

Rotor Airloads Prediction Using Unstructured Meshes and Loose CFD/CSD Coupling

Robert T. Biedron* Elizabeth M. Lee-Rausch†
NASA Langley Research Center, Hampton, VA 23681

Abstract

The FUN3D unsteady Reynolds-averaged Navier-Stokes solver for unstructured grids has been modified to allow prediction of trimmed rotorcraft airloads. The trim of the rotorcraft and the aeroelastic deformation of the rotor blades are accounted for via loose coupling with the CAMRAD II rotorcraft computational structural dynamics code. The set of codes is used to analyze the HART-II Baseline, Minimum Noise and Minimum Vibration test conditions. The loose coupling approach is found to be stable and convergent for the cases considered. Comparison of the resulting airloads and structural deformations with experimentally measured data is presented. The effect of grid resolution and temporal accuracy is examined.

Nomenclature

F/M^{CFD}	Forces/Moments from CFD solution	r	radial position, [m]
F/M^{LL}	Forces/Moments from lifting line solution	AFDD	Aeroflightdynamics Directorate (U.S. Army)
T_m	4x4 transform Matrix	BDF	Backward Differentiation Formulae
μ	advance ratio	BL	Baseline
ψ	azimuthal position, [°]	BVI	Blade Vortex Interaction
θ_0	collective pitch, [°]	CAMRAD	Comprehensive Analytical Model of Rotorcraft Aerodynamics and Dynamics
θ_c	(-) lateral cyclic pitch, [°]	CFD	Computational Fluid Dynamics
θ_s	(-) longitudinal cyclic pitch, [°]	CSD	Computational Structural Dynamics
c	local blade chord [m]	FUN3D	Fully Unstructured 3-Dimensional
$C_m M^2$	sectional pitching moment coefficient ($\frac{M_y}{\frac{1}{2}\rho a^2 c^2}$)	HHC	Higher Harmonic pitch Control
$C_n M^2$	sectional normal force coefficient ($\frac{F_n}{\frac{1}{2}\rho a^2 c}$)	MN	Minimum Noise
Q	second invariant of the velocity-gradient tensor	MV	Minimum Vibration
R	rotor radius, [m]		

Introduction

Rotorcraft airloads prediction presents a very substantial challenge for Computational Fluid Dynamics (CFD). Not only must the unsteady nature of the flow be accurately modeled, but since most rotorcraft blades are not structurally stiff, an accurate simulation must account for the blade structural dynamics. In addition, trim of the rotorcraft to desired thrust and moment targets depends on both aerodynamic loads and structural deformation, and vice versa. Further, interaction of the fuselage with the rotor flow field can be important, so that relative motion between the blades and the fuselage must be accommodated. Thus a complete simulation requires coupled aerodynamics, structures and trim, with the ability to model geometrically complex configurations.

NASA has recently initiated a Subsonic Rotary Wing (SRW) Project under the overall Fundamental Aeronautics Program. Within the context of SRW are efforts aimed at furthering the state of the art of high-fidelity rotorcraft flow simulations, using both structured and unstructured meshes. Structured-mesh solvers have an advantage in computation speed, but even though remarkably complex configurations may be accommodated using the overset grid approach, generation of complex structured-mesh systems can require months to set up. As a result, many rotorcraft simulations using

*Senior Research Scientist, Member AIAA

†Research Engineer, Senior Member AIAA

This material is declared a work of the U.S. Government and is not subject to copyright protection in the United States.

structured-grid CFD neglect the fuselage. On the other hand, unstructured-mesh solvers are easily able to handle complex geometries, but suffer from slower execution speed. However, advances in both computer hardware and CFD algorithms have made previously state-of-the-art computations routine for unstructured-mesh solvers, so that rotorcraft simulations using unstructured grids are now viable. The aim of the present work is to develop a “first principles” rotorcraft simulation tool based on an unstructured CFD solver.

CFD Solver

Baseline Solver

The unstructured-grid flow solver used for this study is FUN3D.¹ Within FUN3D, the Navier-Stokes equations are discretized over the median dual volume surrounding each mesh point, balancing the time rate of change of the averaged conserved variables in each dual volume with the flux of mass, momentum, and energy through the instantaneous surface of the control volume. The convective fluxes are computed with a flux-splitting scheme, and for second-order accuracy the values at dual-cell interfaces are reconstructed using gradients at mesh nodes that are computed using a least-squares technique. Limiting of the reconstructed values may be employed for flows with strong shocks. For all results presented in this paper, the convective flux scheme used is Roe’s flux difference splitting² with second-order unlimited reconstruction. For tetrahedral meshes, the full viscous fluxes are discretized using a finite-volume formulation in which the required velocity gradients on the dual faces are computed using the Green-Gauss theorem. For non-tetrahedral meshes, the Green-Gauss gradients are combined with edge-based gradients to avoid odd-even decoupling. All results presented here employ purely tetrahedral meshes. For turbulent flows, both the one-equation model of Spalart and Allmaras³ (SA) and the two-equation SST model of Menter⁴ are available. The SA model may be solved loosely coupled to the mean-flow equations or tightly coupled to the mean-flow equations. For all results presented in this paper, the one equation SA model is employed, solved in a loosely coupled fashion.

To advance the equations in time, several schemes based on backward differentiation formulae (BDF) are available. In addition to the well-known first-, second- and third-order BDF schemes (BDF1, BDF2, and BDF3, respectively), a blended scheme, referred to as BDF2_{opt},⁵ and a fourth order, modified extended backward differentiation formula (MEBDF4)⁶ scheme are also available. For most results presented here, the BDF2_{opt} is used. This scheme is a blend of second- and third-order backward-difference schemes that, unlike BDF3, is guaranteed to be stable. Although the BDF2_{opt} scheme is formally second-order accurate, it has a smaller coefficient for the error term as compared to the standard BDF2 scheme. Furthermore, the computational cost of BDF2_{opt} is the same as BDF2 and is roughly one-third that of MEBDF4. Regardless of which BDF variant is chosen, a dual-time stepping scheme is used, wherein the flow equations are advanced towards a steady-state in pseudo-time between each physical time step. As described in References 6 and 7, a temporal-error controller is used to provide automated exit from the pseudo-time loop if the residual within the time step drops below a user-specified percentage of the estimated temporal error. The exit criterion is applied to both the mean-flow equations and the turbulence equation(s). If the specified criterion is not met, a user-specified maximum number of pseudo-time steps is taken. For all cases presented here, the exit criterion was taken as 10% of the estimated temporal error, with a maximum of 15 pseudo-time steps. Following Reference 8, a source term derived from the Geometric Conservation Law⁹ is used to ensure that the scheme is free-stream preserving with moving meshes.

Mesh Motion

FUN3D allows for rigid and deforming mesh motion; both are needed for rotorcraft applications. Given a surface displacement, mesh deformation is accomplished by solving the linear elasticity equations for the mesh point displacements throughout the field. The elasticity equations contain material property coefficients; for mesh deformation, these coefficients are set from grid characteristics. In the particular formulation currently used in FUN3D, the material properties are the Poisson ratio and the Young’s modulus. For physical materials, Young’s modulus is a positive quantity that indicates the stiffness of the material; larger values indicate a stiffer material. Poisson’s ratio indicates the relative amount that a material shrinks in the transverse direction as it is extended in the axial direction. The Poisson ratio is set to zero throughout the mesh, while the Young’s modulus is varied from point to point based on either the inverse of the distance from the point to the nearest solid surface or inversely proportional to the dual-cell volume. For all results presented here, the former is used. In either case the effect is that small cells near solid surfaces are essentially rigid, and hence move in concert with the surface, without deformation, while larger cells further removed from the surface are deformed to accommodate the motion. The elasticity equations are solved as a separate steady-state problem at each time step, using the generalized minimum residual (GMRES) method.¹⁰ Typically a reduction of six orders of magnitude in the L_2 norm of the residual is utilized as a convergence criterion for the elasticity equations.

