

Rotor Loads Prediction Using Helios: A Multi-Solver Framework for Rotorcraft CFD/CSD Analysis

Jayanarayanan Sitaraman*¹, Mark Potsdam†², Buvanewari Jayaraman‡³, Anubhav Datta§³, Andrew Wissink¶², Dimitri Mavriplis||¹, and Hossein Saberi**⁴

¹*University of Wyoming, Laramie, WY*

²*U.S. Army Aeroflightdynamics Directorate (AMRDEC), Moffett Field, CA*

³*NASA Ames Research Center, Moffett Field, CA*

⁴*Advanced Rotorcraft Technologies, Sunnyvale, CA*

We explore the use of the Helios high-fidelity rotorcraft simulation software for forward flight CFD/CSD simulation of the UH-60A rotorcraft, comparing computed results for three critical flight conditions. The approach used for moving-body CFD/CSD analysis in Helios applies Cartesian-based unsteady adaptive mesh refinement (AMR) for the off-body wake solution, and results with the enhanced wake solution are compared to traditional fixed off-body refinement. Results show airload predictions that are generally comparable to existing state-of-the-art CFD/CSD analysis codes. The enhanced wake resolution from AMR provides some improvement to vibratory airload predictions but the improvements are marginal.

I. Introduction

This paper presents a validation of the aerodynamic loading predictions of the UH-60A Black Hawk rotorcraft using the multi-solver rotorcraft CFD/CSD analysis framework – Helios. Helios¹ is being designed and developed by the DOD-sponsored High Performance Computing (HPC) institute for Advanced Rotorcraft Modeling and Simulation (HIARMS). Improvement of usability, accuracy and efficiency were the principal metrics utilized in the design and development of Helios. To this end, Helios provides several novel features that are different from existing state-of-the-art. These include (1) a multi-code unstructured/Cartesian dual-mesh CFD solver architecture,² (2) an automated and parallel overset grid assembling capability with implicit hole-cutting³ (3) the capability for solution adaption and high-order accuracy in the off-body solver⁴ (4) an unstructured patch force based fluid-structure interface⁵ and (5) a Python-based software integration framework that supports seamless and efficient data exchange between component modules.⁶

The fluid dynamics model in Helios differs from traditional monolithic CFD code architecture owing to its utilization of a heterogeneous meshing paradigm (Fig. 1). The meshing paradigm in Helios combines the advantages of unstructured grids (their ability to model complex geometry, resolving boundary layers using anisotropic meshes) and Cartesian adaptive grids (efficient high order algorithms, efficient storage, ease of solution based adaption). The use of multiple mesh paradigms provides the potential for optimizing the gridding strategy on a local basis. For the rotorcraft problem in particular, it is apparent that one should employ body fitted unstructured grids around the solid boundaries and then transition the gridding paradigm to solution adaptive Cartesian grids for wake capturing. The use of a heterogeneous meshing paradigm does involve several implementation challenges. For example, integrating different meshing paradigms into a

*Assistant Professor, Dept. Mechanical Engineering, AIAA Member

†Aerospace Engineer, Research Development and Engineering Command, AIAA Member

‡Research Scientist, Science and Technology Corp.,

§Research Scientist, Science and Technology Corp.,

¶Aerospace Engineer, Research Development and Engineering Command, AIAA Member

||Professor, Dept. Mechanical Engineering, AIAA Member

**Vice President




Figure 1. Overset near/off-body gridding paradigm used in Helios. Unstructured or curvilinear grids to capture geometric features and boundary layer near body surface, adaptive block-structured Cartesian grids to capture far-field flow features.

single large monolithic code is complex and has the disadvantage of forcing at least one of the gridding paradigms to be less optimized and flexible than the original stand alone solver. In Helios, we overcome this challenge by utilizing a multiple-mesh strategy that is implemented through the use of *multiple* CFD codes, each optimized for a particular mesh type. In addition to supporting heterogeneous grids and solvers, the software was also developed in a manner to minimize the analysis burden on the user. While an interface procedure is required for each specific solver, this interface must only be written once. Once the interface exists, the domain connectivity software retrieves the information it needs directly from the flow solvers and the grids themselves, so the human analyst does not need to provide input specific to each problem. The design paradigm used in Helios automates most analysis procedures, thus considerably simplifying its routine usage for production needs. For example, for the four-bladed UH-60A simulation, the only grid generation required is for a single blade (i.e. the user needs to create an unstructured grid that extends about one chord from the blade). The creation and partitioning of the four-bladed grid system, the generation/adaption of the off-body Cartesian grid system, the execution of the overset grid assembly, the execution of CFD/CSD coupling and data exchange are completely automated from then on and can proceed seamlessly without any user intervention.

The focus of this paper concerns primarily the prediction of aerodynamic loading and resolution of the vortex wake structure. We utilize the UH-60A flight test data obtained as part of the UH-60A airloads program (jointly sponsored by NASA and the U.S. Army) for all the validation and verification purposes. The flight test data set is quite comprehensive and contains measurements of rotor performance, sectional aerodynamic loading, structural loads, blade motion, control loads and aircraft states. It also spans an extensive range of flight conditions, including level flight and maneuvers. We choose three critical flight conditions for validation/verification purposes in this paper. They are (1) steady high speed forward flight (noted henceforth as C8534 based on the test data counter label) (2) low speed transition flight (C8513) and (3) high-altitude stall flight (C9017). Coupled fluid/structure simulations for these flight conditions have been performed for these flight conditions by several research groups with varying degrees of success (see references^{7,9,10}). The objectives of this paper are (1) to verify that the prediction capabilities of Helios are on par with the existing state-of-the-art (2) demonstrate solution/geometry based dynamic off-body mesh adaption capability in the complex operating condition involving unsteady flows and elastically deforming meshes and (3) characterize the improvements in efficiency and accuracy brought forth by the dynamic mesh adaption.

II. Technical Approach

Full fluid/structure coupled analysis requires the invocation of all the participant modules in **Helios**. These are (a) the parallel NSU3D code¹¹ as the unstructured near-body solver (b) the SAMARC code⁴ which combines the SAMRAI framework¹² and high-order ARC3DC code as the off-body solver (c) the parallel domain connectivity module PUNDIT³ (d) a patch force based fluid/structure interface (RFSI⁵) and (e) the RCAS¹³ rotorcraft comprehensive analysis which performs structural dynamic analysis and trim. The coupling of all these components is accomplished through a Python-based infrastructure with emphasis on preserving the modularity of the participating solvers. In addition to the advantages in efficiency and ease in code development, coupling existing mature simulation codes through a common high-level infrastructure provides a natural way to reduce the complexity of the coupling task and to leverage the large amount of verification, validation, and user experience that typically go into the development of each separate model. Brief descriptions of the implementations of all the component modules and the python infrastructure are outlined below. More detailed descriptions and validations of the Python-based infrastructure can be found in Refs.^{2,6}

II.A. Python Infrastructure

Python-based computational frameworks have been developed previously by other researchers¹⁴ as a means of coupling together existing legacy codes or modules. Such a framework has a number of advantages over a traditional monolithic code structure: (1) it is easier to incorporate well-tested and validated legacy codes rather than to build the capabilities into an entirely new code, (2) there is less code complexity in the infrastructure itself, so maintenance and modification costs are less, and (3) it is easier to test and optimize the performance of each module separately, often yielding better performance for the code as a whole. Essentially, Python enables the legacy solvers to execute independently of one another and reference each other's data without memory copies or file I/O. Further, the Python-wrapped code may be run in parallel using pyMPI or myMPI, with each of the solvers following its native parallel implementations.

II.B. Flow Solver

The flow solution methodology in **Helios** utilizes a heterogeneous meshing paradigm consisting of unstructured near-body grids and Cartesian off-body grids. Three distinct modules operate synchronously to facilitate the flow solution procedure. The modules are termed as the near-body-solver, the off-body solver and the domain connectivity module, respectively. In the following sections, we provide brief descriptions of the meshing paradigm and each of the participating solvers.

II.B.1. Meshing Paradigm

The meshing paradigm consists of separate near-body and off-body grid systems. The near-body grid typically extends a short distance from the body, sufficient to contain the boundary layer. This grid can be an unstructured tetrahedral or prismatic grid that has been extracted from a standard unstructured volume grid or generated directly from a surface triangulation using hyperbolic marching. The reason for using unstructured grids in the near-body region is to properly capture the geometry and viscous boundary layer effects, which are difficult or impossible to capture with Cartesian grids alone. We further note that either structured or unstructured grids will work equally well in the near-body region from the point of view of our infrastructure. Away from the body the near-body grid solution is interpolated onto a Cartesian background mesh with the aid of the domain connectivity algorithm (see section II.B.4). This transition normally occurs at a distance wherein the sizing of the near-body grid cells is approximately commensurate with the sizing of the Cartesian mesh in the off-body region.

The off-body grid system consists of a hierarchy of nested refinement levels, generated from coarsest to finest. In the Structured Adaptive Mesh Refinement (SAMR) paradigm,¹⁵ the coarsest level defines the physical extents of the computational domain. Each finer level is formed by selecting cells on the coarser level and then clustering the marked cells together to form the regions that will constitute the new finer level. Solution-based refinement progresses as follows: physical quantities and gradients are computed at each Cartesian cell using the latest available solution and those cells that hold values deemed to require refinement are marked. The marked cells are then clustered to form the new set of patches that constitute the next finer level. The process is repeated at each new grid level until the geometry and solution features




Figure 2. Block structured AMR grid composed of a hierarchy of nested levels of refinement. Each level contains uniformly-spaced logically-rectangular regions, defined over a global index space.

are adequately resolved. We note that this entire procedure is automated within the software and no user intervention is required. The procedure of adaptive mesh refinement is graphically illustrated in Figure 2.

An example of the overset near-body/off-body meshing strategy is given in Figure 1, which shows the meshes for flow computations over a helicopter fuselage and blade. Here, the mixed-element unstructured near-body grid envelops the surface of the fuselage and blade, while a multi-level Cartesian off-body grid extends from the outer boundary of the near-body grid to the far-field boundary. The two sets of meshes overlap in the so-called fringe region, where data are exchanged between the grids.

II.B.2. Near-body solver

For the unstructured solver in the near-body region, we employ the NSU3D¹¹ code, which is an implicit node-centered Reynolds-Averaged Navier-Stokes (RANS) code capable of handling arbitrary unstructured mesh elements. NSU3D utilizes a node-centered finite-volume scheme, that is spatially second-order accurate. Time-accurate computations utilize second-order accurate backwards-Euler time-stepping along with dual-time-stepping for iterative convergence of the non-linear problem at each physical time-step. The solution scheme is comprised of a non-linear multigrid algorithm that employs a three-stage Runge-Kutta scheme with line-Jacobi relaxation as a smoother at each grid level. The code employs an edge-based data structure, which facilitates flux computations on the edges of the median-dual control volume. NSU3D has been extensively tested and benchmarked in stand-alone mode (see refs^{16,17}).

II.B.3. Off-Body Flow Solver

The off-body solver is an Euler equation derivative of the ARC3D code,¹⁸ that is optimized for isotropic Cartesian grids. It uses an explicit 3rd-order Runge-Kutta time-integration scheme and has spatial differencing options ranging from second- to fifth-order accurate. As mentioned earlier, the code is compiled as a library that is callable by SAMRAI, which is responsible for the overall management of the off-body grid generation, adaption, parallel domain decomposition and load balancing. The combined SAMRAI/ARC3D package is referred to as SAMARC.⁴

II.B.4. Overset Methodology and domain connectivity module

The overall solution procedure for the overset meshes is as follows. At each iteration step, the solutions of the fluid equations in each mesh system are obtained independently with the solution in each fringe region being specified by interpolation from the overlapping “donor” mesh as a Dirichlet boundary condition. At the end of the iteration, the fringe data are exchanged between the solvers so that the evolution of the global solution is faithfully represented in the overset methodology. Further, if one or more of the meshes is moving or changing, the fringe regions and the interpolation weights are recalculated at the beginning of

the time step. For a steady-state calculation, this procedure is repeated for each iteration step until solution convergence is attained in both the near- and off-body meshes.

The domain connectivity module (PUNDIT³) is responsible for performing all the overset grid assembly operations. These include determination of donor-recipient relationships and appropriate interpolation schedules between multiple overlapping meshes. For static meshes, these operations are done once, at the beginning of the computation, while, for the more general case of moving or adapting meshes, the determinations of donors and weights has to be repeated within the time-stepping or iteration loop.

All the overset grid assembly operations are fully automated and do not require any geometry input from the user. In addition, PUNDIT performs all the domain connectivity operations on the partitioned grid system in a fully parallel environment. PUNDIT adopts the implicit hole cutting procedure followed by NAVAIR-IHC.¹⁹ The core idea of this approach is to retain the grids with the finest resolution at any location in space as part of the computational domain and interpolate data at all coarser grids in this region from the solution on the fine grid. This results in the automatic generation of optimal holes without any user specification as in the case of explicit hole-cutting. Moreover, the implicit hole-cutting procedure produces an arbitrary number of fringe points based on mesh density compared to traditional methods which use a fixed fringe width (usually single or double). The critical parameter that quantifies the quality of a grid cell or node is termed as “resolution capacity”. PUNDIT nominally uses the cell volume as the resolution capacity for a grid cell and the average of cell volumes of all associated grid cells as resolution capacity for a grid node. The near-body to near-body and near-body to off-body domain connectivity procedures are separated in PUNDIT to facilitate automatic off-body grid generation and to improve efficiency (Ref.²⁰).

