

Recent Advances in High-Fidelity Multidisciplinary Adjoint-Based Optimization with the NSU3D Flow Solver Framework

Dimitri Mavriplis *

Enrico Fabiano †

Evan Anderson ‡

Department of Mechanical Engineering, University of Wyoming, Laramie, WY 82071

This paper describes the various components in the development of a coupled multidisciplinary optimization capability for rotorcraft problems. The underlying approach is based on the use of a time dependent discrete adjoint capability that is extended to multiple disciplines, including aerodynamic, structural dynamic and acoustic disciplines. Examples of optimizations for rigid and flexible rotors are given based on aerodynamic performance and aeroacoustic signature objectives. The extension of the structural model to a full geometrically non-linear brick and shell finite-element model is also described and preliminary results are shown.

I. Introduction

Over the last decade, substantial progress has been made in the development and application of sensitivity analysis methods for rotorcraft optimization problems. Both analysis and optimization of rotorcraft configurations are particularly challenging simulation problems. Rotorcraft configurations generally involve the simulation of bodies in relative motion which requires the use of time dependent approaches. This is in contrast to many fixed wing aircraft cruise optimization problems which can be formulated as steady-state problems. At the same time, high-fidelity analysis and optimization of rotorcraft problems must necessarily be multidisciplinary, involving aeroelastic effects and time dependent control inputs, since these effects are highly coupled and strongly influence rotorcraft aeromechanic performance. Due to the unsteady nature and tight multidisciplinary coupling required for these problems, rotorcraft applications have become a driver for the development and application of more advanced high-fidelity analysis and optimization techniques, most of which are only now starting to be considered for more challenging fixed wing applications at the edges of the flight envelope such as flutter and maneuvering flight.

Gradient-based optimization methods are well accepted as the most effective strategies for high-fidelity aerodynamic optimization, due to the much lower number of analysis runs or function evaluations required with these methods compared to non-local techniques such as genetic algorithms. Adjoint methods are now well established as the most efficient technique for obtaining the required design sensitivities, since these methods allow the computation of the sensitivities of an arbitrary number of design parameters with respect to an objective or constraint at the cost of the solution of a single adjoint problem, which is generally equivalent to the cost of a flow analysis solution. However, for time-dependent problems, the solution of the corresponding adjoint problem requires a backward integration in time which in turn requires storage of the complete time history of the analysis solution. One decade ago, when investigations into time-dependent adjoint formulations for aerospace applications were first considered, these were thought to be impractical, although the numerical weather prediction community was already employing such strategies for data assimilation purposes. Initial efforts were first demonstrated in two dimensions, and later in three dimensions, for reasonably complex configurations including overlapping mesh problems, and it is now acknowledged that time dependent adjoint problems, including the storage to disk of the analysis time history, are both practical and timely.

More recently, effort has been devoted to incorporating additional disciplines into the design optimization process in a high fidelity and tightly coupled manner, in order to enable more realistic and useful design optimization tools.

*Professor, AIAA Associate Fellow; email: mavriplis@uwyo.edu

†PhD Candidate, Member AIAA; email: efabiano@uwyo.edu

‡PhD Candidate, Member AIAA; email: evanski@uwyo.edu

This includes the incorporation of structural dynamics effects, control inputs, and farfield aeroacoustic signatures. The principal challenges in this area center on the development of the corresponding disciplinary adjoint sensitivities and the tight coupling and solution of the coupled multidisciplinary adjoint problems. For such problems, the discrete adjoint is highly favored since the discrete approach enables a straightforward and methodical (albeit tedious) approach for obtaining sensitivities of arbitrarily complex analysis procedures.

In this paper, we outline the principal techniques developed over the last decade in our group for developing a time-dependent multidisciplinary discrete adjoint capability for realistic rotorcraft problems which include aerodynamic, structural dynamic, aeroacoustic and control effects. Our approach is based on a time-domain simulation and (discrete) adjoint approach. This is in contrast to many efforts in rotorcraft applications which have focused on time-spectral or frequency domain approaches. Time-spectral methods offer the possibility of treating periodic problems as steady-state problems, and thus greatly simplify the cost and complications of the time-dependent adjoint solution process. However, time-spectral methods have other drawbacks, such as difficulties extending these methods to overset meshes, and to multidisciplinary problems. Furthermore, since the bulk of rotorcraft analysis tools are based on time-domain methods, it is most desirable to develop analogue sensitivity analysis methods which can be used directly with existing well validated analysis tools.

II. Aerostructural Analysis and Sensitivity Formulation

A. Flow Solver Analysis Formulation

The flow solver used in this work is the Reynolds-averaged “Navier-Stokes Unstructured 3D” (NSU3D)¹ mesh solver. NSU3D has been widely validated for steady-state and time-dependent flows and contains a discrete tangent and adjoint sensitivity capability which has been demonstrated previously for optimization of steady-state and time-dependent flow problems. As such, only a concise description of these formulations will be given in this paper, with additional details available in previous references.¹⁻³ The flow solver is based on the conservative form of the Navier-Stokes equations which may be written as:

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

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

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

Here V refers to the volume of the control volume bounded by a control surface $B(t)$, $\dot{\mathbf{x}}$ is the vector of mesh face or edge velocities, and \mathbf{n} is the unit normal of the face or edge. The state vector \mathbf{u} consists of the conserved variables and the Cartesian flux vector $\mathbf{F} = (\mathbf{F}_x, \mathbf{F}_y, \mathbf{F}_z)$ contains both inviscid and viscous fluxes. The equations are closed with the perfect gas equation of state and the Spalart-Allmaras turbulent eddy viscosity model⁴ for all cases presented in this work. The time derivative term is discretized using a second-order accurate backward-difference formula (BDF2) scheme, leading to the implicit system of equations at each time step given as:

$$\frac{3}{2\Delta t} \mathbf{V}^n \mathbf{u}^n - \frac{2}{\Delta t} \mathbf{V}^{n-1} \mathbf{u}^{n-1} + \frac{1}{2\Delta t} \mathbf{V}^{n-2} \mathbf{u}^{n-2} + \mathbf{S}^n(\mathbf{u}^n, \mathbf{x}^n, \dot{\mathbf{x}}^n) = 0 \quad (3)$$

where $\mathbf{V}^n = \mathbf{V}(x^n)$ represents the mesh control volumes and $\mathbf{S}^n(\mathbf{u}^n, \mathbf{x}^n, \dot{\mathbf{x}}^n)$ represents the spatial discretization terms at the n^{th} time step. The functional dependence of the implicit system to be solved at each time step can be written in residual form as:

$$\mathbf{R}^n(\mathbf{u}^n, \mathbf{u}^{n-1}, \mathbf{u}^{n-2}, \mathbf{x}^n, \mathbf{x}^{n-1}, \mathbf{x}^{n-2}) = 0, \quad n = 2, 3, \dots, N \quad (4)$$

where the initial conditions are given by \mathbf{u}^0 and \mathbf{x}^0 , and noting that a BDF1 time discretization is used for the first time step.

At each time step, the implicit residual is solved using Newton’s method. The Jacobian matrix is inverted iteratively using a line-implicit agglomeration multigrid scheme that can also be used as a preconditioner for a GMRES Krylov solver.⁵ To simplify the notation, we drop the time level superscripts, and represent the general flow equations to be solved at all time levels as the system given by:

$$\mathbf{R}(\mathbf{u}, \mathbf{x}) = 0 \quad (5)$$

where the vector \mathbf{U} denotes the flow values over all time steps, \mathbf{x} represents the CFD mesh coordinates over all time steps, and equation (5) denotes the simultaneous solution of all time steps. Given the functional dependence of

equation (4), it is evident that equation (5), when written in matrix form, takes on a lower triangular form, and can be solved by forward substitution, which corresponds to a forward integration in time.

B. Structural Model Formulation

A non-linear bend-twist beam model is a suitable and widely utilized structural model for slender fixed and rotary wing aircraft structures within the context of an aeroelastic problem. A bend-twist beam model has previously been developed and coupled to the NSU3D unstructured mesh Reynolds-averaged Navier-Stokes solver for steady^{6,7} and time-dependent problems.⁸ The non-linear governing equations of a slender beam are discretized using the finite-element method (FEM) in space using Hamilton's principle as described in detail in the reference by Hodges and Dowell.⁹ Figure 1(a) shows a typical beam with 15 degrees of freedom for each element to accommodate bend wise, lag wise, axial and torsional displacements. The elastic twist ϕ refers to the rotation about the blade elastic axis.

The ability to employ higher fidelity finite-element based structural models has also been pursued through the development of an in-house finite-element code. This code supports the basic commonly-used brick and shell element types and can import FEM models constructed within commercial software codes such as Abaqus and NASTRAN. Although implementing all features supported by commercial structural analysis codes is not feasible, the motivation for constructing an in-house capability is to enable tight coupling between the computational fluid dynamic (CFD) and computational structural dynamic (CSD) disciplines and also to enable implementation of the discrete adjoint of the structural analysis discretization for use in computing sensitivities. Figure 1(b) provides an illustration of a shell-based FEM model for a wind turbine blade constructed inside Abaqus that has been used in our coupled aero-structural analysis capability for fatigue loading predictions.¹⁰




Figure 1. (a) Beam finite element model with flap, lag, torsional and axial (total 15) degrees of freedom; (b) Shell finite-element model of wind turbine blade

The structural dynamic equations of motion can be expressed as: $M\ddot{\mathbf{q}} + C\dot{\mathbf{q}} + K\mathbf{q} = \mathbf{F}$, where $[M]$, $[C]$ and $[K]$ are mass, damping and stiffness matrices assembled based on the finite-element representation. Vector $\mathbf{F} = \mathbf{F}(t)$ is the forcing vector. Vector \mathbf{q} represents the displacements along all degrees of freedom. This set of equations can be reduced to a first-order system as:

$$[I]\dot{\mathbf{Q}} + [A]\mathbf{Q} = \mathbf{F} \quad (6)$$

where $[I]$ is the identity matrix, $\mathbf{Q} = [\mathbf{q}, \dot{\mathbf{q}}]^T$, $\mathbf{F} = [0, [M]^{-1}\mathbf{F}]^T$ and $[A] = \begin{bmatrix} 0 & -[I] \\ [M]^{-1}[K] & [M]^{-1}[C] \end{bmatrix}$. Similarly to the flow equations, a second-order backwards difference time-integration scheme (BDF2) is used to advance the structural equations in time. This results in the system of equations to be solved at each time step given as:

$$\frac{3}{2\Delta t}\mathbf{Q}^n - \frac{2}{\Delta t}\mathbf{Q}^{n-1} + \frac{1}{2\Delta t}\mathbf{Q}^{n-2} + [A]^n\mathbf{Q}^n = \mathbf{F}^n \quad (7)$$

Alternatively, a beta-Newmark time-integration approach¹¹ can be applied directly to equation (6). In either case, the functional dependence of the structural dynamic equations to be solved at each time step can be written in residual form as:

$$\mathbf{J}(\mathbf{Q}^n, \mathbf{Q}^{n-1}, \mathbf{Q}^{n-2}, \mathbf{F}^n) = 0, \quad n = 2, 3, \dots, N \quad (8)$$

As with the flow equations, to simplify the notation, we drop the time level superscripts, and represent the structural dynamic equations to be solved at all time levels as the system given by:

$$\mathbf{J}(\mathbf{Q}, \mathbf{F}) = \mathbf{0} \quad (9)$$

where \mathbf{Q} and \mathbf{F} represent the structural states and forces, respectively, over all time steps. Written in matrix form, equation (9) exhibits a lower triangular form, and can be solved by forward substitution in time, which corresponds to a forward time-integration procedure.

1. Forward sensitivity formulation of structural dynamic equations

The structural model tangent (forward sensitivity) linearization is similar to the analysis problem. For a given functional, $L(\mathbf{D}, \mathbf{Q}(\mathbf{D}))$, its sensitivity with respect to a blade design parameter, \mathbf{D} can be written as:

$$\frac{dL}{d\mathbf{D}} = \frac{\partial L}{\partial \mathbf{D}} + \frac{\partial L}{\partial \mathbf{Q}} \frac{d\mathbf{Q}}{d\mathbf{D}} \quad (10)$$

This requires solving for the sensitivity of the structural state \mathbf{Q} , which can be obtained by differentiating Eqn. (9) with respect to the design variable \mathbf{D} and rearranging as:

$$\left[\frac{\partial \mathbf{J}}{\partial \mathbf{Q}} \right] \frac{\partial \mathbf{Q}}{\partial \mathbf{D}} = - \frac{\partial \mathbf{J}}{\partial \mathbf{F}} \frac{d\mathbf{F}}{d\mathbf{D}} - \frac{\partial \mathbf{J}}{\partial \mathbf{D}} \quad (11)$$

The last term on the right hand side is non zero for structural design parameters such as structural element stiffnesses, in which case the applied force does not change with the design parameter, making the first term on the right hand side zero. In the coupled aeroelastic case, using aerodynamic shape parameters that primarily affect the airloads on the structure, the first term on the right-hand side is non-zero while the second term vanishes. Solving for $\frac{\partial \mathbf{Q}}{\partial \mathbf{D}}$ in equation (11), and substituting this into equation (10), the forward sensitivity of the objective function $\frac{dL}{d\mathbf{D}}$ can be obtained.

2. Adjoint formulation of structural model

The adjoint formulation of the structural model can be derived by approaching the tangent formulation in the reverse (transpose) direction. Taking the transpose of the objective functional sensitivity yields: $\frac{dL^T}{d\mathbf{D}} = \frac{\partial L^T}{\partial \mathbf{D}} + \frac{d\mathbf{Q}^T}{d\mathbf{D}} \frac{\partial L^T}{\partial \mathbf{Q}}$. This requires solving for the transpose sensitivity of the structural state \mathbf{Q} , which leads to solving for an adjoint vector $\Lambda_{\mathbf{Q}}$ defined as:

$$\left[\frac{\partial \mathbf{J}}{\partial \mathbf{Q}} \right]^T \Lambda_{\mathbf{Q}} = \frac{\partial L^T}{\partial \mathbf{Q}} \quad (12)$$

The final objective sensitivities can be obtained by:

$$\frac{dL^T}{d\mathbf{D}} = \frac{\partial L^T}{\partial \mathbf{D}} + \left[- \frac{\partial \mathbf{J}^T}{\partial \mathbf{D}} - \frac{d\mathbf{F}^T}{d\mathbf{D}} \frac{\partial \mathbf{J}^T}{\partial \mathbf{F}} \right] \Lambda_{\mathbf{Q}} \quad (13)$$

The above forms the adjoint formulation of the structural model. In this work we consider only shape design variables that affect the outer mold line of the CFD geometry without modifying the structural model. Thus, the influence of these aerodynamic design variables is only felt through changes in the aerodynamic forces applied to the structural model, which correspond to the second term in the brackets on the right-hand side of the above equation. The first term in the brackets vanishes in this case, but would be non-zero when considering structural design parameters such as material properties. As previously, the left-hand-side Jacobian term of the structural adjoint equation corresponds to the transpose of the Jacobian in the forward linearization.