Motion is accomplished via application of a 4x4 transform matrix that contains both translation and orthonormal rotation components. Given a point at an initial position $(x, y, z)^T$, application of the transform matrix moves the point to its new position $(x', y', z')^T$

$$\begin{bmatrix} x' \\ y' \\ z' \\ 1 \end{bmatrix} = \begin{bmatrix} R_{11} & R_{12} & R_{13} & T_x \\ R_{21} & R_{22} & R_{23} & T_y \\ R_{31} & R_{32} & R_{33} & T_z \\ 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ z \\ 1 \end{bmatrix} = [\mathbf{T}_m] \begin{bmatrix} x \\ y \\ z \\ 1 \end{bmatrix}$$

The transform matrix $[\mathbf{T}_m]$ may be applied to the entire mesh to achieve rigid mesh motion, or to a collection of surface points to facilitate rigid body motion within a deforming mesh, or to individual points on a surface to accommodate deforming body motion within a deforming mesh. Application of the inverse transform negates the transformation and returns the point to its original position. An extremely useful feature of the transform matrix approach is that multiple transformations telescope via matrix multiplication. Thus, rotation about a point $(x_0, y_0, z_0)^T$, by an angle θ , in the direction \hat{n} is effected by first translating to the origin by application of a pure translation matrix $[\mathbf{T}_0]$, performing the rotation via a pure rotation matrix $[\mathbf{R}_0]$, and then translating back to $(x_0, y_0, z_0)^T$ by application of $[\mathbf{T}_0]^{-1}$. The motion can be accomplished in a single step by the application of the composite transform matrix $[\mathbf{T}_m]$

$$[\mathbf{T}_m] = [\mathbf{T}_0][\mathbf{R}_0][\mathbf{T}_0]^{-1}$$

This telescoping property is used extensively for rotorcraft simulations to facilitate transformation from one coordinate system to the next. For example, by convention the reference coordinate system for the rotor blade is taken with the x-axis running spanwise, the y-axis in the anti-chordwise direction, and the z-axis pointing up. The blade is moved from this reference system to its current position in the rotorcraft coordinate system via a composite transform that first translates to the shaft origin, tilts the blade to the shaft axis, and then rotates about the shaft axis to the proper azimuthal position.

Overset Grids

Deforming meshes can accommodate moderately large surface motions within a single mesh system, but if the motion becomes too large, negative cell volumes can result. To overcome this, overset meshes can be used for applications involving large motions; rotorcraft and store separation are two examples. The overset method was first implemented in the FUN3D solver by O'Brien.¹¹ The implementation uses the Donor Interpolation/Receptor Transaction library¹² (DiRTlib) to facilitate the use of overset grids in a parallel environment without extensive modification to the flow solver. As for non-overset meshes, the flow solver continues to operate on a single mesh (partitioned for multiple processors). For overset meshes, points are flagged with an identifier for the particular component mesh with which they are associated. With a few simple calls within the flow solver, DiRTlib handles the equation blanking and solution interpolation required for the overset method. Linear interpolation of the solution between points associated with different component meshes is used with two layers of donor points, consistent with the underlying second-order spatial accuracy of the baseline solver. Points within holes (blanked regions) are assigned solution values by averaging the solution at neighboring points. This also helps ensure that for moving mesh problems, points that were blanked at previous time steps do not suddenly become unblanked with initial freestream values. Orphan points (if any) are assigned solution values in the same manner as hole points.

DiRTlib does not perform composite grid assembly, cut holes (establish blanking), or determine the requisite interpolation coefficients. For that, the Structured, Unstructured, and Generalized overset Grid Assembler¹³ (SUGGAR) program is employed. SUGGAR may be compiled as a stand-alone executable or as a callable library. In a preprocessing step prior to the initial flow-solver execution, SUGGAR reads two or more component meshes and creates a single composite mesh with the configuration in the initial position, along with a file (a DCI file in SUGGAR/DiRTlib parlance) identifying points corresponding to each component mesh, blanked points, and interpolation coefficients. This composite mesh is then partitioned for execution on multiple processors, after which the flow-solver execution may begin. The current usage of SUGGAR within the FUN3D solver is as a library, with one processor devoted to the SUGGAR task. When called upon, this processor receives updated grid information from the other processors, computes the new overset connectivity data, and then sends that data back to the flow-solve processors. A new DCI file containing the updated connectivity data is also written out for subsequent reuse, if desired. Currently, only a single processor runs the SUGGAR task, since SUGGAR does not parallelize well. An updated version of SUGGAR, SUGGAR++, is due to be released soon and is expected to parallelize significantly better than its predecessor. The current FUN3D parallel-processing implementation supports multiple MPI (Message Passing Interface) communicators, and will be able to take advantage of multiple processors running SUGGAR++ when that code is available.

In applications involving overset deforming meshes the situation may occur where large parts of the mesh undergo no deformation. In this case some efficiency is lost because the linear elasticity solver, by default, operates on all points of the mesh. Thus, for overset deforming meshes the non-deforming points are masked from solution in the linear elasticity equations. In rotorcraft applications, this means that points in the fuselage/background mesh are masked; these points may in fact be the majority of the points in the mesh. Simply masking the points is typically not sufficient to achieve

increased efficiency. During the preprocessing step of mesh partitioning for parallel execution, the deforming points need to be weighted so that an equal number of deforming points are given to each processor.

CSD Solver

Although FUN3D has a linear structures, mode-shape-based aeroelastic capability, the unique aspects of rotorcraft problems must be handled by a CSD code specifically tailored for rotorcraft. The rotorcraft CSD solver used for this study is the Comprehensive Analytical Model of Rotorcraft Aerodynamics and Dynamics II (CAMRAD II) code.¹⁴ As the name suggests, CAMRAD II models both the aerodynamics and structural dynamics of rotorcraft systems. However, the aerodynamics modules within CAMRAD II are based on lifting-line models utilizing airfoil tables, coupled with wake models. Although such aerodynamic models can provide reasonable results for many flight conditions, "tuning" is often required to produce accurate results. For other flight conditions, the predictions of the airloads can be inaccurate because of limitations of the relatively low-order aerodynamic modeling. The goal of the loose coupling approach described below is to replace the low-order lifting-line aerodynamics of the CSD code with the higher-fidelity Navier-Stokes aerodynamics of the CFD code.

Within CAMRAD II, each blade is modeled as a set of nonlinear beam elements. Since rotor blades are typically long and slender, the use of a beam model is a good approximation. In addition to the structural dynamics modelling, CAMRAD II offers a sophisticated trim capability. Rotor system dynamics depend on the aerodynamics, structural response, and trim settings needed to achieve the thrust-and-moment values appropriate to the specific flight conditions. Indeed, all three - aerodynamics, structural dynamics and trim - depend very strongly on each other. Thus trim is an important element of a rotorcraft simulation.