II.C. Comprehensive Analysis and Fluid/Structure Interface

Accurate modeling of the UH-60 requires establishing a trim condition where all rotor aerodynamic forces are balanced. A structural model is also needed to properly model the structural dynamics of the flexible rotor. These structural dynamics and trim operations are provided by a rotorcraft comprehensive analysis package that is called by Helios. That is, at each timestep Helios computes the aerodynamic loads using the dual-mesh CFD analysis and then calls the comprehensive analysis package which sets the rotor pitch angle, structural deflections, etc. Aerodynamic loads computed by CFD are transmitted to the comprehensive package through a fluid/structural interface. This section discusses details of this procedure.

II.C.1. Rotorcraft Comprehensive Analysis System

Rotor structural dynamics and trim coupling are handled by the Rotorcraft Comprehensive Analysis System (RCAS).¹³ RCAS is a rotorcraft comprehensive analysis (CA) platform developed and maintained by the U.S. Army Aeroflightdynamics Directorate (AFDD) in partnership with Advanced Rotorcraft Technology, Inc. over the past ten years. In general, RCAS contains state-of-the-art multibody, large deformation, composite nonlinear beam finite-element modeling (FEM), full aircraft trim and vehicle dynamics, and lower-order unsteady lifting-line aerodynamic models. It also contains a suite of solution procedures for level flight trim solution, flutter and aeroelastic stability, ground/air resonance capabilities, unsteady maneuvers, and external slung load/missile release capabilities. In **Helios**, the RCAS component supplies the CSD and trim modules. Importantly, it provides the delta-coupling capability that is central to multi-disciplinary CFD-CA coupled analysis. Additional capabilities such as tight-coupling, multiple rotors and rotor-fuselage configurations will be included in future versions of **Helios**. RCAS uses a Hilbert-Hughes-Taylor time marching with an azimuthal time step of $\Delta\psi = 5^\circ$. The RCAS trim model used for these calculations is a shaft fixed, moment trim, targeting the estimated thrust and measured hub moments from the standard Airloads database, to calculate the rotor control pitch angles.

II.C.2. Fluid/structure interface

The rotor fluid-structure interface (RFSI) transfers surface forces from the near-body CFD solver to the structural dynamic model of the rotor. The implementation is based on the exact coupling method described in Ref.⁵ The method consists of a patch force based interface and a modified delta procedure. The interface does not require 2-D sectional airloads from CFD, and instead, imposes each and every surface force (integrated fluid stresses over surface elements) on the structure by an exact calculation of virtual work. Thus, it is both exact (conservative in energy and preservative in total airloads) and generic (applicable

to arbitrary grids, geometries, and structural models). The modified delta procedure also bypasses the requirement for 2-D sectional airloads and implements the delta-quantities directly at the finite element level (of the structure). The patch force interface can be collapsed, as an option, on to a reduced set of span-wise points to re-create sectional airloads. In this form, the interface (no longer conservative but still preservative) allows the implementation of the conventional delta procedure. It is this form which is used here for coupling with RCAS.

The reduced interface provides an efficient sectional airloads calculation, without the need for this calculation in the CFD code itself. An example is shown in Figure 3 (from Ref.⁵). Figure 3(a) shows typical patch forces at zero azimuth for a high-speed forward flight case from a structured CFD solver. The beam excitations from each patch force – accumulated on a limited number of span-wise points (say x) on the beam axis (Figure 3(b)) by simply distributing them proportionally among two neighboring points around the span-wise station to which it is assigned – if divided by Δx provide an estimate of sectional airloads. The only difference is that whereas the sectional airloads contain only the sectional pressure/shear solution, the reduced airloads contain all patch forces within a domain (Figure 3(b)). As the number of span-wise points increase, the reduced airloads approach sectional airloads. An example is shown in Figure 3(c) where the patch forces have been accumulated on a set of 27 span-wise points (with a point at 96.5% R). The isolated example shown here uses a structured near-body CFD to serve the purpose of a consistent verification. The coupled solutions in the rest of the paper use unstructured near-body CFD.

RFSI ensures consistency of geometry and coordinate system frames between CFD and CSD solvers. Currently, the module works for conventional single rotor configurations, although the capability to handle multi-element rotors (advanced geometry and flaps) and multiple rotor configurations is under development.

II.C.3. Mesh Motion Module (MMM)

The final component that plugs into *Helios* are the mesh motion and deformation modules. There are two specific sub-modules: *pydeform* which is used for prescribed motion cases, and *rdeform* which is used for full aero-elastic coupling. In both instances, the module uses the original mesh coordinates and moves and deforms the mesh to account for rotor rotation and blade surface deformations based on the azimuthal and radial location of a grid point.

II.C.4. CFD-CA Delta Coupling Analysis

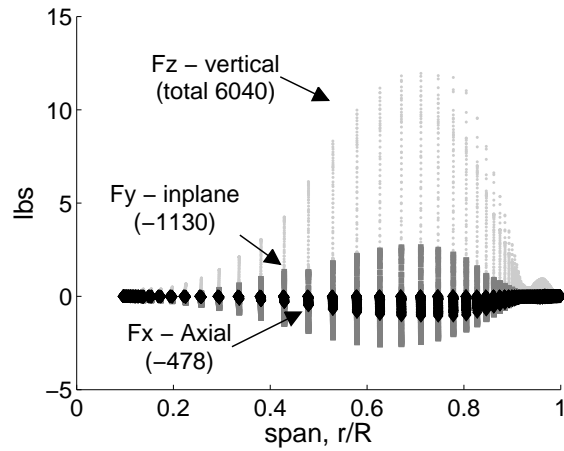
The delta-coupling or loose-coupling formulation²³ takes advantage of the periodic nature of the rotor dynamics. In this scenario, the aerodynamic forces on the blade are obtained over a period and input into the comprehensive analysis code. In turn, the periodic deflections are used to deform and move the mesh for the aerodynamic simulations over the next period. This sequence of CFD and CA simulations is repeated until the load and deflections converge to a periodic steady-state, thereby ensuring time-accuracy of the response harmonics and, simultaneously, trim. In practice, the coupling is performed every N time steps of the CFD time-marching solution, where N is sufficient to construct airloads over one entire rotor revolution using the solution from all blades. These deflections are used to move the grid over the next period of rotor revolution and so on until the loads and deflections converge.

III. UH-60 Black Hawk


III.A. Flight test data

A unique and extensive flight test database exists for a UH-60A Black Hawk helicopter in level flight and transient maneuvers.²² The data were obtained during the NASA/Army UH-60A Airloads Program. The database provides aerodynamic pressures, structural loads, control positions, and rotor forces and moments, allowing for the validation of both aerodynamic and structural models. The test matrix contains a range of advance ratios and gross weight coefficients, as shown in Figure 4, with the test points investigated here indicated. Absolute pressures were measured at $r/R = 0.225, 0.40, 0.55, 0.675, 0.775, 0.865, 0.92, 0.965,$ and 0.99 (Figure 5) along the blade chord and integrated to obtain normal force, pitching moment, and chord force.

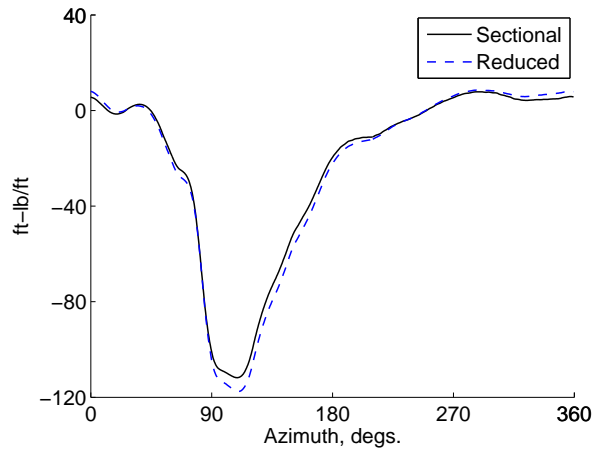
The data have undergone a significant amount of careful investigation,²¹ however, some discrepancies have not yet been resolved. Measured rotor thrust was determined from the gross weight of the helicopter



(a) CFD surface patch forces on a blade at $\psi = 0^\circ$; F_x is to tip, F_z is in direction of thrust; F_y is to leading edge (for counter clock-wise rotation).



(b) A limited set of span-wise points on CSD beam axis.



(c) Accumulated airloads from reduced interface (patch forces collapsed on to a set of CSD beam axis points) compared with sectionally integrated 2-D airloads; pitching moment (1/4-c) at 96.5% R.

Figure 3. Calculation of 1-D beam excitation from 3-D CFD surface forces; Verification of an exact patch force interface by reduction on to a limited set of span-wise points and comparison with 2-D sectional airloads.

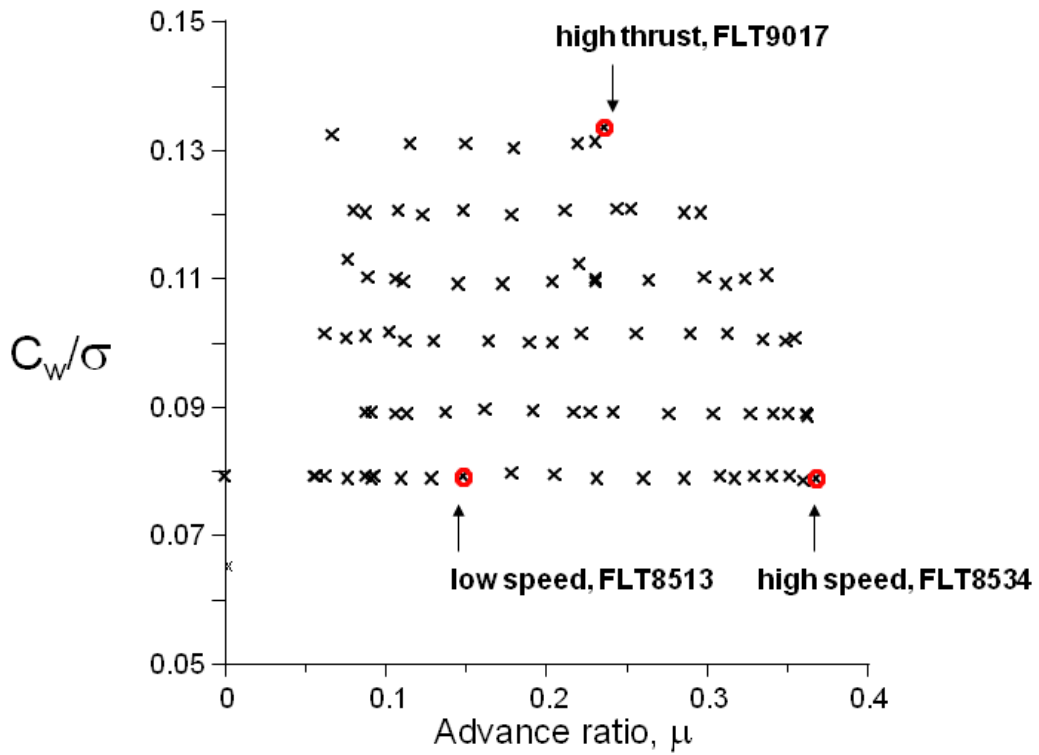


Figure 4. UH-60 test matrix




Figure 5. UH-60 blade with pressure sensor locations

plus estimates for the fuselage and tail rotor loads. Measured hub moments, roll and pitch, were determined from an upper shaft bending moment gauge. However, integration of the measured pressures over the rotor results in poor agreement with measured thrust and moments. For example, for the high speed test point, the integrated thrust is 10% higher and the total integrated hub moment is 50% larger with an 80-degree phase difference compared to the measured values. Consequently, it is clear that there is some uncertainty in the aircraft trim condition, and there will be some discrepancy in comparison of mean airloads values. Errors in the blade pressures can have large effects on integrated section pitching moments. Faulty trailing edge pressure taps have been discovered in this dataset that considerably skew the pitching moment mean values. For this reason, all plots of pitching moment have the mean removed.

The UH-60A master input database, available to approved researchers, has been used to define the elastic UH-60A four-bladed rotor model. The database contains geometric, aerodynamic, and structural material properties. Figure 5 shows the blade planform and pressure transducer locations. The blade has a radius of 322 inches, and the 20 deg swept tip begins at $r/R = 0.929$. The solidity (blade area/disk area) σ is .0826 with a root cut-out of 14% and there is about 16 degs of nonlinear twist. The chord of the unswept portion of the blade is 20.76 inches. The nominal rotor speed is 258 RPM. Further details of the blade can be found in Ref.²¹ The geometry and grids have been developed from CAD-based surface grid definitions in PLOT3D format distributed to the UH-60 Airloads Workshop participants by Sikorsky Aircraft in 2010. In comparison with previous theoretical definitions of the UH-60 blade,²² the CAD-based definition corrects the tab location, strings the airfoils along the span more correctly, and modifies a small region at the tip. None of the major blade characteristics (twist, chord, airfoil sections, planform) are changed from that used in a large range of previous investigations. Three flight conditions are analyzed in the paper, shown in Table 1.