C. Fluid-structure interface (FSI)

In addition to the solution of the aerodynamic problem and the structural dynamics problem, the solution of the fully coupled time-dependent aeroelastic problem requires the exchange of aerodynamic loads from the CFD solver to the structural solver, which in turn returns surface displacements to the fluid flow solver. In practice, the FSI computes forces at each CFD surface mesh point $\mathbf{F}(\mathbf{x}, \mathbf{u})$ by integrating the surface stresses (pressure and shear) over the surface area associated with each surface grid point. These forces are then projected onto the finite-element basis functions of

the beam or brick and shell model where they are assembled in the form of forces on the finite-element nodal locations denoted as \mathbf{F}_b . Conversely, once the structural deflections have been computed, they are transferred back to the surface CFD mesh in a similar manner. The governing equations for the FSI can be written in residual form as:

$$\mathbf{S}^{\text{FSI}}(\mathbf{F}_b, \mathbf{x}, \mathbf{u}) = \mathbf{F}_b - [T(\mathbf{Q})]\mathbf{F}(\mathbf{x}, \mathbf{u}) = 0 \quad (14)$$

$$\mathbf{S}^{\text{SFI}}(\mathbf{x}_s, \mathbf{Q}) = \mathbf{x}_s - [T(\mathbf{Q})]^T \mathbf{Q} = 0 \quad (15)$$

where \mathbf{S}^{FSI} represents the force transfer from CFD to CSD, and \mathbf{S}^{SFI} denotes the displacement transfer from CSD to CFD. In these equations, $[T]$ represents the rectangular transfer matrix which projects point-wise CFD surface forces $\mathbf{F}(\mathbf{x}, \mathbf{u})$ onto the individual structural model elements resulting in the forces \mathbf{F}_b . The transpose of this matrix is used to obtain the CFD surface displacements \mathbf{x}_s from the structural model degrees of freedom \mathbf{Q} , thus ensuring the principle of conservation of virtual work. The interpolation patterns which define the $[T]$ matrix are computed by locating the perpendicular projection of each point of the surface CFD mesh on the structural model beam elements or brick and shell surface faces, as shown in Figure 2. This approach allows for non-matching surfaces between the CFD and CSD elements, even in the case of the full brick/shell finite-element structural model, since these may contain widely differing discretizations and are often built using different outer mold line definitions.




Figure 2. Illustration of fluid-structure interface for (a) beam model (b) brick and shell finite element model

D. Control input facility

Control inputs for defining rotor blade motion must be applied in order to simulate a rotor in forward flight. For a fixed hub, the overall rotor motion can be broken down into two parts: the cyclic pitching of the blades defined by a pitch angle (θ), and the azimuthal rotation of the blades about the hub axis (ψ), as shown in Figure 3. The time-varying blade pitch angle can be represented by a combination of mean pitch angle (θ_0) and several harmonic components of pitch angles (e.g. θ_{c_i} and θ_{s_i} for i 'th harmonic): $\theta = \theta_0 + \theta_{c_i} \cos(i\psi) + \theta_{s_i} \sin(i\psi)$, for harmonics of $i = [1, 2, \dots]$. This introduces the set of control parameters, i.e. $\mathbf{D} = [\theta_0, \theta_{c_1}, \theta_{s_1}]$, which define the prescribed blade cyclic pitching motion. In this work, three pitch parameters are used, namely collective θ_0 , and two cyclics ($\theta_{c_1}, \theta_{s_1}$) corresponding to a single harmonic $i = 1$. At each time step, the individual blade pitch angles are computed using the above control inputs, and used to displace all CFD surface mesh points by the required angular rotation θ . Thus, the pitching motion results in a new set of surface grid points defined by the solution of the blade pitching equation, in residual form:

$$\mathbf{S}^\theta(\mathbf{x}_{\text{sp}}, \mathbf{x}_s, \mathbf{D}) = \mathbf{0} \quad (16)$$

where \mathbf{x}_{sp} represents the new pitched blade surface mesh points, \mathbf{x}_s represents the original surface mesh points, and \mathbf{D} represents the control inputs, which may also be used as design parameters for optimization runs.

Once the surface mesh points have been displaced, the interior mesh must be deformed to avoid negative cells in neighboring regions. This is accomplished using the mesh deformation facility described subsequently in the **Mesh deformation capability** section, resulting in a new set of interior mesh points denoted as \mathbf{x}_o . The azimuthal rotation is next applied by rotating the entire grid as a solid body. This is achieved by solving the prescribed mesh rotation equations, written in residual form as:

$$\mathbf{S}^\Psi(\mathbf{x}, \mathbf{x}_o) = \mathbf{0} \quad (17)$$

where \mathbf{x} denotes the final deformed and rotated mesh coordinates.




Figure 3. Illustration of prescribed pitch and azimuthal rotation for rotor in forward flight.

E. Geometry parameterization facility

In order to perform shape optimization of a multi-bladed rotor, a shape parametrization capability must be implemented. The specific parametrization used in this work is described in section 1. Mathematically, the blade design methodology defines a blade surface, which in turn defines the coordinates of the CFD surface mesh points on each blade in a fixed reference frame as a function of a set of shape design variables \mathbf{D} . Thus, the blade surface mesh points \mathbf{x}_{so} are obtained as the solution of the parametrization equations, in residual form:

$$\mathbf{B}_{cad}(\mathbf{x}_{so}, \mathbf{D}) = \mathbf{0} \quad (18)$$

where \mathbf{B}_{cad} denotes a possible computer-aided design (CAD) parametrization of the blade, and \mathbf{D} refers to the blade shape and pitching parameters.

F. Mesh deformation capability

A mesh deformation capability is required to deform the volume mesh in response to prescribed displacements of the surface CFD mesh. Surface displacements can be produced in three manners: 1) through blade geometry shape changes resulting during the design process, 2) through blade deflections generated by the structural model, and 3) due to the prescribed blade cyclic pitching, defined by the three pitch control parameters. In practice, a single mesh deformation calculation is performed at each coupling cycle. Thus, the surface displacements are first determined by adding together the displacements from the three sources when present, and the interior mesh point coordinates are obtained by solving the mesh deformation equations, which can be written in residual form as:

$$\mathbf{G}(\mathbf{x}_o, \mathbf{x}_{sp}) = \mathbf{0} \quad (19)$$

where \mathbf{x}_o denotes the deformed interior mesh coordinates in response to the displaced surface mesh coordinates \mathbf{x}_{sp} . Note that the mesh deformation calculation is carried out prior to the azimuthal mesh rotation phase, which operates on the deformed coordinates \mathbf{x}_o before rotation to produce the final mesh coordinates \mathbf{x} used in the CFD simulation, as described by equation (17). The mesh deformation approach is based on a linear elastic analogy. In this approach, the mesh is modeled as a linear elastic solid with a variable modulus of elasticity that can be prescribed either as inversely proportional to the cell volume or to the distance of each cell from the nearest wall (and thus is a function of current grid). The resulting equations are discretized and solved on the mesh in its original undeformed configuration in response to surface displacements using a line-implicit multigrid algorithm analogous to that used for the flow equations. Detailed explanation of the method can be found in references.^{6,12} An equivalent adjoint solver¹² for the mesh deformation problem is also implemented for use in the coupled aeroelastic adjoint capability.

G. General solution procedure

The general analysis problem consists of multiple coupled sets of equations namely, the blade design equations, the prescribed motion equations, the mesh deformation equations, the flow solution equations, the structural model equations, and the fluid-structure interface. All these equations are coupled and must be solved simultaneously at each time

step. The coupled system of equations to be solved can be written as:

$$\mathbf{B}_{\text{cad}}(\mathbf{x}_{\text{so}}, \mathbf{D}) = \mathbf{0} \quad (20)$$

$$\mathbf{S}^{\theta}(\mathbf{x}_{\text{sp}}, \mathbf{x}_{\text{so}}, \mathbf{x}_{\text{s}}, \mathbf{D}) = \mathbf{0} \quad (21)$$

$$\mathbf{G}(\mathbf{x}_{\text{o}}, \mathbf{x}_{\text{sp}}) = \mathbf{0} \quad (22)$$

$$\mathbf{S}^{\psi}(\mathbf{x}, \mathbf{x}_{\text{o}}) = \mathbf{0} \quad (23)$$

$$\mathbf{R}(\mathbf{u}, \mathbf{x}) = \mathbf{0} \quad (24)$$

$$\mathbf{S}^{\text{FSI}}(\mathbf{F}_{\text{b}}, \mathbf{x}, \mathbf{u}) = \mathbf{0} \quad (25)$$

$$\mathbf{J}(\mathbf{F}_{\text{b}}, \mathbf{Q}) = \mathbf{0} \quad (26)$$

$$\mathbf{S}^{\text{SFI}}(\mathbf{x}_{\text{s}}, \mathbf{Q}) = \mathbf{0} \quad (27)$$

In equation (21), the cyclic pitching acts directly on the deflected blade shape coordinates \mathbf{x}_{s} , although these depend on the original undeflected blade surface coordinates \mathbf{x}_{so} and reduce to these values at the beginning of the solution procedure prior to the solution of the structural equations. The above equations denote the system of equations to be solved over all time steps. However, since each residual equation at a given time level n depends on (at most) the values from time levels n , $n - 1$, and $n - 2$, the entire system can be solved by forward substitution in time (i.e. forward time integration). In this case, the equations to be solved at a given time level n can be written more precisely by reintroducing the time level superscripts on the residual operators and their corresponding arguments.

Within each physical time step, solution of the fully coupled fluid structure problem consists of performing multiple coupling iterations on each discipline using the latest available values from the other disciplines. Thus, a typical solution procedure starts by generating the blade surface mesh coordinates and prescribing the required pitch angles to these points, as determined by the first two equations, which can be evaluated explicitly. Next, the mesh deformation equations (i.e. equation (22)) must be solved iteratively, after which the deformed mesh is rotated through explicit evaluation of equation (23). The flow equations (i.e. equation (24)) can then be solved iteratively using the updated grid coordinates \mathbf{x} and the resulting flow variables \mathbf{u} are then used to compute forces which are applied to the structural model through explicit evaluation of the FSI (i.e. equation (25)). Using these forces, the structural model is solved directly (i.e. equation (26)) and the resulting displacements are transferred back to the CFD mesh as determined by equation (27). Since the surface mesh coordinates are now modified, the entire process must be repeated, starting at the second equation, until convergence is achieved.

H. Sensitivity Analysis for Coupled Aeroelastic Problem

In this work we consider only aerodynamic performance objectives which depend only on the flow solution and geometric shape. Starting with the forward sensitivity problem, the sensitivity of such an objective L can be written as:

$$\frac{dL}{d\mathbf{D}} = \begin{bmatrix} \frac{\partial L}{\partial \mathbf{x}} & \frac{\partial L}{\partial \mathbf{u}} \end{bmatrix} \begin{bmatrix} \frac{d\mathbf{x}}{d\mathbf{D}} \\ \frac{d\mathbf{u}}{d\mathbf{D}} \end{bmatrix} \quad (28)$$

As previously, the variables \mathbf{u} and \mathbf{x} represent the values over all time steps, and the inner products are over all space and time. The individual disciplinary sensitivities are given as the solution of the coupled system, which is obtained by differentiating the coupled analysis equations:

$$\begin{bmatrix}
\frac{\partial \mathbf{B}_{\text{cad}}}{\partial \mathbf{x}_{\text{so}}} & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
\frac{\partial \mathbf{S}^\theta}{\partial \mathbf{x}_{\text{so}}} & \frac{\partial \mathbf{S}^\theta}{\partial \mathbf{x}_{\text{sp}}} & 0 & 0 & 0 & 0 & 0 & \frac{\partial \mathbf{S}^\theta}{\partial \mathbf{x}_s} \\
0 & \frac{\partial \mathbf{G}}{\partial \mathbf{x}_{\text{sp}}} & \frac{\partial \mathbf{G}}{\partial \mathbf{x}_o} & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & \frac{\partial \mathbf{S}^\psi}{\partial \mathbf{x}_o} & \frac{\partial \mathbf{S}^\psi}{\partial \mathbf{x}} & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & \frac{\partial \mathbf{R}}{\partial \mathbf{x}} & \frac{\partial \mathbf{R}}{\partial \mathbf{u}} & 0 & 0 & 0 \\
0 & 0 & 0 & \frac{\partial \mathbf{S}^{\text{FSI}}}{\partial \mathbf{x}} & \frac{\partial \mathbf{S}^{\text{FSI}}}{\partial \mathbf{u}} & \frac{\partial \mathbf{S}^{\text{FSI}}}{\partial \mathbf{F}_b} & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & \frac{\partial \mathbf{J}}{\partial \mathbf{F}_b} & \frac{\partial \mathbf{J}}{\partial \mathbf{Q}} & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & \frac{\partial \mathbf{S}^{\text{SFI}}}{\partial \mathbf{Q}} & \frac{\partial \mathbf{S}^{\text{SFI}}}{\partial \mathbf{x}_s}
\end{bmatrix}
\begin{bmatrix}
\frac{d\mathbf{x}_{\text{so}}}{d\mathbf{D}} \\
\frac{d\mathbf{x}_{\text{sp}}}{d\mathbf{D}} \\
\frac{d\mathbf{x}_o}{d\mathbf{D}} \\
\frac{d\mathbf{x}}{d\mathbf{D}} \\
\frac{d\mathbf{u}}{d\mathbf{D}} \\
\frac{d\mathbf{F}_b}{d\mathbf{D}} \\
\frac{d\mathbf{Q}}{d\mathbf{D}} \\
\frac{d\mathbf{x}_s}{d\mathbf{D}}
\end{bmatrix}
=
\begin{bmatrix}
-\frac{\partial \mathbf{B}_{\text{cad}}}{\partial \mathbf{D}} \\
-\frac{\partial \mathbf{S}^\theta}{\partial \mathbf{D}} \\
0 \\
0 \\
0 \\
0 \\
0 \\
0
\end{bmatrix}$$

The first equation corresponds to equations for surface mesh point sensitivities with respect to the shape design variables, while the second equation rotates these sensitivities through the prescribed pitch angle, and adds the sensitivities with respect to the pitch control inputs, which are also used as design parameters. The third equation propagates the surface mesh sensitivities to the interior mesh points through the mesh deformation equations, while the fourth equation corresponds to the azimuthal rotation of these mesh sensitivities. The fifth equation generates the flow sensitivities based on the mesh sensitivities and the solution of the linearized flow problem. The flow sensitivities are then used to construct force sensitivities for the structural model using the FSI, which in turn generate structural sensitivities in the seventh equation. These are passed back to the CFD surface mesh in the last equation, to be reused at the next coupling iteration. As can be seen, each disciplinary solution procedure requires the inversion of the same Jacobian matrix as the corresponding analysis problem, which is done using the same solution technique.