Loose Coupling Methodology

The idea for loose coupling of a rotorcraft CSD code with a CFD code was first suggested in Reference 15. In the loose coupling approach, airloads data from the CFD solver and blade motion data from the CFD solver are exchanged at relatively infrequent periodic intervals, for example once per revolution or once per blade passage ($1/N_{blade}$). This approach has subsequently been adopted in numerous rotorcraft studies,¹⁶⁻²⁰ and has proven to be stable and convergent for conditions in which there is a periodic state - i.e. hover and steady flight. For conditions that are fundamentally transient, for example in maneuvering flight, an exchange of data each time step is more appropriate. As only steady flight conditions are considered here, the more computationally efficient loose-coupling approach is adopted.

The flow of a loosely coupled CFD/CSD rotorcraft simulation is shown in Figure 1. The simulation is begun by first running the CSD code without input from CFD, using the forces and moments obtained from the internal lifting line model, F/M^{LL} . The CSD code determines the structural response and trim settings along with the resulting blade motion. The blade motion is output as a set of three displacements from the reference blade quarter-chord position, plus a set of three rotations (Euler angles), at a number of radial (r) positions along the quarter-chord. The motion data is output at a number of azimuthal (ψ) stations around the rotor disk for one complete revolution of the shaft. This rotation and displacement data encompasses all motion of the blade exclusive of the overall rotation at the shaft speed.

Before the CFD code is executed, a pre-processing code is used to parse the blade motion data from the CSD output and create a file that is read by the flow solver. Within the CFD code, these displacements and rotations are first fit with a two-dimensional spline in (r, ψ) . This gives a set of six continuous functions for the quarter-chord displacements and rotations that can be queried at any point in (r, ψ) . The quarter-chord data is applied to all points on the blade having the same radius - i.e. sectional bending is neglected. The displacements and Euler angles are used to form a local transform matrix T_m at each point on the airfoil surface, which is then applied to the coordinates of the point in the reference position. When all points on the blade surface have been processed in this manner, the blade has the correctly deformed shape, but is still in the reference coordinate system. A final composite transform is then used to translate the deformed blade to the hub center, tilt to the shaft angle, and rotate to the correct azimuthal position. With all blades in the new position, the volume mesh is deformed to accommodate the new blade shapes, and then the flow equations are advanced in time. At each time step, the pressure and three components of skin friction are extracted at spanwise stations along the blades. The resulting pressure and skin friction distributions at each radial station are then integrated along the chord to give sectional force and moment coefficients, F/M^{CFD} , which are output for the next step of the coupling process. The CFD solution is advanced until at least one complete revolution of F/M^{CFD} data is obtained. For an N-bladed rotor, this requires only $360/N$ degrees of shaft rotation. However, during the initial coupling cycle, since the CFD solution starts from uniform flow, several (typically 3) complete shaft revolutions are utilized to establish the flow field and a (nearly) periodic solution. For subsequent coupling cycles the CFD solver is run for $2 \times 360/N$ degrees of shaft rotation, rather than the minimum $360/N$ required, to remove transients that occur after each CSD/CFD coupling. The transients are reduced in magnitude with each coupling cycle, eventually disappearing altogether.

Following Reference 16, at the start of subsequent coupling cycles, $i = 1, 2, 3 \dots$, a second pre-processing code is used to compute a delta airloads file containing

$$\Delta \mathbf{F}/\mathbf{M}_i = \Delta \mathbf{F}/\mathbf{M}_{i-1} + (\mathbf{F}/\mathbf{M}_{i-1}^{CFD} - \mathbf{F}/\mathbf{M}_{i-1})$$

where $\Delta \mathbf{F}/\mathbf{M}_0 = 0$. CAMRAD II reads this file and adds the delta airloads to its internal lifting line airloads to give the total

$$\mathbf{F}/\mathbf{M}_i = \mathbf{F}/\mathbf{M}_i^{LL} + \Delta \mathbf{F}/\mathbf{M}_i$$

Starting with $\Delta \mathbf{F}/\mathbf{M}_0 = 0$, i.e. $\mathbf{F}/\mathbf{M}_0 = \mathbf{F}/\mathbf{M}_0^{LL}$, successive substitution for $i = 1, 2, 3 \dots$ into the above equations leads to the relation

$$\mathbf{F}/\mathbf{M}_i = (\mathbf{F}/\mathbf{M}_i^{LL} - \mathbf{F}/\mathbf{M}_{i-1}^{LL}) + \mathbf{F}/\mathbf{M}_{i-1}^{CFD}$$

Thus, provided the lifting-line aerodynamics converge between successive iterations in CAMRAD II, so that $\mathbf{F}/\mathbf{M}_{i-1}^{LL} \Rightarrow \mathbf{F}/\mathbf{M}_i^{LL}$, then $\mathbf{F}/\mathbf{M}_i \Rightarrow \mathbf{F}/\mathbf{M}_{i-1}^{CFD}$, and the total airloads within CAMRAD II will be the CFD airloads. In the current work, the coupling process is repeated until the thrust computed by the CFD solver is within 0.1% of the thrust computed by the CSD code. In addition, the computed trim control settings are monitored until there is no change to within plotting accuracy from one coupling cycle to the next. The computed structural responses and airloads at selected stations are also monitored, with convergence deemed to have occurred when there are no changes (to plotting accuracy) from one trim cycle to the next.

HART-II Test Description

The Higher harmonic control Aeroacoustics Rotor Test II (HART-II) is part of an ongoing international cooperative program between the German Aerospace Center (DLR), the German-Dutch Wind Tunnels (DNW), the French Office National D'Etudes et de Recherches Aerospatiales (ONERA), the United States NASA Langley Research Center and the United States Army Aero-Flight Dynamics Directorate (AFDD) research organizations. The program focuses on a geometrically and aeroelastically scaled model of a BO-105 main rotor that was tested in an open-jet anechoic test section of the German-Dutch Windtunnel (DNW) in October 2001.²¹ The test model installed in the DNW is shown in Figure 2. In the test program, detailed acoustic, aerodynamic, wake, and blade deformation data were collected with pressure instrumented blades. A set of pressure transducers, of sufficient density to provide aerodynamic loading analysis, was installed at the 87% span location. Blade motion data were obtained at a large number of spanwise locations using a Stereo Pattern Recognition system in which white markers were attached to the undersurface of the blade and photographed during the test.

The rotor is a 40% scale model of the hingeless BO-105 main rotor, with a 2m radius and a 0.121m chord. The blade section consists of a modified NACA 23012 section wherein the trailing edge is modified to have a 5.4mm long, 0.8mm thick tab. Except for a small section near the root, the blade has a rectangular planform, with a solidity of 0.077. The blade has a built in -8° linear twist (wash out) and a square tip. The rotor has a pre-cone angle of 2.5° . The model rotor is mounted on a fuselage-shaped fairing supported by a sting, as seen in Figure 2.