Table 1. UH-60 flight conditions. RM=rolling moment, PM=pitching moment, AOA=angle of attack.

Flight Cond	Advance Ratio	Hover Tip Mach	AOA (deg)	Density (slug/ft ³)	Weight coeff/ σ	Thrust coeff/ σ	Torque coeff/ σ	Hub RM (ft-lb)	Hub PM (ft-lb)
8534	0.368	0.642	-7.31	0.0020823	0.079	0.084	0.00905	1510	-1042
8513	0.153	0.638	+0.75	0.0021717	0.079	0.076	0.00384	239	-1367
9017	0.237	0.666	-0.19	0.0013242	0.135	0.126	0.01058	97	34


IV. Computational Details

IV.A. Near-body Unstructured Grid Generation


Two unstructured near-body grids of the UH-60 blade have been used in this work, a coarse grid and fine grid. The relevant sizes of the two grids are detailed in Table 2. Both grids were generated in the same manner. A surface grid was developed using the commercial GRIDGEN/Pointwise software, and a volume grid based on this surface grid was created using AFLR3 from Mississippi State University. The coarse surface mesh uses a structured mesh on the main part of the blade with fully unstructured tip caps (Figure 6(a)). The fine surface mesh uses a quasi-structured approach (Figure 6(b)), where structured cells are generated from the leading edge aft and the trailing edge forward. They are then transitioned to triangles after some distance where the cells become isotropic. In both cases, the volume grid generation process converts the quadrilateral surface cells into two triangles. The volume meshes are comprised of prismatic cells in the boundary layer and tetrahedral cells in the farfield, with some pyramidal cells for transition (Figures 6(c) and (d)). The purpose of initially developing structured grids on the surface is to allow for anisotropic cells and stretching, which is not generally available from unstructured surface mesh generators. For the viscous boundary layer, the initial wall normal spacing on both grids is 5e-05 inches ($y^+ < 1$). The outer boundary spacing and distance for each grid relative to the rotor chord are also shown in Table 3. For the fine mesh, a utility has been used to trim the original outer boundary back (Figure 6(f)), resulting in an irregular outer boundary compared with the coarse mesh outer boundary box (Figure 6(e)).




(a) coarse surface




(b) fine surface




(c) coarse volume



(d) fine volume



(e) coarse outer boundary



(f) fine outer boundary

Figure 6. UH-60 unstructured near-body rotor unstructured meshes, coarse (left) and fine (right). (a)–(b) surface meshes; (c)–(d) cut through volume mesh; (e)–(f) outer boundary (overlaps with Cartesian off-body).

Table 2. UH-60 unstructured near-body rotor meshes. Per-blade statistics shown, full computational mesh consists of 4-blades.

	Nodes	Cells	Prisms	Tets	Blade Surf Faces	Outer Bdry Faces
Coarse	4.5M	11.3M	7.9M	3.4M	48486	16044
Fine	15.4M	36.5M	27.4M	9.0M	122314	61338

Table 3. UH-60 surface grid parameters, near-body blade mesh. (OB=outer boundary, LE=leading edge, TE=trailing edge, AR=aspect ratio)

	OB Spacing	OB Dist	LE Spacing	TE Spacing	Max Cell AR	# Cells across TE
Coarse	0.22c	1.6c	0.0029c	0.014c	37	2
Fine	0.09c	0.7c (cut)	0.0067c	0.0067c	75	10

IV.B. Off-body Cartesian Structured Grids

The off-body grids are detailed in Table 4. Both fixed refinement and geometry/solution adaption cases are run for comparison. Fixed refinement specifies the extent of the finest level grid domain at the start of the calculation: $\pm 1.16R$ in X and Y and $\pm 4.3c$ in Z is used. The finest off-body grid spacing used with the two near-body unstructured grids is typically designed to match the outer boundary spacing, in order to obtain reasonable and consistent overlap. This finest spacing is determined in SAMARC obliquely by the spacing specified at the outer boundary (level 1) and the total number of levels to be used. For the coarse grid calculations the outer boundary spacing is 5.8 chords (120 in.), while for the fine mesh calculations the outer boundary spacing is 6.2 chords (128 in.). The nomenclature of the differing off-body grids is then designated by levels 6, 7, and 8 (L6, L7, L8), corresponding to finest mesh spacings of 0.18, 0.10, and 0.05 chords (3.75, 2.0, and 1.0 inches), respectively.

Table 4. UH-60 Off-body Cartesian meshes.

	Refinement levels	Finest Mesh Spacing	Nodes non-overlap	Nodes overlap	Max # Blocks
Coarse L6	6	0.18c	4.3M	9.1M	495
Fine L7	7	0.10c	22.1M	39.1M	1208
Fine L8	8	0.05c	145.7M	218.0M	2488
Fine L8 adapt	8	0.05c	24.0-125.0M	48.0-208.0M	max 5200

IV.C. Adaptive Mesh Refinement

Adaptive mesh refinement (AMR) is performed based on geometry proximity and solution features. Geometry-based refinement is performed to encompass the near-body grids and then, based on an input tag buffer size, extends a certain distance beyond this minimum extent (see, for example, the fine grid region surrounding the blade in Figs. 7(a) and (c)). In this case, the geometry-refinement buffer is 8 cells. Vorticity is used as the solution adaption criteria. Lower thrust cases use a non-dimensional value of vorticity of 0.0045 (based on speed of sound and a grid reference length of one inch). The high thrust coefficient case increases this to 0.005 due to the increased levels of vorticity in the flow. The solution threshold value must be tuned to each particular flight condition and geometry/grid. Future Helios developments will automate this for the user with a normalized function based on the Q criteria.²⁴ For the cases with adaption, the point count continuously increases as the wake is convected downstream and resolved. The number of points for the adapted case evolves over time and depends on the length of the calculation and solution adaption criteria. There is no adaption cut-off. Figure 8 shows a plot of the grid point count in the L8 adapt cases over the course of the simulation. In all cases the L8 adapt calculations were initiated from the L7 fixed refinement

solution. Even after 144 adaption cycles (one rotor revolution) the point count is less than that obtained with L8 fixed refinement, which is more limited in the extent of wake capture downstream of the rotor. Specifically, the adapted grids for the high speed, low speed, and high thrust coefficient cases have generated 67, 59, and 125 million grid points, respectively, after one rotor revolution (only half a rotor revolution for the high speed case). This is in comparison to the L7 fixed grid with 22 million points, which was used as the starting point, and the L8 fixed grid with 146 million points.

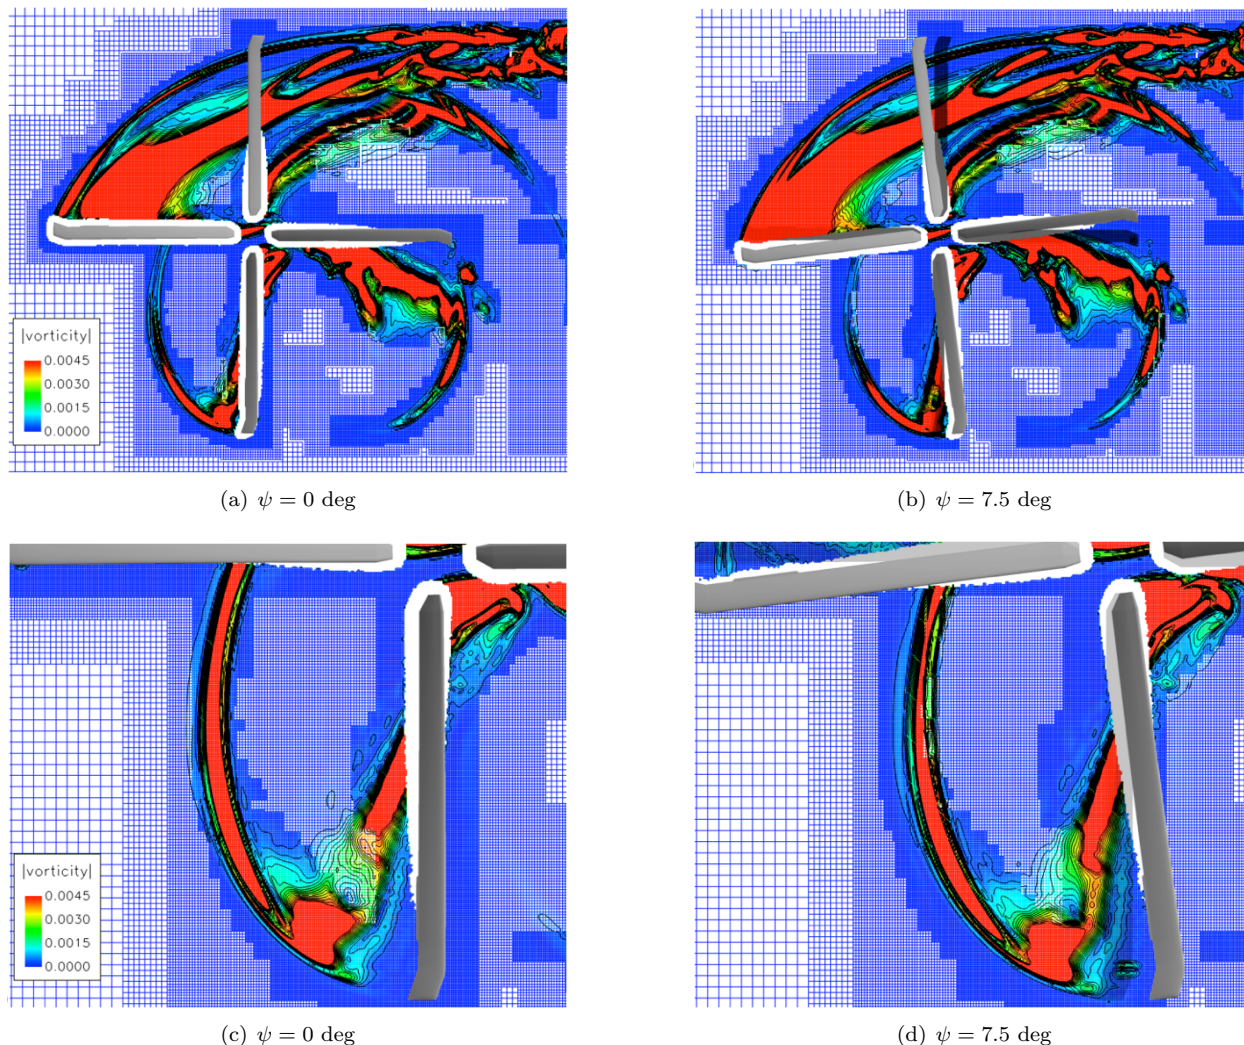


Figure 7. High speed wake. Vorticity magnitude overlaid on AMR mesh at (a) $\psi = 0$ deg and (b) $\psi = 7.5$ deg; (c)–(d) close up of third-quadrant tip vortex at these azimuthal locations.

Geometry and solution-based refinement are performed concurrently every 2.5 degs of rotor rotation. This is problem dependent, but for conventional rotorcraft the adapt frequency is driven by the rotor rotation speed and the movement of the rotor tip (advance ratio < 1). Figure 7(a/c) shows an AMR mesh at a particular time ($\psi = 0$ deg) and the corresponding vorticity magnitude contours for the high speed case. Also shown are the vorticity contours 7.5 degs later (Figure 7(b/d)) on the same mesh. The maximum vorticity value shown (red) is the refinement criteria (0.0045). It is seen that based on the size of the default solution-based buffer region, the high vorticity regions (tip vortices) are just exiting the refined region. However, the blade geometry has moved well beyond the fine grid geometry-refinement boundaries, and is, therefore, the limiting factor setting the required frequency of adaption. This could be remedied by increasing the geometry-based tag buffer region, and would be particularly beneficial for low speed cases in order to significantly reduce the adaption frequency. However, the cost of an adaption is roughly on the order of 5 SAMARC steps, or one total Helios step, and is, therefore, relatively efficient given the significant

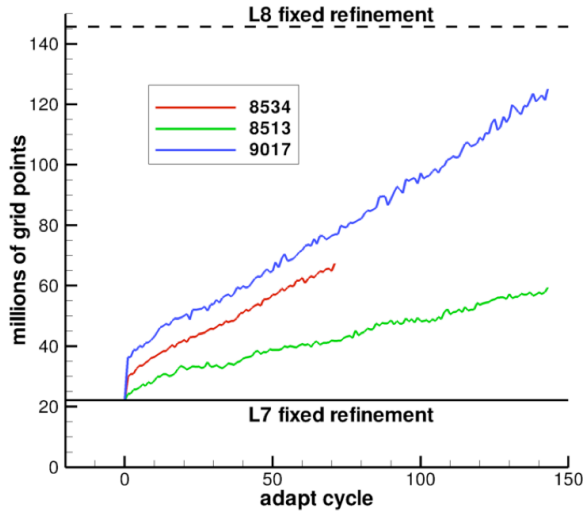


Figure 8. L8 adaptive grid point count for the three UH-60A flight conditions, compared with L7 and L8 fixed refinement.

regridding and data communication that is being performed. This amounts to 1% of the total calculation cost.