Technically, the above system represents the forward sensitivity equations to be solved over all time steps. However, as mentioned previously, due to the dependence on at most the previous two time levels in all the individual disciplinary equations, the above matrix equation can be solved by forward substitution in time. The left-hand side matrix can be rearranged into a lower triangular form, by reordering of the equations grouping together all disciplinary equations at a given time step, which are then solved simultaneously prior to advancing to the next time step, as described in references.^{3,13} However, for convenience we retain the global temporal notation with the understanding that the above system can also be interpreted as the equations to be solved at a given time step, simply by adding the appropriate time level superscripts to the corresponding operators and variables.

The corresponding adjoint problem can be obtained by pre-multiplying the above matrix equation by the inverse of the large coupling matrix and substituting this into equation (28), transposing the entire system, and defining adjoint variables as solutions to the following coupled system:

$$\begin{bmatrix}
\frac{\partial \mathbf{B}_{\text{cad}}^T}{\partial \mathbf{x}_{\text{so}}} & \frac{\partial \mathbf{S}^{\theta T}}{\partial \mathbf{x}_{\text{so}}} & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & \frac{\partial \mathbf{S}^{\theta T}}{\partial \mathbf{x}_{\text{sp}}} & \frac{\partial \mathbf{G}^T}{\partial \mathbf{x}_{\text{sp}}} & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & \frac{\partial \mathbf{G}^T}{\partial \mathbf{x}_o} & \frac{\partial \mathbf{S}^{\psi T}}{\partial \mathbf{x}_o} & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & \frac{\partial \mathbf{S}^{\psi T}}{\partial \mathbf{x}} & \frac{\partial \mathbf{R}^T}{\partial \mathbf{x}} & \frac{\partial \mathbf{S}^{\text{FSI}^T}}{\partial \mathbf{x}} & 0 & 0 \\
0 & 0 & 0 & 0 & \frac{\partial \mathbf{R}^T}{\partial \mathbf{u}} & \frac{\partial \mathbf{S}^{\text{FSI}^T}}{\partial \mathbf{u}} & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & \frac{\partial \mathbf{S}^{\text{FSI}^T}}{\partial \mathbf{F}_b} & \frac{\partial \mathbf{J}^T}{\partial \mathbf{F}_b} & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & \frac{\partial \mathbf{J}^T}{\partial \mathbf{Q}} & \frac{\partial \mathbf{S}^{\text{SFI}^T}}{\partial \mathbf{Q}} \\
0 & \frac{\partial \mathbf{S}^{\theta T}}{\partial \mathbf{x}_s} & 0 & 0 & 0 & 0 & 0 & \frac{\partial \mathbf{S}^{\text{SFI}^T}}{\partial \mathbf{x}_s}
\end{bmatrix}
\begin{bmatrix}
\Lambda_{\mathbf{x}_{\text{so}}} \\
\Lambda_{\mathbf{x}_{\text{sp}}} \\
\Lambda_{\mathbf{x}_o} \\
\Lambda_{\mathbf{x}} \\
\Lambda_{\mathbf{u}} \\
\Lambda_{\mathbf{F}_b} \\
\Lambda_{\mathbf{Q}} \\
\Lambda_{\mathbf{x}_s}
\end{bmatrix}
=
\begin{bmatrix}
0 \\
0 \\
0 \\
\frac{\partial \mathbf{L}^T}{\partial \mathbf{x}} \\
\frac{\partial \mathbf{L}^T}{\partial \mathbf{u}} \\
0 \\
0 \\
0
\end{bmatrix}$$

Once again, the solution of the various disciplinary adjoint equations requires the inversion of the corresponding disciplinary Jacobians (transposed in this case) which can be accomplished using the same iterative solvers as for the analysis and forward sensitivity problems, but proceeding in the reverse order. By analogy with the forward sensitivity problem, this coupled adjoint system over all time steps can be written in upper triangular form and solved by back substitution, which corresponds to a backwards integration in time.^{3,13} The specific coupled system to be solved at each time step can be recovered by specifying the time level superscripts on all operators and variables.

Once the solution of the coupled disciplinary adjoint problems has been obtained, these can be used to compute the objective sensitivities as:

$$\frac{dL}{d\mathbf{D}} = \begin{bmatrix} -\frac{\partial \mathbf{B}_{\text{cad}}}{\partial \mathbf{D}} & -\frac{\partial \mathbf{S}^{\theta}}{\partial \mathbf{D}} & 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} \Lambda_{\mathbf{x}_{\text{so}}} \\ \Lambda_{\mathbf{x}_{\text{sp}}} \\ \Lambda_{\mathbf{x}_o} \\ \Lambda_{\mathbf{x}} \\ \Lambda_{\mathbf{u}} \\ \Lambda_{\mathbf{F}_b} \\ \Lambda_{\mathbf{Q}} \\ \Lambda_{\mathbf{x}_s} \end{bmatrix} \quad (29)$$

Figure 4 illustrates the flow of information between fluid and structural disciplines for the analysis, tangent, and adjoint solvers. In the analysis problem, the CFD solver sends a force vector to the structural solver and receives a surface displacement vector, while in the tangent problem, these are replaced with a force sensitivity and displacement sensitivity vector, respectively. For the adjoint problem, the information proceeds in reverse order with the CFD solver receiving a force adjoint vector from the structural model, and sending a displacement adjoint to the structural model. In all cases, these vectors have the same dimension as those used in the analysis problem and make use of the same

data structures for inter-disciplinary coupling. The unsteady adjoint requires the solution state from each of the time steps, which is easily read from hard disk saved during the forward analysis step. This avoids additional memory requirements.

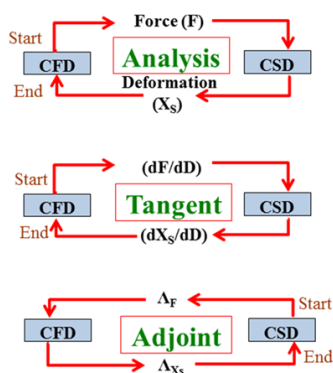


Figure 4. Corresponding flow of information for analysis, tangent and adjoint solution processes.

The forward and adjoint sensitivities for the coupled aeroelastic problem are verified using the complex step method.¹⁴ As in the case of finite-differencing, the complex step-based differentiation also requires a step size. However, unlike finite-differencing the complex step method is insensitive to small step sizes since no differencing is required. In theory it is possible to verify forward and adjoint-based gradients using the complex step method to machine precision. With this in mind, a complex version of the complete coupled aero-structural analysis code has been constructed through scripting of the original source code to redefine variables from real to complex types and to overload a small number of functions for use with complex variables. Agreement between the sensitivities obtained with the complex step method, the forward sensitivity and the adjoint implementations has been established to within 11 significant digits as reported in references.^{15, 16} Furthermore, this work has demonstrated that the agreement of these sensitivities does not degrade as the number of time steps is increased, for up to several hundred time steps, provided the equations are converged to machine precision at each time step. On the other hand, algebraic error accumulation has been found to contaminate the precision of the computed gradients and grow with the number of time steps, suggesting that deeper convergence is required for cases with larger numbers of time steps.¹⁶

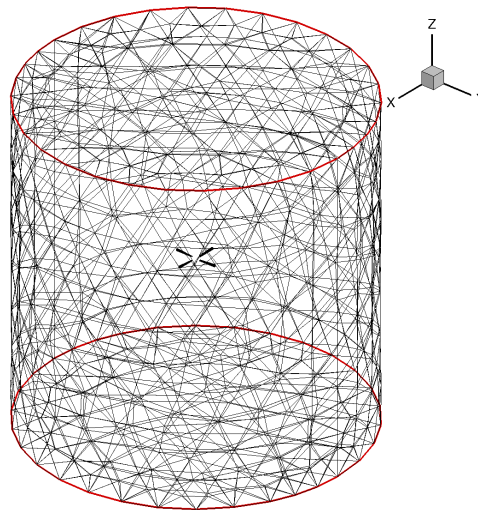
III. Aerostructural analysis and design

A. Time dependent analysis problem

The chosen test case is a four bladed HART2 rotor in a forward flight condition. The flight conditions correspond to a wind tunnel test¹⁷ with a freestream Mach number of $M_\infty = 0.095$, a rotor tip Mach number of $M_{tip} = 0.638$, and a shaft tilt angle towards freestream of $\alpha_{shaft} = 5.4^\circ$. The corresponding rotor rotational speed is $\Omega = 1041$ RPM. The blade pitch angle actuation is prescribed as: $\theta = \theta_0 + \theta_{1s}\sin(\psi) + \theta_{1c}\cos(\psi)$, with $\theta_0 = 5.0^\circ$, $\theta_{1s} = -1.1^\circ$, $\theta_{1c} = 2.0^\circ$, where ψ corresponds to the azimuthal angle of rotation about the hub axis. The rotor is impulsively started from rest, in an initially quiescent flow field, and rotated with the mesh as a solid body for a fixed number of revolutions. This problem is solved both for a rigid blade model (using no structural model), as well as for a flexible blade model (using the beam structural model). For the latter, the flow is solved in tight coupling mode with the beam solver as described previously.

The baseline simulation (coarse mesh) makes use of a mixed element mesh made up of prisms, pyramids and tetrahedra consisting of approximately 2.32 million grid points and is shown in Figure 5(a). The simulations are run for 5 rotor revolutions using a 2 degree time-step, i.e. for 900 time-steps starting from freestream initialization. For the rigid blade simulation, the time-dependent mesh motion is determined by first pitching the blade about the blade axis followed by solving the mesh deformation equations and then rotating the entire mesh as a solid body at each time step. The unsteady Reynolds-averaged Navier-Stokes equations are solved at each time step in ALE form, using the Spalart-Allmaras turbulence model.

The coupled CFD/CSD simulation is run in a similar manner. However, the flow solution (CFD) is coupled with the beam solver (CSD) at every time step by appropriately exchanging, a) airloads information from the flow domain to the beam and b) blade deformation information from the beam to the flow domain, through the fluid-structure




(a) Computational Domain: 2.32 million nodes

Figure 5. Computational polyhedra mesh

interface (i.e. blade surface). In this coupled simulation, the mesh is moved according to the deformations dictated by the new flexed blade coordinates determined from the structural beam code after the combined kinematics of pitch actuation followed by solid body rotation of the entire mesh have been performed. This coupled fluid-structure interaction problem needs to be iterated until satisfactory convergence (up to more than two orders of magnitude drop) is achieved on flow, structure and mesh deformation problems within each time step. The flow and mesh deformation problems are converged up to more than two orders of magnitude, and the structure problem up to machine precision (since it uses direct matrix inversion) values of the respective residuals.

The simulations were performed on the Yellowstone supercomputer at the NCAR-Wyoming Supercomputing Center (NWSC), with the analysis problem running in parallel on 1024 cores. Each time step used 6 coupling iterations, and each coupling iteration used 10 non-linear flow iterations with each non-linear iteration consisting of a three-level line-implicit multigrid cycle. The typical simulation at this level of resolution requires approximately 40 minutes of wall clock time per rotor revolution.



(a) Blade Deformation (C_p contours)

(b) Blade-1 tip displacement vs time

Figure 6. HART2 blade deformation

The effect of the CFD/CSD aeroelastic coupling is clearly demonstrated in Figure 6(a), which compares the

deformed blade shape and its corresponding C_p surface contours from the coupled simulation with that from the rigid blade simulation. From the figure it is noted that all four blades show different deformation characteristics due to corresponding different aerodynamic environment they experience in forward flight. The blade attains the largest flap displacement at azimuth, $\psi = 180^\circ$ and the smallest flap displacement at azimuth, $\psi = 0^\circ$. For both the flexible and rigid blades, the pressure contours demonstrate that the advancing side blades experience higher loading (larger pressure gradients near the rotor tips) than the retreating side blades. The flexible blade tip vertical displacement time history shown in Figure 6(b) demonstrates $1/rev$ behavior of blade flapping, as expected.

1. Geometry Parameterization

In order to obtain sensitivities with respect to a set of shape parameters that are well suited for design optimization purposes, a baseline blade is constructed by stacking 11 airfoil section along the span. Each airfoil contains 10 Hicks-Henne bump functions, 5 on the upper surface, and 5 on the lower surface, that can be used to modify the airfoil shape. Additionally, the twist values of the blade at the root and tip airfoil sections are also used as design variables resulting in a total of 112 design variables. Figure 7(a) provides an illustration of the baseline blade design setup. A high density structured mesh is generated about this blade geometry, which is then rotated and translated to match each individual blade in the CFD mesh, as shown in Figure 7(b). Interpolation patterns between each unstructured mesh surface point and the baseline structured mesh are determined in a preprocessing phase. These interpolation patterns are then used to interpolate shape changes from the baseline blade to all four blades in the CFD mesh (as determined by changes in the design variables) and to transfer sensitivities from the surface CFD mesh points to the design variables using the chain rule of differentiation.