The tests were conducted with a rotor speed of approximately 1040 rpm. The nominal values of hover tip Mach number and tunnel Mach number were 0.64 and 0.096, respectively, yielding an advance ratio $\mu=0.15$. A shaft tilt of 5.3° (aft) was used to test descending flight conditions at which Blade Vortex Interaction (BVI) noise radiation is significant. Subsequent wind-tunnel corrections reduced the effective shaft tilt to approximately 4.5° ; this corrected value was used for the computational results presented here. For all cases the rotor was trimmed to a nominal thrust of $3300N$, with nominally zero rolling and pitching moments (on the order of $20 - 30Nm$). Three cases involving differing amounts of higher-harmonic pitch control (HHC) were investigated. The baseline (BL) case had no HHC. The minimum noise (MN) and minimum vibration (MV) cases had 3/rev HHC control each with different phase lags. The commanded controls may be described via

$$\begin{aligned} \theta &= \theta_0 + A_1 \cos(\psi - \delta_1) + A_3 \cos(3\psi - \delta_3) \\ &= \theta_0 + \theta_{1c} \cos(\psi) + \theta_{1s} \sin(\psi) + \theta_{3c} \cos(3\psi) + \theta_{3s} \sin(3\psi) \end{aligned}$$

The following table summarizes the pitch control settings for the three cases considered here. Rotor azimuthal angles ψ and phase angles δ are measured in degrees from 0° , corresponding to blade 1 of the rotor pointing aft, over the tunnel mounting sting.

	θ_0	A_1	δ_1	A_3	δ_3	θ_{1c}	θ_{1s}	θ_{3c}	θ_{3s}
BL	3.80	2.34	325	0	0	1.92	-1.34	0	0
MN	3.91	2.41	326	0.81	300	2.00	-1.35	0.41	-0.70
MV	3.80	2.51	323	0.79	180	2.01	-1.51	-0.79	0

Table 1 Experimental control settings for HART-II descent cases (degrees, measured relative to $r/R = 0.70$).

Computational Results

For the simulations presented here, the blades and the fairing/sting assembly are modeled. The shaft and linkages between the shaft and the blades are ignored. These geometric details are situated in an area of low speed flow and are assumed to have negligible contributions to the overall airloads. The wind tunnel walls are neglected as well, although as mentioned, a wind tunnel correction to the shaft angle is used.

A series of three unstructured, overset grids consisting of tetrahedral cells were generated, containing 3.6, 6.9 and 13.6 million nodes. All three were generated using the VGRID²² mesh generation software. Overset structured grids for rotorcraft applications are often categorized not only by the number of points in the mesh, but also by the size of the cells just beyond the outer boundary of blade mesh. A similar categorization of unstructured meshes is considerably less precise, owing to the indirect means of specifying cell sizes and the relatively wide variations in size from one cell to the next. Very roughly speaking, the cell sizes in background meshes near the blades in the three unstructured grids used here correspond to 0.40c, 0.15c, and 0.10c for the 3.6, 6.9 and 13.6 million node meshes, respectively. For all the computational meshes, wall spacings for the fuselage and rotor blades were chosen so that $y^+ \approx 1$ at the first grid point off the surface. Figure 3 shows a slice taken through the volume mesh as well as the surface meshes of the 13.6 million node grid. The single background mesh (blue) contains the fairing/sting and the majority of the grid points. The densely-clustered region near the blade is created using a new volume sourcing capability in VGRID 4.1. Each blade mesh (red) extends only a short distance (roughly one chord-length) from the blade surface. The composite mesh is formed within SUGGAR from a single blade mesh (replicated and rotated to form the four-bladed rotor system) and the fairing/sting mesh. Holes in the background mesh (light red region in the sliced plane) are cut by SUGGAR, with an overlap region remaining (red/blue region).

Within the CSD code, the HART-II blade is modeled using seven non-linear beam elements, with a higher concentration of elements near the root. Each element has eleven degrees of freedom (4 axial, 2 lag bending, 2 flap bending, and 3 torsional). A time step corresponding to a 15° increment in azimuth was used for the harmonic balance method employed by the CSD code, along with sixteen elements for the linear aerodynamics model in CAMRAD II. Blade motion is output at the same azimuthal increments, and at radial increments of $0.01R$. Likewise, delta airloads from the CFD solver are input with radial and azimuthal resolutions of $0.01R$ and 15° , respectively. These resolution levels are typical of other published results that have utilized the CAMRAD II CSD code.¹⁷ For all cases, trim targets within CAMRAD II were set to a thrust level of $3300N$, with zero rolling and pitching moment.

Uncoupled AFDD Motion

As part of the HART-II program, a number of workshops have been held to provide information exchange between researchers in computational and experimental methods. For the 4th International HART II Workshop, a file containing the blade motions (exclusive of the overall rotation at the shaft speed) for the BL case was made available to workshop participants. This blade-motion file was the result of a loose coupling between the OVERFLOW-2 structured-grid CFD code and CAMRAD II,¹⁷ but without accounting for the presence of the fairing/sting in the computations - i.e. rotor only. Hereafter this motion file is referred to as the AFDD motion file, since it was generated by researchers at the U.S. Army Aeroflightdynamics Directorate. This motion file (and similar ones for the MN and MV cases) was used to obtain the results presented in Reference 17 wherein meshes of up to 113 million nodes were used. The results of Reference 17 are considered state-of-the-art for structured meshes. Having access to the AFDD motion file was invaluable to the development of the loose coupling between FUN3D and CAMRAD II; it allowed independent testing of half the coupling procedure without having to perform the entire coupling process.

Before comparing computed airloads with data it is useful to have an understanding of the rotor wake structure. Figure 4 illustrates the wake structure in the vicinity of the rotor by way of computed isosurfaces of the second invariant of the velocity-gradient tensor, Q

$$Q = \frac{1}{2}(P^2 - S_{ij}S_{ji} - W_{ij}W_{ji})$$

where P , S and W are the divergence $\partial u_i/\partial x_i$, the rate-of-strain tensor $(\partial u_i/\partial x_j + \partial u_j/\partial x_i)/2$ and the rate-of-rotation tensor $(\partial u_i/\partial x_j - \partial u_j/\partial x_i)/2$, respectively. In the figure, the isosurfaces correspond to $Q = 0.0075$ and are colored by

the magnitude of vorticity. To orient the discussion that follows, rotor-disk quadrants 1 (advancing blade) and 4 (retreating blade) are labeled in Figure 4. The first quadrant corresponds to rotor shaft angles $0^\circ \leq \psi \leq 90^\circ$ while the fourth corresponds to rotor shaft angles $270^\circ \leq \psi \leq 360^\circ$, where $\psi = 0^\circ$ corresponds to a rotor blade located over, and aligned with, the wind-tunnel support sting. The blades rotate counterclockwise as seen from above. Because all three HART II cases correspond to descending forward flight, the rotor wakes are convected downstream, but tend to stay above the rotor disk, at least near the rotor. A rotor blade entering the fourth quadrant encounters relatively "young" vortices originating in quadrant 4, and slightly older ones from quadrant 3. In contrast a blade entering the first quadrant initially encounters much older vortices from quadrants 4 and 3. In addition, the vortices generated in quadrants 1 and 2 (advancing side) are inherently weaker than those in generated quadrants 3 and 4 (retreating side), since the angle of attack is smaller on the advancing side. Furthermore, because of numerical dissipation, it is much more difficult to accurately capture older wakes with CFD. Thus, given the inherently weaker vortices generated in quadrants 1 and 2, the relatively older vortices from quadrants 3 and 4, and numerical dissipation, it is anticipated that the effects of BVI on the rotor airloads will be more difficult to capture in the first quadrant as compared to the fourth quadrant. This is evident in Figure 4, where the vortex cores in quadrant 1 are less well-defined than those in quadrant 4.

