, a library of time integration schemes is available including standard first-order and second-order backward difference schemes termed as BDF1 and BDF2 schemes, respectively, as well as a blended second- and third-order backward difference scheme, referred to as BDF2opt [36]. The BDF2opt scheme is second-order accurate in time, but has a smaller error than the BDF2 scheme. For overset meshes, the DiRTlib [29] and SUGGAR++ [30] codes are used to facilitate communications between components of the mesh.

The CA solver used in this study is DYMORE [1], a nonlinear flexible multibody dynamics analysis model that provides static, dynamic, stability, and trim analyses of rotorcraft configurations. DYMORE contains libraries of primitive elements such as rigid bodies, mechanical joints, elastic springs, dampers, beams, shells and plates. For computational structural dynamics analysis, DYMORE uses a high-fidelity finite-element method in the time domain without relying on a low-fidelity model reduction method. A generalized alpha discretization [37] is used in DYMORE for time advancement. Although DYMORE has not been developed specifically for rotorcraft applications, it has been widely used in this area due to its capabilities of handling nonlinear flexible systems with arbitrary topology. The internal aerodynamics model in DYMORE is a low fidelity approximation based on lifting line theory and vortex wake models. For this study, a complex-variable version of DYMORE has been developed and verified to enable accurate computations of the sensitivity of surface displacements with respect to aerodynamic loads.

B. Multidisciplinary Analysis

A FUN3D/DYMORE interface has been previously developed to exchange aerodynamic loads and structural responses between high-fidelity aero and structural dynamics models [38]. In the current work, this interface has been updated and extended to enable multidisciplinary sensitivity analysis as well as to handle mixed-mode communications between real-valued FUN3D and complex-valued DYMORE solutions. The extended multidisciplinary system can conduct loose or tight coupling CFD/CA analysis [26].

In the loose coupling approach [5], the CA solution obtained with an internal low-fidelity aerodynamics

model provides initial elastic blade deformations for the CFD solver. In particular, the two-dimensional unsteady airfoil theory of Peters [39] is employed in DYMORE for obtaining the internal aerodynamic loads at the airload stations. CFD computations are conducted over a characteristic period of time corresponding to a coupling iteration. The first coupling iteration starts from a freestream condition and proceeds for the time period corresponding to one full revolution. The subsequent coupling iterations use the time period corresponding to the $2/N_b$ fraction of a revolution, where N_b is the number of rotor blades. The extended time period used in the first coupling iteration helps mitigate the transient effects of initial conditions. The differences between airloads computed by the CFD solver in a coupling iteration and the corresponding airloads computed through the internal, low-fidelity aerodynamics of the CA model are called delta airloads. The delta airloads are collected over the coupling iteration time period. After completing the CFD computations, the delta airloads are passed to the CA solver to compute trimmed elastic motions. The latter are transferred back to the CFD solver to update the blade surface geometries and motions. This process is repeated until converged periodic solution and trim control angles are achieved.

In the tight coupling approach, the CFD and CA models are still executed separately, but the data exchange occurs once per time step. In the present implementation of the tight coupling procedures, the data exchange between the CFD and CA models is performed in memory by linking FUN3D with a prebuilt DYMORE library, which results in direct, fast data access. Fig. 1 shows a diagram illustrating the CFD/CA tight coupling procedures, in which FUN3D functions involving DYMORE structural solutions are highlighted.

At each tight coupling iteration, DYMORE performs a single time step advancement by using sectional airloads as inputs; and the internal aerodynamics model in DYMORE is disabled. Initial sectional airloads can be computed from a given initial flow condition. A suitable initial condition can be obtained from a loose coupling CFD/CA solution. The linear and angular deflections at the quarter-chord line of rotor blades computed by DYMORE are sent to FUN3D. Within FUN3D, high-order surface spline functions are generated to distribute deflections over the entire blade surface, averaged rigid motions are extracted from the CA deflections, and the CFD surface grid is updated. The reference CFD volume grid is rigidly moved using the averaged rigid motions and elastically deformed to accommodate the new surface mesh. The CFD solution at the current time is iterated on the updated mesh. Using the updated CFD solution, a boundary slicing procedure is performed to compute sectional airloads (forces and moments) at specified

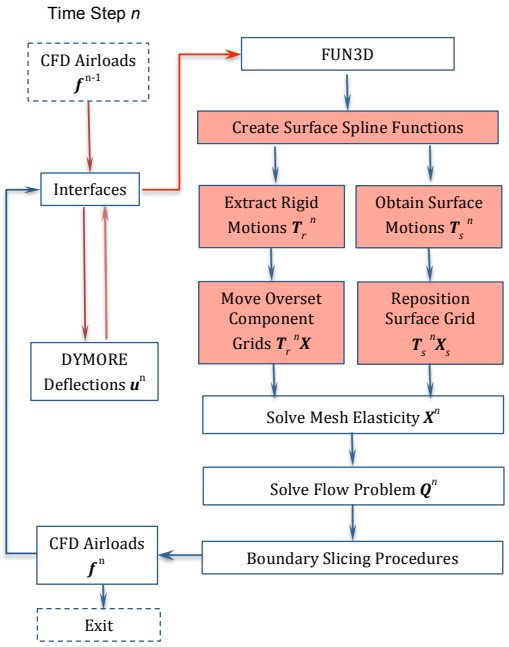

Fig. 1 Diagram illustrating the CFD/CA tight coupling procedures.

station locations. The new CFD airloads are transferred to DYMORE to start the next coupling iteration. In Section III A, a study is conducted to verify the accuracy of the CFD/CA loose and tight coupling procedures for a HART-II rotorcraft configuration [8].

C. Complex-Variable Method for Sensitivity Analysis

In this work, the complex-variable finite-difference method [18] is employed by the multidisciplinary CFD/CA sensitivity analysis to compute sensitivities of structural responses with respect to aerodynamic loads and to verify multidisciplinary sensitivities. A brief description of the method is given as follows.

For a real-valued function, $f(x)$, of a real variable x, the complex-variable method allows simultaneous evaluation of the function and its derivative, $\frac{df}{dx}$. For an input, $x+ih$, with an imaginary perturbation, ih, the complex-valued function can be expanded in the Taylor series as:

$$f(x+ih) = f(x) + ih\frac{df}{dx} - \frac{h^2}{2}\frac{d^2f}{dx^2} - i\frac{h^3}{6}\frac{d^3f}{dx^3} + \cdots . \tag{1}$$

From this equation, the real part of the complex-valued expansion gives the function itself,

$$f(x) \approx Re[f(x+ih)], \tag{2}$$

and the imaginary part yields the derivative

$$\frac{df}{dx} \approx \frac{Im[f(x+ih)]}{h}, \tag{3}$$

with a truncation error $O(h^2)$.

Unlike a finite-difference method with a real-valued perturbation, the complex-variable approach does not suffer from subtraction error cancellations, thereby providing increased accuracy and robustness. Within machine accuracy, derivatives computed by the complex-variable method with perturbation $h = 10^{-50}$ can be considered as the true derivatives with zero truncation error.

The standard FUN3D software suite allows for complex-mode operations to verify hand-differentiated sensitivities, but a similar capability was not available in DYMORE. For the present study, a complex mode has been developed in DYMORE by replacing all double and float variables declared in the original (i.e., real-valued) DYMORE code with complex variables of double precision and redefining operators needed for complex-valued operations. The complex mode of DYMORE enables the evaluation of CA sensitivities via the complex-variable approach. The FUN3D/DYMORE interfaces have been extended to operate in the complex mode, supporting the exchange of complex-valued aerodynamic forces, moments and surface displacements.

D. Mathematical Formulation for Tight Coupling Analysis

A mathematical formulation for the tight coupling multidisciplinary CFD/CA analysis is given in this section. A CA surface-displacement vector at time level n, denoted as $\mathbf{u}^n$, is provided at specified stations on the surface and used to update the CFD surface grid, $\mathbf{X}_s^n$. A vector of CFD solutions, $\mathbf{Q}^n$, is computed on the overset volume grid, $\mathbf{X}^n$, at time level n. A vector of airloads, $\mathbf{f}^n$, is computed by the CFD code at time level n at specified stations and transferred to the CA model to compute structural responses. The dimension for $\mathbf{u}^n$ and $\mathbf{f}^n$ is equal to the number of airload stations from all rotor blades multiplied by the degrees of freedom (i.e., three translations and three rotations for $\mathbf{u}^n$, and three forces and three moments for $\mathbf{f}^n$) at each station.

An overset CFD volume grid is composed of a stationary background grid and component grids corresponding to moving bodies. Each body and the corresponding component grid can experience an independent motion. A combination of rotational and translational motions applied to a point in space can be compactly

represented by a 4×4 transform matrix, $\mathbf{T}$, in the following form

$$\mathbf{T} = \begin{bmatrix} W_{11} & W_{12} & W_{13} & t_x \\ W_{21} & W_{22} & W_{23} & t_y \\ W_{31} & W_{32} & W_{33} & t_z \\ 0 & 0 & 0 & 1 \end{bmatrix}. \tag{4}$$

Here, the upper-left 3×3 block is the prescribed rotation matrix, $\mathbf{W}$, and the vector $\mathbf{t} = (t_x, t_y, t_z)^T$ appearing in the fourth column defines the translation. A rotation matrix describing rotation around a unit vector $\mathbf{n} = (n_x, n_y, n_z)$ on angle θ can be represented as

$$\mathbf{W}(\theta) = \cos(\theta)\,\mathbf{I} + \sin(\theta)\,[\mathbf{n}]_{\mathbf{x}} + (1 - \cos(\theta))\,\mathbf{n}^T\mathbf{n}. \tag{5}$$

Here $\mathbf{I}$ is the 3×3 identity matrix, $[\mathbf{n}]_{\mathbf{x}}$ is the cross product of vector $\mathbf{n}$ defined as

$$[\mathbf{n}]_{\mathbf{x}} = \begin{bmatrix} 0 & n_z & -n_y \\ -n_z & 0 & n_x \\ n_y & -n_x & 0 \end{bmatrix}, \tag{6}$$

and $\mathbf{n}^T\mathbf{n}$ is the tensor product

$$\mathbf{n}^T\mathbf{n} = \begin{bmatrix} n_x^2 & n_x n_y & n_x n_z \\ n_x n_y & n_y^2 & n_y n_z \\ n_x n_z & n_y n_z & n_z^2 \end{bmatrix}. \tag{7}$$

A motion $\mathbf{T}$ of a point $\mathbf{r} = (x, y, z)^T$ to the new position, $\tilde{\mathbf{r}} = (\tilde{x}, \tilde{y}, \tilde{z})^T$, is formally expressed as

$$\begin{bmatrix} \tilde{\mathbf{r}} \\ 1 \end{bmatrix} = \mathbf{T} \begin{bmatrix} \mathbf{r} \\ 1 \end{bmatrix}. \tag{8}$$

In an expanded form, the same motion can be described as

$$\tilde{\mathbf{r}} = \mathbf{W}\mathbf{r} + \mathbf{t}. \tag{9}$$

A rigid motion applied to a component grid implies that all points of the grid undergo the same motion.

A transform matrix can be formed to represent a sequence of motions by multiplying individual transform matrices corresponding to each motion. For example, the composite transform matrix of a rigidly moving body experiencing pitch, flap and lead-lag motions is obtained as

$$\mathbf{T}_r(\delta, \beta, \theta) = \mathbf{T}(\delta)\mathbf{T}(\beta)\mathbf{T}(\theta), \tag{10}$$

where δ, β, and θ are lead-lag, flap, and pitch control angles, respectively.

A component surface grid, $\mathbf{X}_s^n$, at time level n is derived from a reference surface grid, $\hat{\mathbf{X}}_s(\mathbf{D})$. The reference surface grid is based on surface parameterization and explicitly depends on the shape design parameters, $\mathbf{D}$. The reference surface grid does not include elastic deformations or rigid body motions. The corresponding reference volume grid, $\hat{\mathbf{X}}$, is obtained in the following steps. An input volume grid, $\overline{\mathbf{X}}$, is read from the disk and assumed independent of design parameters. The corresponding input surface grid, $\overline{\mathbf{X}}_s$, is a subset of $\overline{\mathbf{X}}$ that contains coordinates of the surface points. The input surface grid has to have the same degrees of freedom and the same surface-grid connectivity as the reference surface grid, but coordinates of surface points may not match. If the surface point coordinates match, then the input volume grid becomes the reference volume grid. This straightforward situation is typical for a coupled CFD/CA analysis, where the input volume grid matching a given surface grid is obtained elsewhere, e.g., from a grid generator, and the surface parameterization is not changed during simulations. In general, a mismatch may occur between surface grids $\hat{\mathbf{X}}_s$ and $\overline{\mathbf{X}}_s$. Then, the input volume grid can be elastically deformed to match $\hat{\mathbf{X}}_s$. The elasticity matrix, $\overline{\mathbf{K}}$, is computed based on the input volume grid, $\overline{\mathbf{K}} = \mathbf{K}(\overline{\mathbf{X}})$, and is considered independent of design parameters. The reference volume grid, $\hat{\mathbf{X}}$, is obtained by solving the following mesh elasticity equation

$$\overline{\mathbf{K}}\left(\hat{\mathbf{X}} - \overline{\mathbf{X}}\right) = \hat{\mathbf{X}}_s - \overline{\mathbf{X}}_s. \tag{11}$$

Here, the surface grids are extended to the dimension of the volume grids by assigning zeros to the positions corresponding to coordinates of the interior points. Note that if the reference and input surface grids match, i.e., $\hat{\mathbf{X}}_s = \overline{\mathbf{X}}_s$, then the reference and input volume grids match too, $\hat{\mathbf{X}} = \overline{\mathbf{X}}$. Thus, Eq. (11) describes both matched and mismatched cases.