In addition to the aforementioned design variables, inclusion of control parameters for pitching in forward flight introduces a new set of three design parameters, i.e. $\mathbf{D} = [\theta_0, \theta_{c1}, \theta_{s1}]$. Design variable θ_0 is known as rotor collective pitch angle and θ_{c1} and θ_{s1} are rotor cyclic pitch angles. This results in a total of 115 design variables. These sensitivities, in addition to the above mentioned geometric sensitivities, are transferred onto the CFD surface mesh points from the master blade shape through the interpolation method described earlier.

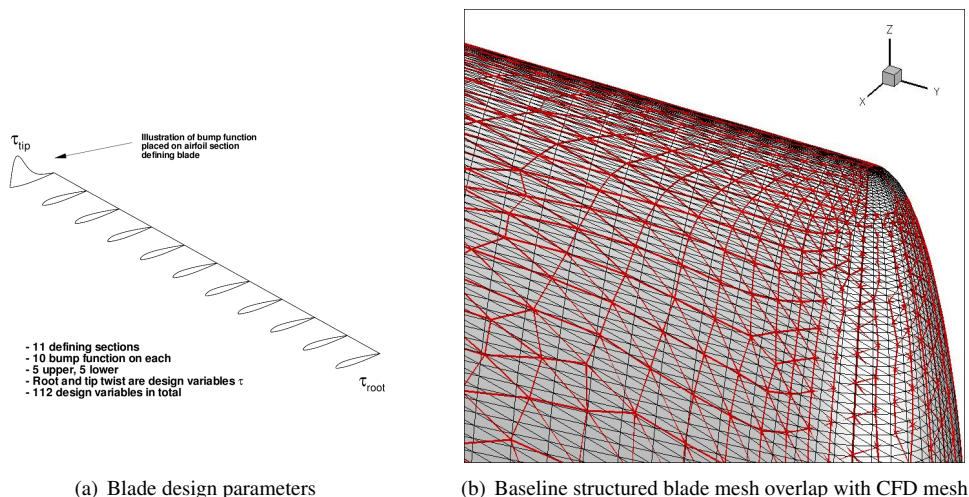


Figure 7. Illustration of (a) baseline blade with design parameters and (b) overlap in tip region between baseline blade structured mesh and CFD surface unstructured mesh.

2. Unsteady Objective Function Formulation

A time-integrated objective function based on the time variation of the thrust (C_T), torque (C_Q) and moment coefficients (C_{M_x} and C_{M_y}) is used for this test case. The goal of the optimization is to reduce the time-integrated torque coefficient while constraining the time-integrated thrust coefficient to the baseline rotor performance as well as constraining the moment coefficients along roll (C_{M_x}) and pitch (C_{M_y}) axes of the rotor to a trimmed value, i.e. zero average moment values. The objective function is based on the summation of the differences between a target and a computed objective

value at each time step n . Mathematically the global objective function is defined as:

$$L^g = L_{Shape} + L_{Trim} \quad (30)$$

$$L_{Shape} = \omega_1 \frac{1}{T} \sum_{n=1}^{n=N} \Delta t [\delta C_Q^n]^2 \quad (31)$$

$$L_{Trim} = \omega_2 \frac{1}{T} \sum_{n=1}^{n=N} \Delta t [\delta C_T^n]^2 + \omega_3 [\delta \bar{C}_{M_X}]^2 + \omega_4 [\delta \bar{C}_{M_Y}]^2 \quad (32)$$

$$\delta C_Q^n = (C_Q^n - C_{Q_{target}}^n) \quad (33)$$

$$\delta C_T^n = (C_T^n - \bar{C}_{T_{target}}) \quad (34)$$

$$\delta \bar{C}_{M_X} = \frac{1}{T} \sum_{n=1}^{n=N} \Delta t (C_{M_X}^n - C_{M_X_{target}}^n) \quad (35)$$

$$\delta \bar{C}_{M_Y} = \frac{1}{T} \sum_{n=1}^{n=N} \Delta t (C_{M_Y}^n - C_{M_Y_{target}}^n) \quad (36)$$

where the mean target thrust coefficient value is specified to the Hart2 baseline mean trim value of $\bar{C}_{T_{target}} = 0.0044$, and the target torque and moment values are set to zero. The weights ($\omega_i, i = [1, 2, 3, 4]$) are included to equalize the difference in orders-of-magnitude between the thrust, torque and moment coefficients. Use of pitch control parameters ($[\theta_0, \theta_{c_i}, \theta_{s_i}]$) as design variables and use of moment penalty terms ensure that the optimized rotor shape with optimized control parameters tend towards a final trimmed state when the rotor design cycles converge. Ideally, this penalty approach would be expected to be computationally more efficient than the use of a hard constraint formulation, since this latter approach would require the computation of multiple adjoint problems at each design cycle, as opposed to the single adjoint required in the current formulation. However, since the trim constraint may not be satisfied exactly in penalty approach, additional computational cost may be incurred in trimming the rotor (using a trim optimization step) before shape design is carried out. Given only few number of iterations required for trim optimization, the overall penalty approach can still be more efficient. Although, some references¹⁸ claimed to have achieved hard constraint design optimization solving multiple adjoints with only an incremental cost, it remains unclear if this additional cost using their suggested approach would be any less than that required in trim step.

B. Optimization of Helicopter Rotor in Forward Flight

For the optimization runs, the simulation is run for two full rotor revolutions, starting impulsively from rest in quiescent flow. The objective consists of the time-integrated torque with thrust and moment penalty, as described previously. However, the objective to be minimized is only integrated over the last rotor revolution, in order to avoid the optimization process from focusing on start-up transients. The optimization problem is solved on a relatively coarse grid (2.3 million points) using a time step size of 2.0 degrees. The optimization procedure used is the L-BFGS-B bounded reduced Hessian algorithm.¹⁹ Each request by the optimization driver for a function and gradient value results in a single forward time-integration of the analysis solver and a single backward integration in time of the adjoint solver. A bound of $\pm 5\%$ chord for each defining airfoil section was set on the Hicks-Henne²⁰ bump functions, a bound of $\pm 1.0^\circ$ of twist was set on the root and tip twist definitions, and a bound of $\pm 5.0^\circ$ of pitch angle was set on all the pitch parameters (collective and cyclics). The optimizations were performed on the Yellowstone supercomputer at the NCAR-Wyoming Supercomputing Center (NWSC) with the simulations (analysis/adjoint) running in parallel on 1024 cores. Each time step in the analysis problem employed 6 coupling cycles. Each coupling cycle used 10 nonlinear iterations. A typical coupled functional gradient (analysis/adjoint) computation step requires approximately 70 minutes when run on 1024 cores.

The performance optimization consists of three main stages: i) Trim, ii) Shape/performance optimization and, iii) re-trim. The 'Trim' step involves trimming the rotor to a target wind tunnel rotor thrust value of $C_T = 0.0044$ and zero longitudinal and lateral moments ($C_{M_Y}, C_{M_X} = 0$). The objective function used in this step consists only of the thrust and moment terms or simply L_{Trim} in equation (30) and the objective minimization is performed using only the three pitch parameters, namely, collective and cyclics as design variables.

In the second stage, blade shape optimization is performed by including the performance objective, i.e., C_Q term into the time dependent objective function to be minimized. Appropriate weights ($\omega_i, i = 1, 2, 3, 4$) are used to maintain the rotor trim state through a penalty function while the blade shape is optimized to obtain minimum rotor power. In total 115 design parameters were used in this stage, including 112 blade shape parameters and 3 pitch control

parameters. However, even after optimization convergence in stage two, the exact trim state is not maintained. This is because the trim objective components in the objective function were used only as weak constraint terms, i.e. as penalty terms and not as hard constraints. Therefore, the last stage involves trimming the rotor back to the target thrust and moment values, once again using only the three control pitch parameters as design inputs. This stage is otherwise referred to as the 're-trim' stage in this paper.

Figure 8 illustrates the trim optimization of the flexible HART2 rotor to the trim target thrust ($C_T = 0.0044$) and moment values ($C_{M_x}, C_{M_y} = 0$). Fig. 9(a) shows that the trim objective gradient drops by more than 2 orders of magnitude, while the objective functional achieves a minimum over 19 design iterations. Figure 9(b) further shows consistent convergence of all the three pitch parameters.




Figure 8. HART2 flexible blade trim




Figure 9. HART2 flexible blade trim convergence

Figure 10 summarizes the flexible blade shape optimization as well as re-trimmed performance plots. Figure 10(a) shows the baseline trimmed flexible rotor is shape optimized to achieve a significant power reduction of approximately 5.0% with a thrust loss of approximately 2.6%. Similar to the rigid blade re-trim, for the flexible blade some of this gain is lost upon re-trim. However, the final shape optimized and re-trimmed rotor achieves an overall 3.1% reduction in power compared to the baseline blade. Figure 10(c) shows one and half orders of magnitude gradient drop and consistent functional convergence over 21 design iterations for the blade shape optimization stage. Figure 11 shows the shape optimized blade sections at 9 stations. The optimized blade shapes show similar trends to those observed in the rigid blade shape optimization, i.e. thicker inboard and thinner outboard stations.

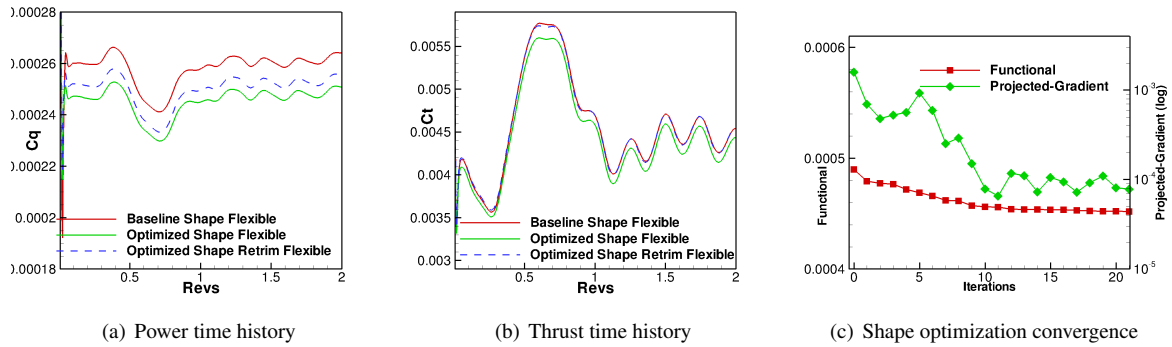


Figure 10. HART2 flexible blade shape optimization and re-trim

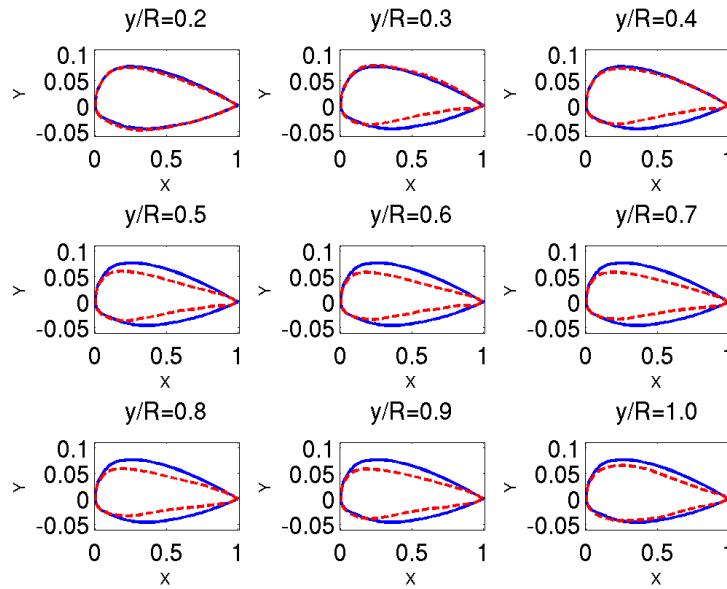


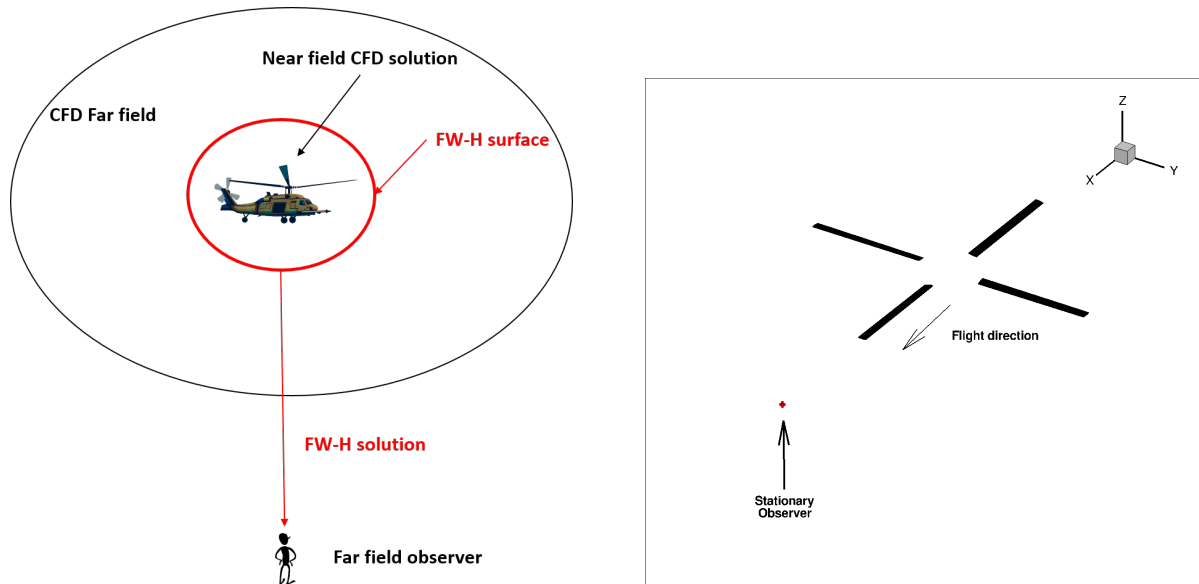
Figure 11. HART2 flexible blade optimized blade sections; blue solid-baseline, red dashed-optimized

IV. Acoustic formulation

Despite the continuous increase in computational resources, numerical simulations that resolve wave propagation from the nearfield to a farfield observer are still impractical, hence a viable approach to predicting farfield noise level is the use of hybrid methods that have been developed over the years²¹ and are now fairly well established.