Figure 5 shows the variation of $C_n M^2$ for the 13.6 million node mesh at $r/R = 0.87$, the single radial station where experimental data are available. Also shown in the figure are the experimentally measured data and the OVERFLOW-2 results^{17,23} on a somewhat larger mesh (19.6 million nodes) with roughly comparable off-body spacing ($0.1c$). The experimental data were taken for 80 consecutive revolutions, and then conditionally averaged to eliminate data scatter.²⁴ The OVERFLOW-2 results were obtained using fourth-order differencing in the off-body grids and second-order differencing within the rotor blade grids, whereas FUN3D is limited to second-order differencing everywhere. For reasons of numerical stability, the OVERFLOW-2 results were generated using first-order backward differencing in time with a very small time step, corresponding to 7200 steps/revolution (0.05° per step).¹⁷ The present results were computed using a step size corresponding to 360 steps/revolution (1° per step); no stability issues were encountered during the present study. The underlying 3/rev variation in normal force is well captured by both CFD simulations. In both cases there is a similar offset from the experimental data. Both CFD simulations capture the BVI in the fourth quadrant (retreating blade) fairly well, less so in the first quadrant (advancing blade), with the finer structured mesh showing somewhat better peak-to-peak variation of the BVI in both quadrants. Figure 6 shows the corresponding comparisons for the pitching moment coefficient, taken about the quarter chord. Again, there is an offset between the experimental data and both computations, though the offset is slightly larger for the unstructured-grid results. Apart from the offset, the predicted variation in pitching moment is quite similar between the two codes.

Before tackling fully-coupled CFD/CSD simulations, a limited study was conducted to assess the effect of mesh size and time step on the computed results. It is quite difficult to generate a sequence of unstructured grids that represent a consistent spatial refinement in the sense of Reference 25. The three grids generated for this study certainly do not meet that definition. Nevertheless, Figure 7 shows the effect of mesh refinement on $C_n M^2$ using these grids, with the 2^{nd} -order accurate BDF2_{opt} scheme for time advancement and a fixed time step corresponding to 1° per step. Apart from the larger peak-to-peak variation of the BVI, there is little difference in the overall loading variation with ψ between the 6.9 million node mesh and the 13.6 million node mesh. Greater differences are seen in the loading with the coarsest mesh, and the BVI is missed entirely.

Figure 8 shows the effect of the temporal accuracy on the BL case. For this study, a time step corresponding to 1° per step was held fixed on the 13.6 million node mesh and the temporal accuracy was increased from 1^{st} to 4^{th} order. The 1^{st} -order results on the 13.6 million node mesh are similar to the 2^{nd} -order results on the 3.7 million node mesh (see Figure 7) - the BVI is not captured at all. Very little difference is observed between the 2^{nd} - and 4^{th} -order results. It should be noted that the 4^{th} -order MEBDF scheme consists of three stages for each step, and is therefore roughly three times as expensive as either of the two lower-order single-stage schemes. The 1^{st} - and 2^{nd} -order schemes are both roughly the same computational cost. Although not shown, an additional time-step refinement study was conducted wherein the time step was halved (using the BDF2_{opt} scheme). The differences between 1° per step and 0.5° per step were quite small, comparable to the differences between 2^{nd} and 4^{th} order at 1° per step.

Coupled CFD/CSD Motion

Based on the grid-refinement and time-step refinement studies on the uncoupled AFDD Motion results, the coupled unstructured-grid CFD/CSD simulations are carried out using the 6.9 million node mesh, utilizing the BDF2_{opt} scheme for time advancement with a time step corresponding to 1° per step. This combination of mesh and time step/temporal advancement scheme is capable of capturing the overall airload variations as well as the finer mesh with a higher order temporal scheme, although the BVI is not nearly as well resolved. The relatively smaller loading oscillations from the BVI are expected to have negligible impact on the rotorcraft trim and blade response, so the extra computational cost to capture them is not warranted in the coupling computations. Recently, Lim²⁶ has studied the effect of a low pass filter applied to CFD data prior to use in the CAMRAD II CSD code, thereby removing the high-frequency BVI component. Negligible

differences were observed in the final airloads distributions as compared to using unfiltered CFD data, further strengthening the argument that the capture of BVI during the coupling process is not essential. Once the coupled FUN3D/CAMRAD II blade motions have been determined, the finer mesh is used to obtain CFD solutions with improved resolution of the BVI.

To illustrate the convergence of the CFD/CSD coupling methodology, results for the MN case are considered; convergence of the BL and MV cases are very similar to the MN case. The variation of the coupled results for $C_n M^2$ and $C_m M^2$ with trim cycle are shown in Figures 9 and 10, respectively. As expected, the BVI is not captured on the coarser mesh used for the coupled results. The largest changes in $C_n M^2$ and $C_m M^2$ with trim cycle occur within the first three coupling steps, with minimal changes between coupling cycle thereafter. Between the seventh and eighth coupling cycles, no perceptible change in the airloads occurs. The fact that the level of $C_n M^2$ at $r/R = 0.87$ is lower than measured might suggest that the thrust is underpredicted, but this is not the case, as the coupled solutions are trimmed to match the specified thrust. Thus, there must be compensating differences in $C_n M^2$ at other radial stations; however, without measured $C_n M^2$ data at other locations it is impossible to verify that conjecture. In addition, the experimental data for $C_n M^2$ is obtained by integration of only 17 unsteady pressure taps (11 upper, 6 lower),²⁴ so there may be some error in the experimental values of $C_n M^2$ as a result. The convergence of the collective and cyclic pitch settings with trim cycle is shown in Figure 11. After approximately six trim cycles the pitch settings show negligible change. The trimmed settings are within $1/2^\circ$ of the experimental values.

The HART-II database contains extensive measurements of blade motion, allowing for assessment of the structural response computed by the coupled CFD/CSD solvers. Data are available for elastic flap displacement (local flap displacement with pre-cone angle removed), elastic torsion (local blade pitch with collective pitch, cyclic pitch, HHC, and built-in twist removed), and lead-lag displacement. Figure 12 shows the elastic torsion response at the blade tip. The phase of the torsional response is in excellent agreement with the measured data, although the peak-to-peak variation is smaller in the computed results. In Figure 12 and all subsequent figures showing experimentally-measured tip-motion data, the blade-to-blade variations are shown as error bars plotted on the mean (over blades 1-4) motion. The blade-to-blade variation can be considerable. Two blades were instrumented with pressure transducers, while the other two were not; the two uninstrumented blades tended to track closely with each other, while the instrumented blades typically tracked differently from the uninstrumented blades, and differently from each other.²⁶ Figure 13 illustrates the convergence of tip elastic flap with trim cycle for the MN case, and again, convergence with trim cycle is found to be fairly rapid. The agreement with data is reasonably good, with excellent agreement of the flap phase angle with moderate differences in amplitude, typically less than $0.005R$ ($< 8\%$ chord). Finally, Figure 14 shows the convergence of the tip lead-lag deflections. In this case, convergence is quite rapid, within a couple of trim cycles. However, there is a large difference between the computed and measured displacements. The measured motion is positive (lead) whereas most of the predicted motion is negative (lag), although the overall shape of the waveform is well predicted in the simulation. It should be noted that the discrepancy between computed and measured lead-lag motion is not limited to the current simulations; other CFD/CSD couplings show nearly identical discrepancies. Reference 27 presents results from three different simulation methodologies employed by three different teams of researchers, and all show similar offsets. Lim²⁶ recently speculated that the blade lead-lag stiffness used in the CSD model, especially for the inboard section, may not match the actual blade stiffness - blade properties near the root cutout being especially difficult to accurately determine.