IV.D. Flow Solver Inputs

The 5th-order ARC3D spatial scheme is employed for off-body calculations. Time integration in the off-body solver uses an explicit time-dependent 3-stage Runge-Kutta scheme which has a CFL limitation on the size of the timestep for stability. Finer off-body grids require smaller timesteps. For the unsteady coarse grid calculations with 6 off-body levels, a physical time step of $\Delta\psi = 0.25$ degs is used. This corresponds to an off-body CFL=1.67, which is slightly high given the target stable CFL is 1.0, but the solution remained stable with this setting. Enforcing the CFL limitation with finer off-body meshes requires either 1) reducing the physical timestep $\Delta\psi$, or 2) applying sub-steps within each physical timestep. We chose to use a combination of these two options. For the medium (7 level) and fine (8 level) off-body grid systems a $\Delta\psi = 0.1$ degs is employed, and sub-steps are applied to maintain a CFL < 1. Table 5 shows the parameters used. During the off-body substepping, the overset boundary conditions within the off-body solver are updated, however, the near-body overset boundary points cannot be updated due to inconsistent solution methodologies during sub-stepping/sub-iterating. Since the most significant cost of the Helios solution is the near-body solver, the off-body substepping provides an efficient alternative to unnecessarily reducing the global time step below that required to capture the physics.

Near-body NSU3D calculations use an iterative scheme to converge the solution at each physical time step. An implicit line solver is applied in the prismatic boundary layer part of the grid, along with multigrid acceleration. The number of subiterations are chosen based on the drop in subiteration residual, convergence of the global rotor forces (C_T , C_Q) at each time step (Figure 9), and convergence of the section airloads from revolution to revolution. Such convergence criteria typically require at least one order of magnitude reduction in the subiteration residual. Table 5 lists the near-body subiterations and multigrid levels used for the UH-60 calculations.

Table 5. Unsteady CFD time-stepping parameters.

Off-body Levels	Time step (deg Azimuth)	Off-body Substeps	CFL	Near-body Subiters	Near-body MG Levels
6	0.25	1	1.67	25	1
7	0.10	2	0.66	15	3
8	0.10	3	0.88	15	3

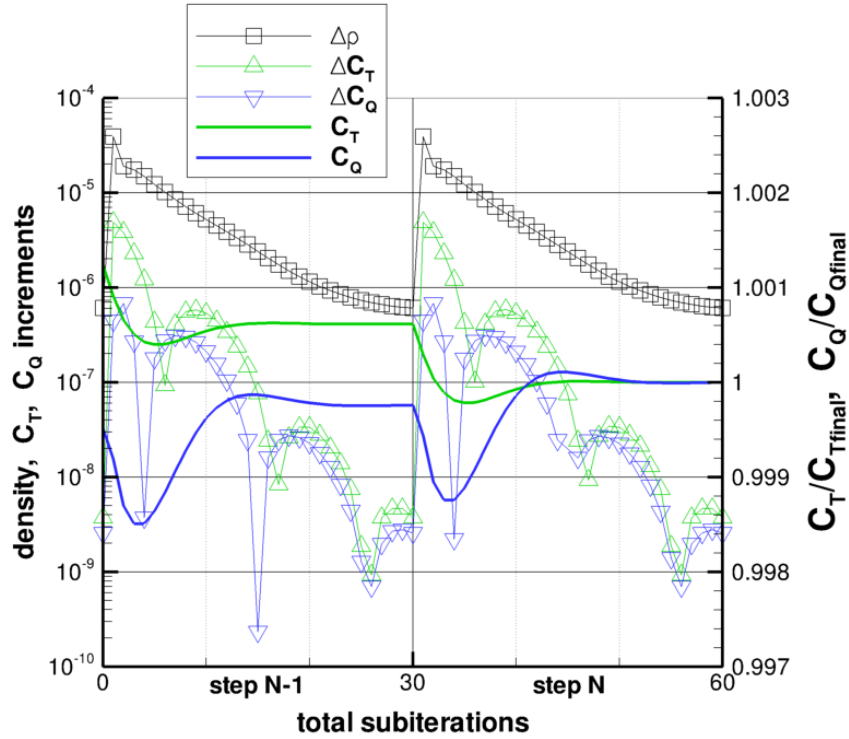


Figure 9. NSU3D fine grid subiteration convergence. *left axis: density ρ and force C_T , C_Q residual increments; right axis: normalized C_T and C_Q convergence (local iteration value divided by final value).*

IV.E. CFD/CSD Coupling

For the both the coarse and fine grids fully coupled CFD/CSD results are obtained. For the coarse grid, results are obtained using fixed L6 refinement. On the fine grid, both the low speed transition and high speed (C8513 and C8534) cases have been fully coupled using L8 adapt. All three cases on the fine grid have, however, been computed using prescribed motions from coarse grid coupled calculations. The trim conditions (thrust, hub moments, and angle of attack) are taken from the UH-60 Airloads database with corrections made for fuselage, empennage, and tail rotor lift. The trim conditions are shown in Table 1. For the coarse grid coupling is performed every 180 degs of rotor rotation. For the fine grid the frequency is increased to every 90 degs. All cases are run for at least 6 coupling iterations, which is generally enough for engineering convergence of the coupling procedure.

IV.F. Solution Timing

All cases were run on 256 processors of a Dell PowerEdge M610 (DOD MHPCC mana). Solution times broken down by Helios module are shown in Table 6. The near-body (NSU3D) cost is fixed since the near-body grid is unchanged. The off-body solution time varies with the number of grid points. In all cases, the timings are the averaged values for all the flight conditions. Due to the nature of the adaption, the timings for these cases are approximate. The near-body and off-body timings are generally very repeatable between runs. It is seen that the near-body flow solver dominates the cost of the solution, even for the L8 grid which has 9.5 times as many off-body as near-body points, making the off-body solver almost 40 times more efficient on this grid. It should be noted, however, that the off-body scheme is structured, inviscid and explicit, while the near-body is unstructured, viscous and implicit. The domain connectivity cost is quite variable (min/max times are quoted), particularly with finer off-body grids. At best, PUNDIT requires only 7% of the time step, although at worst the cost is upwards of 40%. At present we are unclear of the cause. but we have found disk access speeds on the mana system are highly variable and can at times be very slow. Although communication in the domain connectivity software are not file based, the inter-processor

communication takes place in core memory using MPI, information is written to a centralized log file and the consequent disk access may be introducing the variability in timing. The issue is under investigation. The total times quoted in the table are based on minimum domain connectivity times. It is seen that on average the cost of an adapted solution is on par with a fixed refinement case, with the extra cost due not to adaption but domain connectivity.

Table 6. Breakdown of time per timestep in the CFD solution (seconds/percentage), computed on 256 cores of a Dell PowerEdge M610 system.

	Near-body	Off-body	Adaption Overhead	Domain Conn (min/max)	Total
Coarse L6	13.1 (82%)	0.31 (1%)	–	1.28 (7%)	15.8
Fine L7	27.5 (86%)	1.3 (4%)	–	2.2/5.3 (7%)	31.8
Fine L8	27.8 (71%)	6.8 (17%)	–	3.7/14.8 (10%)	38.9
Fine L8 adapt	27.0 (54%)	4.0 (17%)	0.8 (2%)	8.9/21.3 (18%)	50.0

V. Results and Discussions

The results are organized as follows. We first present the convergence of the fluid/structure simulation strategy. Following that more detailed comparisons of sectional aerodynamic loading (both time history as well as harmonics) with flight test data are demonstrated to verify the prediction capability of Helios. Finally, we illustrate comparisons of wake visualizations with and without solution-based AMR.

V.A. Convergence of Fluid/Structure coupling

Fluid/structure coupling involves exchange of loads/displacements according to a prescribed schedule until convergence of the aerodynamic/structural dynamic loading is attained based on the periodicity of the rotor. In Fig. V.A we compare the evolution of steady sectional aerodynamic normal force variation over the span of the rotor blade with coupling iteration. It can be noticed that the aerodynamic loading approaches the flight data with coupling iterations and finally shows fair correlation with the measured flight data. Aerodynamic loading (normal force) history at a radial location close to the tip of the blade (96.5% span) is also shown in Fig. V.A. The high speed forward flight condition shows the best convergence and requires only six coupling iterations for the variations to be within plotting accuracy. The low speed and stall cases require further iterations for tight convergence – though the general pattern of vibratory and oscillatory airloads are already in place.

V.B. Airload Comparisons

Airloads comparisons are made in this section between the CFD/CSD coupled coarse near-body/L6 off-body and fine near-body/L8 adaptive grid systems. As described previously, prescribed motion cases were also run for all the flight conditions using L7 fixed refinement, L8 fixed refinement, and L8 adaptive with the fine near-body grid. The differing wake resolutions had minor impact on section airloads and global forces and moments (rotor thrust and torque). As would be expected, torque was most noticeably dependent on near-body grid resolution.

V.B.1. High speed forward flight condition (C8534)

The high speed forward flight condition is characterized by high vibratory loads owing to returning wake interactions on the advancing side. Results for this case shown in Fig. 11 indicate good prediction of the advancing blade interaction compared to other state-of-the-art computations.⁷ Note the lift and pitching moment impulses around 90 degs blade azimuth, indicating resolution of some probable wake-induced vibratory loads on the advancing side. This phenomena is a key contributor to the vibratory lift phase and 3/rev flap bending moment.

The results for the L6 and L8 adaptive grids, which differ significantly in both near-body and off-body grid resolution, are quite similar in their ability to capture the salient features of the aerodynamic loading

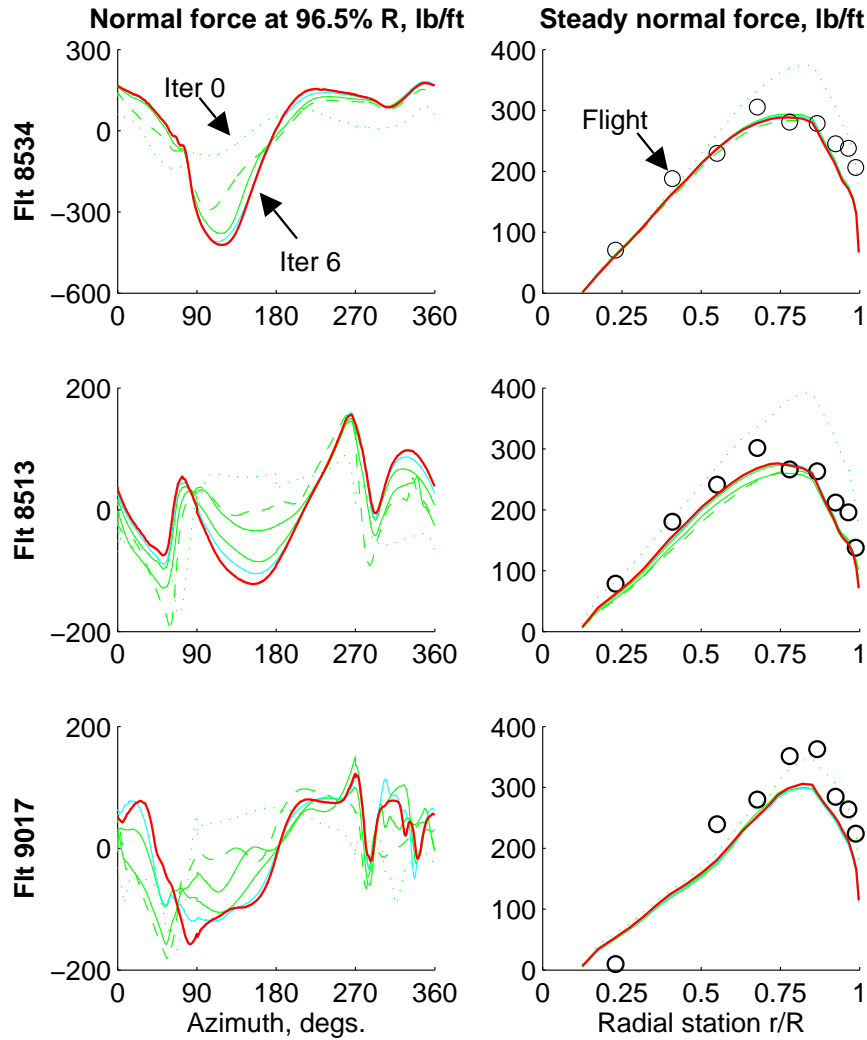


Figure 10. Convergence of oscillatory (mean removed) and mean normal forces.

waveform. It appears the L8 adapt grids do show more vibratory characteristics, likely owing to better resolution of the wake, but the difference is slight. Other studies⁷ have indicated that improved resolution of both the near- and off-body grids is required for optimal load prediction.

From basic rotor dynamics it is possible to infer that the first two harmonics (of the rotor rotating frequency) of the sectional aerodynamic loading do not contribute to fixed frame vibration, if the number of blades in the rotor is greater than or equal to four. Therefore it is a common practice in the rotary wing community to characterize the vibratory load prediction by plotting sectional aerodynamic loading after filtering the first two harmonics of variation. The vibratory part of the aerodynamic loading (3/rev and higher) for the UH-60A high speed forward flight condition is shown in Fig. 12. The overall waveform of the vibratory loading is captured quite well by both mesh systems. Minor improvement in the magnitude of vibratory lift can be noticed from the adaptive solution (denoted by L8a in plots). It is worth noting that although the magnitude of vibratory lift from Helios does show improvement from state-of-the-art predictions, it still remains underpredicted compared to the flight test data.