In hybrid methods the finely resolved nearfield flow time history is used as input to an acoustic formulation that predicts the noise radiated to a given observer. The acoustic formulations are often based on Lighthill's acoustic analogy, in particular the Ffowcs Williams-Hawkings (FW-H) approach, which is the one used in this work.²²

Here the NSU3D CFD flow solver provides the nearfield flow time history to a newly developed FW-H acoustic integration module that propagates the acoustic pressure at a farfield observer. A schematic of the hybrid approach is shown in Figure 12(a). The location of the FW-H acoustic integration surface depends on the noise problem being investigated. When significant non-linear flow effects are present, as in the case of high speed helicopter forward flight, an off-body permeable integration surface should be used.^{21,23} Noise prediction with an off-body permeable surface is only accurate when the flow field between the body and the integration surface is finely resolved, resulting in a significant increase in the computational cost of the CFD time-integration process. However, the case considered in this work consists of a HART-II¹⁷ four-bladed rotor in forward flight with a freestream mach number 0.095 and a tip Mach number of 0.638. At these conditions nonlinear flow effects are expected to be negligible so that an on-body impermeable integration surface can be used. Figure 12(b) shows the FW-H acoustic integration surface and



(a) Schematic of the hybrid approach for the aeroacoustic computation

(b) Acoustic integration surface and observer location

Figure 12. Schematics of the hybrid acoustic approach (a) and acoustic problem formulation (b): the observer is stationary in the plane of the rotor two radii from the rotor hub at $\psi = 180$ deg.

the observer location used for validation purposes. The acoustic surface coincides with the CFD rotor surface mesh and every node on the surface is an acoustic source that produces an acoustic pressure at the observer location via the FW-H integration process.

1. Acoustic Analysis Formulation: the FW-H equation

The FW-H equation can be expressed in differential form as²¹

$$\left(\frac{\partial^2}{\partial t^2} - c_0^2 \frac{\partial^2}{\partial x_i \partial x_j} \right) (H(f)\rho') = -\frac{\partial F_i \delta(f)}{\partial x_i} + \frac{\partial Q \delta(f)}{\partial t} \quad (37)$$

where

$$Q = (\rho_o v_i + \rho (u_i - v_i)) \frac{\partial f}{\partial x_i} \quad (38)$$

and

$$F_i = (p\delta_{ij} + \rho u_i (u_j - v_j)) \frac{\partial f}{\partial x_j} \quad (39)$$

and the Lighthill's stress tensor, the quadrupole term, has been omitted since it is not used in this work. Equation (38) gives rise to an unsteady monopole-type contribution that can be associated with mass addition, while the dipole term, equation (39), involves an unsteady force. The function $f(x_i, t) = 0$ defines the surface of integration outside of which the solution is sought. Total density and pressure are ρ and p respectively, the fluid velocities are u_i , while v_i are the surface velocities and c_o is the freestream speed of sound. The prime denotes perturbation relative to the freestream which itself is denoted with the subscript o . x_i and t are Cartesian coordinates and time respectively. $H(f)$ is the Heaviside function while $\delta(f)$ is the Dirac function. A time domain solution to equation (37) can be obtained from the derivation of Farassat²⁴ using the variables

$$\begin{aligned} U_i &= \left(1 - \frac{\rho}{\rho_o}\right) v_i + \frac{\rho u_i}{\rho_o} \\ L_i &= p' n_j + \rho u_i (u_n - v_n) \end{aligned} \quad (40)$$

as proposed by Di Francescantonio,²⁵ where $p' = p - p_o$, $u_n = u_i n_i$, $v_n = v_i n_i$ and n_i is the unit normal of the acoustic surface. The integral solution is then given, neglecting the quadrupole term, by equation (41)

$$\begin{aligned} 4\pi p'_T(\mathbf{y}, t) &= \int_{f=0} \left[\frac{\rho_o (\dot{U}_n + U_{\dot{n}})}{r(1 - M_r)^2} \right]_{ret} dS \\ &+ \int_{f=0} \left[\frac{\rho_o U_n K}{r^2(1 - M_r)^3} \right]_{ret} dS \\ 4\pi p'_L &= \frac{1}{c_o} \int_{f=0} \left[\frac{\dot{L}_r}{r(1 - M_r)^2} \right]_{ret} dS \\ &+ \int_{f=0} \left[\frac{L_r - L_M}{r^2(1 - M_r)^2} \right]_{ret} dS \\ &+ \frac{1}{c_o} \int_{f=0} \left[\frac{L_r K}{r^2(1 - M_r)^3} \right]_{ret} dS \\ 4\pi p'(\mathbf{y}, t) &= 4\pi p'_T(\mathbf{y}, t) + 4\pi p'_L(\mathbf{y}, t) \end{aligned} \quad (41)$$

where \mathbf{y} is the observer location, t is the observer time, r is the distance between the source \mathbf{x} and the observer \mathbf{y} , M is the surface Mach number and

$$\begin{aligned} U_n &= U_i n_i \\ M_r &= M_i r_i \\ L_r &= L_i r_i \\ L_M &= L_i M_i \\ K &= r \dot{M}_r + c_o (M_r - M^2) \end{aligned} \quad (42)$$

Equation (41) requires the evaluation of the integrals at the emission or retarded time τ : for a given observer location \mathbf{y} and time t the retarded time must be computed via the solution of the nonlinear equation (43)

$$\tau = t - r(\mathbf{x}(\tau), \mathbf{y}(t)) / c_o \quad (43)$$

where $r(\mathbf{x}(\tau), \mathbf{y}(t))$ is the distance between the source \mathbf{x} at the emission time τ and the observer \mathbf{y} at the observer time t . This approach has been followed in most acoustic-analogy-based codes.²¹ However, by regarding the source time as the primary time one can choose the source time for an acoustic source and determine when the signal will reach the observer using equation (44)

$$t = \tau + r(\mathbf{x}(\tau), \mathbf{y}(t)) / c_o \quad (44)$$

where τ is again the source time and t is the reception time, i.e. the time at which the acoustic disturbance reaches the observer. Equation (44), a rearrangement of equation (43), is easier to solve than its retarded time counterpart as the observer motion is usually simpler than the motion of the acoustic integration surface. For each acoustic source a sequence of uniformly spaced source times leads to a sequence of unequally spaced observer times as each source has a different source-to-observer distance, hence the acoustic pressure time history for each source must be interpolated at the desired observer time to determine the final acoustic pressure at the observer location. A comparison of the retarded time and the source-time-dominant algorithm is given in.²⁶

In this work we follow this so-called source-time-dominant algorithm as it allows a seamless integration of the acoustic module with the CFD solver. The source time is the CFD time and the integrals in equation (41) are evaluated at every time step during the CFD time integration process. The acoustic integration surface corresponds to the rotor's unstructured surface mesh and each node of the CFD grid corresponds to an acoustic source. All the time derivative terms in equation (41) are approximated with a BDF2 time discretization for consistency with the CFD time integration with the exception of the source acceleration term that is approximated with a second order accurate central difference scheme. The final observer pressure time history is built via linear interpolation of each source time history at the desired observer time after the CFD integration process is completed. The observer acoustic pressure time history can finally be processed to compute the aeroacoustic objective to be minimized. The aeroacoustic objective used in this work is the root mean square of the acoustic pressure time history

$$L_{FWH} = p'_{RMS} = \sqrt{\frac{\sum_{i=1}^{N_{sample}} p'^2(\mathbf{D})}{N_{sample}}} \quad (45)$$

where N_{sample} is the number of samples in the observer acoustic pressure time history and \mathbf{D} is the vector of design variables. The current implementation of the FW-H integral equation has been validated against the PSU-WOPWOP²⁶ acoustic code for the HART-II flexible rotor in trimmed forward flight as shown in Figure 13. The flow conditions are described in section III. The CFD simulation has been carried out for two rotor revolutions on a mesh consisting of approximately 2.32 millions nodes with a 2-degree timestep, 6 CFD/CSD coupling iterations per time step and a stationary in-plane observer located two radii from the rotor hub at an azimuthal angle $\psi = 180$ deg as shown in Figure 12(b) and described in Table 1. Two different observer time windows are shown in Figure 13: an extended time window, Figure 13(a), and the observer time window targeted during the optimization, Figure 13(b). In both cases agreement between the new FW-H implementation and the legacy PSU-WOPWOP code²⁶ is excellent. In Figure 14 we investigate the effect that blade flexibility has on the predicted total acoustic pressure: the positive pressure peak is more pronounced in the flexible case while the negative pressure peak is significantly reduced as a consequence of blade deformation. A comparison of the rigid and flexible blades is shown in Figure 15 together with a surface contour plot of the pressure coefficient. All four blades show different deformation characteristics due to the corresponding different aerodynamic environment they experience during trimmed forward flight: the largest flap displacement is attained at $\psi = 180$ deg while the smallest happens at $\psi = 0$ deg.


Table 1. Observer location for the acoustic objective function with respect to the rotor hub, R being the rotor radius

\mathbf{x}	\mathbf{y}	\mathbf{z}
2R	0	0

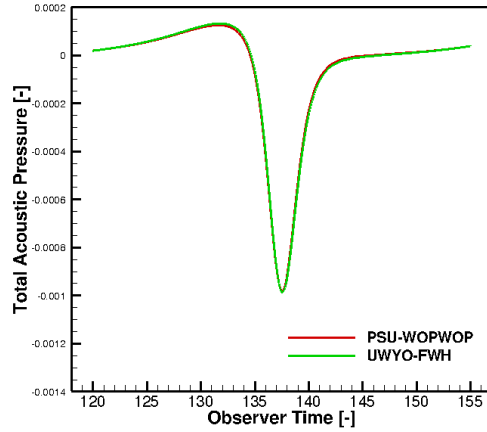
V. Sensitivity Analysis for Aero-structural-acoustic Problem

A. Sensitivity formulation for the integral FW-H equation

Computing the sensitivities of the coupled flexible aeroacoustic problem requires the linearization of the acoustic module with respect to the design variables. The sensitivity formulation for the current implementation of the integral FW-H equation closely mimics the linearization of the CFD code as detailed in references.^{1,2} First the tangent linearization of the acoustic code has been developed by exact hand-differentiation of the discretized integral FW-H equation. The tangent linearization has then been transposed and applied in reverse order to obtain the adjoint sensitivity with respect to the full vector of design variables.



(a) Validation of the current FW-H implementation with the PSU-WOPWOP acoustic code over an extended observer time window.



(b) Comparison between the current FW-H implementation and PSU-WOPWOP over the observer time window used for optimization purposes

Figure 13. Validation of the current FW-H implementation over an extended (a) and the optimization (b) time window

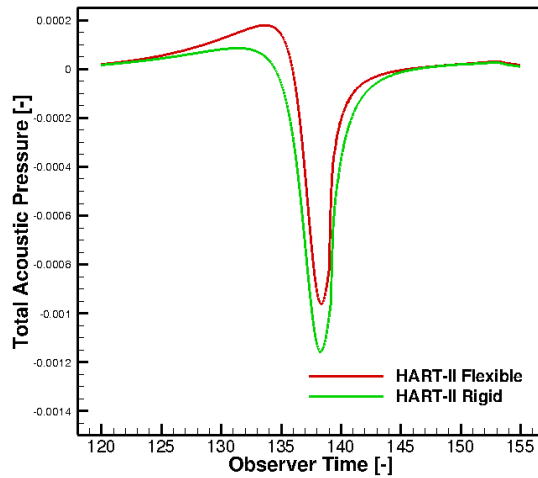


Figure 14. Acoustic pressure at the observer for the baseline rigid and flexible HART-II rotor.




Figure 15. Comparison between the rigid and the flexible HART-II rotor.

Tangent sensitivity formulation

The tangent linearization of the acoustic objective function allows the computation of the objective function sensitivity with respect to a single design variable. We express the acoustic pressure at the observer location and time as

$$p'(\mathbf{y}, t, \mathbf{D}) = \mathbf{FWH}(\mathbf{U}(\mathbf{D}), \mathbf{x}(\mathbf{D})) \quad (46)$$

where \mathbf{D} is the vector of design variables and $\mathbf{FWH}(\mathbf{U}(\mathbf{D}), \mathbf{x}(\mathbf{D}))$ represents all the discrete operations necessary to evaluate equation (41) numerically. In this expression $\mathbf{U}(\mathbf{D})$ and $\mathbf{x}(\mathbf{D})$ are the flow and mesh solution at every time step of the time-integration process at the acoustic integration surface, after the aeroelastic coupling has converged. The acoustic pressure tangent sensitivity time history can be expressed as

$$\frac{dp'(\mathbf{y}, t, \mathbf{D})}{dD} = \sum_n \frac{\partial \mathbf{FWH}}{\partial U_{FWH}^n} \frac{dU_{FWH}^n}{dD} + \frac{\partial \mathbf{FWH}}{\partial x_{FWH}^n} \frac{dx_{FWH}^n}{dD} \quad (47)$$

so that the tangent linearization of the acoustic objective function p'_{RMS} defined in equation (45) becomes:

$$\begin{aligned} \frac{dL_{FWH}}{dD} &= \sum_n \frac{\partial L_{FWH}}{\partial U_{FWH}^n} \frac{dU_{FWH}^n}{dD} + \frac{\partial L_{FWH}}{\partial x_{FWH}^n} \frac{dx_{FWH}^n}{dD} \\ &= \frac{\partial p'_{RMS}}{\partial p'} \left[\sum_n \frac{\partial \mathbf{FWH}}{\partial U_{FWH}^n} \frac{dU_{FWH}^n}{dD} + \frac{\partial \mathbf{FWH}}{\partial x_{FWH}^n} \frac{dx_{FWH}^n}{dD} \right] \\ &= \frac{\partial p'_{RMS}}{\partial p'} \frac{dp'}{dD} \end{aligned} \quad (48)$$

In equation (48) the terms $\frac{dU_{FWH}^n}{dD}$ and $\frac{dx_{FWH}^n}{dD}$ are the aeroelastically-converged flow and mesh tangent sensitivities at every time step evaluated at the acoustic integration surface that are computed via the forward time integration described by equation (29). The terms $\frac{\partial \mathbf{FWH}}{\partial U_{FWH}^n}$ and $\frac{\partial \mathbf{FWH}}{\partial x_{FWH}^n}$ are the tangent linearizations of equation (41). The tangent acoustic problem proceeds in analogy with the acoustic analysis problem. The tangent flow solution and the tangent acoustic solution are carried out simultaneously: the aeroelastically-converged tangent flow solution is used to assemble the tangent sensitivities of all the terms in equation (42) which in turn are used to evaluate the tangent sensitivity of the integral equation (41) at every timestep, resulting in an unequally spaced acoustic pressure sensitivity at the observer location for every acoustic source. The final observer acoustic pressure sensitivity time history is built via linearization of the time interpolation algorithm at the end of the time-integration process allowing for the computation of the acoustic objective function sensitivity, equation (48), with respect to one design variable.