Having obtained coupled CFD/CSD solutions for all three pitch-control cases using the medium (6.9M node) mesh, the resulting quarter-chord displacement and rotation data from the final coupling cycle were extracted and used for CFD simulations on the finest (13.6M node) mesh. First, comparison is made between the previous BL results using the AFDD motion data and the loosely coupled CFD/CSD results, in Figures 15 and 16. Little overall difference is seen between the the AFDD motion solutions and the coupled solutions, except that coupling has generally moved the normal force distribution slightly closer to the data, particularly near $\psi = 0^\circ/360^\circ$ and near $\psi = 180^\circ$. This is the region where the presence of the fuselage is expected to have the greatest influence, and as mentioned earlier, the AFDD motion data was obtained in the absence of the fuselage. Figure 17 shows the computed variation of $C_n M^2$ along with the experimental data for all three cases at $r/R = 0.87$. All three results show the offset to a lower mean value of $C_n M^2$ than exhibited by the experimental data, as was the case when using the AFDD motion data. The MN case shows generally good agreement with the data, apart from the offset from the measured data. The BVI on the fourth quadrant is captured on this mesh; azimuthal positions of the BVI is very well predicted, although the amplitude is substantially reduced compared to experiment. On the advancing side BVI is barely perceptible in the load distribution. The MV case shows the poorest correlation with data; the amplitude of the basic 3/rev distribution is poorly predicted. Again, BVI is captured on the retreating side, but not at all on the advancing side. The same results are replotted in Figure 18, but with the mean value of $C_n M^2$ removed from both the computation and experiment. With the means removed, it is clear that the overall normal force distribution is well predicted by the current method, with the exception of the MV case. It should be noted that the overall agreement between the current results and the measured data is generally quite comparable to those of Reference 17, with the exception that the first-quadrant BVI for the MV case is captured in OVERFLOW-2 simulations.

Figure 19 shows the comparison of computed and measured pitching moment for all three pitch-control cases. As with

the AFDD motion data, the BL case exhibits a shift towards an more negative (pitch down) mean value of the pitching moment. The MN and MV cases also exhibit mean offsets from the measured data, but toward a more positive (pitch up) mean value. The pitching moment results are replotted in Figure 20 with the mean removed from both the experimental and computed data. With the means removed, the comparison of the coupled results with data is generally quite good, especially for the MN case. The pitching moment variation for the MV is also well captured, even though the normal force was not.

Having looked at the airloads results, the tip-motion predictions are now examined. Figure 21 shows the tip elastic torsion for all three cases. Because the pitch control rotations and built-in blade twist are removed, the elastic torsion indicates aeroelastic response of the blade. The tip elastic torsion is well predicted by the computations for both the BL and MN cases, but the MV case is seen to be poorly predicted, especially in the first quadrant. Therefore the torsion, and thus angle of attack, slightly inboard at $r/R = 0.87$ where airloads were measured is also likely to be in error for the MV case. This error in torsion for the MV case is the most likely cause of the relatively poorer prediction of $C_n M^2$ for that case compared to the BL and MN cases, where tip torsion is reasonably well predicted. Since the pitching moment $C_m M^2$ is measured about the quarter chord (approximately the aerodynamic center, where pitching moment is insensitive to angle of attack), the errors in torsion do not taint the pitching moment prediction as they do the normal force prediction. Turning next to the tip elastic flapping motion, Figure 22, it is observed that all three cases are reasonably well predicted in the current simulations, with the largest discrepancies with measured data occurring in the fourth quadrant of the MV case. Overall, flap displacement errors are no larger than $0.005R$, which corresponds to approximately 8% chord. It has recently been suggested²⁸ that even better agreement of the flap displacement with data might have been obtained if the experimentally-measured pitch and roll moments ($O(20Nm)$) had been used as moment trim targets, rather than assuming these moments to be zero. The tip elastic lead-lag displacements for all three cases are shown in Figure 23. As observed when discussing Figure 14, a large offset, on the order of $0.02R$ or roughly 33% chord is observed for all three cases, with blade lag predicted when blade lead is measured. However, the overall waveform is reasonably well predicted, as are subtle changes in the shape of the predicted response from one case to the next (e.g. the slight flattening of the waveform of the MV case as compared to the BL case). Finally, the computed and measured control angles for all three pitch-control cases are listed below in Table 2, along with computed control angles from OVERFLOW-2/CAMRAD II for the BL case. In all cases the computed control angles are within 0.5° of the experimentally measured values.

	θ_0	θ_{1c}	θ_{1s}	θ_{3c}	θ_{3s}
BL (measured)	3.80	1.92	-1.34	0	0
BL (computed)	3.47	1.77	-0.97	0	0
BL (Ref. 23)	3.31	1.53	-0.95	0	0
MN (measured)	3.91	2.00	-1.35	0.41	-0.70
MN (computed)	3.57	1.81	-0.89	0.38	-0.67
MV (measured)	3.80	2.01	-1.51	-0.79	0
MV (computed)	3.40	1.73	-1.03	-0.75	-0.0002

Table 2 Comparison of measured and computed control settings for HART-II descent cases (degrees).

Summary

A Navier-Stokes solver for unstructured grids has been coupled to a rotorcraft CSD code to allow computation of rotor airloads including the effects of structural deformation and rotorcraft trim. The use of unstructured grids is expected to make the inclusion of rotorcraft fuselages and other components in computational analysis much easier than current structured-grid approaches. Computed results for both aerodynamic loading and blade structural response were obtained for the HART-II wind tunnel model for three descending flight cases, including effects of the fuselage (fairing/sting). Generally good agreement was obtained with measured airload and blade deflection data, typically comparable to the state-of-the-art structured mesh solver OVERFLOW-2 on meshes of comparable resolution. Although target thrust targets were met in the loose-coupling procedure, the mean value $C_n M^2$ at the $r/R = 0.87$ station where measured data were available showed an offset from the experimental data. However, the same offset was observed in the OVERFLOW-2 results. Limited time step studies indicated that time steps on the order of 1° provided sufficient temporal resolution in conjunction with the BDF2_{opt} scheme. Solutions were found to be more sensitive to grid refinement. The prominent BVI associated with the descending flight conditions of the HART II test suite were largely absent in the computed results until the mesh was sufficiently refined (approximately $0.1c$, with roughly 14 million nodes). It is expected that increased mesh resolution will lead to further improvement in the the BVI predictions; such computations are planned, pending parallelization of the computation of overset connectivity data.

Acknowledgment

The authors would like to thank Dr. Douglas Boyd of the Aeroacoustics Branch at NASA Langley Research Center for providing the CAMRAD II input deck for the HART II configuration, for guidance in the use of the CAMRAD II comprehensive analysis code, and for providing the OVERFLOW-2/CAMRAD interface routines. The efforts of Mr. William Jones of the Advanced Engineering Environments Branch at NASA Langley Research Center to develop the HART-II geometry model are greatly appreciated. The authors would like to thank Dr. Shahyar Pirzadeh of the Configuration Aerodynamics Branch at NASA Langley Research Center for providing an advance copy of the VGRID 4.1 software. Finally, the many helpful discussions, structured-grid numerical results, and encouragement of Dr. Joon Lim of the U.S. Army Aeroflightdynamics Directorate, Moffet Field, CA, is gratefully acknowledged.