A component surface grid, $\mathbf{X}_s^n$, represents at time level n the actual shape and position of a solid body such as a blade or a fuselage. The surface grid of a flexible body can experience elastic deformations that are determined by the displacements $\mathbf{u}^n$ computed by the CA model. The surface deformations are computed with respect to the corresponding reference surface grid, $\hat{\mathbf{X}}_s$. Different moving bodies of the same type, e.g., different rotor blades, can use the same reference surface grid. The displacements are available at specified stations of the reference surface grid. The displacement interpolation in the spanwise direction is based on high-order spline functions. The spanwise displacements are treated as invariant along the cross-sections of

the blades. The relation between $\mathbf{X}_s^n$ and $\hat{\mathbf{X}}_s$ can be expressed as

$$\mathbf{X}_s^n = \mathbf{T}_s^n(\mathbf{u}^n)\hat{\mathbf{X}}_s(\mathbf{D}), \tag{12}$$

$$\mathbf{T}_s^n = \mathbf{T}_{\psi_n}\mathbf{T}(\delta(\mathbf{u}^n))\mathbf{T}(\beta(\mathbf{u}^n))\mathbf{T}(\theta(\mathbf{u}^n))\mathbf{T}_{\phi_0}, \tag{13}$$

where $\mathbf{T}_s$ is a transform matrix denoting elastic motions of the body surface points. The transform $\mathbf{T}_{\phi_0}$ rigidly moves the body from the specific initial position corresponding to the azimuthal angle ϕ_0 to the reference frame, in which the CA displacements are defined; the transforms $\mathbf{T}(\delta(\mathbf{u}^n))$, $\mathbf{T}(\beta(\mathbf{u}^n))$, and $\mathbf{T}(\theta(\mathbf{u}^n))$ describe displacements at individual surface points (in lead-lag angle, flap angle and pitch angle, respectively); the transform $\mathbf{T}_{\psi_n}$ rigidly moves the body from the reference frame to the position corresponding to time level n. Note that, at time level $n = 0$, $\mathbf{T}_{\psi_0} = \mathbf{T}_{\phi_0}^{-1}$. At time levels $n \geq 1$, $\mathbf{T}_{\psi_n}$ depends on kinematic design parameters, such as control angles and rotation rate.

The volume component grid, $\mathbf{X}^n$, corresponding to time level n is obtained by a combination of rigid motions and elastic mesh deformations applied to the reference volume grid. The associated rigid motion is computed as

$$\mathbf{T}_r^n = \mathbf{T}_{\psi_n}\mathbf{T}_r(\bar{\delta}(\mathbf{u}^n), \bar{\beta}(\mathbf{u}^n), \bar{\theta}(\mathbf{u}^n))\mathbf{T}_{\phi_0}, \tag{14}$$

where $\mathbf{T}_r(\bar{\delta}(\mathbf{u}^n), \bar{\beta}(\mathbf{u}^n), \bar{\theta}(\mathbf{u}^n))$ is an averaged rigid motion evaluated by extracting average CA displacements; $\bar{\delta}, \bar{\beta}$, and $\bar{\theta}$ denote averaged lead-lag, flap, and pitch angles, respectively. This averaged rigid motion helps keep the elastically-deflected blade centered within the blade volume mesh, and improves robustness of the mesh deformation process, when large flapping motions are present [38].

The elasticity matrix, $\mathbf{K}^n$, at time level n is determined by the rigidly moved input volume grid. Based on the invariance of the material properties to a rigid motion, $\mathbf{K}^n$ can be derived from $\overline{\mathbf{K}}$ via the following relation

$$\mathbf{K}^n\ \mathbf{W}_r^n = \mathbf{W}_r^n\overline{\mathbf{K}}, \tag{15}$$

where $\mathbf{W}_r^n$ is a square block-diagonal matrix of the elasticity-matrix dimension with the identical diagonal 3×3 blocks representing the rotation part of the rigid motion transform, $\mathbf{T}_r^n$ (see Eq. (14)). A compact form of the grid equations describing elastic grid deformations is the following

$$\mathbf{K}^n\left(\mathbf{X}^n - \mathbf{T}_r^n\hat{\mathbf{X}}\right) - \left(\mathbf{T}_s^n\ \hat{\mathbf{X}}_s - \mathbf{T}_r^n\hat{\mathbf{X}}_s\right) = 0. \tag{16}$$

In an extended form, the same equations are written as follows

$$\mathbf{K}^n \left(\mathbf{X}^n - \mathbf{W}_r^n \hat{\mathbf{X}}\right) - \left(\mathbf{W}_s^n \, \hat{\mathbf{X}}_s - \mathbf{W}_r^n \hat{\mathbf{X}}_s + \tilde{\mathbf{t}}_s^n\right) = 0. \tag{17}$$

Here vector $\tilde{\mathbf{t}}_s^n$ has the translation components of $\mathbf{T}_s^n$ in positions corresponding to the surface points and zeros in the positions corresponding to the interior points. Note that (1) the elasticity matrix multiplied by a constant translation vector is zero in the interior; and (2) the translation contribution due to the rigid motion $\mathbf{T}_r^n$ is the same for the volume and surface reference grids and is canceled on the boundary.

The initial CA surface displacements, $\mathbf{u}^0$, at time level $n = 0$ is assumed independent of design parameters. Specifically, a static (wind-off) CA solution or a loose coupling CFD/CA solution can be used as an initial CA solution. The CFD solution at time level $n = 0$ is specified from a given initial flow condition, $\overline{\mathbf{Q}}(\mathbf{D})$, for example a freestream flow or a loose coupling CFD/CA analysis solution. The initial flow condition is assumed independent of the shape design parameters, but may depend on CFD design parameters such as angle of attack, Mach number, etc.

Given inputs that are independent of design parameters, $\overline{\mathbf{X}}$, $\overline{\mathbf{K}}$, $\mathbf{u}^0$, and $\mathbf{T}_{\phi_0}$, and inputs that may depend on design parameters, $\overline{\mathbf{Q}}(\mathbf{D})$, $\hat{\mathbf{X}}_s(\mathbf{D})$, and $\mathbf{T}_{\psi_n}(\mathbf{D})$, the following equations for a tight coupling formulation are written in the order in which the equations are solved:

$$\text{Reference Grid} : \hat{\mathbf{G}}(\hat{\mathbf{X}}, \hat{\mathbf{X}}_s, \mathbf{D}) \equiv \overline{\mathbf{K}} \, \left(\hat{\mathbf{X}} - \overline{\mathbf{X}}\right) - \left(\hat{\mathbf{X}}_s(\mathbf{D}) - \overline{\mathbf{X}}_s\right) = \mathbf{0}, \tag{18}$$

$$\text{Initial Grid} : \mathbf{G}^0(\mathbf{X}^0, \hat{\mathbf{X}}, \mathbf{T}_r^0(\mathbf{u}^0), \mathbf{D}) \equiv \mathbf{X}^0 - \mathbf{T}_r^0 \, \hat{\mathbf{X}} = \mathbf{0}, \tag{19}$$

$$\text{Initial Flow} : \mathbf{R}^0(\mathbf{Q}^0, \overline{\mathbf{Q}}(\mathbf{D})) \equiv \mathbf{Q}^0 - \overline{\mathbf{Q}}(\mathbf{D}) = \mathbf{0}, \tag{20}$$

$$\text{Initial Load} : \mathbf{F}^0(\mathbf{f}^0, \, \mathbf{Q}^0, \mathbf{X}^0) = \mathbf{0}, \tag{21}$$

$$\text{CA} : \mathbf{C}^n(\mathbf{u}^n, \mathbf{u}^{n-1}, \mathbf{f}^{n-1}, \mathbf{D}) = \mathbf{0}, \quad (1 \le n \le N) \tag{22}$$

$$\text{CFD Grid} : \mathbf{G}^n(\mathbf{X}^n, \hat{\mathbf{X}}_s, \mathbf{W}_r^n(\mathbf{u}^n), \mathbf{W}_s^n(\mathbf{u}^n), \tilde{\mathbf{t}}_s^n, \mathbf{D}) \equiv \mathbf{K}^n \left(\mathbf{X}^n - \mathbf{W}_r^n \hat{\mathbf{X}}\right)$$

$$- \left(\mathbf{W}_s^n \, \hat{\mathbf{X}}_s(\mathbf{D}) - \mathbf{W}_r^n \hat{\mathbf{X}}_s + \tilde{\mathbf{t}}_s^n\right) = \mathbf{0}, \; (1 \le n \le N) \tag{23}$$

$$\text{CFD Flow} : \mathbf{R}^n(\mathbf{Q}^n, \mathbf{Q}^{n-1}, \mathbf{X}^n, \mathbf{X}^{n-1}, \mathbf{D}) = \mathbf{0}, \quad (1 \le n \le N) \tag{24}$$

$$\text{Load} : \mathbf{F}^n(\mathbf{f}^n, \mathbf{Q}^n, \mathbf{X}^n) = \mathbf{f}^n - \mathbf{S}^n(\mathbf{Q}^n, \mathbf{X}^n) = \mathbf{0}. \quad (1 \le n \le N) \tag{25}$$

The vector $\mathbf{S}$ in Eq. (25) denotes a boundary slicing that is performed in FUN3D for obtaining sectional airloads. For simplicity of presentation, the time-dependent equations are assumed to depend on solutions at

the current and one preceding time level only. In the actual flow solver, the time derivatives are discretized through second-order backward difference schemes that use solutions at the current and up to three preceding time levels. The residuals of the CFD flow equations are in the form of an Arbitrary Eulerian-Lagrangian formulation that is suitable for dynamic deformable grids and accounts for grid speeds and geometric conservation law. Details of the CFD equations on overset grids are provided elsewhere [15, 40].

The current implementation of the coupled CFD/CA analysis does not distinguish between the reference and input volume grids. Only one grid is used. Its surface grid serves as the reference surface grid, $\hat{\mathbf{X}}_s$, to compute elastic surface deformations at time level n, Eq. (12). The volume grid itself serves as the input volume grid to compute the elasticity matrix at time level n. The coupled CFD/CA analysis formulation describing the current implementation assumes that input and reference volume grids are always matching. With this assumption, no reference-grid solution is required; Eq. (18) is trivially satisfied. The sensitivity analysis formulation described in Section II E does not make this grid-matching assumption. It requires the reference grid to satisfy Eq. (18).

E. Mathematical Formulation for Sensitivity Analysis

In the sensitivity analysis framework presented here, the CFD flow and grid equations and the load equations use adjoint formulations, whereas sensitivities of the surface-grid elastic motion, $\mathbf{T}_s$, and the averaged rigid motion, $\mathbf{T}_r$, to aerodynamic loads, $\mathbf{f}^n$, are computed by the complex-variable method. For a general objective function $f(\mathbf{Q},\mathbf{u},\mathbf{X},\mathbf{f},\mathbf{D})$, the adjoint formulation is derived via Lagrangian, L, formed as

$$L = f(\mathbf{Q},\mathbf{u},\mathbf{X},\mathbf{f},\mathbf{D}) + \sum_{n=0}^{N} \left[\boldsymbol{\Lambda}_R^n\right]^T \mathbf{R}^n + \sum_{n=0}^{N} \left[\boldsymbol{\Lambda}_G^n\right]^T \mathbf{G}^n + \sum_{n=0}^{N} \left[\boldsymbol{\Lambda}_F^n\right]^T \mathbf{F}^n + \left[\boldsymbol{\Lambda}_{\hat{G}}\right]^T \hat{\mathbf{G}}, \tag{26}$$

where $\mathbf{Q}$, $\mathbf{u}$, $\mathbf{X}$, and $\mathbf{f}$ represent CFD solutions, CA surface displacements, grids, and airloads solutions at all time levels, respectively; $\boldsymbol{\Lambda}_R$, $\boldsymbol{\Lambda}_G$, and $\boldsymbol{\Lambda}_F$ are vectors of the time-dependent adjoint solutions for the flow, grid, and load equations, respectively; $\boldsymbol{\Lambda}_{\hat{G}}$ denotes the adjoint solution for the reference grid; the superscript T represents the transposition operator; N denotes the total number of time steps. Differentiating the Lagrangian with respect to design parameters $\mathbf{D}$ and equating the coefficients of state-variable sensitivities

to zero [41] yields the following adjoint equations:

for $n = N$,

$$\text{Load}: \left[\frac{\partial f}{\partial \mathbf{f}^n}\right] + \left[\mathbf{\Lambda}_F^N\right]^T \left[\frac{\partial \mathbf{F}^N}{\partial \mathbf{f}^n}\right] = \mathbf{0}, \quad (27)$$