Adjoint sensitivity formulation

The adjoint sensitivity can be derived by transposing the tangent sensitivity formulation. Transposing equation (48) yields

$$\begin{aligned} \frac{dL_{FWH}}{dD}{}^T &= \left[\sum_n \frac{d\mathbf{U}_{FWH}^n}{dD}{}^T \frac{\partial FWH}{\partial \mathbf{U}_{FWH}^n}{}^T + \frac{dx_{FWH}^n}{dD}{}^T \frac{\partial FWH}{\partial x_{FWH}^n}{}^T \right] \frac{\partial p'_{RMS}}{\partial p}{}^T \\ &= \sum_n \frac{\partial U^n}{\partial D}{}^T \frac{\partial L_{FWH}}{\partial U^n}{}^T + \frac{\partial x^n}{\partial D}{}^T \frac{\partial L_{FWH}}{\partial x^n}{}^T \end{aligned} \quad (49)$$

The terms $\frac{\partial L_{FWH}}{\partial U^n}{}^T = \frac{\partial FWH}{\partial \mathbf{U}_{FWH}^n}{}^T \frac{\partial p'_{RMS}}{\partial p}{}^T$ and $\frac{\partial L_{FWH}}{\partial x^n}{}^T = \frac{\partial FWH}{\partial x_{FWH}^n}{}^T \frac{\partial p'_{RMS}}{\partial p}{}^T$ represent the right-hand-side of equation (29) for the case of the acoustic objective function and drive the adjoint backward time-integration described in equation (29).

The term $\frac{\partial p'_{RMS}}{\partial p}{}^T$ is the reverse linearization of the observer time interpolation process that needs to be evaluated before the adjoint time-integration can start while the term $\frac{\partial FWH}{\partial \mathbf{U}_{FWH}^n}{}^T$ and $\frac{\partial FWH}{\partial x_{FWH}^n}{}^T$ correspond to the reverse linearization of the acoustic integrals in equation (41). At the end of the backward time-integration process, the full sensitivity vector of the acoustic objective function is recovered.

VI. Aero-structural-acoustic Optimization

The goal of the present work is to reduce the noise signature of the HART-II rotor in trimmed forward flight by means of gradient-based optimization techniques that exploit the multidisciplinary flexible aeroacoustic adjoint method developed so far. As a preliminary step, the flexible HART-II rotor in forward flight needs to be trimmed to establish a starting point for all subsequent optimizations. The trim problem is formulated as an optimization problem described by equation (50)

$$\begin{aligned} &\min L_{THRUST} \\ &\text{subject to} \\ &L_{LATERAL} = 0 \\ &\text{w.r.t. } \mathbf{D}_{pitch} \end{aligned} \quad (50)$$

where the vector of design variables \mathbf{D}_{pitch} is made up of one collective and two cyclics and is described in the **Control input facility** section. One design cycle for this optimization corresponds to one unsteady aeroelastic flow solution and two aeroelastic adjoint solutions, one for the objective and one for the constraint function in equation (50).

Once the Hart-II baseline rotor has been trimmed, it is used as the starting point for a multidisciplinary flexible rotor aeroacoustic optimization which aims at reducing the torque required for trimmed forward flight with a significantly lower noise signature according to equation (51).

$$\begin{aligned} &\min L_{TORQUE} \\ &\text{subject to} \\ &L_{THRUST} = 0 \\ &L_{LATERAL} = 0 \\ &p'_{RMS} = p'_{RMS_{TARGET}} \\ &\text{w.r.t. } \mathbf{D} \end{aligned} \quad (51)$$

The value $p'_{RMS_{TARGET}}$ has been chosen to guarantee a 2dB OSPL noise signature reduction of the baseline rotor. In this case one design cycle consists of one aeroelastic unsteady flow solution, and four aeroelastic unsteady adjoint solutions.

The functionals L_{THRUST} , $L_{LATERAL}$ and L_{TORQUE} are defined as

$$\begin{aligned}
 L_{THRUST} &= \frac{1}{N} \left(\sum_{i=1}^N (C_T^i - C_{T_{AVERAGE}}^i) \right)^2 \\
 L_{LATERAL} &= \frac{1}{N} \left[\left(\sum_{i=1}^N C_{M_x}^i \right)^2 + \left(\sum_{i=1}^N C_{M_y}^i \right)^2 \right] \\
 L_{TORQUE} &= \frac{1}{N} \sum_{i=1}^N (C_Q^i)^2
 \end{aligned} \tag{52}$$

where $C_{T_{AVERAGE}}^i = 0.0044$ is the target thrust coefficient for the baseline Hart-II rotor. The flow conditions are those described previously and the simulations are run for two rotor revolutions using a 2-degree timestep and a computational mesh that consists of approximately 2.32 million nodes. Six CFD/CSD coupling iterations are performed at every time step. The aerodynamic functionals in equation (52) are accumulated only over the second revolution to prevent the optimization algorithm from focusing on the initial transient. The acoustic integration is performed only over the second revolution and the acoustic objective function is evaluated over the observer time window shown in Figure 13(b). For all optimizations the observer is stationary and is placed in front of the rotor as shown in Figure 12(b) and detailed in Table 1.

In this optimization test case we consider only shape design variables that affect the outer mold line of the CFD geometry without modifying the structural model. The influence of these aerodynamic design variables affects the aeroelastic loop only through changes in the aerodynamic forces applied to the structural model. The baseline blade is constructed by stacking 9 airfoils along the span. Each airfoil is parameterized with 10 Hicks-Henne bump functions, 5 for the upper surface and 5 for the lower surface respectively. Blade twist varies linearly between root and tip, while one collective and two cyclics allow the trimming of the rotor, for a total of 95 design variables in the vector of design variables \mathbf{D} in equation (51).

One unsteady flow solution takes approximately 90 minutes of wall-clock time on 1024 cores with each unsteady adjoint solution costing approximately 45 minutes of wall clock time, so that the cost of one design cycle spans from three hours for the trim optimization, equation (50), to five hours for the aeroacoustic optimization, equation (51). Approximately 650 Gb of disk space are stored to disk during the analysis problem and read in by the flexible aeroacoustic adjoint procedure. The SNOPT²⁷ sequential quadratic programming algorithm is used to drive the constrained optimization. The computational cost of the optimization in equation (51) is approximately 96 hours of wall-clock time on 1024 cores on the Yellowstone supercomputer at the NCAR-Wyoming Supercomputing Center (NWSC).


We give convergence results for the optimization problems as a function of nonlinear optimization iterations and of design cycles. While the number of design cycles is a measure of the optimization wall-clock time, the number of nonlinear iterations is the number of iterations used by SNOPT to solve the optimization problem; convergence of the optimization is assessed through feasibility, which is a measure of the constraint violation, and optimality, which is a measure of the satisfaction of the Karush - Kuhn - Tucker condition as described in.²⁷

A. Trim results

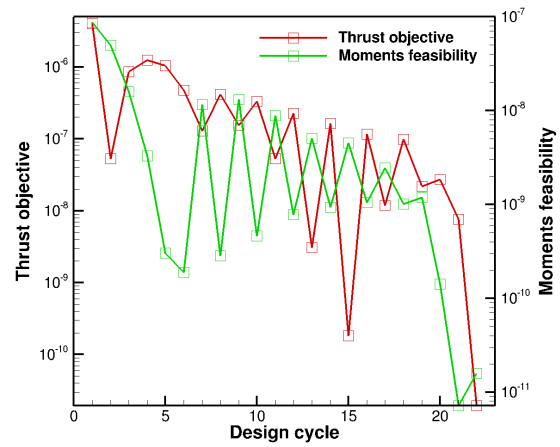
The trim optimization problem is described by equation (50). The optimization convergence is shown in Figure 16. After 15 nonlinear iterations and 22 design cycles the feasibility and optimality of the optimization problem are reduced by approximately 3 orders of magnitude and the baseline flexible HART-II rotor is trimmed as further confirmed by thrust and lateral moments time histories shown in Figure 17.

B. Acoustically-constrained torque minimization

The trimmed HART-II rotor serves as the initial design for the aeroacoustic optimization described by equation (51). In as few as 5 nonlinear iterations feasibility is reduced by more than 2 orders of magnitude while optimality is reduced by approximately 9 orders of magnitude, as shown in Figure 18(a). After 10 design cycles the acoustic constraint is satisfied resulting in a rotor with a 2dB OSPL reduction compared to the baseline, while the required torque is reduced by 2.5% as shown in Figure 18(b). Satisfaction of the trim constraint is shown in Figure 18(c) and confirmed by the thrust and lateral moment time histories shown in Figure 19. A comparison of the torque time histories is shown in Figure 20 highlighting the 2.5% torque reduction achieved for this optimization. The thickness, loading and total acoustic pressure time histories are shown in Figure 21. A comparison of the airfoil shapes for the baseline and the



(a) Feasibility and optimality of the trim problem as a function of nonlinear optimization iterations



(b) Convergence of the thrust objective function and moment constraint as a function of design cycles.

Figure 16. Convergence of the trim optimization problem

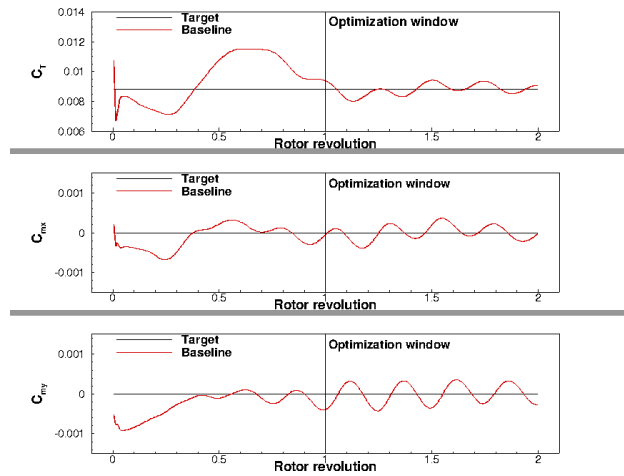



Figure 17. Thrust and moments time histories for the baseline HART-II rotor



(a) Feasibility and optimality for the torque-constrained aeroacoustic optimization problem as a function of nonlinear iterations (b) Convergence of the acoustic objective function and the torque constraint as a function of design cycles. (c) Convergence of the trim feasibility as a function of design cycles.

Figure 18. Convergence of the acoustically constrained torque optimization

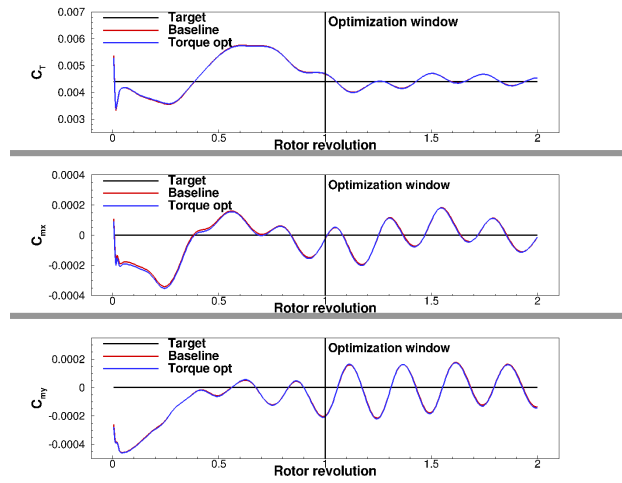


Figure 19. Thrust and moments time histories for the baseline rotor and the acoustically constrained torque optimized rotor




Figure 20. Torque time histories for the baseline rotor and the acoustically constrained torque optimized rotor




Figure 21. Thickness, loading and total acoustic pressures at the observer for the acoustically constrained torque optimized rotor.

optimized rotor is shown in Figure 22 highlighting that the acoustically constrained torque minimization results in reduced blade thickness particularly at the outboard sections of the rotor.

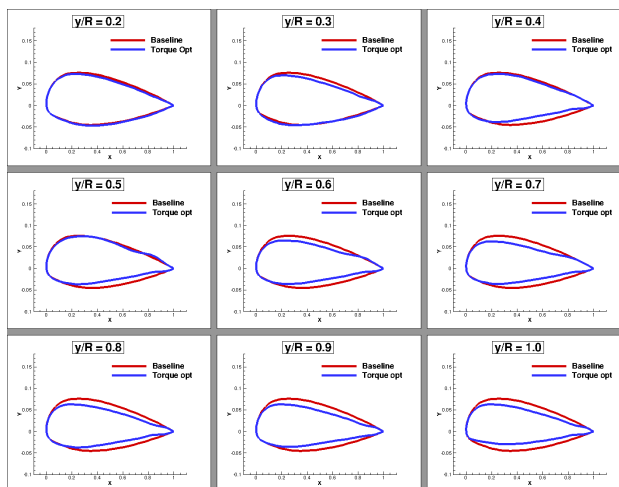



Figure 22. Acoustically constrained torque optimized blade shapes

Previous noise mitigation studies²⁸ have shown that optimizing for one observer location can result in higher noise levels at different observer locations. A directivity study for the optimized rotor reveals that the current design is capable of reducing noise at in-plane observer locations that have not been targeted by the optimization as shown for two different observers placed at $\psi = 135$ and $\psi = 315$ degrees as shown in Figure 23. Additionally, in reference²⁹ it is shown that even though the optimizations targeted only the second rotor revolution, the optimized rotor performance is improved over multiple rotor revolutions. Thus, the resulting optimal rotors can be used as initial guesses for more expensive optimizations that target multiple rotor revolutions.