References

- ¹Anderson, W. K. and Bonhaus, D. L., "An Implicit Upwind Algorithm for Computing Turbulent Flows on Unstructured Grids," *Computers and Fluids*, Vol. 23, No. 1, 1994, pp. 1–22.
- ²Roe, P. L., "Approximate Riemann Solvers, Parameter Vectors, and Difference Schemes," *Journal of Computational Physics*, Vol. 43, 1981, pp. 357–372.
- ³Spalart, P. R. and Allmaras, S. R., "A One-Equation Turbulence Model for Aerodynamic Flows," *La Recherche Aerospaciale*, No. 1, 1994, pp. 5–21.
- ⁴Menter, F. R., "Two-Equation Eddy-Viscosity Turbulence Models for Engineering Applications," *AIAA Journal*, Vol. 32, No. 8, 1994, pp. 1598–1605.
- ⁵Nyukhtikov, M., Smelova, N., Mitchell, B. E., and Holmes, D. G., "Optimized Dual-Time Stepping Technique For Time-Accurate Navier-Stokes Calculation," *Proceedings of the 10th International Symposium on Unsteady Aerodynamics, Aeroacoustics and Aeroelasticity of Turbomachines*, 2003.
- ⁶Vatsa, V. and Carpenter, M. H., "Higher-Order Temporal Schemes with Error Controllers for Unsteady Navier-Stokes Equations," AIAA Paper 2005-5245, June 2005.
- ⁷Biedron, R. T., Vatsa, V. N., and Atkins, H. L., "Simulation of Unsteady Flows Using an Unstructured Navier-Stokes Solver on Moving and Stationary Grids," AIAA Paper 2005-5093, June 2005.
- ⁸Morton, S. A., Melville, R. B., and Visbal, M. R., "Accuracy and Coupling Issues of Aeroelastic Navier-Stokes Solutions on Deforming Meshes," AIAA Paper 1997-1085, April 1997.
- ⁹Thomas, P. D. and Lombard, C. K., "Geometrical Conservation Law and Its Application," *AIAA J.*, Vol. 17, No. 10, Oct. 1978, pp. 1030–1037.
- ¹⁰Saad, Y. and Schultz, M. H., "GMRES: A Generalized Minimum Residual Algorithm for Solving Nonsymmetric Linear Systems," *SIAM Journal of Scientific and Statistical Computing*, Vol. 7, 1986, pp. 856–869.
- ¹¹O'Brien, D., *Analysis of Computational Modeling Techniques for Complete Rotorcraft Configurations*, Ph.D. thesis, Georgia Institute of Technology, 2006.
- ¹²Noack, R. W., "DIRTlib: A Library to Add an Overset Capability to Your Flow Solver," AIAA Paper 2005-5116, June 2005.
- ¹³Noack, R. W., "SUGGAR: a General Capability for Moving Body Overset Grid Assembly," AIAA Paper 2005-5117, June 2005.
- ¹⁴Johnson, W., "Rotorcraft Aerodynamics Models for a Comprehensive Analysis," *American Helicopter Society 54th Annual Forum Proceedings*, 1998.
- ¹⁵Tung, C., Caradonna, F., and Johnson, W., "The Prediction of Transonic Flows on an Advancing Rotor," *American Helicopter Society 40th Annual Forum Proceedings*, 1984.
- ¹⁶Potsdam, M., Yeo, H., and Johnson, W., "Rotor Airloads Prediction Using Aerodynamic/Structural Coupling," *American Helicopter Society 60th Annual Forum Proceedings*, 2004.
- ¹⁷Lim, J. W. and Strawn, R. C., "Prediction of HART II Rotor BVI Loading and Wake System Using CFD/CSD Loose Coupling," AIAA Paper 2007-1281, Jan. 2007.
- ¹⁸Pahlke, K. and van der Wall, B., "Chimera simulations of multibladed rotors in high-speed forward flight with weak fluid-structure-coupling," *Aerospace Science and Technology*, Vol. 9, No. 5, 2005, pp. 379–389.
- ¹⁹Dietz, M., KeBler, M., Krämer, E., and Wagner, S., "Tip vortex conservation on a main rotor in slow descent flight using vortex-adapted Chimera grids," *AIAA J.*, Vol. 45, No. 8, 2007, pp. 2062–2074.
- ²⁰Gopalan, G., Sitaraman, J., Baeder, J. D., and Schmitz, F. H., "Aerodynamic and Aeroacoustic Prediction Methodologies with Application to the HART II Model Rotor," *American Helicopter Society 62nd Annual Forum Proceedings*, 2006.
- ²¹Yu, Y. H., Tung, C., van der Wall, B., Pausder, H., 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," *American Helicopter Society 58th Annual Forum Proceedings*, 2002.
- ²²Prizadeh, S., "Three-Dimensional Unstructured Viscous Grids by the Advancing Front Method," *AIAA Journal*, Vol. 34, No. 1, Jan. 1996, pp. 43–49.
- ²³Lim, J., Private Communication, U.S. Army Aeroflightdynamics, Moffett Field, CA.
- ²⁴Schneider, O. and van der Wall, B., *Conditionally Averaging of Pressure Data from HART II*, DLR Institute Report IB 111-2006/08, Jan. 2006.
- ²⁵Thomas, J. L., Diskin, B., and Rumsey, C. L., "Towards Verification of Unstructured-Grid Solvers," AIAA Paper 2008-0666, Jan. 2008.
- ²⁶Lim, J., "An assessment of Rotor Dynamics Correlation for Descending Flight Using CFD/CSD Coupled Analysis," *American Helicopter Society 64th Annual Forum Proceedings*, 2008.
- ²⁷Lim, J., Tung, C., Yu, Y. H., Burley, C. L., Brooks, T., Boyd, D., van der Wall, B., Schneider, O., Richard, H., Raffel, M., Beaumier, P., Bailly, J., Delrieux, Y., Pengel, K., and Mercker, E., "HART-II: Prediction of Blade-Vortex Interaction Loading," *Proceedings of the 29th European Rotorcraft Forum*, 2003.
- ²⁸van der Wall, B., Private Communication, German Aerospace Center, Institute of Flight Systems, Braunschweig, Germany.

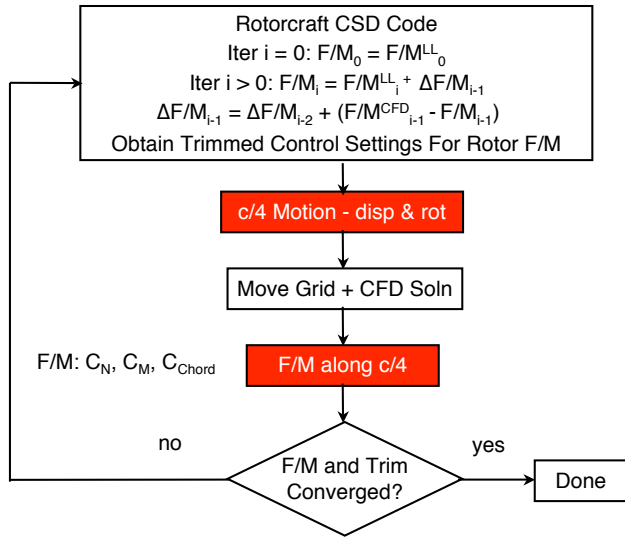


Fig. 1 Loose coupling methodology, adapted from Reference 16.




Fig. 2 HART-II model installed in the DNW wind tunnel.




Fig. 3 HART-II overset unstructured grid (13.6 million nodes).