$$\text{Flow}: \left[\frac{\partial f}{\partial \mathbf{Q}^N}\right] + \left[\mathbf{\Lambda}_R^N\right]^T \left[\frac{\partial \mathbf{R}^N}{\partial \mathbf{Q}^N}\right] + \left[\mathbf{\Lambda}_F^N\right]^T \left[\frac{\partial \mathbf{F}^N}{\partial \mathbf{Q}^N}\right] = \mathbf{0}, \quad (28)$$

$$\text{Grid}: \left[\frac{\partial f}{\partial \mathbf{X}^N}\right] + \left[\mathbf{\Lambda}_G^N\right]^T \left[\frac{\partial \mathbf{G}^N}{\partial \mathbf{X}^N}\right] + \left[\mathbf{\Lambda}_R^N\right]^T \left[\frac{\partial \mathbf{R}^N}{\partial \mathbf{X}^N}\right] + \left[\mathbf{\Lambda}_F^N\right]^T \left[\frac{\partial \mathbf{F}^N}{\partial \mathbf{X}^N}\right] = \mathbf{0}, \quad (29)$$

for $1 \le n < N$,

$$\text{Load}: \left[\frac{\partial f}{\partial \mathbf{f}^n}\right] + \left[\mathbf{\Lambda}_F^n\right]^T \left[\frac{\partial \mathbf{F}^n}{\partial \mathbf{f}^n}\right] + \sum_{k=n+1}^{N} \left[\mathbf{\Lambda}_G^k\right]^T \left(\left[\frac{\partial \mathbf{G}^k}{\partial \mathbf{T}_r^k}\right] \left[\frac{\partial \mathbf{T}_r^k}{\partial \mathbf{f}^n}\right] + \left[\frac{\partial \mathbf{G}^k}{\partial \mathbf{X}_s^k}\right] \left[\frac{\partial \mathbf{X}_s^k}{\partial \mathbf{f}^n}\right] \right) = \mathbf{0}, \quad (30)$$

$$\text{Flow}: \left[\frac{\partial f}{\partial \mathbf{Q}^n}\right] + \left[\mathbf{\Lambda}_R^n\right]^T \left[\frac{\partial \mathbf{R}^n}{\partial \mathbf{Q}^n}\right] + \left[\mathbf{\Lambda}_R^{n+1}\right]^T \left[\frac{\partial \mathbf{R}^{n+1}}{\partial \mathbf{Q}^n}\right] + \left[\mathbf{\Lambda}_F^n\right]^T \left[\frac{\partial \mathbf{F}^n}{\partial \mathbf{Q}^n}\right] = \mathbf{0}, \quad (31)$$

$$\text{Grid}: \left[\frac{\partial f}{\partial \mathbf{X}^n}\right] + \left[\mathbf{\Lambda}_G^n\right]^T \left[\frac{\partial \mathbf{G}^n}{\partial \mathbf{X}^n}\right] + \left[\mathbf{\Lambda}_R^n\right]^T \left[\frac{\partial \mathbf{R}^n}{\partial \mathbf{X}^n}\right] + \left[\mathbf{\Lambda}_R^{n+1}\right]^T \left[\frac{\partial \mathbf{R}^{n+1}}{\partial \mathbf{X}^n}\right] + \left[\mathbf{\Lambda}_F^n\right]^T \left[\frac{\partial \mathbf{F}^n}{\partial \mathbf{X}^n}\right] = \mathbf{0}, \quad (32)$$

and for $n = 0$,

$$\text{Load}: \left[\frac{\partial f}{\partial \mathbf{f}^0}\right] + \left[\mathbf{\Lambda}_F^0\right]^T \left[\frac{\partial \mathbf{F}^0}{\partial \mathbf{f}^0}\right] + \sum_{k=1}^{N} \left[\mathbf{\Lambda}_G^k\right]^T \left(\left[\frac{\partial \mathbf{G}^k}{\partial \mathbf{T}_r^k}\right] \left[\frac{\partial \mathbf{T}_r^k}{\partial \mathbf{f}^0}\right] + \left[\frac{\partial \mathbf{G}^k}{\partial \mathbf{X}_s^k}\right] \left[\frac{\partial \mathbf{X}_s^k}{\partial \mathbf{f}^0}\right] \right) = \mathbf{0}, \quad (33)$$

$$\text{Flow}: \left[\frac{\partial f}{\partial \mathbf{Q}^0}\right] + \left[\mathbf{\Lambda}_R^0\right]^T \left[\frac{\partial \mathbf{R}^0}{\partial \mathbf{Q}^0}\right] + \left[\mathbf{\Lambda}_R^1\right]^T \left[\frac{\partial \mathbf{R}^1}{\partial \mathbf{Q}^0}\right] + \left[\mathbf{\Lambda}_F^0\right]^T \left[\frac{\partial \mathbf{F}^0}{\partial \mathbf{Q}^0}\right] = \mathbf{0}, \quad (34)$$

$$\text{Grid}: \left[\frac{\partial f}{\partial \mathbf{X}^0}\right] + \left[\mathbf{\Lambda}_G^0\right]^T \left[\frac{\partial \mathbf{G}^0}{\partial \mathbf{X}^0}\right] + \left[\mathbf{\Lambda}_R^0\right]^T \left[\frac{\partial \mathbf{R}^0}{\partial \mathbf{X}^0}\right] + \left[\mathbf{\Lambda}_R^1\right]^T \left[\frac{\partial \mathbf{R}^1}{\partial \mathbf{X}^0}\right] + \left[\mathbf{\Lambda}_F^0\right]^T \left[\frac{\partial \mathbf{F}^0}{\partial \mathbf{X}^0}\right] = \mathbf{0}, \quad (35)$$

$$\text{Reference Grid}: \left[\mathbf{\Lambda}_{\hat{G}}\right]^T \left[\frac{\partial \hat{\mathbf{G}}}{\partial \hat{\mathbf{X}}}\right] + \sum_{n=0}^{N} \left[\mathbf{\Lambda}_G^n\right]^T \left[\frac{\partial \mathbf{G}^n}{\partial \hat{\mathbf{X}}}\right] = \mathbf{0}. \quad (36)$$

The adjoint equations are solved in the reverse order in time. In the current implementation, $[\partial \mathbf{F}^n / \partial \mathbf{f}^n]$ is the identity matrix. The sensitivities of rigid motions and surface grid to sectional airloads, denoted as $[\partial \mathbf{T}_r^k / \partial \mathbf{f}^n]$ and $[\partial \mathbf{X}_s^k / \partial \mathbf{f}^n]$ in Eq. (30), are computed with the complex-variable finite-difference approach by perturbing the imaginary part of individual components of the airloads $\mathbf{f}^n$ at each time level n. For each such perturbation, the complex-variable CA surface displacements, $\mathbf{u}^k$, the transform matrices for rigid motions, $\mathbf{T}_r^k$, and the surface mesh, $\mathbf{X}_s^k$, are computed for time levels k, $n+1 \le k \le N$. In contrast to the CA analysis operating on a single processor, the computations of sensitivities to airloads are distributed to all the processors used in the CFD analysis. The total computational time for evaluating these sensitivities is linearly proportional to the number of airloads, quadratically proportional to the total number of time steps N in the simulation, and inversely proportional to the number of processors used in parallel computations. In the current implementation, due to the quadratic dependence on the number of time steps, computing

the sensitivities of rigid motions and elastic surface motions to airloads is the most expensive part of the multidisciplinary sensitivity analysis. The multidisciplinary sensitivity analysis performance is discussed in Section III B 4.

The CA sensitivities contribute to the CFD flow and grid adjoint equations through the vectors $[\boldsymbol{\Lambda}_F^n]^T [\partial \mathbf{F}^n / \partial \mathbf{Q}^n]$ and $[\boldsymbol{\Lambda}_F^n]^T [\partial \mathbf{F}^n / \partial \mathbf{X}^n]$, respectively. Here, the $[\partial \mathbf{F}^n / \partial \mathbf{Q}^n]$ and $[\partial \mathbf{F}^n / \partial \mathbf{X}^n]$ matrices represent the sensitivities of the sectional airloads to the flow solution and CFD grid, respectively, and are implemented by hand differentiation. A test verifying correctness of the implementation is provided in Section II F. At time level $n = 0$, the matrices $[\partial \mathbf{F}^0 / \partial \mathbf{f}^0]$, $[\partial \mathbf{R}^0 / \partial \mathbf{Q}^0]$, and $[\partial \mathbf{G}^0 / \partial \mathbf{X}^0]$ are identity matrices; and the load, flow, and grid adjoint equations become simple evaluations of algebraic expressions. Finally, the formulation for sensitivity derivatives is given as

$$\begin{aligned} \frac{dL}{d\mathbf{D}} &= \left[\frac{\partial f}{\partial \mathbf{D}}\right] + \left[\boldsymbol{\Lambda}_R^0\right]^T \left[\frac{\partial \mathbf{R}^0}{\partial \overline{\mathbf{Q}}}\right] \left[\frac{\partial \overline{\mathbf{Q}}}{\partial \mathbf{D}}\right] + \sum_{n=1}^{N} \left[\boldsymbol{\Lambda}_R^n\right]^T \left[\frac{\partial \mathbf{R}^n}{\partial \mathbf{D}}\right] + \left[\boldsymbol{\Lambda}_{\hat{G}}\right]^T \left(\left[\frac{\partial \hat{\mathbf{G}}}{\partial \hat{\mathbf{X}}_s}\right] \left[\frac{d\hat{\mathbf{X}}_s}{d\mathbf{D}}\right] \right) \\ &+ \sum_{n=1}^{N} \left[\boldsymbol{\Lambda}_G^n\right]^T \left(\left[\frac{\partial \mathbf{G}^n}{\partial \hat{\mathbf{X}}_s}\right] \left[\frac{d\hat{\mathbf{X}}_s}{d\mathbf{D}}\right] + \left[\frac{\partial \mathbf{G}^n}{\partial \mathbf{T}_r^n}\right] \left[\frac{\partial \mathbf{T}_r^n}{\partial \mathbf{D}}\right] + \left[\frac{\partial \mathbf{G}^n}{\partial \mathbf{T}_s^n}\right] \left[\frac{\partial \mathbf{T}_s^n}{\partial \mathbf{D}}\right] \right). \end{aligned} \tag{37}$$

Specifically, the sensitivity with respect to flow design parameters, such as angle of attack or Mach number, is computed as

$$\frac{dL}{d\mathbf{D}} = \left[\frac{\partial f}{\partial \mathbf{D}}\right] + \left[\boldsymbol{\Lambda}_R^0\right]^T \left[\frac{\partial \mathbf{R}^0}{\partial \overline{\mathbf{Q}}}\right] \left[\frac{\partial \overline{\mathbf{Q}}}{\partial \mathbf{D}}\right] + \sum_{n=1}^{N} \left[\boldsymbol{\Lambda}_R^n\right]^T \left[\frac{\partial \mathbf{R}^n}{\partial \mathbf{D}}\right]; \tag{38}$$

the sensitivity with respect to shape design parameters, such as camber, thickness, or twist, is computed as

$$\frac{dL}{d\mathbf{D}} = \left[\frac{\partial f}{\partial \mathbf{D}}\right] + \sum_{n=1}^{N} \left[\boldsymbol{\Lambda}_G^n\right]^T \left[\frac{\partial \mathbf{G}^n}{\partial \hat{\mathbf{X}}_s}\right] \left[\frac{d\hat{\mathbf{X}}_s}{d\mathbf{D}}\right] + \left[\boldsymbol{\Lambda}_{\hat{G}}\right]^T \left[\frac{\partial \hat{\mathbf{G}}}{\partial \hat{\mathbf{X}}_s}\right] \left[\frac{d\hat{\mathbf{X}}_s}{d\mathbf{D}}\right]; \tag{39}$$

and the sensitivity with respect to kinematics design parameters, such as rotation rate, collective or cyclic control angles, is computed as

$$\frac{dL}{d\mathbf{D}} = \left[\frac{\partial f}{\partial \mathbf{D}}\right] + \sum_{n=1}^{N} \left[\boldsymbol{\Lambda}_G^n\right]^T \left(\left[\frac{\partial \mathbf{G}^n}{\partial \mathbf{T}_r^n}\right] \left[\frac{\partial \mathbf{T}_r^n}{\partial \mathbf{D}}\right] + \left[\frac{\partial \mathbf{G}^n}{\partial \mathbf{T}_s^n}\right] \left[\frac{\partial \mathbf{T}_s^n}{\partial \mathbf{D}}\right] \right). \tag{40}$$

Similar to the approach for computing sensitivities of rigid and elastic surface motions to airloads, the complex-variable method is also employed to evaluate the sensitivities of structural responses to kinematics design parameters, $[\partial \mathbf{T}_r^n / \partial \mathbf{D}]$ and $[\partial \mathbf{T}_s^n / \partial \mathbf{D}]$, shown in Eq. (40). The computational cost of the evaluation of sensitivities to kinematics parameters is linearly proportional to the number of time steps and considered small in comparison to the cost of sensitivities to airloads. The sensitivities of structural responses are computed by the complex-variable finite-difference method; all other derivatives required for the evaluation of

the adjoint solution and the final sensitivity are implemented by hand differentiation of the corresponding routines in the multidisciplinary CFD/CA system.