Basic rotor dynamics also indicates that $N_b - 1$, N_b and $N_b + 1$ (N_b =number of blades) harmonics of rotor frequency in rotating frame loading are the primary contributors to fixed frame vibration. Therefore, for the four-bladed UH-60A rotor, we show the variation of the magnitudes of the 3, 4 and 5/rev rotating

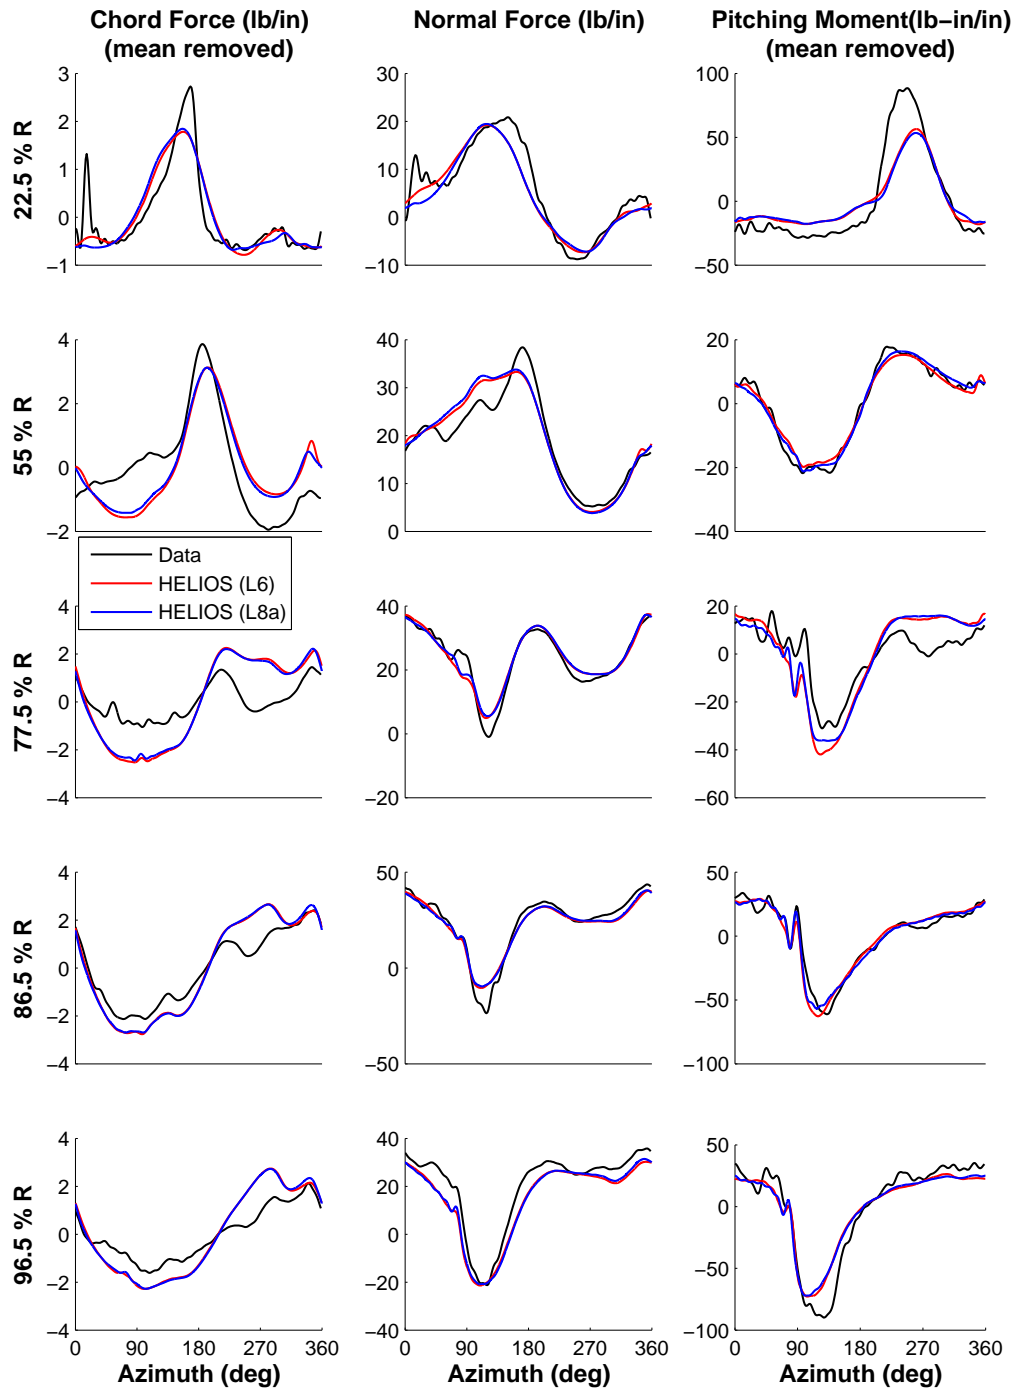


Figure 11. High speed (C8534) airloads. Predicted and measured airloads for high speed flight using coarse (off-body L6) and fine (off-body L8 adapt) grid systems; $\mu = 0.368$; $C_T/\sigma = 0.084$

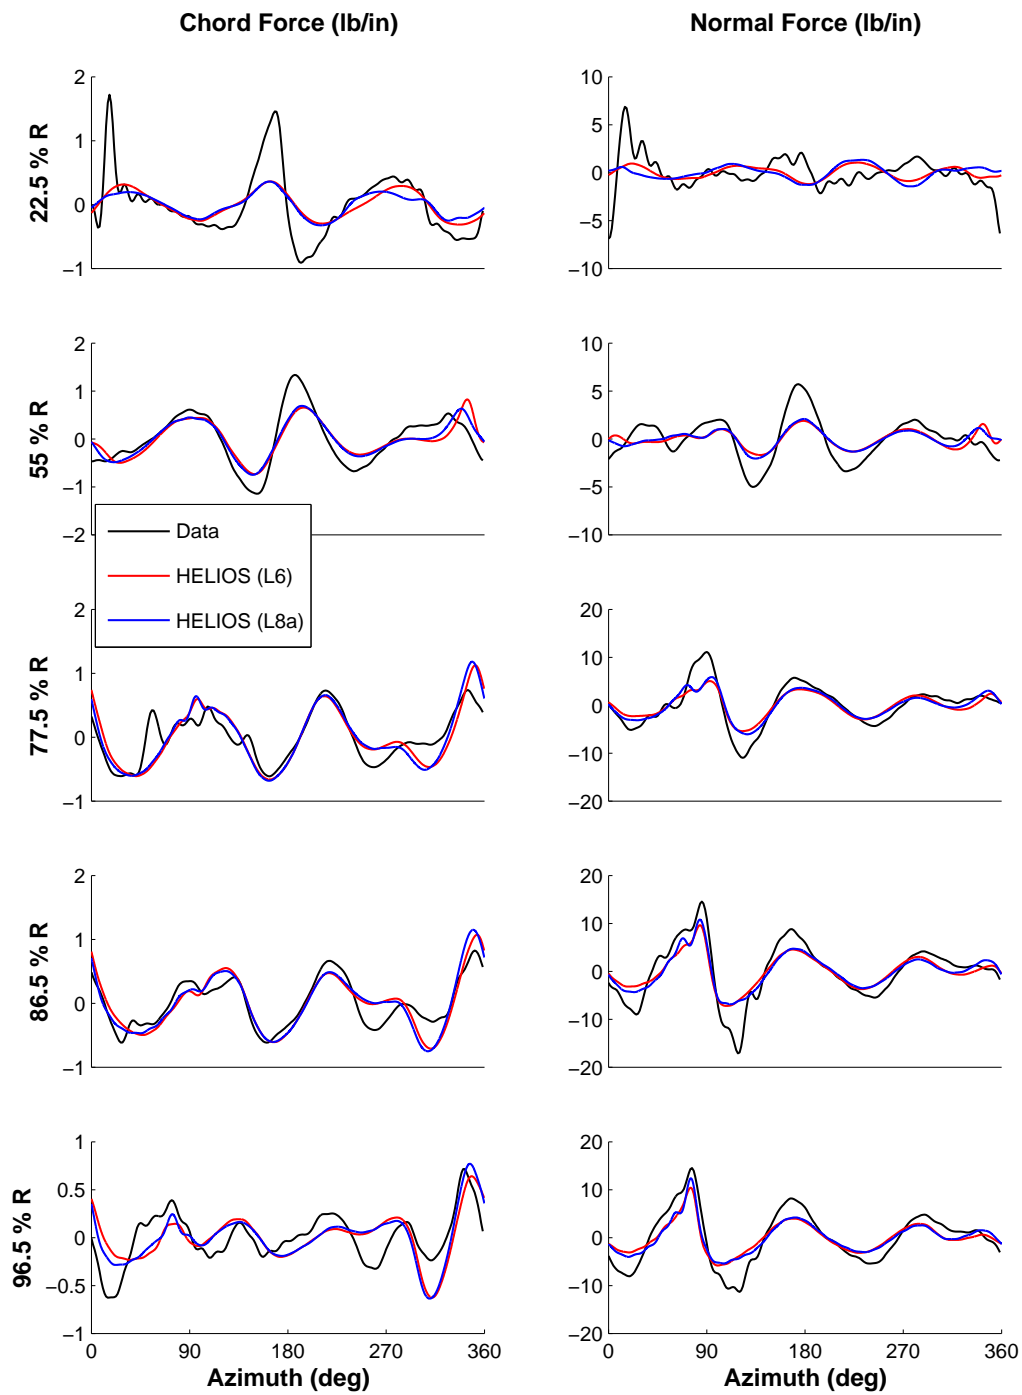


Figure 12. High speed (C8534) vibratory airloads. Predicted and measured vibratory airload (3/rev and higher) time-histories for high speed flight using coarse (off-body L6) and fine (off-body L8 adapt) grid systems; $\mu = 0.368$; $C_T/\sigma = 0.084$

frame aerodynamic loading with non-dimensional spanwise coordinate in Figure 13. It can be noticed that adaptive mesh refinement (denoted by L8a in plots) does produce non-negligible improvement in the 3/rev and 4/rev harmonics of loading in mid-span locations. Indeed, mid-span locations are where returning wake interactions are prominent at this flight condition, and adaptive mesh refinement contributes to the improvement of aerodynamic loading at these locations by better capturing the vortex wake field.

V.B.2. Low speed transition flight condition (C8513)

The low speed transition flight condition is also a high vibration flight condition similar to the high speed flight. The mechanism of vibration is, however, quite different. The aerodynamic loading in this flight condition is dominated by the interactions of the rotor blade with the rolled up vortex structures on the rear of the rotor disk (often termed 'disk-vortices' or 'super vortices'). Helios predictions for this flight condition is shown in Figure 14. The data shows good correlation with the flight test data indicating effective capture of the vortex wake structures. Similar to the high-speed C8534 case, the L6 and L8 adapt grid systems show largely similar characteristics, although the L8 adapt results do show more vibratory characteristics likely from better resolution of the wake.

Figure 15 shows the vibratory aerodynamic loading (3/rev and higher) at this flight condition. The waveforms of the vibratory loads are in good agreement with flight data for both normal and chord forces. The peak magnitude remains underpredicted, although to a lesser extent compared to high speed forward flight condition. The variation on rotor frequency harmonics with spanwise coordinate is compared with test data in Fig. 16. Differences are observed between the L6 and L8 adapt grids but, in general, it is difficult to discern whether the finer L8 adapt off-body grids are leading to improvements in the spanwise harmonic loading as well.

V.B.3. High-altitude stall flight condition (C9017)

The high altitude stall flight condition represents one of the extremes in the level flight operating envelope. The high thrust generated at this condition causes distinct dynamic stall events on the retreating blade. These stall events dominate the high frequency blade pitch excitation causing large torsion moments and consequently large push-rod loads (almost the maximum in level flight).

For this flight condition, the target thrust was arbitrarily reduced by 5% from the previously published $C_T/\sigma = 0.129$ to $C_T/\sigma = 0.126$ due to well known uncertainties associated with this flight. Unlike the other two flight conditions, this flight occurs at the stall boundary of the rotor, and achieving the correct trim state is a critical pre-requisite to any meaningful stall prediction.

The retreating blade stall events in the aerodynamic loading is the key phenomena that needs to be resolved at this flight condition. From Fig. 17 it is evident that the two stall cycles are resolved quite accurately at the stations outboard of 80%. The advancing blade lift waveform, however, is not as well predicted as the other two flight conditions. This is partly due to previously described differences in the trim targets. Coupled solutions were available only using the coarse grid (L6) at the time of preparation of this article. We compare the coarse grid solutions with an adaptive solution obtained by prescribing the blade motions obtained from the coarse grid coupled solution (represented as L8a-prescribed) in the figure. We observe noticeable improvements in prediction of stall patterns (especially in the inboard stations) when AMR is utilized. This improvement is also evident in the harmonic variation with span coordinate (shown in Fig. 18), where the AMR clearly shows improved magnitude of the vibratory airloads. Torsion moments and push-rod loads at this flight conditions are characterized by the presence of large 5/rev magnitude harmonics owing to the frequency content of the two prominent stall cycles in the pitching moment wave form. It is worth noting that the 5/rev harmonic variation across span is well captured by both grid systems, with the AMR grid system showing comparatively better prediction inboard.