VII. Aeroacoustic optimization of the rigid UH-60 rotor

The aeroacoustic optimization of the rigid UH-60 rotor in trimmed forward flight is presented in this section. Following an approach similar to the HART-II aeroacoustic optimization from the previous section, the UH-60 rotor is first trimmed to establish a baseline acoustic signature. Then, an unconstrained aeroacoustic optimization aims at reducing the noise signature at the observer in trimmed forward flight.

The baseline UH-60 geometry has been reconstructed from available literature.^{30,31} The flight condition for the UH-60



(a) Baseline and optimized total acoustic time history at $\psi = 135$ degrees

(b) Baseline and optimized total acoustic time history at $\psi = 315$ degrees

Figure 23. Directivity study for the baseline and optimized geometry

rotor corresponds to counter 8534 in the UH-60 airloads catalog³² and is detailed in Table 2.

Table 2. Operating condition for the UH-60 rotor in trimmed forward flight

Freestream Mach number (velocity)	$M_\infty = 0.236$
Tip Mach Number	$M_{tip} = 0.642$
Advance ratio	$\mu = 0.368$
Rotational speed	$\Omega = 258 \text{ RPM}$
Shaft angle	$\alpha_{shaft} = -7.31 \text{ deg}$
Thrust	16602 lb
Pitching moment	-5025 lb-ft
Rolling moment	5350 lb-ft

The aeroacoustic observer lies in the rotor plane, 2 radii from the rotor hub and travels with the rotor.

The trim optimization problem is the same as in equation (50), but the objective $L_{LATERAL}$ is now modified to account for non-zero pitching and rolling moments as

$$L_{LATERAL} = \frac{1}{N} \left[\left(\sum_{i=1}^N (C_{M_x}^i - C_{M_x-average}^i) \right)^2 + \left(\sum_{i=1}^N (C_{M_y}^i - C_{M_y-average}^i) \right)^2 \right] \quad (53)$$

The trim optimization problem is solved with respect to the same design variables \mathbf{D}_{pitch} described in the **Control input facility** section.

After the baseline UH-60 rotor has been trimmed an unconstrained aeroacoustic optimization is performed to minimize the noise signature at the specified observer location. For this unconstrained optimization the problem statement is given by

$$\begin{aligned} \min L_{UNC} \\ \text{w.r.t. } \mathbf{D} \end{aligned} \quad (54)$$

and the objective function L_{UNC} is defined as

$$L_{UNC} = L_{FWH} + 10(L_{THRUST} + 10L_{LATERAL}) \quad (55)$$

where L_{FWH} is defined in equation (45) and L_{THRUST} and $L_{LATERAL}$ are defined in equations (52) and (53) respectively. For this unconstrained optimization a total of 127 design variables has been used. The blade is built by stacking 11 airfoils along the span with 10 Hicks-Henne bump functions per airfoil. Twist is defined at 11 spanwise sections and linearly interpolated between two consecutive sections. Tip taper, sweep and droop allow for the modification of the tip planform shape. Finally, one collective and two cyclics allow for the trimming of the rotor.

At the end of the unconstrained optimization, since the penalty function formulation in equation (55) doesn't satisfy the trim constraint exactly, the optimized rotor needs to be retrimmed to the thrust and lateral moments values in Table 2 according to equation (50). This final optimization is once again solved with respect to the three design variables in \mathbf{D}_{pitch} .

The CFD mesh used in this work, shown in Figure 24, consists of 5.1 million nodes and has been built according to the AIAA Drag Prediction Workshop guidelines.³³ The rotor is simulated for two rotor revolutions with a 2-degree

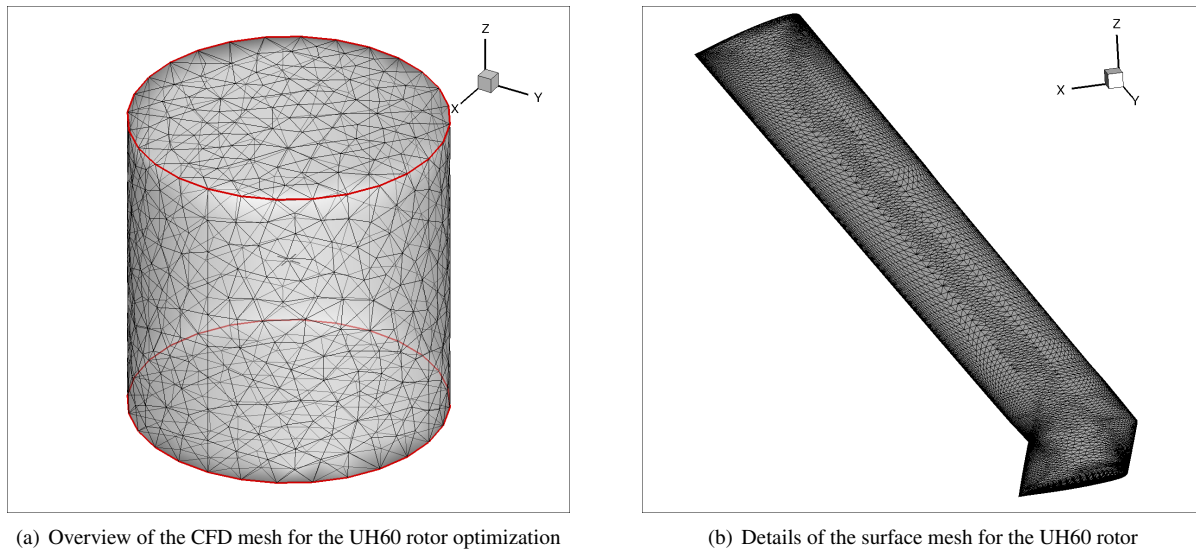


Figure 24. CFD mesh for the UH60 rotor. The mesh consists of 5.1 Mln nodes and is made up of prisms, pyramids and tetrahedra

timestep and the objective functional is accumulated only over the second rotor revolution.

For the unconstrained optimization problem in equation (54) 1 design cycle corresponds to 1 unsteady flow solution and one unsteady adjoint solution. The unsteady flow solution takes approximately 60 minutes of wall-clock time on 3008 cores with 2Tb of disk storage, while 1 unsteady adjoint solution takes approximately 30 minutes.

The SNOPT optimization algorithm is used to drive both the trim optimization problem and the unconstrained aeroacoustic optimization problem. The unconstrained optimization problem, including the final retrim, took approximately 66 hours of wall-clock time on 3008 cores on the Yellowstone supercomputer of the NCAR Wyoming Supercomputing Center (NWSC).

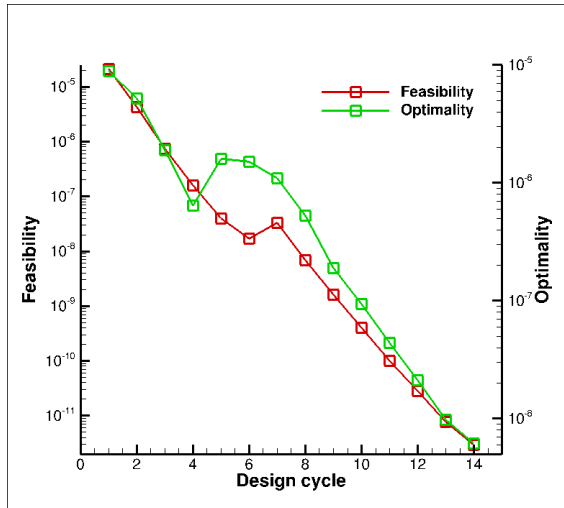
UH-60 trim results

Convergence results for the trim optimization problem of the baseline UH-60 rigid rotor are shown in Figure 25. In 14 design cycles optimality, as computed by SNOPT, is reduced by 3 orders of magnitude while feasibility is reduced by almost 7 orders of magnitude.

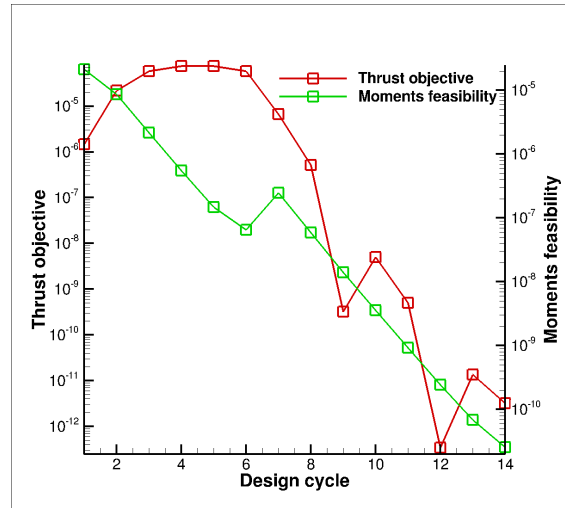
This optimization establishes a baseline result for the rigid UH-60 rotor that will be used to assess the improvement gained with the unconstrained aeroacoustic optimization in equation (55).

Unconstrained aeroacoustic optimization with retrim

The noise signature of the trimmed UH-60 rotor is now minimized using the unconstrained minimization approach described in equation 54. At the end of the unconstrained optimization a final retrim of the optimized rotor guarantees



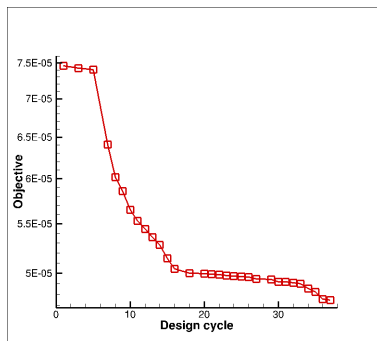
(a) Feasibility and optimality of the trim problem for the UH-60 rigid rotor as a function of design cycles




(b) Convergence of the thrust objective function and moment constraint as a function of design cycles.

Figure 25. Convergence of the trim optimization problem for the rigid UH-60 rotor

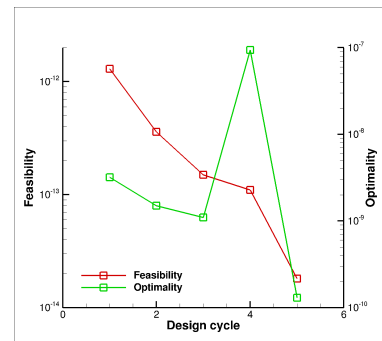
the exact satisfaction of the trim constraint. Convergence results for the unconstrained optimization problem and for the final retrim are shown in Figure 26. Figure 26(a) shows that after a significant initial gain the objective function in



(a) Convergence of the unconstrained objective function equation (55).



(b) Optimality for the unconstrained minimization problem in equation (54).



(c) Feasibility and optimality for the final retrim of the optimized rotor.

Figure 26. Convergence of the trim optimization problem for the rigid UH-60 rotor

equation (55) starts to settle down after approximately 20 design cycles, while Figure 26(b) shows that optimality is reduced by approximately 1 order of magnitude after 37 design cycles. Figure 26(c) shows feasibility and optimality for the final trim problem highlighting the satisfaction of the trim constraint for the aeroacoustically optimized rotor as further confirmed by the thrust and lateral moments time histories shown in Figure 27. The noise signatures for the baseline and optimized rotors for this moving-observer acoustic setup are shown in Figure 28. The unconstrained optimization yields a 3.9dB reduction in Overall Sound Pressure Level after retrim. The geometry of the optimized rotor is shown in Figure 29. Figure 29(a) shows thicker airfoils at the inboard sections and thinner airfoils at the outboard sections, together with tapering of the tip airfoil. Figure 29(b) shows the blades' planform shape, particularly in the tip area, highlighting both anhedral and forward sweep of the tip section as opposed to an aft-swept tip of the baseline UH-60 rotor.

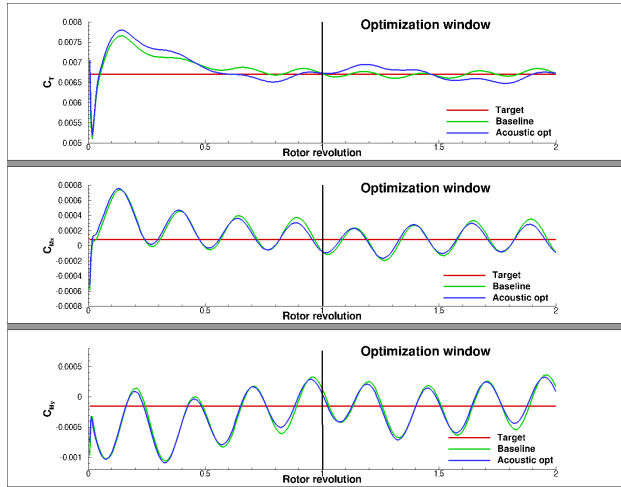


Figure 27. Thrust, pitching and rolling moment time histories for the rigid UH-60 and the optimized rotors.