F. Sensitivity Analysis Verification

This section presents verification studies for implementations of the interfaces used in the CFD/CA sensitivity analysis and the entire multidisciplinary CFD/CA sensitivity analysis. The complex-valued sensitivities of the entire multidisciplinary CFD/CA system, representing the true sensitivities, are compared with the sensitivities computed either by hand-differentiation routines or by a multidisciplinary sensitivity analysis that involves adjoint-based sensitivities integrated with complex-variable finite-difference sensitivities. The verification computations are performed for the HART-II baseline rotorcraft configuration [8, 42] modeled with a composite grid containing a total of 463,340 nodes and 2,691,423 tetrahedral elements and including four moving component grids for blades and a stationary background grid. The time-dependent solutions of the complex-valued multidisciplinary analysis and the solutions of the adjoint equations are converged to machine zero residuals. All real- and complex-valued variables are declared and operated in double precision. The perturbation is selected to be $i10^{-50}$, which is sufficiently small to keep the real parts of solutions unaffected by the imaginary-part perturbations.

1. True Sensitivities of Coupled CFD/CA Analysis

True sensitivities for the coupled CFD/CA analysis can be obtained by operating both models in fully complex-variable modes and perturbing the imaginary part of a design parameter. One such computation can be used to obtain the sensitivities of multiple objective functions with respect to a single design parameter. To ensure that the complex-valued DYMORE model functions properly, real parts of the complex-valued DYMORE solutions describing the linear and angular displacements at the quarter-chord line of a rotor blade in the HART-II baseline rotorcraft test model are examined. The examination shows that the real parts of the DYMORE solutions are identical (matching all significant digits) to those obtained by the real-valued DYMORE model.

Table 1 lists aerodynamic lift at time level $n = 3$ computed by the real-valued tight coupling procedures and by the complex-valued tight coupling procedures with several perturbed design parameters including angle of attack and shape design parameters on different blade surfaces of the HART-II rotorcraft config-

Table 1 Real-valued and complex-valued lift (at time level 3) computed by coupled CFD/CA analysis.

Analysis	Perturbed parameters	Real part of solution	Sensitivity
Real-valued		0.000798918217192	
	Angle of attack	0.000798918217192	0.000024054292434
	Shape variable, blade 1	0.000798918217192	−0.000044924191748
Complex-valued	Shape variable, blade 2	0.000798918217192	0.000009594940189
	Shape variable, blade 3	0.000798918217192	−0.000000225139305
	Shape variable, blade 4	0.000798918217192	0.000005817185338

uration. In this test, the initial freestream flow condition and the BDF1 scheme are employed. As shown in the table, all digits of the real parts of the lift evaluated by the complex-valued analyses agree with the lift evaluated by the real-valued analysis. The imaginary parts of the lift evaluations by the complex-valued analyses correspond to the lift sensitivities with respect to the specified design parameters and can be used for verification of the multidisciplinary CFD/CA sensitivity analysis.

2. Verification of Multidisciplinary CFD/CA Sensitivity Computation

First, the implementation of hand-differentiation sensitivities of sectional airloads with respect to flow and grid solutions, $[\partial \mathbf{F}^n/\partial \mathbf{Q}^n]$ and $[\partial \mathbf{F}^n/\partial \mathbf{X}^n]$, is examined for accuracy. These sensitivities, as shown in Eqs. (31)–(32) for instance, contribute to the flow and grid adjoint equations. Due to the viscous force contributions, the sensitivities must take into account not only flow solutions and mesh coordinates of surface points, but also flow solutions and mesh coordinates of the interior points, which are neighbors of surface points. Table 2 compares representative aerodynamic airload derivatives for a given state of viscous flow and grid solutions. The derivatives of the three components of the aerodynamic load at the 70% radial station of the blade are computed with respect to the y-coordinate of a surface point located at this station and with respect to the flow solution representing the x-momentum (flow density multiplied by the x-directional velocity) evaluated at the point. The hand-differentiation derivatives are compared with the derivatives computed by the complex-variable approach. As observed, the hand-differentiated sensitivities agree well with the complex-variable sensitivities; at least 12-digit agreement is obtained, which is considered as a sufficiently good agreement [15].

Next, the multidisciplinary CFD/CA sensitivity analysis integrating adjoint-based sensitivities of

Table 2 Sensitivities of airloads at 70% radial station with respect to grid coordinates and flow solutions.

Method	$\partial F_x/\partial Q$	$\partial F_y/\partial Q$	$\partial F_z/\partial Q$
Hand differentiation	0.000041123560446	0.000126949885201	-0.006184933414864
Complex variable	0.000041123560446	0.000126949885201	-0.006184933414864

Method	$\partial F_{mx}/\partial Q$	$\partial F_{my}/\partial Q$	$\partial F_{mz}/\partial Q$
Hand differentiation	-0.000035570041307	-0.000005662910700	-0.000000350458662
Complex variable	-0.000035570041307	-0.000005662910700	-0.000000350458662

Method	$\partial F_x/\partial X$	$\partial F_y/\partial X$	$\partial F_z/\partial X$
Hand differentiation	-1.66200987502660	1.97008265283363	0.520855236517486
Complex variable	-1.66200987502690	1.97008265283398	0.520855236517574

Method	$\partial F_{mx}/\partial X$	$\partial F_{my}/\partial X$	$\partial F_{mz}/\partial X$
Hand differentiation	-0.015116700097615	-0.015298667011763	0.007316138436884
Complex variable	-0.015116700097618	-0.015298667011765	0.007316138436883

FUN3D with complex-variable sensitivities of DYMORE is verified. In this verification test, the initial grid and initial conditions for the CFD and CA solutions are assumed given and independent of design parameters. In particular, the term $[\partial \mathbf{G}^0/\partial \hat{\mathbf{X}}]$ in Eq. (36) and the term $[\partial \overline{\mathbf{Q}}/\partial \mathbf{D}]$ in Eq. (38) vanish. Specifically, a grid corresponding to the baseline geometry is used as the initial grid, and a loose coupling CFD/CA solution is used as the initial condition for the CFD and CA solvers to start a tight coupling procedure. Although only flow adjoint solutions are required to evaluate the lift sensitivity to a flow design parameter, load and grid adjoint solutions must also be computed due to the coupling between the flow, grid, and structure disciplines. The verification test is performed in the following three steps.

1. The real-valued CFD/CA tight coupling analysis equations, Eqs. (19)–(25), are solved to machine-zero residuals. The grid for the baseline configuration obtained from a grid generation tool serves as the input and reference grid. The time-dependent CFD flow, CFD grid, and CA solutions, including the reference grid, $\hat{\mathbf{X}}$, are saved to disk.

2. The adjoint equations, Eqs. (27)–(36), use the grids and solutions saved in disk and are solved to

machine-zero residuals. Recall that the adjoint formulation relies on Eq. (15), which has been derived under the assumption that the input grid and the elasticity matrices do not depend on shape design parameters.

3. The design parameters are perturbed, one at a time, by the complex-valued perturbation. The complex-valued solutions of the CFD/CA tight coupling analysis equations, Eqs. (19)–(25), are converged to machine-zero residuals. If the perturbed design parameter is a shape parameter, and Eq. (18) has been iterated, the imaginary perturbation propagates to the reference grid and introduces an imaginary component to the elasticity matrix. This violates the sensitivity analysis assumption that the elasticity matrices do not depend on shape design parameters. To avoid this violation, the imaginary parts of the elasticity matrix computed by the input (or reference) grid are set to be zero.

Table 3 lists lift sensitivities evaluated by the multidisciplinary CFD/CA sensitivity analysis and the complex-valued CFD/CA approach for the BDF1, BDF2, and BDF2opt schemes. The design parameters consist of the angle-of-attack, several representative shape design parameters placed on different blades, and kinematics design parameters including the collective and cyclic control angles. Sensitivities computed

Table 3 Sensitivities of lift to design parameters evaluated by multidisciplinary CFD/CA sensitivity analysis (M) and complex-valued approach (C). Loose coupling solution is given as initial condition.

Design parameters	BDF1	BDF2	BDF2opt
Angle of attack	M: −0.000001105132519	M: −0.000001100711231	M: −0.000001092344481
	C: −0.000001105132519	C: −0.000001100711231	C: −0.000001092344481
Shape, Blade 1	M: −0.000002557422**216**	M: −0.000005642808**476**	M: −0.000017109969**9901**
	C: −0.000002557422**563**	C: −0.000005642808**946**	C: −0.000017109977**7709**
Shape, Blade 3	M: 0.000009998583796	M: 0.000009722985571	M: −0.000003815029**580**
	C: 0.000009998583811	C: 0.000009722985594	C: −0.000003815029**368**
Trim, Collective	M: 0.000108820672956	M: 0.000148871328612	M: 0.000449424709552
	C: 0.000108820680183	C: 0.000148871337842	C: 0.000449424871529
Trim, Cyclic	M: −0.000000765860088	M: −0.000000562605238	M: 0.000001298910374
	C: −0.000000765860584	C: −0.000000562606052	C: 0.000001298921008

with the BDF1 and BDF2 schemes correspond to the lift coefficient evaluated at three time steps; five time steps are used for sensitivities of the BDF2opt scheme. The multidisciplinary and complex-valued CFD/CA sensitivities are in a good agreement for all temporal schemes; all digits agree for the sensitivities to the angle of attack; 11–13 digits after the decimal point agree for the sensitivities to the shape design parameters; and 9–12 digits agree for the sensitivities to the trim design parameters. The larger-than-usual discrepancies between the multidisciplinary CFD/CA sensitivities and complex-variable sensitivities for trim parameters are hypothesized as a result of poor scaling of the elasticity matrix. In particular, because of grid anisotropy, the diagonal coefficients of the elasticity matrix corresponding to different control volumes vary by many orders of magnitude. This matrix appears in terms $[\partial \mathbf{G}^n / \partial \mathbf{X}^n]$ in Eqs. (29) and (32) and affects accuracy of the grid adjoint solution.

G. Geometry Parameterization

In the current study, a geometry parameterization tool, MASSOUD (Multidisciplinary Aerodynamic-Structural Shape Optimization Using Deformation) software package [43], is employed for surface mesh deformation. The parameterization tool uses a set of aircraft-centric shape design parameters such as planform, twist, shear, thickness, and camber parameters to realize general configuration shape control. Figure 2 shows an example of blade planform geometry of a HART-II rotor and locations of twist, thickness, and camber shape design parameters. Analytical sensitivities of the CFD surface grid to various types of shape design parameters (c.f., $[d\hat{\mathbf{X}}_s / d\mathbf{D}]$ in Eq. (39)) are provided for sensitivity derivative computations.

H. Design Optimization Cycle

A sequential quadratic programming optimizer, SNOPT [44, 45], is used to drive the gradient-based rotorcraft design optimization. The design optimization cycle is illustrated in Fig. 3 and summarized as the following steps.

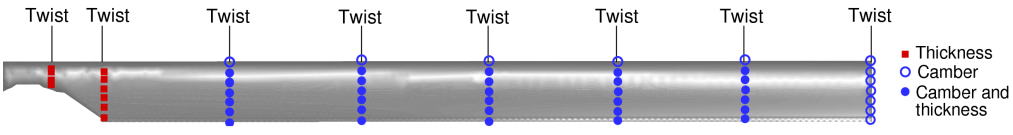

Fig. 2 Blade planform geometry and shape parameter locations.

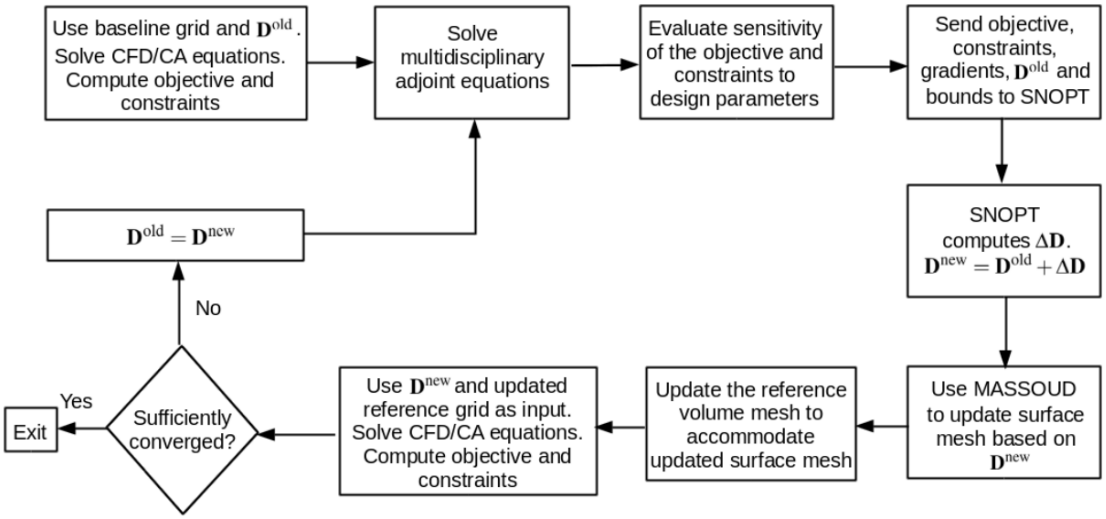

Fig. 3 Design optimization cycle.