V.C. Wake Comparisons

Figure 19 show a top view comparison of the wakes for the various grid combinations – L6, L7, L8, and L8 adapt – for the high speed (C8534) case. For the fixed refinement calculations ((a)–(c)) the wake detail is increasingly well captured as indicated by the iso-surface of the Q criteria ($Q * R^2 = 0.3$). The grid shows the extent of the fixed refinement region. For the L8 off-body grids ((c)–(d)) only every other point is shown. The adapted grid (d) maintains clustering near the tip and root vortices and vortices emanating from the

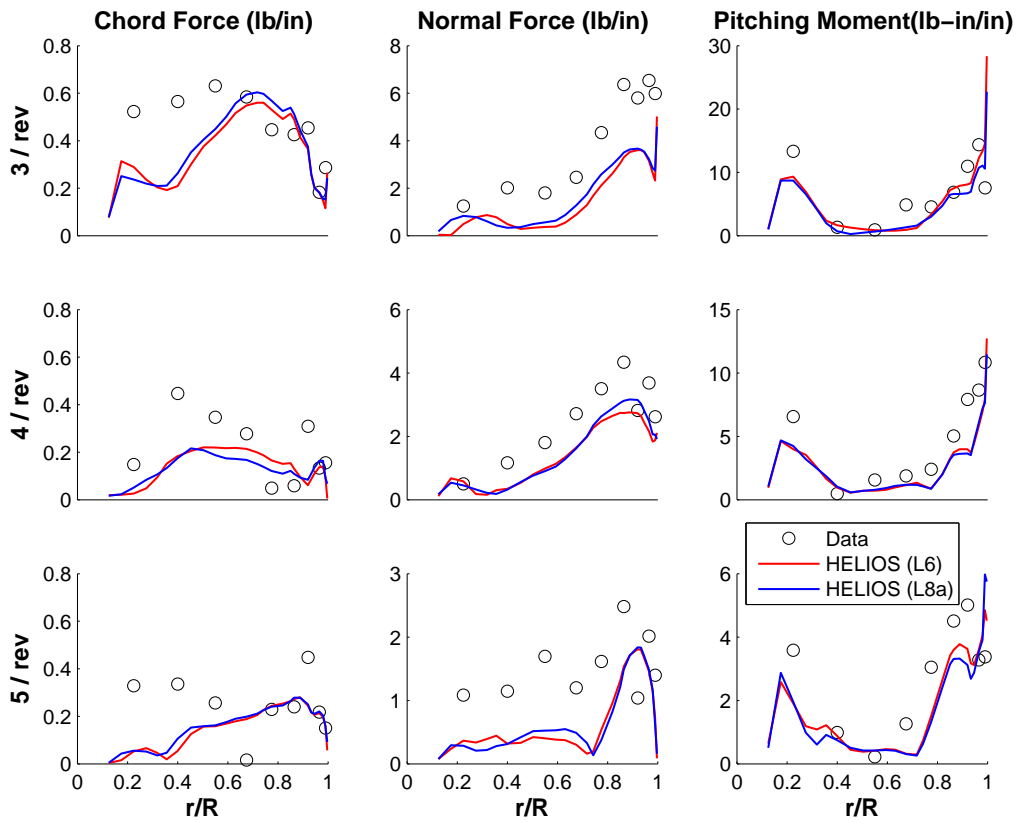


Figure 13. High speed (C8534) harmonics. Comparison of harmonics of sectional aerodynamic loading along the spanwise direction for high speed flight using coarse (off-body L6) and fine (off-body L8 adapt) grid stem's; $\mu = 0.368$; $C_T/\sigma = 0.084$

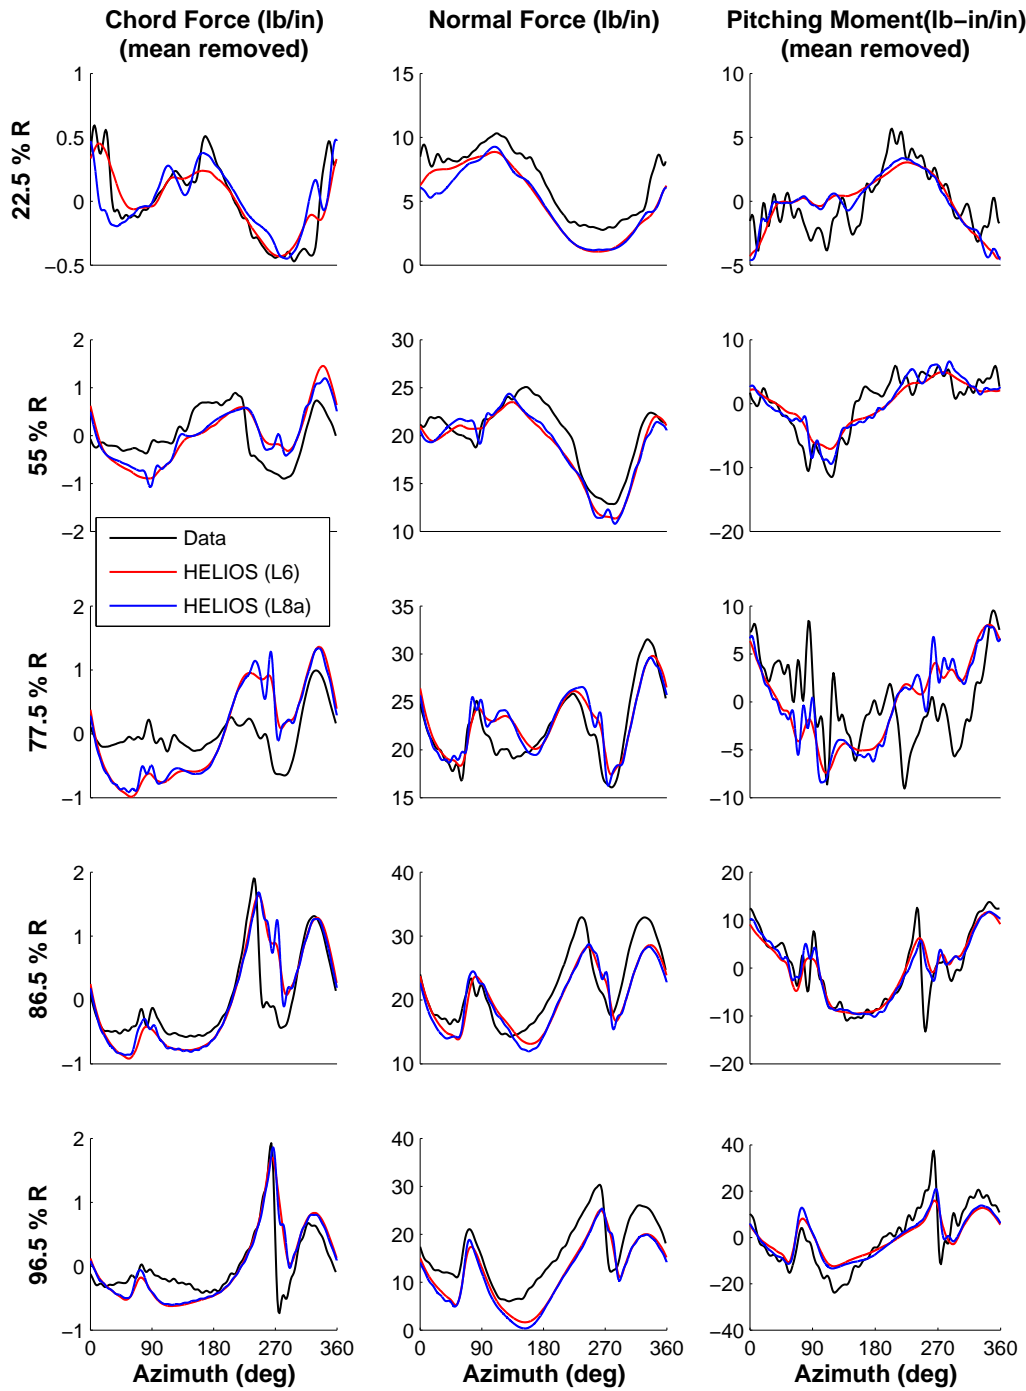


Figure 14. Low speed (C8513) airloads. Predicted and measured airloads for low speed transition flight using coarse (off-body L6) and fine (off-body L8 adapt) grid systems; $\mu = 0.153; C_T/\sigma = 0.076$

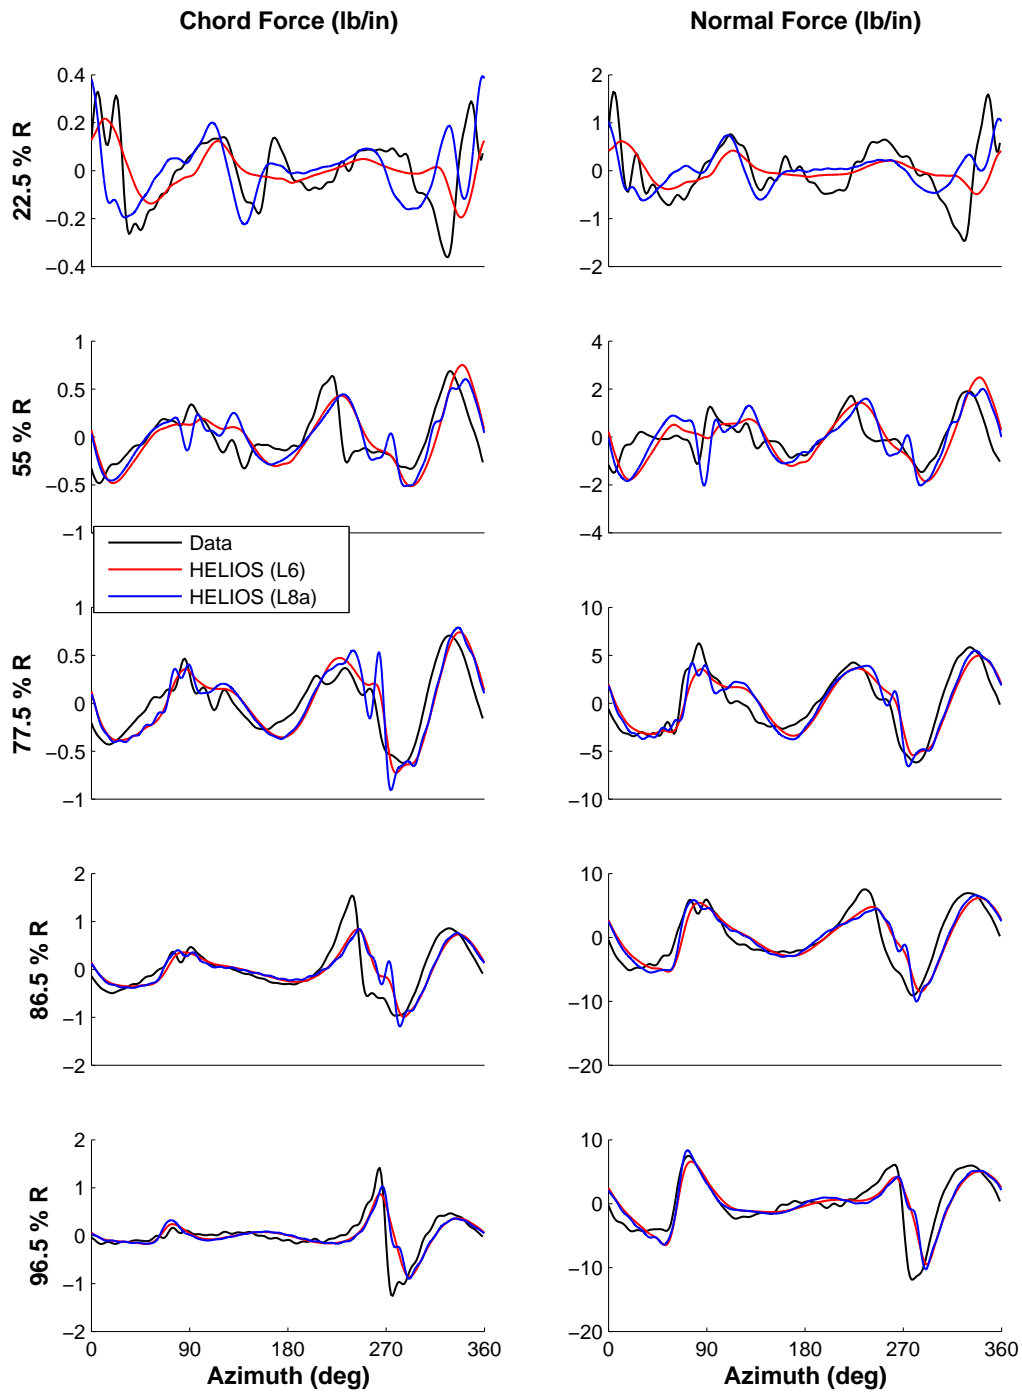


Figure 15. Low speed (C8513) vibratory airloads. Predicted and measured vibratory airload (3/rev and higher) time-histories for low speed transition flight using coarse (off-body L6) and fine (off-body L8 adapt) grid systems; $\mu = 0.153$; $C_T/\sigma = 0.076$