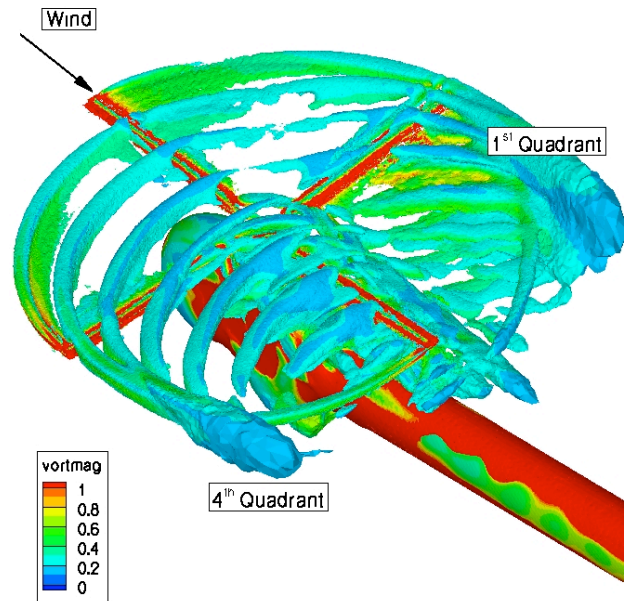


Fig. 4 Computed vortex structure near the rotor, HART-II BL case, AFDD motion data. Vortex visualization via isosurfaces of the second invariant of the velocity-gradient tensor ($Q = 0.0075$), colored by vorticity magnitude. Blade rotation counterclockwise viewed from above.




Fig. 5 Section normal force coefficient at $r/R = 0.87$, HART-II BL case; FUN3D results obtained with AFDD motion file. OVERFLOW-2 results and measured data shown for comparison.




Fig. 6 Section pitching moment coefficient at $r/R = 0.87$, HART-II BL case; FUN3D results obtained with AFDD motion file. OVERFLOW-2 results and measured data shown for comparison.




Fig. 7 Effect of mesh refinement on section chord force coefficient at $r/R = 0.87$, HART-II BL case, $BDF2_{opt}$ scheme.




Fig. 8 Effect of temporal order of accuracy on section chord force coefficient at $r/R = 0.87$, HART-II BL case, 13.6M node mesh.




Fig. 9 Convergence of section normal force coefficient at $r/R = 0.87$ with trim cycle for the HART-II MN case, 6.9 million node mesh.




Fig. 10 Convergence of section pitching coefficient at $r/R = 0.87$ with trim cycle for the HART-II MN case, 6.9 million node mesh.




Fig. 11 Convergence of collective and cyclic pitch settings with trim cycle for the HART-II MN case, 6.9 million node mesh.




Fig. 12 Convergence of tip elastic torsion with trim cycle for the HART-II MN case, 6.9 million node mesh. Error bars indicate blade-to-blade variations in the measured data.




Fig. 13 Convergence of tip elastic flap with trim cycle for the HART-II MN case, 6.9 million node mesh. Error bars indicate blade-to-blade variations in the measured data.




Fig. 14 Convergence of tip elastic lead-lag with trim cycle for the HART-II MN case, 6.9 million node mesh. Error bars indicate blade-to-blade variations in the measured data.




Fig. 15 Comparison of FUN3D results for section normal force coefficient using the AFDD motion data with motion data resulting from loose CFD/CSD coupling, along with measured data for reference.




Fig. 16 Comparison of FUN3D results for section pitching moment coefficient using the AFDD motion data with motion data resulting from loose CFD/CSD coupling, along with measured data for reference.

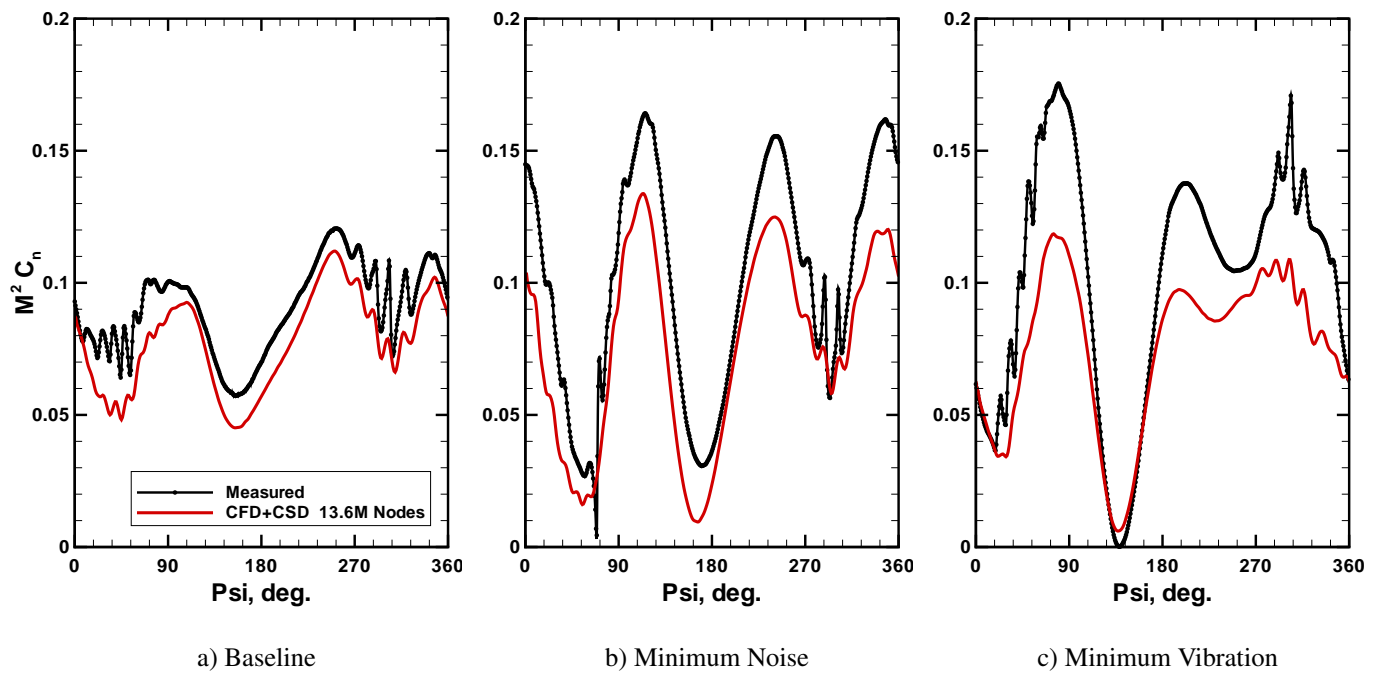


Fig. 17 Coupled section normal force coefficient at $r/R = 0.87$, 13.6 million node mesh.




Fig. 18 Coupled section normal force coefficient (mean removed) at $r/R = 0.87$, 13.6 million node mesh.




Fig. 19 Coupled section pitching moment coefficient at $r/R = 0.87$, 13.6 million node mesh.




Fig. 20 Coupled section pitching moment coefficient (mean removed) at $r/R = 0.87$, 13.6 million node mesh.




Fig. 21 Coupled elastic tip torsion $r/R = 1.0$, 6.9 million node mesh. Error bars on measured data represent blade-to-blade variation




Fig. 22 Coupled elastic tip flap $r/R = 1.0$, 6.9 million node mesh. Error bars on measured data represent blade-to-blade variation




Fig. 23 Coupled elastic tip lead-lag $r/R = 1.0$, 6.9 million node mesh. Error bars on measured data represent blade-to-blade variation