1. *Startup phase.* Ensure that the input surface grid matches the surface grid of the baseline configuration, $\hat{\mathbf{X}}_s = \overline{\mathbf{X}}_s$. Solve the tight coupling CFD/CA analysis equations, Eqs. (19)–(25), for the baseline configuration corresponding to the initial design parameters, $\mathbf{D}^{\text{old}}$. Save to disk the time-dependent CFD flow, CFD grid, and CA solutions, including the reference grid, $\hat{\mathbf{X}}$. Evaluate the values of the objective function and constraints corresponding to $\mathbf{D}^{\text{old}}$.

2. Solve the adjoint equations, Eqs. (27)–(36), by marching in the reverse time. Use the latest available design parameters, $\mathbf{D}^{\text{old}}$, and the corresponding CFD flow, CFD grid, and CA solutions available on disk. Use the rigidly moved reference grid from disk to compute the elasticity matrix at time level n. This ensures that $\hat{\mathbf{X}} = \overline{\mathbf{X}}$ and provides consistency with the multidisciplinary analysis solutions computed in the previous optimization cycle.

3. Evaluate the sensitivities of the objective and constraints to design parameters by using Eqs. (38)–(40).

4. Provide the optimizer with the current values of objective and constraint functions, the current design parameters, the sensitivities of the objective and constraints to the design parameters, as well as lower and upper bounds of the design parameters.

5. SNOPT provides corrections for the design parameters, $\Delta\mathbf{D}$. SNOPT may iterate with the solver to compute acceptable $\Delta\mathbf{D}$. Apply those corrections to compute new design parameters as $\mathbf{D}^{\text{new}} = \mathbf{D}^{\text{old}} + \Delta\mathbf{D}$.

6. Use MASSOUD to update the surface grid $\hat{\mathbf{X}}_s$ based on the new (shape) design parameters $\mathbf{D}^{\text{new}}$.

7. Update the reference volume grid, $\hat{\mathbf{X}}$, by using Eq. (18).

8. Use the updated reference grid as the input grid, $\overline{\mathbf{X}} = \hat{\mathbf{X}}$. Solve the tight coupling CFD/CA analysis equations, Eqs. (19)–(25), using the new design parameters, $\mathbf{D}^{\text{new}}$. Save to disk the time-dependent CFD flow, CFD grid, and CA solutions including the reference grid. Evaluate the values of the objective function and constraints.

9. Set $\mathbf{D}^{\text{old}} = \mathbf{D}^{\text{new}}$.

10. Repeat the steps of 2–9 until the objective function is sufficiently minimized and constraints are converged to a satisfactory level.

III. Numerical Results

In this section, the current implementations of the loose and tight coupling analyses are validated for a benchmark HART-II model corresponding to a descending flight. A constrained gradient-based design optimization is demonstrated for the HART-II rotorcraft configuration.

A. HART-II Rotorcraft Analysis

1. CFD/CA Loose Coupling Simulations

The flow conditions in the HART-II baseline case include a tip Mach number of 0.6387, an advance ratio of 0.151, and a shaft tilt angle of 4.5° (adjusted after wind tunnel wall correction). This case represents a descending flight, where blade-vortex interactions are significant. The trim targets for the rotor thrust and the rolling and pitching moments in this test are 3300 N, 20 N·m, and -20 N·m, respectively. Other important parameters of the blade geometry and test conditions are detailed elsewhere [8]. Simulations for the HART-II baseline case conducted with the current system are compared to FUN3D/CAMRAD-II (NL-2) simulations [8] and simulations using a previous implementation of FUN3D/DYMORE coupling (GIT-1) [8] as well as experimental data. Note that while the structural model parameters used in the CAMRAD and DYMORE solvers are quite similar, small differences in common parameters were observed. In addition, each CA solver has specific input parameters, which can lead to slightly different structural dynamics modeling. The

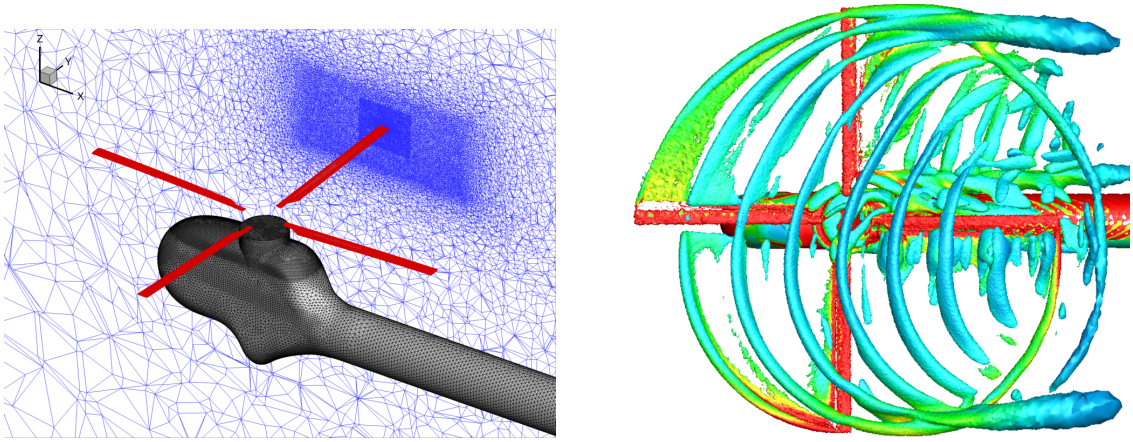

(a) Computational grid

(b) Isosurfaces of Q-criterion

Fig. 4 Computational overset grid and vortex visualization via isosurfaces of Q-criterion ($Q = 0.0075$) simulated by coupled CFD/CA analysis.

unstructured mixed-element overset computational grid, shown in Fig. 4(a), consists of four blade component grids, one fuselage component grid, and a stationary background grid, resulting in a total of 13,615,983 nodes, and 66,660,091 unstructured mixed elements consisting of 60,093,307 tetrahedrons, 44,647 pyramids, and 6,522,137 prisms. In the present study, the BDF2opt temporal scheme is used to advance the flow computation in time with a time step corresponding to the increment of 1° azimuth angle. In ten coupling cycles that involve 5.5 revolutions, the variations in the control angles and thrust deltas are within 0.01° and 0.1%, respectively. Figure 4(b) illustrates vortex structures near the HART-II rotor, where the blade-vortex interaction phenomena and the formation of blade tip vortices are clearly observed.

Figure 5 demonstrates the normal force and pitching moment at the 87% blade radial station evaluated from the converged loose coupling solutions, and the measurement data. In Fig. 5(a), the magnitudes and shapes of the normal force predicted by the current FUN3D/DYMORE solution are in good agreement with those of the NL-2 solutions. The maximum difference between the normal forces computed with the current and NL-2 systems is observed close to 116° and does not exceed 0.005. Furthermore, the blade-vortex interaction at the advancing side and particularly the retreating side are captured well by the current FUN3D/DYMORE analysis. The underprediction between 5° and 40° and the overprediction between 230° and 270° appearing in the GIT-1 solution are either greatly reduced or absent in the current simulation. The differences between the present implementation and GIT-1 solutions observed here are pos-

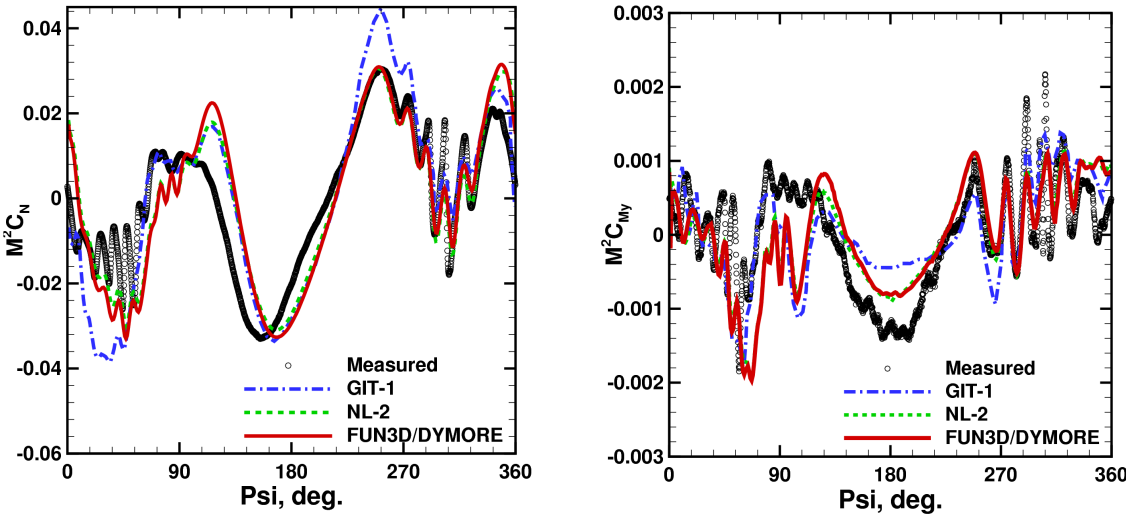

(a) Normal force

(b) Pitching moment

Fig. 5 Comparisons of normal force and pitching moment predictions at 87% radial station with means removed.

sibly due to the fact that the present implementation of FUN3D/DYMORE coupling follows closely with the FUN3D/CAMRAD-II development to date. The GIT-1 solution was obtained with an earlier version of FUN3D, where differences exist in the global slicing procedures and the conversions of elastic motions to the blade surface geometries. A similar comparison for the aerodynamic pitching moment is shown in Fig. 5(b), where all the solutions of the current FUN3D/DYMORE, NL-2 and GIT-1 methods are generally in good agreement, particularly in the retreating side. The maximum difference between the pitching moments computed with the current FUN3D/DYMORE and NL-2 systems is observed close to 126° and does not exceed 0.0003.

Figure 6 compares the azimuthal variation of elastic motions including flap, lead-lag and torsion deformations at the blade tip. The flap deflections are obtained by removing the precone angle (2.5°) from the vertical displacements, while the elastic torsion is obtained by excluding the pretwist angle and the pitch control inputs from the total geometric pitch angle. As demonstrated in Fig. 6(a), the overall characteristics of the predicted blade tip flap deflections in the current FUN3D/DYMORE solutions agree well with those of the NL-2 solution, both in magnitude and phase. Compared to the GIT-1 solution, some differences can be observed, especially in the first and third quadrants of the rotor disk. These discrepancies are likely caused by the differences in the airload predictions as observed in Fig. 5. All the computational results exhibit a

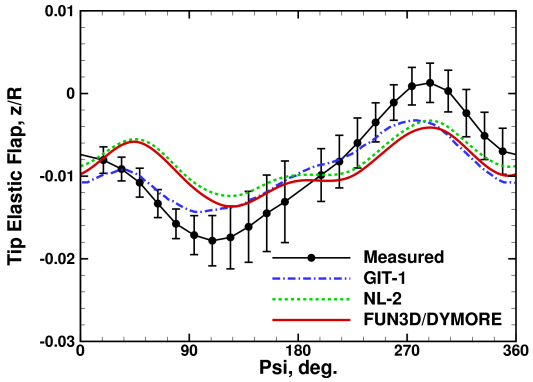

(a) Flap deflections relative to precone (2.5°)

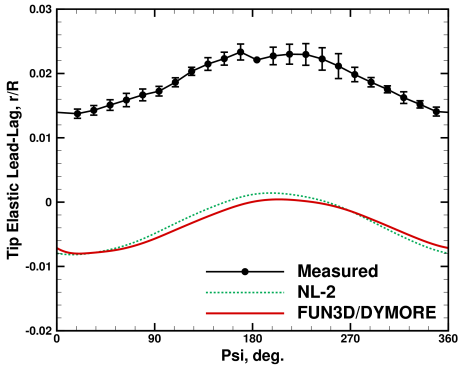

(b) Lead-lag deflections

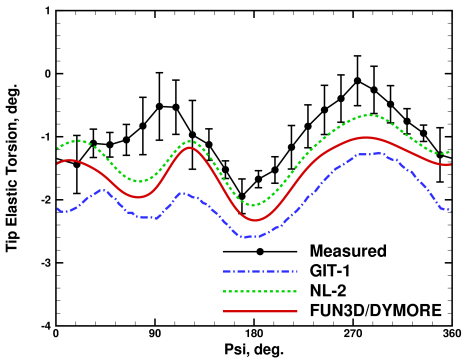

(c) Torsion deflections relative to pretwist and control angles

Fig. 6 Blade tip flap, lead-lag and torsion deflections.

secondary elastic flap maximum at approximately 40°–50° azimuth. In this azimuth range, the magnitude of the current solution is higher than that of the GIT-1 solution, but quite similar to that of the NL-2 solution. Blade tip lead-lag deflections are evaluated in Fig. 6(b), where a consistent agreement between the current FUN3D/DYMORE and NL-2 solutions is again observed. The maximum difference between the normalized lead-lag deflections computed by the current and NL-2 systems occurs between 120° and 190° and is less than 0.0012. Compared to the measurement data, a constant offset roughly amounting to 1/5 of the blade chord in the elastic lead-lag deflections is also observed; this translational offset behavior is commonly seen in other work [46, 47]. Fig. 6(c) compares the elastic torsion at the blade tip, where the 2/rev behavior of the elastic twist is predicted well in the current FUN3D/DYMORE solution. In addition, the overall torsion levels are more accurately captured by the current FUN3D/DYMORE solution as compared to the GIT-1 solution. The current FUN3D/DYMORE solution shows approximately 0.5° underpredictions across

all azimuth locations, while the GIT-1 solution underpredicts by approximately 1°.