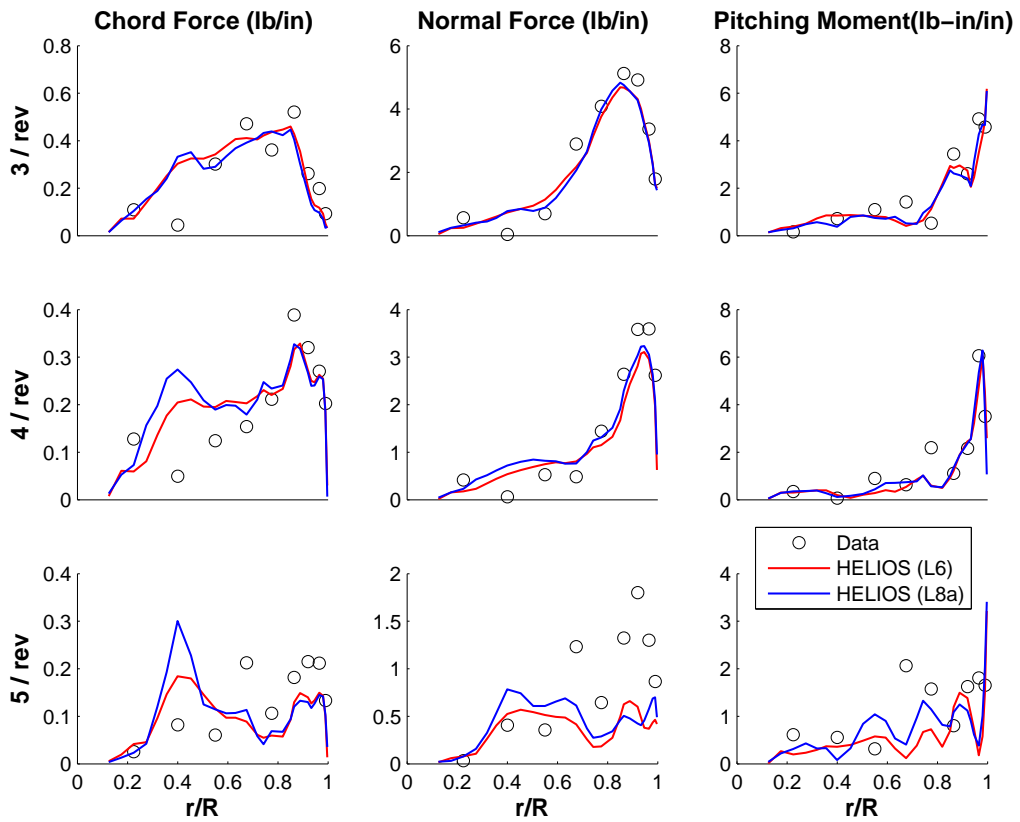


Figure 16. Low speed (C8513) harmonics. Comparison of harmonics of sectional aerodynamic loading along the spanwise direction for low speed transition flight using coarse (off-body L6) and fine (off-body L8 adapt) grid systems; $\mu = 0.153$; $C_T/\sigma = 0.076$

swept tip break. The wake sheet in the 2nd quadrant is also picked up by the adaption. Comparing the L8 fixed refinement and adapted cases indicates very little difference in the structure. The most noticeable difference is in the 4th quadrant. The adapted grid does a much better job at resolving the region outside of the fixed refinement box. For clarity the near-body grids have not been shown, although the hole cuts in the off-body meshes indicate their appropriate location.

Figure 20 shows a comparison of the vorticity contours in a streamwise cut at 62% span on the retreating side for the high speed (C8534) case. The maximum vorticity shown is the refinement criteria (0.0045 non-dimensional). Again the wake is increasingly well captured as the off-body grid is refined ((a)-(c)). The cut off and dissipation of the vortex at the downstream boundary of the refinement region is drastic. For this flight condition the refined region is poorly placed and should be moved downward, but this is flight condition dependent. The adapted solution (d) removes the requirement to know the location of the wake a priori. It does an excellent job of capturing the tip vortices with the finest refinement level. The vortices are effectively tracked continuously downstream. Unlike the L8 fixed refinement case (c) the shear layers are not fully refined so that the vorticity contours differ visually. The location of the blade is indicated by the blanked out region and is seen to have geometry adaption (Fig. 20(d)).

At the start of the adaption calculation the L8 adapted grid has only 29 million grid points used to refine to the geometry and solution in the fixed refinement region. This is a significant savings over the L8 fixed refinement box which required 146 million points and is comparable to the 22 million required for the L7 fixed refinement case but with better resolution. At the end of the adaption portion of the calculation, i.e. half a rotor revolution for C8534, the L8 adapted grid (Fig. 20(d)) has half the number of grid points as the L8 fixed grid (Fig. 20(c)) – 67 million compared with 146 million. Overall, there is a much more efficient use of grid points (computational resources) with the adaption scheme throughout the calculation.

Figures 21 and 22 shown Q iso-surfaces (top view) and vorticity contours (side view, retreating side, 62% span) for the high thrust coefficient case (C9017). The top view (Figure 21) clearly shows the wake and shed vorticity being increasingly well captured with increasing grid density. The L8 grids (fixed and adapted) capture the same wake structures (Fig. 21(c-d)). This includes what appears to be a "bursting" of the tip vortex early in the 4th quadrant and the entrainment of this structure into the retreating-side super-vortex. A significant amount of turbulent wake is also generated from the hub vortices and their interaction with the stall cycle late in the 4th quadrant. This detailed type of turbulent (i.e. chaotic) wake structure has not been seen previously in UH-60 high thrust coefficient calculations and deserves further investigation as to the causes and effects of these instabilities.

In the vorticity contours (Fig. 22) it is seen that the L8 fixed refinement actually better captures some of the wake sheets and maintains the older vortices longer compared with the L8 adapted grid after 1 rotor revolution (Fig. 22(d)). The older vortices (from the front of the rotor disk) are lower in the side view. The vorticity magnitude of these older structures that originated in the non-refined L7 fixed grid starting solution were not strong enough to be identified by the feature-based refinement procedure. In this sense, the fixed L8 refinement would seem to be better capturing the wake, and the reduction in vorticity of the older vortices would appear to be an artifact of the adaption. However, additional visualization of the adapted solution after 2 rotor revolutions (Fig. 22(e)) indicates that vortices which originate in the L8 refined mesh are properly maintained. This delay in adapting to the vortex is due to 1) the choice of adaption criteria (i.e. vorticity) and 2) the physical time it takes for the highly resolved adaption regions to be "convected" with the flow. This second point is a function of the starting conditions for the AMR calculations (L7 fixed in this case). Additional study of the effects of time-dependent AMR on rotorcraft flowfields, including adaption frequency, extent of buffer regions, and grid interpolation on vortex capturing is required to understand this phenomenon. Note that the Cartesian "holes" in the L8 off-body grid that are not associated with the blade hole cuts (e.g. regions between refinement levels) are an artifact of plotting every other point.

Due to a combination of lower advance ratio producing more wake, increased thrust producing higher vorticity levels, and shed vorticity from the stall regions, this case generates the most grid points from solution adaption. The vorticity adaption criteria was only increased slightly from 0.0045 to 0.0050 to account for the increased levels of vorticity. At the adaption initiation from L7 fixed, 36 million grid points are generated. The adaption portion of the calculation results in 125 million grid points after one rotor revolution, and over 240 million points after 2.5 rotor revolutions. For one rotor revolution, the L8 adapt point count is, therefore, still less than the L8 fixed refinement (145 million) case but with noticeable differences in wake capturing as previously described. The cost of refinement and the associated domain connectivity make this calculation approximately 12% more expensive on average.

Finally, for completeness, Figs. 23 and 24 shown Q iso-surfaces and vorticity contours for the low speed case (C8513). The trends from the previous cases are maintained. Even though the advance ratio is low, which maintains more revolutions of the wake near the plane of the rotor, (compare Figs. 19 and 23), the L8 adaption grid point count is not as large as would be expected. In fact, the starting and ending point count values (after one rotor revolution) of 24 and 59 million are the lowest of the three flight conditions. In this case the L8 adapt cost is 12% less than the fixed refinement. However, as in the high thrust coefficient case, the wake sheets and older wake ages are not always well captured or maintained (Fig. 24(d)) unless additional rotor revolutions are performed (Fig. 24(e)).

VI. Concluding Observations

This article describes the CFD/CSD implementation of the `Helios` software framework and validates its prediction capability using experimental flight data from the UH-60 rotorcraft in three flight conditions. The computations exercise all of the capabilities of `Helios` including adaptive mesh refinement, automated over-set domain connectivity and seamless in-core fluid/structure coupling. Furthermore, this is perhaps the first successful demonstration of solution based adaptive mesh refinement for the complex coupled fluid/structure problem encountered in rotorcraft aeromechanics, particularly in forward flight. Key observations and conclusions are summarized below.

1. The CFD flow solution process involving a dual-mesh paradigm – unstructured near-body solver (NSU3D), Cartesian-adaptive off-body solver (SAMARC), and automated parallel domain connectivity (PUNDIT) – is verified to be robust for application in rotary-wing systems with relative motion and elasticity.
2. Some outstanding issues were identified with respect to solution timings. In particular, the cost of domain connectivity was quite variable, sometimes as low as 7% while at other times as high as 40% of the total cost per step, with some dependency on AMR (i.e. grid point and grid block count). Also, the cost of the unstructured near-body solution at each unsteady timestep remains high relative to calculations with structured-grid solvers.
3. The fluid/structure interface process used by `Helios` is verified to be seamless, consistent, and accurate, hence providing accurate and convergent solutions to the coupled fluid/structure/vehicle trim problem.
4. The off-body AMR scheme generally achieves equal accuracy in the wake as compared to a finely-resolved fixed refinement off-body grid with equal resolution but with significantly fewer gridpoints. It also aids in automatically preserving flow features for much larger distances/times downstream of the rotor, which removes from the user the burden of having to define, a priori, the condition/problem dependent wake location. In certain cases, older vortices from the front of the rotor disk were not always well captured by vorticity-based adaption, thus motivating the development of more sophisticated feature detection criteria and schemes.
5. The improved wake resolution with AMR is found to have minimal impact on aerodynamic sectional loading predictions (normal force, chord force, and pitching moment). Furthermore, global force prediction (thrust and torque coefficient) showed little sensitivity to off-body wake resolution, based on prescribed motion calculations with various levels of off-body spacing using fixed refinement and feature-based adaption. This suggests that the fine wake resolution from AMR is not absolutely required for obtaining reasonable aerodynamic loading predictions for the three level forward flight conditions studied. This would not be a likely conclusion for blade-vortex interaction (BVI) analysis.
6. Although the AMR scheme reduces the off-body grid point count, the cost of solutions are not significantly changed compared with equivalently fine fixed refinement meshes ($\pm 12\%$). One reason is that AMR maintains the wake for much larger downstream distances and, in some cases, creates comparable amounts of grid points as a case with fixed refinement. This effect can be mitigated by limiting the extent of the wake tracking and adaption, although maintaining refinement of the wake may be preferable for modeling rotor/fuselage or main-rotor/tail-rotor interactions. Another is that AMR introduces many more grid blocks which potentially increases domain connectivity cost. Nonetheless, the off-body solution component is quite fast compared with the near-body module, on average about

15% of the total cost. Even significant efficiency improvements to the off-body solution will thus have a lesser impact on the overall cost.

7. Further study of the effects of AMR on a range of rotorcraft flowfields (i.e. BVI and hover), performance prediction, and wake capturing is required. This includes automation of the feature detection and optimization of the adaption parameters, such as extent and frequency.

VII. Acknowledgments

Material presented in this paper is a product of the CREATE-AV (Air Vehicles) Element of the Computational Research and Engineering for Acquisition Tools and Environments (CREATE) Program sponsored by the U.S. Department of Defense HPC Modernization Program Office. This work was conducted at the High Performance Computing Institute for Advanced Rotorcraft Modeling and Simulation (HIARMS). The authors gratefully acknowledge additional contributions by Dr. Venkateswaran Sankaran, Dr. Roger Strawn, and Prof. Dimitri Mavriplis. The authors gratefully acknowledge the use of the DoD High Performance Computing and Modernization Office (HPCMO) DOD Shared Resource Centers (DSRC).