Figure 28. Baseline and optimized acoustic pressure time histories

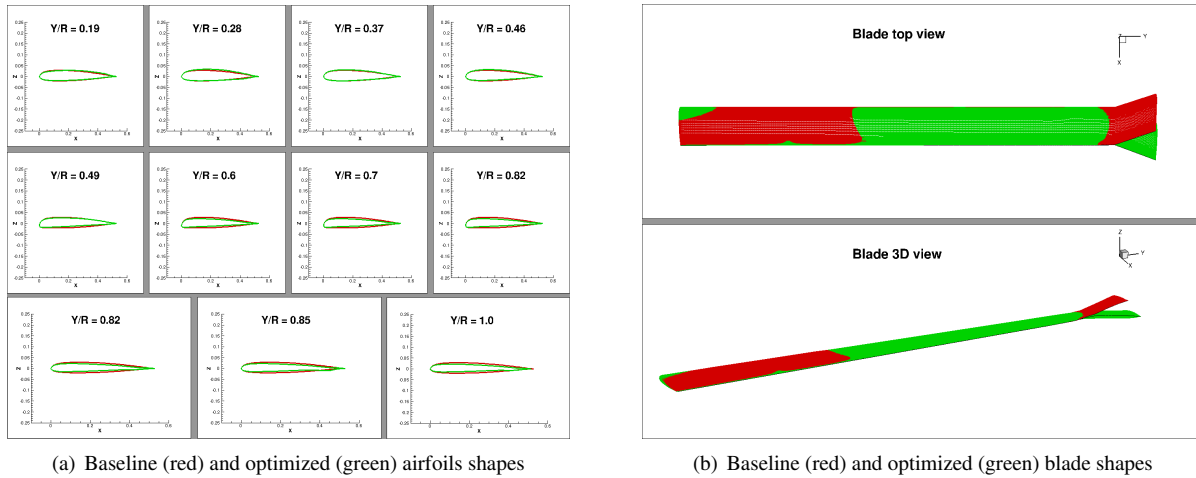


Figure 29. Baseline and optimized blade shapes

VIII. Incorporating Higher fidelity

The next logical step in the drive to higher fidelity aeroelastic effects for both fixed wing and rotary wing problems involves the incorporation of a high-fidelity finite-element structural model based on brick and shell elements, as is used in standard aerospace structural design practices. While such models are common in many commercial simulation tools, the need to obtain adjoint-derived sensitivities for the individual contributing disciplines results in the requirement of having complete source code access, since adjoint linearization is an inherently intrusive process. Therefore, the most practical approach consists of developing an in-house structural model that can be tailored for sensitivity analysis.

In previous work we have outlined the development of a brick and shell finite-element structural model that supports FEM models based on the most commonly used element types built in the NASTRAN and/or Abaqus commercial code frameworks.¹¹ This structural dynamics code has been tightly coupled to the NSU3D CFD solver and used for static and dynamic aeroelastic analysis problems. The discrete adjoint formulation of the structural model has also been implemented, and the adjoint derived sensitivities have been verified to machine precision using the complex step method for both static and dynamic cases. In Table 3 the adjoint-derived sensitivities of the time-integrated norm of displacement with respect to various structural design parameters are verified for the case of a cantilevered shell-based HART-2 blade model undergoing free vibration, showing agreement with the forward and complex step methods to near machine precision.

Table 3. Verification of time-dependent adjoint sensitivities for Hart-2 shell model

Design Variable : D	Adjoint	Tangent	Complex
Stiffness	-5.93308711120498e-3	-5.93308711120513e-3	-5.93308711122332e-3
Density	1.17662420188903e-3	1.17662420188909e-3	1.17662420188903e-3
Thickness	-1.65506902436382e-2	-1.65506902436360e-2	-1.65506902436435e-2
Applied Force	4.75646290931996e-3	4.75646290931999-3	4.75646290932002e-3

Figure 30 depicts results for a time-dependent structural optimization using a HART-2 blade model based on linear brick elements. As before, we clamp the root end of the blade and model it in free vibration, but we now take one step further and optimize the stiffness, density and cross-sectional dimension of the blade to conform the response to a target periodic response, with a chosen amplitude and frequency. The figure shows the time history of the tip displacement over one period of the target displacement for various design optimization cycles, as well as the convergence of the objective and objective gradient. The figure shows how, after seven design cycles, the time history of the blade tip displacement closely follows that of the target response, and the optimization procedure converges monotonically both in terms of the reduction of the objective and the gradients.




Figure 30. (a) Design evolution of time history of blade tip displacement with optimization steps versus target time history; (b) Convergence history of structural optimization problem in terms of objective and norm of sensitivity vector

Recently the structural dynamics capability has been extended to account for geometric nonlinearity, to model structures with displacements and rotations that are large compared to structural dimensions, such as blades on helicopter rotors and wind turbines. This has been achieved by implementing Green-Lagrange strain³⁴ as a deformation metric, a well-established approach in the field of nonlinear finite element modeling. Green-Lagrange strain effectively removes rigid body rotation modes from the deformation of a point in an element, and can be expressed

$$\boldsymbol{\varepsilon} = \frac{1}{2} [\nabla \mathbf{u} + \nabla \mathbf{u}^T + \nabla \mathbf{u} \cdot \nabla \mathbf{u}^T] \quad (56)$$

where $\nabla \mathbf{u} = \frac{\partial u_i}{\partial x_j}$. As can be seen, for small strain, the second nonlinear term vanishes and the usual linear strain form is recovered. The current nonlinear formulation has been verified by computing a large displacement problem for a brick-based mode of a HART-2 blade, as shown in Figure 31. In this case, the prescribed clamped end boundary conditions are such that the final solution produces a 45° solid body rotation of the blade, in addition to significant bending. The results are validated by comparing with displacements produced by the Abaqus code for this problem, and by verifying that the computed displacements converge to the values obtained for the linear formulation when the rotation mode is extracted and bending displacements are small.




Figure 31. Illustration of geometrically non-linear finite-element solution for large deflection of blade brick-based FEM model of HART2 rotor blade

IX. Conclusions and Future Work

This paper has described the steady progression in the development of a coupled multidisciplinary optimization capability for rotorcraft problems. The underlying approach is based on the use of a time dependent discrete adjoint capability that is extended to multiple disciplines. The acoustic optimization test cases have shown that optimizations can be performed using either constrained or unconstrained (penalty-based) approaches, with the latter technique

requiring the solution of only a single adjoint problem. However, more experience is needed to determine the best optimization approach as the problem formulation becomes more complex, including the extension to multipoint optimization. The drive to higher fidelity has motivated the development of a geometrically non-linear finite-element based structural dynamics solver with discrete adjoint capability. In future work, this will be coupled with the NSU3D flow solver to improve the overall fidelity of the aero-structural-acoustic analysis and optimization capability.

X. Acknowledgements

This work was partly funded by AFOSR Contract FA9550-14-C-0035 and by the Alfred Gessow Rotorcraft Center of Excellence through a subcontract with the University of Maryland. Computer resources were provided by the University of Wyoming Advanced Research Computing Center and by the NCAR-Wyoming Supercomputer Center. Special thanks Dr. Brentner for making PSU-WOPWOP available to us for validation purposes.

References

- ¹Mavriplis, D. J., "Discrete Adjoint-Based Approach for Optimization Problems on Three-Dimensional Unstructured Meshes," *AIAA Journal*, Vol. 45-4, April 2007, pp. 741–750.
- ²Mavriplis, D. J., "Solution of the Unsteady Discrete Adjoint for Three-Dimensional Problems on Dynamically Deforming Unstructured Meshes," *Proceedings of the 46th Aerospace Sciences Meeting and Exhibit, Reno NV*, 2008, AIAA Paper 2008-0727.
- ³Mani, K. and Mavriplis, D. J., "Geometry Optimization in Three-Dimensional Unsteady Flow Problems using the Discrete Adjoint," *51st AIAA Aerospace Sciences Meeting, Grapevine, TX*, January 2013, AIAA Paper 2013-0662.
- ⁴Spalart, P. R. and Allmaras, S. R., "A One-equation Turbulence Model for Aerodynamic Flows," *La Recherche Aéropatiale*, Vol. 1, 1994, pp. 5–21.
- ⁵Mavriplis, D. J., "Multigrid Strategies for Viscous Flow Solvers on Anisotropic Unstructured Meshes," *Journal of Computational Physics*, Vol. 145, No. 1, Sept. 1998, pp. 141–165.
- ⁶Mavriplis, D. J., Yang, Z., and Long, M., "Results using NSU3D for the first Aeroelastic Prediction Workshop," *Proceedings of the 51st Aerospace Sciences Meeting and Exhibit, Grapevine TX*, 2013, AIAA Paper 2013-0786.
- ⁷Schuster, D. M., Chwalowski, P., Heeg, J., and Wieseman, C. D., "Summary of Data and Findings from the First Aeroelastic Prediction Workshop," *Seventh International Conference on Computational Fluid Dynamics (ICCFD7)*, ICCFD, Big Island, HI, July 9-13 2012.
- ⁸Mishra, A., Mani, K., Mavriplis, D., and Sitaraman, J., "Time dependent adjoint-based optimization for coupled fluid–structure problems," *Journal of Computational Physics*, Vol. 292, 2015, pp. 253–271.
- ⁹Hodges, D. H. and Dowell, E., "Nonlinear equations of motion for the elastic bending and torsion of twisted nonuniform rotor blades," 1974.
- ¹⁰Bhuiyan, F., Mavriplis, D., and Fertig, R., "Predicting Composite Fatigue Life of Wind Turbine Blades Using Constituent-Level Physics and Realistic Aerodynamic Load," AIAA Paper 2016-0988, 54th AIAA Aerospace Sciences Meeting, San Diego CA.
- ¹¹Mavriplis, D., Anderson, E., Fertig, R., and Garnich, M., "Development of a High-Fidelity Time-Dependent Aero-Structural Capability for Analysis and Design," AIAA Paper 2016-1175, 54th AIAA Aerospace Sciences Meeting, San Diego CA.
- ¹²Yang, Z. and Mavriplis, D. J., "A Mesh Deformation Strategy Optimized by the Adjoint Method on Unstructured Meshes," *AIAA Journal*, Vol. 45, No. 12, 2007, pp. 2885–2896.
- ¹³Mavriplis, D. J. and Mani, K., "Unstructured Mesh Solution Technique Using the NSU3D Solver," *Proceedings of the 52nd AIAA Aerospace Sciences Conference, National Harbor, MD*, 2014, AIAA Paper 2014-0081.
- ¹⁴Nielsen, E. J. and Kleb, W. L., "Efficient Construction of Discrete Adjoint Operators on Unstructured Grids by Using Complex Variables," *AIAA Journal*, Vol. 44-4, January 2005, pp. 827–836.
- ¹⁵Mishra, A., Mavriplis, D. J., and Sitaraman, J., "Time-dependent Aero-elastic Adjoint-based Aerodynamic Shape Optimization of Helicopter Rotors in Forward Flight," *56th AIAA Aerospace Sciences Meeting, Kissimmee, FL*, January 5–9 2015, AIAA Paper 2015-1130.
- ¹⁶Mishra, A., Mavriplis, D. J., and Sitaraman, J., "Multipoint Time-Dependent Aero-elastic Adjoint-based Aerodynamic Shape Optimization of Helicopter Rotors," *71st American Helicopter Society Annual Forum, Virginia Beach, VA*, May 2015.
- ¹⁷Yu, Y. H., Tung, C., van der Wall, B., Pausder, H.-J., Burley, C., Brooks, T., Beaumier, P., Delrieux, Y., Mercker, E., and Pengel, K., "The HART-II Test: Rotor Wakes and Aeroacoustics with Higher-Harmonic Pitch Control (HHC) Inputs -The Joint German/French/Dutch/US Project-," *58th American Helicopter Society Annual Forum, Montreal, Canada*, June 11–13 2002.
- ¹⁸Nielsen, E. J. and Diskin, B., "Discrete adjoint-based design for unsteady turbulent flows on dynamic overset unstructured grids," *AIAA journal*, Vol. 51, No. 6, 2013, pp. 1355–1373.
- ¹⁹Zhu, C., Byrd, R. H., and Nocedal, J., "L-BFGS-B: Algorithm 778: L-BFGS-B, FORTRAN routines for large scale bound constrained optimization," *ACM Transactions on Mathematical Software*, Vol. 23-4, 1997, pp. 550–560.
- ²⁰Hicks, R. and Henne, P., "Wing Design by Numerical Optimization," *Journal of Aircraft*, Vol. 15-7, 1978, pp. 407–412.
- ²¹Brentner, K. S. and Farassat, F., "Modeling aerodynamically generated sound of helicopter rotors," *Progress in aerospace sciences*, Vol. 39, 2003, pp. 83–120.
- ²²Crighton, D. G., Dowling, A. P., and Ffowcs Williams, J. E., *Modern Methods in Analytical Acoustics*, Springer-Verlag, 1992.
- ²³Brentner, K. S. and Farassat, F., "An Analytical Comparison of The Acoustic Analogy And Kirchoff Formulation For Moving Surfaces," *AIAA Journal*, Vol. 36-8, 1998, pp. 1379–1386.
- ²⁴Farassat, F. and Succi, G. P., "The prediction of helicopter discrete frequency noise," *Vertica*, Vol. 7-4, 1983, pp. 309–320.
- ²⁵di Francescantonio, P., "A New Boundary Integral Formulation for the Prediction of Sound Radiation," *Vertica*, Vol. 7-4, 1983, pp. 309–320.
- ²⁶Bres, G. A., Brentner, K. S., Perez, G., and Jones, H. E., "Maneuvering rotorcraft noise prediction," *Journal of sound and vibration*, Vol. 275, 2004, pp. 719–738.

- ²⁷Gill, P. E., Murray, W., and Saunders, M. A., "SNOPT: An SQP Algorithm for Large-Scale Constrained Optimization," *SIAM review*, Vol. 47-1, 2005, pp. 99–131.
- ²⁸Sim, B. W., JanakiRam, R. D., Barbely, N. L., and Solis, E., "Reduced In-Plane, Low Frequency Noise of an Active Flap Rotor," *American Helicopter Society 65th Annual Forum*, Grapevine, TX, May 2009.
- ²⁹Fabiano, E. and Mavriplis, D., "Adjoint-based Aeroacoustics Design-Optimization of Flexible Rotors in Forward Flight," AHS 72 Forum paper, West Palm Beach, FL.
- ³⁰Yeager, W. T., Mantay, W. R., Wilbur, M. L., Cramer, R. G., and Singleton, D. J., "Wind-Tunnel Evaluation of an Advanced Main-Rotor Blade Design for a Utility-Class Helicopter," Tech. Rep. TM-89129, NASA, September 1987.
- ³¹Bousman, W. G., "Aerodynamics Characteristics of SC1095 and SC1094R8 Airfoils," Tech. Rep. TP-2003-212265, NASA, December 2003.
- ³²Bousman, W. G. and Kufeld, R. M., "UH-60A Airloads Catalog," Tech. Rep. TM-2005-212827, NASA, August 2005.
- ³³<https://aiaa-dpw.larc.nasa.gov/> .
- ³⁴Bathe, K.-J., *Finite element procedures*, Klaus-Jurgen Bathe, 2006.