2. CFD/CA Tight Coupling Simulations

The tight coupling CFD/CA algorithm is verified in this section by comparing tight and loose coupling simulations conducted with the same trim control angles. Trim computations are often performed with a loose coupling CFD/CA approach. To achieve converged trim solutions, the CA model simulates many revolutions per each coupling iteration, which becomes expensive for a tight coupling procedure. Here, the trim control angles are taken from a converged loose coupling solution. The tight coupling simulations then start from the loose coupling solution, and the collective and cyclic pitch control angles prescribed by the loose coupling trimmed solutions are kept throughout the tight coupling computation. Such a tight coupling procedure should result in minimal changes in the solutions.

Figure 7 compares the loose and tight coupling solutions for predictions of the normal force and pitching moment at 87% of the radial station, and the tip flap deflections in the HART-II baseline rotorcraft case. The solutions correspond to the fourth quadrant of the second rotor revolution. The tight and loose coupling solutions are almost identical. The tight coupling FUN3D/DYMORE solver appears to function correctly, and the rotor maintains the trimmed condition.

B. Design Optimization of HART-II Rotorcraft

In this section, the multidisciplinary sensitivity analysis is applied to gradient-based design optimization of the baseline HART-II rotorcraft configuration. In this demonstration, the computational grid is composed of four identical blade component grids, a single fuselage component grid, and an outer background grid. Each of the blade component grids contains 900,723 nodes and 2,882,711 mixed elements of tetrahedrons, pyramids and prisms, while the fuselage component grid contains 3,455,726 nodes and 17,500,042 mixed elements. The resulting composite overset mesh consists of a total of 7,058,618 nodes and 29,030,886 elements consisting of 22,799,206 tetrahedrons, 44,874 pyramids, and 6,186,806 prisms. The BDF2opt time integration scheme is used in FUN3D with 20 nonlinear subiterations and a time step corresponding to 1° azimuth angle of rotor rotation. A tolerance of 10^{-13} is set for the mesh elasticity solver. To avoid nonphysical deflections of the blades in the initial transient stage, FUN3D and DYMORE are initialized with a given loose coupling solution. DYMORE initial conditions are based on the blade motions and trim control

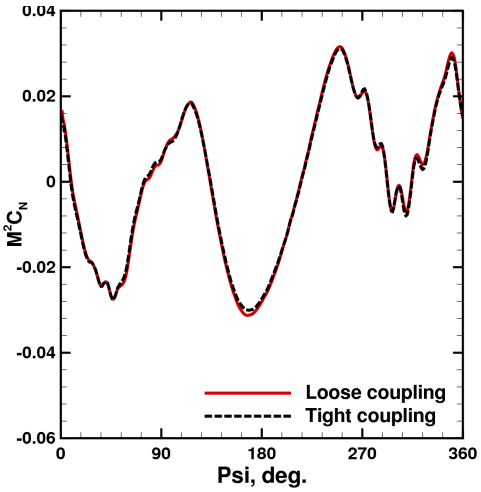

(a) Normal force

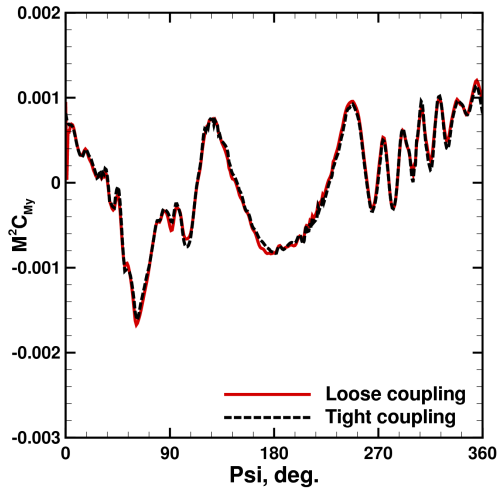

(b) Pitching moment

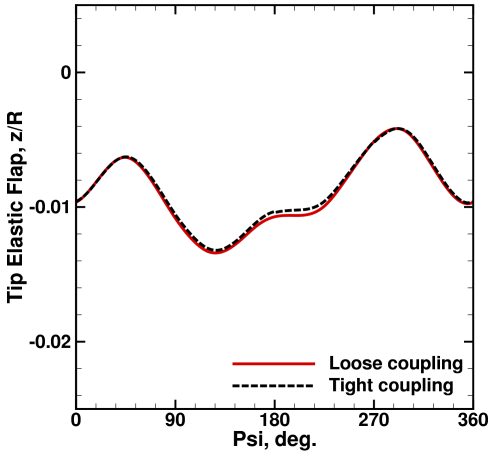

(c) Blade tip deflections

Fig. 7 Normal force and pitching moment (with means removed) at 87% radial station and blade tip flap deflections computed by FUN3D/DYMORE system with tight and loose coupling approaches on 7M node unstructured mesh.

angles obtained after 11 loose coupling iterations. The initial flow condition is the flow obtained after three-revolution simulations with the blade motions and trim angles from the DYMORE initial conditions. Note that these initial conditions correspond to a loose coupling solution that is not fully converged. The same initial conditions are used in every design cycle during the optimization process.

1. Objective and Constraints

The objective for the current design optimization case is to reduce the time-averaged torque coefficient, while explicitly constraining the thrust and the rolling and pitching moments to the trimmed values [8] of the baseline rotor performance. Specifically, the objective function to be minimized is based on a weighted difference between a target torque coefficient, C_Q^*, and a time-averaged objective value over a time period, expressed as

$$f_{obj} = w_{obj}\left[\left(\frac{1}{N_2 - N_1 + 1}\sum_{n=N_1}^{N_2} C_Q^n\right) - C_Q^*\right]^2 \Delta t. \tag{41}$$

For the purpose of minimization of rotor torque, C_Q^* is set to be zero. A weighting parameter, w_{obj}, is used to scale the objective function to be one for the baseline case. Note that due to the relation between the rotor shaft power and torque, $P = \Omega Q$ (Ω denotes the angular speed of rotor), the current objective can also be used for minimization of the rotor power. Parameters N_1 and N_2 represent the optimization interval and are set to $N_1 = 271$ and $N_2 = 360$, respectively. Therefore, the optimization window in the current design demonstration corresponds to the fourth quadrant of the first rotor revolution. Although the optimization window can be increased to include more rotor revolutions, the current selection is made based on the following considerations. (1) Comparing to the CFD/CA tight coupling simulations with freestream initial conditions, the initial transients in the simulations that start from a loose coupling solution and trimmed control angles are considerably smaller. (2) This short optimization interval is sufficient for demonstrating viability of the multidisciplinary CFD/CA sensitivity analysis for rotorcraft design optimization. (3) The cost of the present CFD/CA sensitivity analysis grows quadratically with the increase in the number of time steps. Therefore, limiting the number of time steps alleviates the computational cost of each design cycle.

The explicit constraints on the rotor thrust and the two moments are also defined based on time-averaged functions over the same optimization interval:

$$g_1 = w_1\left[\left(\frac{1}{N_2 - N_1 + 1}\sum_{n=N_1}^{N_2} C_T^n\right) - C_T^*\right]^2 \Delta t, \tag{42}$$

$$g_2 = w_2\left[\left(\frac{1}{N_2 - N_1 + 1}\sum_{n=N_1}^{N_2} C_{Mx}^n\right) - C_{Mx}^*\right]^2 \Delta t, \tag{43}$$

$$g_3 = w_3\left[\left(\frac{1}{N_2 - N_1 + 1}\sum_{n=N_1}^{N_2} C_{My}^n\right) - C_{My}^*\right]^2 \Delta t. \tag{44}$$

Here, C_T^*, C_{Mx}^*, and C_{My}^* are target values of the thrust coefficient and the rolling- and pitching-moment coefficients, respectively, and are set to be $C_T^* = 0.00457$, $C_{Mx}^* = 0.0000139$, and $C_{My}^* = -0.0000139$ based on the dimensional quantities [8] of the baseline configuration. Similar to the definition of the objective function, weighting parameters, w_1, w_2, and w_3, are used to scale the constraints to be one for the baseline case. Note that proper scaling of objective functions and constraints is important as it allows the optimizer to explore the design space more effectively. Note also that material properties, such as bending stiffness, mass, moment of inertia, etc., have been kept invariant throughout the design optimization process, although a substantial change in blade shapes would require a re-evaluation of blade material properties.

2. Design Parameters

The design parameters in this case, as displayed in Fig. 2, consist of 92 shape design parameters on the rotor blades including 8 twist parameters, 42 thickness parameters, and 42 camber parameters. In this demonstration, shape design parameters are not used to modify the fuselage geometry. Specifically, thickness of the blade leading edge is not allowed to be modified and camber variables are activated within 65% of the blade span while being deactivated in the blade root areas. In addition, three control angles describing the collective control input, θ_0, the lateral and longitudinal cyclic control inputs, θ_{1c} and θ_{1s}, respectively, are also used as design parameters and are allowed to vary as much as ± 5 degrees. The collective and cyclic control parameters are used to define the imposed time-dependent pitching motions of the blades, governed by the relation

$$\theta = \theta_0 + \theta_{1c}\cos\psi + \theta_{1s}\sin\psi, \tag{45}$$

where ψ denotes the blade azimuth position. The cumulative pitching motion is composed of the imposed pitching motion, pretwist of the blades, and structural responses due to fluid-structure interactions. Recall that the initial tight coupling solution corresponds to a loose coupling trimmed solution with the same computational grid system. Therefore, the present optimization procedure aims to maintain the trimmed conditions throughout the design optimization process.

3. Optimization Results

The optimization demonstration has been performed on the NASA Advanced Supercomputing (NAS) facility at NASA Ames Research Center; the execution has been conducted in parallel on 2089 processing cores for the CFD/CA analysis and on 2088 cores for the sensitivity analysis. Approximately 251 gigabytes of disk space have been used to store the complete time-dependent flowfield solution and the associated over-set domain-connectivity data. This optimization demonstration took about 8 days for completing 9 design cycles corresponding to 9 tight coupling CFD/CA solutions and 36 multidisciplinary adjoint solutions. In the NAS computation environment, execution of the coupled CFD/CA and multidisciplinary sensitivity analyses corresponding to one design cycle requires 2.5 and 19.4 hours of wall-clock time, respectively. The time required for computing the multidisciplinary sensitivities exceeds the time required for the CFD/CA analysis by a factor of 7.8. Recall that the primal analysis and the corresponding adjoint based sensitivity analysis are expected to have approximately the same computational cost. The observed large disparity in computational costs is caused by the high cost of sensitivities of structural responses with respect to airloads as well as the need for four sets of adjoint solutions corresponding to the objective and three constraints. More detailed discussions of the performance and a breakdown of the computational cost are given in Section III B 4.

Figures 8(a) and 8(b) show the convergence of torque, thrust, and moment coefficients. As shown in Fig. 8(a), a relatively rapid torque decrease is observed during the first three design cycles, while a total of 4.6% reduction in the rotor torque is achieved after nine design cycles. As displayed in Fig. 8(b), the rotor thrust and the rolling- and pitching-moment coefficients have adequately converged to the target values. Table 4 lists the absolute and relative differences between the thrust and the rolling- and pitching-moment coefficients for the optimized configuration and the corresponding targets. The rolling-moment coefficient is within the 2% margin of the target value; all other constraints are well within the 1% margin. The initial values of the constraints are farther from the targets. These initial deviations are the transient effects of initiating the tight coupling analysis with a loose coupling solution; but the absolute differences especially for thrust are considered small.

Figure 8(c) depicts the convergence of the collective and cyclic pitch control angles. Variations of the control angles from the initial values corresponding to a loose coupling solution are less than 1° during the entire optimization process. Table 5 summarizes the converged control angles for the optimized HART-II

Table 4 Deviation of the thrust and the rolling and pitching moments of the optimized configuration from targets.

Constraint functions	Target	Absolute diff.	Relative diff.
Thrust coefficient C_T	4.57×10^{-3}	1.1×10^{-6}	0.02%
Rolling-moment coefficient C_{Mx}	1.39×10^{-5}	2.7×10^{-7}	2.0%
Pitching-moment coefficient C_{My}	-1.39×10^{-5}	9.3×10^{-8}	0.67%

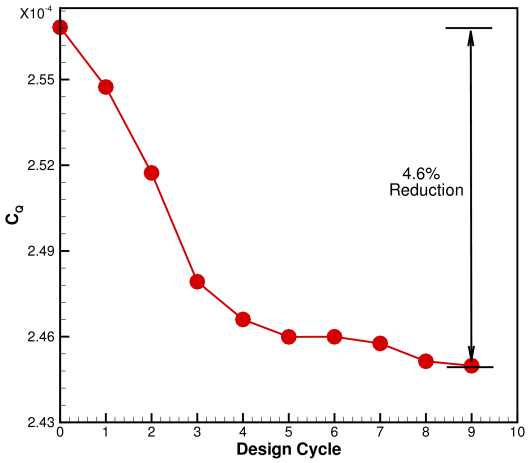

(a) Convergence of torque

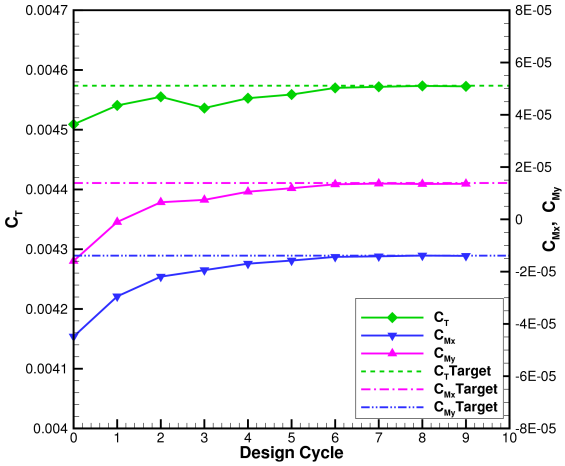

(b) Convergence of thrust and rolling- and pitch-moment coefficients

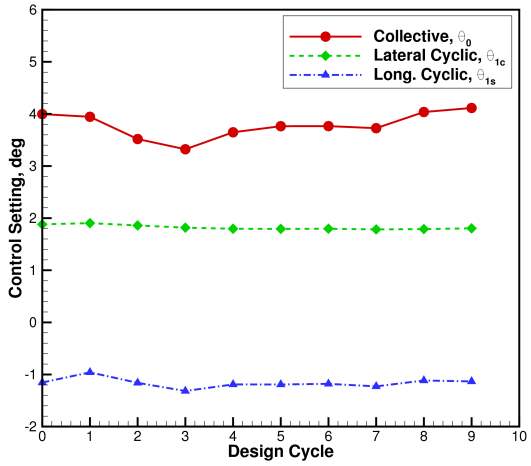

(c) Convergence of control angles

Fig. 8 Convergence of torque, thrust, moments, and control angles.

rotorcraft configuration as compared to the experimental trim results and the baseline case; the latter corresponds to the loose coupling solution. Specifically, the collective control angle receives a 0.12° increase, thereby inclining to create a small production of lift, while the lateral and longitudinal cyclic control param-

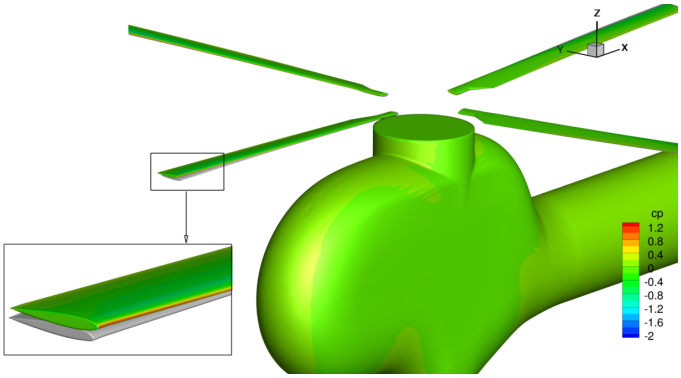

Fig. 9 Blade deflections for baseline (shaded) and optimized (colored with surface pressure coefficient levels) configurations.

eters receive a decrease of 0.08° and an increase of 0.02°, respectively, to generate small additional pitching and rolling motions of the tip-path plane. Overall, the computational controls are close to the experimental trim results.

Figure 9 displays a snapshot of the elastic blades after one revolution for the baseline and optimized configurations. The blades corresponding to the optimized geometries are shown on the top and colored with the surface pressure coefficient levels, while the baseline blades are shaded and shown at the bottom for comparison. The observed changes in blade deflections are small. Note that the constraints of the optimization problem penalize changes in the trim parameters.

Airfoil cross sections of the initial and optimized blade geometries are shown in Fig. 10 for spanwise stations ranging from 25% to 98% of the blade radius. Note that the observed variations in the airfoil cross sections are combined results of changes in many shape design parameters. Compared to the baseline geometry, the optimized shape has more negative twist across most of the blade. Reduced thickness is observed in outboard and aft regions. Negative camber is observed in aft regions corresponding to the middle span. The optimized shape shown in Fig. 10 is in an exaggerated vertical scale; the actual changes in the blade

Table 5 Collective and cyclic control angles for the baseline and optimized configurations. Experimental trim control angles are provided for reference.

Control angles	Experiment	Baseline	Optimized
Collective θ_0	3.80°	3.99°	4.11°
Lateral cyclic θ_{1c}	1.92°	1.88°	1.80°
Longitudinal cyclic θ_{1s}	−1.34°	−1.15°	−1.13°

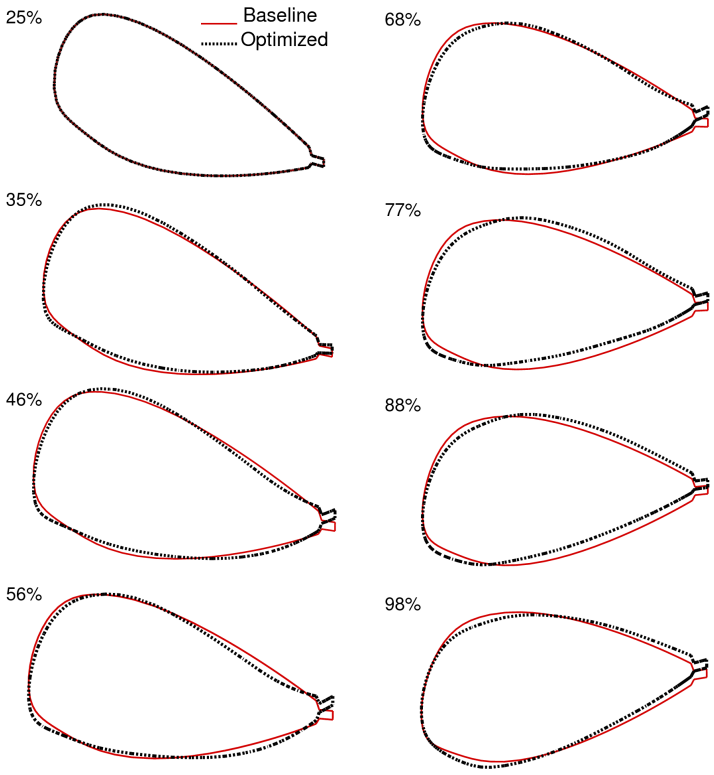

Fig. 10 Baseline and optimized blade section geometries in exaggerated vertical scales ($y:z = 0.23:1$)**.**

shape are not large, and therefore, the invariant material property approximation is valid. Figure 11 shows time histories of the torque, lift, and rolling and pitching moments for the baseline and optimized configurations. The simulations extended beyond the optimization time interval exhibit the same mean values of lift and moments, and trimmed conditions are preserved. The observed mean torque value is slightly varying in both optimized and baseline simulations. This behavior is caused by the initial conditions that correspond to a not-fully-converged loose coupling solution. The reduction of the mean torque value observed for the optimized configuration is maintained and appears more significant in the extended simulations.

4. Performance Analysis

For the design optimization demonstration, a breakdown of the time usage within a design cycle is provided in Table 6. The wall-clock times spent to solve the tightly coupled CFD/CA analysis equations, Eqs. (18)–(25); the load adjoint equations, Eqs. (27), (30), and (33); the flow adjoint equations, Eqs. (28), (31), and (34); and the grid adjoint equations, Eqs. (29), (32), and (35), are shown separately in the table as well as the percentage of time spent on solutions of the particular set of equations. The execution time for

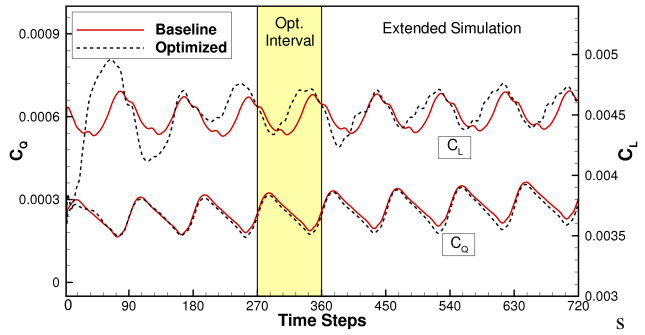

(a) Torque and lift coefficients

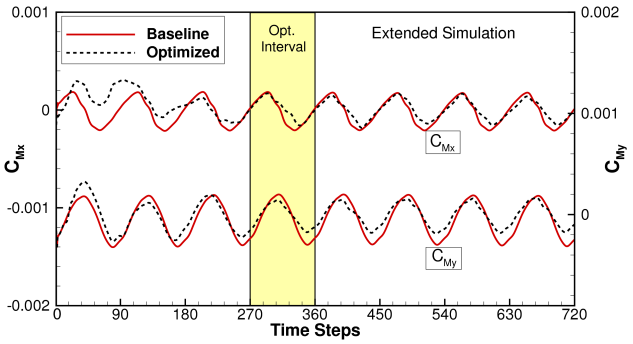

(b) Rolling- and pitching-moment coefficients

Fig. 11 Time histories of torque, lift, and rolling and pitching moments.

solving the grid adjoint equations exceeds the time for solving the flow adjoint equations by approximately 32%. The increased execution time is caused by two factors. (1) The grid adjoint solver uses a tighter convergence tolerance than the flow adjoint solver. (2) The flow adjoint solver takes advantage of computing simultaneously adjoint solutions for all objectives and constraints, while the grid adjoint solver repeats the solution process for every objective/constraint function. The total time required for computing the multidisciplinary sensitivities exceeds the time required for the CFD/CA analysis by a factor of 7.8. As mentioned earlier, one reason for the high sensitivity-analysis cost is the need to compute four adjoint solutions corresponding to one objective and three constraint functions. However, the dominant cost is the cost of evaluating sensitivities of DYMORE solutions to the aerodynamic loads. This cost is reflected in the large time spent on the load-adjoint solution because DYMORE sensitivities contribute to the right-hand side of Eq. (30).

Evaluation of the DYMORE sensitivities at a single time step is not expensive. Specifically, 87 stations are prescribed on each blade; total is 348 stations on four blades. Each station has a specified airload defined by six parameters. Thus, 2088 complex-variable DYMORE solutions are required to compute sensitivities

Table 6 Breakdown cost per design cycle (2088 processors, 21.9 hours).

	CFD/CA analysis	Load adjoint	Flow adjoint	Grid adjoint
Wall-clock time (hours)	2.5	13.7	2.4	3.3
Percent (%)	11.4	62.5	11.0	15.1

to all airloads at a given time. All these computations take only 0.763 seconds. However, the number of DYMORE solutions required to evaluate the sensitivities contributing to Eq. (30) grows quadratically with the number of time steps; nearly 135 million DYMORE solutions are required for solving the time-dependent load-adjoint equations for the simulation involving 360 time steps. The total cost for the entire DYMORE sensitivity evaluation rises to 49,305 seconds or 13.7 hours in wall-clock time.

To improve the efficiency of computing DYMORE sensitivities, the global surface mesh is assembled and stored by each processor to more efficiently compute the elastic surface sensitivities, i.e., $[\partial \mathbf{X}_s^k / \partial \mathbf{f}^n]$ in Eq. (30). This approach accelerates parallel computations by minimizing communications and ensuring that the surface mesh sensitivity computations are distributed among all available processors, but requires excessive memory. Fig. 12 depicts the scalability of computing DYMORE sensitivities in the current implementation. The computational mesh employed in this scalability test corresponds to the mesh employed in the optimization problem in Section III B 3, and 720, 1080 and 2088 processors are used.

A few approaches can be implemented to further accelerate computations. (1) The current DYMORE model can be replaced by a newer version, DYMORE 5, in which domain decomposition and parallelization have been implemented, together with a more efficient formulation based on the motion formalism [48–50]. (2) The number of operations for computing surface elastic sensitivities can be further reduced by separating the surface grid sensitivities $[\partial \mathbf{X}_s^k / \partial \mathbf{f}^n]$ from DYMORE sensitivities by using a chain rule $[\partial \mathbf{X}_s^k / \partial \mathbf{f}^n] = [\partial \mathbf{X}_s^k / \partial \mathbf{u}^k][\partial \mathbf{u}^k / \partial \mathbf{f}^n]$ and performing hand differentiation for the term $[\partial \mathbf{X}_s^k / \partial \mathbf{u}^k]$. As a consequence, the matrix-vector product $[\mathbf{\Lambda}_G^k]^T [\partial \mathbf{X}_s^k / \partial \mathbf{u}^k]$ used in Eq. (30) can be precomputed and used for all sensitivities corresponding to DYMORE solutions at time level k. (3) DYMORE sensitivities can be computed via a real-valued finite-difference approach, which would lead to an additional speedup factor for DYMORE sensitivities. (4) The grid adjoint solver can be optimized to compute simultaneously mesh adjoint solutions for multiple objectives and constraints. Although various strategies can be implemented to increase the efficiency of the present multidisciplinary sensitivity analysis relying on finite-difference sensitivities for

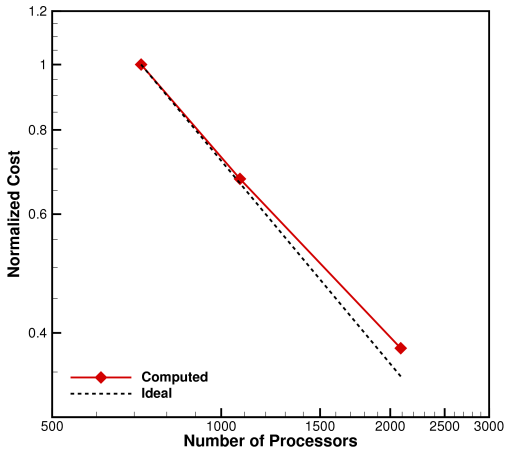

Fig. 12 Scalability of computing DYMORE sensitivities to airloads.

DYMORE, the cost of the analysis grows quadratically with the number of time steps. Elimination of this quadratic dependence requires a discretely consistent adjoint formulation for DYMORE sensitivities.

IV. Conclusions

This paper has presented a multidisciplinary design and optimization framework for high-fidelity rotorcraft computations conducted by a tightly coupled multidisciplinary system including a computational fluid dynamics code, FUN3D, and a rotorcraft comprehensive analysis code, DYMORE. The implementation of the loose and tight coupling procedures for the FUN3D/DYMORE analysis has been validated by simulating the HART-II model in a descending flight condition that exhibits significant blade-vortex interactions. The computed solutions show good agreement with established predictions of unsteady aerodynamic loads and various elastic deflections. A formulation for the tight coupling multidisciplinary sensitivity analysis has been developed, implemented, and verified. The formulation integrates the adjoint-based sensitivities of flow and grid solutions available in FUN3D with DYMORE sensitivities computed by a complex-variable finite-difference approach. A constrained gradient-based design optimization procedure has been formulated and successfully applied to optimization of the HART-II rotorcraft configuration. This demonstration confirms feasibility of high-fidelity multidisciplinary sensitivity analysis for rotorcraft applications. It also highlights the need to reduce the cost of computing sensitivities of structural responses to aerodynamic loads. This cost grows quadratically with the number of time steps and represents the dominant part of the total cost of the sensitivity analysis. Completion of nine optimization cycles took more than 8 days of high-performance

computations involving over 2000 processors. Implementation of a discretely consistent adjoint formulation for DYMORE will eliminate this quadratic dependence and dramatically decrease the computational time required for a high-fidelity rotorcraft design optimization.

Acknowledgments

The work was supported by the NASA Revolutionary Vertical Lift Technology Project contract NNL15AB93T. The authors would like to thank William Jones at NASA Langley Research Center for assistance with geometry parameterization and Prof. Marilyn Smith at Georgia Institute of Technology for providing the initial FUN3D/DYMORE interface routines as well as the DYMORE input deck for the HART-II configuration. The authors would also like to acknowledge the support of the HART II partners from the U.S. Army Aero-Flight Dynamics Directorate (AFDD), German Aerospace Center (DLR), German-Dutch Windtunnel (DNW), and French Aeronautics and Space Research Center (ONERA). Resources supporting this work were provided by the NASA High-End Computing (HEC) Program through the NASA Advanced Supercomputing (NAS) Division at Ames Research Center.

References

[1] O. A. Bauchau, C. L. Bottasso, Y. G. Nikishkov, Modeling rotorcraft dynamics with finite element multibody procedures, Mathematical and Computer Modeling 33 (2001) 1113–1137.

[2] W. Johnson, Rotorcraft Aeromechanics, Cambridge Aerospace Series, Cambridge University Press, 2013.

[3] H. Saberi, M. Khoshlahjeh, R. A. Ormiston, M. J. Rutkowski, Overview of RCAS and application to advanced rotorcraft problems, American Helicopter Society 4th Decennial Specialists Conference on Aeromechanics (Jan. 2004).

[4] W. Johnson, A history of rotorcraft comprehensive analyses, NASA/TP-2012-216012, NASA Ames Research Center, Moffett Field, CA (April 2012).

[5] M. Potsdam, H. Yeo, W. Johnson, Rotor airloads prediction using loose aerodynamic/structural coupling, Journal of Aircraft 43 (3) (2006) 732–742.

[6] A. Altmikus, S. Wagner, On the timewise accuracy of staggered aeroelastic simulations of rotary wings, Aerodynamics, Acoustics, and Test and Evaluation Specialist Meeting, American Helicopter Society, San Francisco, CA (Jan. 2002).

[7] R. T. Biedron, E. M. Lee-Rausch, Rotor airloads prediction using unstructured meshes and loose CFD/CSD coupling, AIAA Paper 2008–7341 (June 2008).

[8] M. J. Smith, J. W. Lim, B. G. van der Wall, J. D. Baeder, R. T. Biedron, D. D. Boyd Jr., B. Jayaraman, S. N. Jung, B.-Y. Min, The HART II international workshop: An assessment of the state of the art in CFD/CSD prediction, CEAS Aeronautical Journal 4 (2013) 345–372.

[9] R. Celi, Recent applications of design optimization to rotorcraft - a survey, Journal of Aircraft 36 (1999) 176–189.

[10] J. L. Walsh, K. C. Young, F. J. Tarzanin, J. E. Hirsh, D. K. Young, Optimization issues with complex rotorcraft comprehensive analysis, AIAA Paper 1998–4889 (Sept. 1998).

[11] T. S. Murthy, Design sensitivity analysis of rotorcraft airframe structures for vibration reduction, NASA CP 2457, NASA Langley Research Center Sensitivity Analysis in Engineering, 1986, pp. 299–318.

[12] L. Wang, D. J. Mavriplis, W. K. Anderson, Adjoint sensitivity formulation for discontinuous Galerkin discretizations in unsteady inviscid flow problems, AIAA Journal 48 (12) (2010) 2867–2883.

[13] E. J. Nielsen, B. Diskin, N. K. Yamaleev, Discrete adjoint-based design optimization of unsteady turbulent flows on dynamic unstructured grids, AIAA Journal 48 (6) (2010) 1195–1206.

[14] L. Wang, W. K. Anderson, Shape sensitivity analysis for the compressible Navier-Stokes equations via discontinuous Galerkin methods, Computers & Fluids 69 (2012) 93–107.

[15] E. J. Nielsen, B. Diskin, Discrete adjoint-based design for unsteady turbulent flows on dynamic overset unstructured grids, AIAA Journal 51 (6) (2013) 1355–1373.

[16] A. Mishra, K. Mani, D. Mavriplis, J. Sitaraman, Time dependent adjoint-based optimization for coupled fluid-structure problems, Journal of Computational Physics 292 (2015) 253–271.

[17] E. J. Nielsen, Adjoint-based aerodynamic design of complex aerospace configurations, ASME 2016 Fluids Engineering Division Summer Meeting, FEDSM 2016–7573 (July 2016).

[18] J. C. Newman III, W. K. Anderson, D. L. Whitfield, Multidisciplinary sensitivity derivatives using complex variables, Computational Fluid Dynamics Laboratory, NSF Engineering Research Center for Computational Field Simulation, MSSU-COE-ERC-98-08, 1998 (July 1998).

[19] W. T. Jones, E. J. Nielsen, E. M. Lee-Rausch, J. Cecil W. Acree, Multi-point adjoint-based design of tilt-rotors in a noninertial reference frame, AIAA Paper 2014–0290 (Jan. 2014).

[20] E. Fabiano, D. Mavriplis, J. Sitaraman, Adjoint-based aeroacoustics design optimization for blade vortex interaction noise, AIAA Paper 2015–1801 (Jan. 2015).

[21] E. Fabiano, D. Mavriplis, Adjoint-based aeroacoustics design optimization of flexible rotors in forward flight, 72nd Forum, American Helicopter Society (May 2016).

[22] G. Corliss, C. Faure, A. Griewantk, L. Hascoet, U. Naumann, Automatic differentiation of algorithms: From simu-

lation to optimization, Springer-Verlag, 2002.

[23] P. Aubert, N. Cesare, O. Pironneau, Automatic differentiation in C++ using expression templates and application to a flow control problem, Computing and Visualization in Science 3 (2001) 197–208.

[24] M. Sagebaum, E. Özkaya, N. R. Gauger, J. Backhaus, C. Frey, S. Mann, M. Nagel, Efficient algorithmic differentiation techniques for turbo-machinery design, AIAA Paper 2017–3998 (Jan. 2017).

[25] J. R. R. A. Martins, J. T. Hwang, Review and unification of methods for computing derivatives of multidisciplinary computational models, AIAA Journal 51 (11) (2013) 2582–2599.

[26] L. Wang, B. Diskin, R. T. Biedron, E. J. Nielsen, O. A. Bauchau, Sensitivity analysis of multidisciplinary rotorcraft simulations, AIAA Paper 2017–1670 (Jan. 2017).

[27] W. K. Anderson, D. L. Bonhaus, An implicit upwind algorithm for computing turbulent flows on unstructured grids, Computers & Fluids 23 (1) (1994) 1–21.

[28] R. T. Biedron, J.-R. Carlson, J. M. Derlaga, P. A. Gnoffo, D. P. Hammond, W. T. Jones, B. Kleb, E. M. Lee-Rausch, E. J. Nielsen, M. A. Park, C. L. Rumsey, J. L. Thomas, W. A. Wood, FUN3D Manual: 13.1, NASA-TM-2017-219580, NASA Langley Research Center, Hampton, VA (2017).

[29] R. Noack, DiRTlib: A library to add an overset capability to your flow solver, AIAA Paper 2005–5116 (June 2005).

[30] R. Noack, D. Bogar, R. Kunz, P. Carrica, SUGGAR++: An improved general overset grid assembly capability, AIAA Paper 2009–3992 (June 2009).

[31] P. Spalart, S. Allmaras, A one-equation turbulence model for aerodynamic flows, Le Recherche Aerospatiale 1 (1994) 5–21.

[32] E. J. Nielsen, B. Diskin, High-performance aerodynamic computations for aerospace applications, Parallel Computing 64 (2017) 20–32.

[33] P. L. Roe, Approximate Riemann solvers, parameter vectors, and difference schemes, Journal of Computational Physics 43 (2) (1981) 357–372.

[34] C. O. E. Burg, Higher order variable extrapolation for unstructured finite volume RANS flow solvers, AIAA Paper 2005–4999 (June 2005).

[35] B. van Leer, Towards the ultimate conservative difference scheme, V. a second order sequel to Godunov's method, Journal of Computational Physics 32 (1) (1979) 101–136.

[36] V. Vatsa, M. H. Carpenter, D. Lockard, Re-evaluation of an optimized second order backward difference (BDF2OPT) scheme for unsteady flow applications, AIAA Paper 2010–0122 (Jan. 2010).

[37] M. Arnold, O. Bruls, Convergence of the generalized-α scheme for constrained mechanical systems, Multibody System Dynamics 18 (2) (2007) 185–202.

[38] J. N. Abras, C. E. Lynch, M. J. Smith, Computational fluid dynamics-computational structural dynamics rotor cou-

pling using an unstructured Reynolds-averaged Navier-Stokes methodology, Journal of The American Helicopter Society 57 (012001) (2012) 1–14.

[39] D. A. Peters, M. J. Johnson, Finite-state airloads for deformable airfoils on fixed and rotating wings, In Proceedings of the Symposium on Aeroelasticity and Fluid/Structure Interaction. American Society of Mechanical Engineers Winter Annual Meeting (Nov. 1994).

[40] R. T. Biedron, J. L. Thomas, Recent enhancements to the FUN3D flow solver for moving-mesh applications, AIAA Paper 2009–1360 (Jan. 2009).

[41] E. J. Nielsen, W. K. Anderson, Recent improvements in aerodynamic design optimization on unstructured meshes, AIAA Journal 40 (6) (2002) 1155–1163.

[42] B. G. van der Wall, A comprehensive rotary-wing data base for code validation: The HART II international workshop, The Aeronautical Journal 115 (1164) (2011) 91–102.

[43] J. A. Samareh, Novel multidisciplinary shape parameterization approach, Journal of Aircraft 38 (6) (2001) 1015–1024.

[44] P. E. Gill, W. Murray, M. A. Saunders, SNOPT: An SQP algorithm for large-scale constrained optimization, SIAM Journal on Optimization 12 (4) (2002) 979–1006.

[45] P. E. Gill, W. Murray, M. A. Saunders, SNOPT: An SQP algorithm for large-scale constrained optimization, SIAM Review 47 (1) (2005) 19–131.

[46] A. A. Kumar, S. R. Viswamurthy, R. Ganguli, Correlation of helicopter rotor aeroelastic response with HART II wind tunnel test, Aircraft Engineering and Aerospace Technology 82 (4) (2010) 237–248.

[47] J. W. Lim, An assessment of rotor dynamics correlation for descending flight using CSD/CFD coupled analysis, 64th Annual Forum, American Helicopter Society, Montreal, Canada (April–May 2008).

[48] V. Sonneville, A. Cardona, O. Brüls, Geometrically exact beam finite element formulated on the special Euclidean group SE(3), Computer Methods in Applied Mechanics and Engineering 268 (1) (2014) 451–474.

[49] V. Sonneville, O. Brüls, A formulation on the special Euclidean group for dynamic analysis of multibody systems, Journal of Computational and Nonlinear Dynamics 9 (4) (2014) 041002.

[50] V. Sonneville, O. Brüls, O. Bauchau, Interpolation schemes for geometrically exact beams: A motion approach, International Journal of Numerical Methods in Engineering. To appear.