References

- ¹Sankaran, V., J. Sitaraman, A. Wissink, A. Datta, B. Jayaraman, M. Potsdam, D. Mavriplis, Z. Yang, D. O'Brien, H. Saberi, R. Cheng, N. Hariharan, and R. Strawn, "Application of the Helios Computational Platform to Rotorcraft Flowfields," AIAA-2010-1230, 48th AIAA Aerospace Sciences Meeting, Orlando FL, Jan 2010.
- ²J. Sitaraman, A. Katz, B. Jayaraman, A. Wissink, V. Sankaran, "Evaluation of a multi-solver paradigm for CFD using overset unstructured and structured adaptive Cartesian grids," AIAA-2008-0660, AIAA 48th Aerospace Sciences Meeting, Reno NV, Jan. 2008.
- ³J. Sitaraman, M. Floros, A. Wissink, M. Potsdam, "Parallel Domain Connectivity Algorithm For Unsteady Flow Computations Using Overlapping And Adaptive Grids," *Journal of Computational Physics*, Volume 229, Issue 12, p. 4703-4723.
- ⁴A. Wissink, S. Kamkar, T. Pulliam, J. Sitaraman, and V. Sankaran, "Cartesian Adaptive Mesh Refinement for Rotorcraft Wake Resolution," AIAA-2010-4554, AIAA 28th Applied Aerodynamics Conference, Chicago IL, June 2010.
- ⁵S. Choi, and A. Datta, "CFD Prediction of Main Rotor Vibratory Loads Using Time-Spectral Method," AIAA Paper 2008-7325, 26th AIAA Applied Aerodynamics Conference, Honolulu, Hawaii, Aug. 18-21, 2008.
- ⁶A. Wissink, J. Sitaraman, D. Mavriplis, T. Pulliam, V. Sankaran, "A python-based infrastructure for overset CFD with adaptive Cartesian grids," AIAA-2008-0927, AIAA 48th Aerospace Sciences Meeting, Reno NV, Jan. 2008.
- ⁷M. Potsdam, H. Yeo, and W. Johnson, "Rotor Airloads Prediction Using Loose Aerodynamic/Structural Dynamic Coupling," Presented at the 60th Forum of American Helicopter Society, Baltimore, MD, May 2004.
- ⁸Wissink, A.M., M. Potsdam, V. Sankaran, J. Sitaraman, Z. Yang, and D. Mavriplis, "A Coupled Unstructured-Adaptive Cartesian CFD Approach for Hover Prediction," Presented at the 66th Forum of the American Helicopter Society, Phoenix AZ, May 11-16, 2010.
- ⁹J. Sitaraman, J. D. Baeder, "Evaluation of the wake prediction methodologies used in CFD based rotor airload computations," AIAA-2006-3472, AIAA 24th Conference on Applied Aerodynamics, Washington, DC, 2006.
- ¹⁰R. Biedron and E. Lee-Rausch, "Rotor Airloads Prediction Using Unstructured Meshes and Loose CFD/CSD Coupling", Paper AIAA-2008-7341, 26th AIAA Applied Aerodynamics Conference, Honolulu, Hawaii, Aug. 18-21, 2008.
- ¹¹Z. Yang, D. Mavriplis, "Higher-order time integration schemes for aeroelastic applications on unstructured meshes," AIAA-2006-0441, AIAA 44th Aerospace Sciences Meeting, Reno NV, Jan. 2006.
- ¹²R. D. Hornung, A. M. Wissink, S. R. Kohn, "Managing complex data and geometry in parallel structured AMR applications," *Engineering with Computers*, Vol. 22, No. 3, 2006, pp. 181-195.
- ¹³Saberi, H., Khoshlahjeh, M., Ormiston, R., and Rutkowski, M. J., "Overview of RCAS and Application to Advanced Rotorcraft Problems," 4th AHS Decennial Specialist's Conference on Aeromechanics, San Francisco, CA, January 21-23, 2004.
- ¹⁴J. J. Alonso, P. LeGresley, E. Van Der Weide, "pyMDO: A framework for high-fidelity multi-disciplinary optimization," AIAA-2004-4480, AIAA/ISSMO 10th Conference on Multidisciplinary Analysis and Optimization, Washington, DC, 2004.
- ¹⁵M. J. Berger, P. Colella, "Local adaptive mesh refinement for shock hydrodynamics," *Journal of Computational Physics*, Vol. 82, No. 1, 1989, pp. 65-84.
- ¹⁶Mavriplis, D. and Levy, W., D., "Transonic Drag Prediction Using an Unstructured Multigrid Solver", *Journal of Aircraft*, Vol. 42, No. 4, July-August, 2005.
- ¹⁷Mavriplis, D., "DLRF6 WB and WBF Results using NSU3D," Proceedings of the 3rd Drag Prediction Workshop, San Francisco, CA, June 2006.
- ¹⁸Pulliam, T. H., "Euler and Thin-Layer Navier-Stokes Codes: ARC2D, and ARC3D," Computational Fluid Dynamics Users Workshop, The University of Tennessee Space Institute, Tullahoma, Tennessee, March 12-16, 1984.
- ¹⁹Y. Lee, J. D. Baeder, "Implicit hole cutting - a new approach to overset grid connectivity," AIAA-2003-4128, AIAA 16th Conference on Computational Fluid Dynamics, Washington, DC, 2003.
- ²⁰R. L. Meakin, A. M. Wissink, W. C. Chan, S. A. Pandya, J. Sitaraman, "On strand grids for complex flows," AIAA-2007-3834, AIAA 18th Conference on Computational Fluid Dynamics, Miami FL, June 2007.
- ²¹W. G. Bousman and R. M. Kufeld, "UH-60A Airloads Catalog," NASA TM-2005-212827/AFDD TR-05-003, 2005.

²²R. M. Kufeld, D. L. Balough, J. L. Cross, K. F. Studebaker, C. D. Jennison, and W. G. Bousman, "Flight Testing of the UH-60A Airloads Aircraft," Presented at the 50th Forum of the American Helicopter Society, Washington, DC, May 1994.

²³C. Tung, F. X. Caradonna, and W. R. Johnson, "The Prediction of Transonic Flows on an Advancing Rotor, Presented at the 40th Forum of the American Helicopter Society," Arlington, VA, May 1984.

²⁴Kamkar, S.J., A. Wissink, V. Sankaran, A. Jameson, "Feature-Driven Cartesian Adaptive Mesh Refinement in the Helios Code," AIAA-2010-171, 48th AIAA Aerospace Sciences Meeting, Orlando FL, Jan 2010.

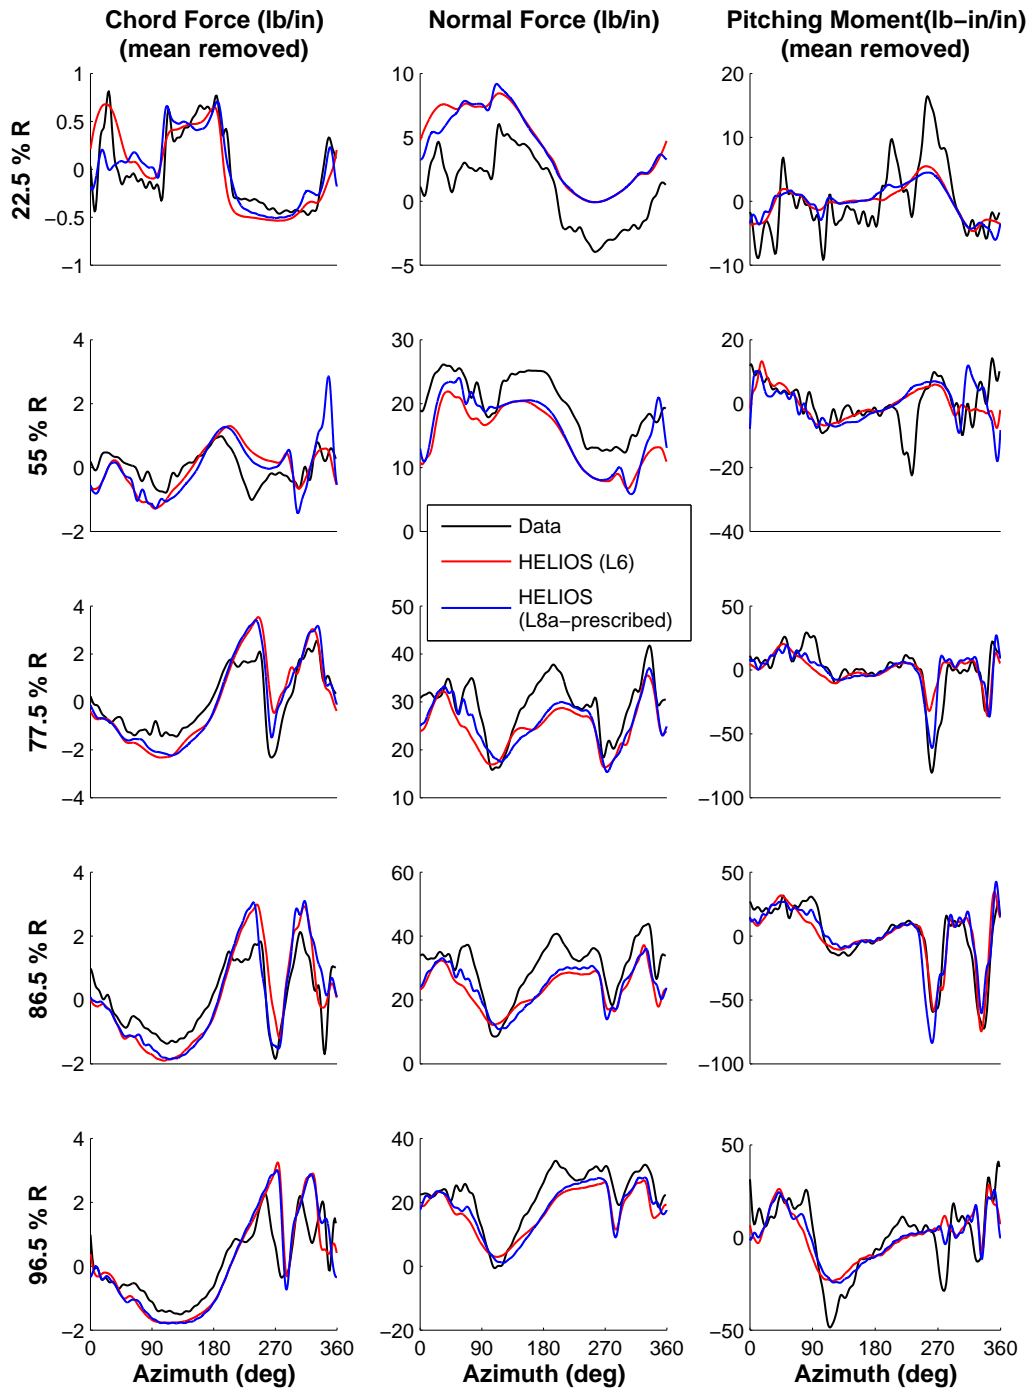


Figure 17. High thrust coefficient (C9017) airloads. Predicted and measured airloads using coarse (off-body L6) and fine (off-body L8 adapt) grid systems; $\mu = 0.237$; $C_T/\sigma = 0.126$

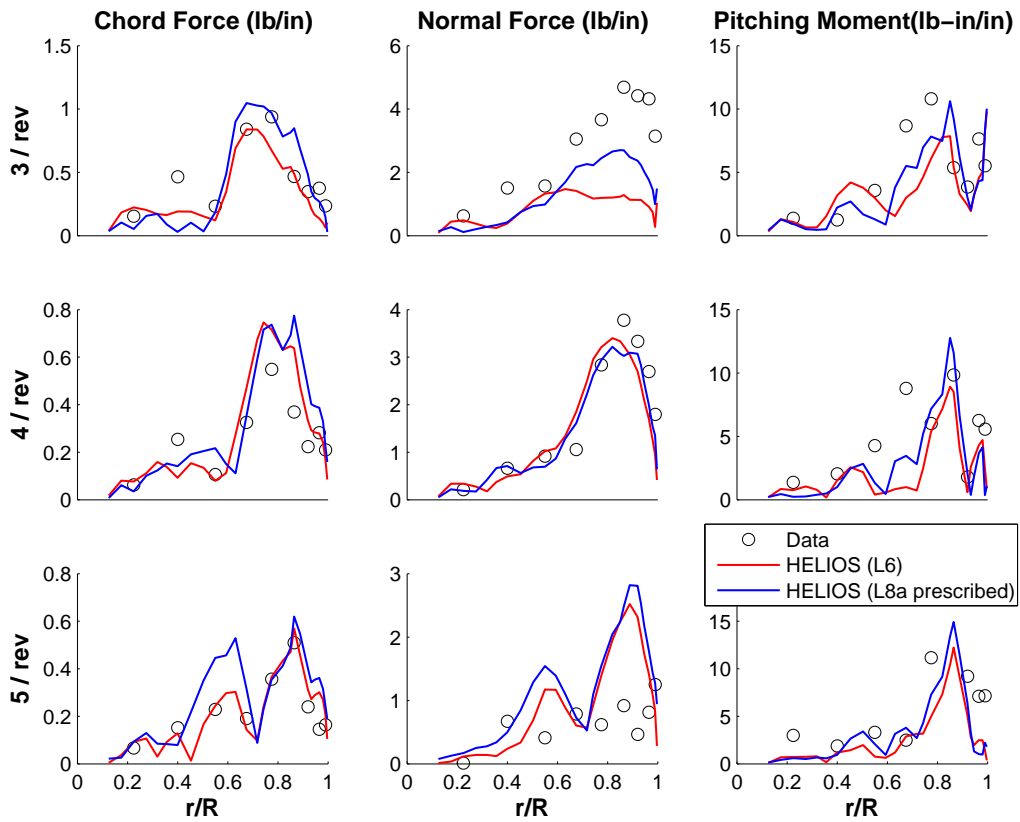





Figure 18. High thrust coefficient (C9017) harmonics. Comparison of various harmonics of sectional aerodynamic loading along the spanwise direction using coarse (off-body L6) and fine (off-body L8 adapt) grid systems; $\mu = 0.237$; $C_T/\sigma = 0.126$




(a) coarse L6



(b) fine L7

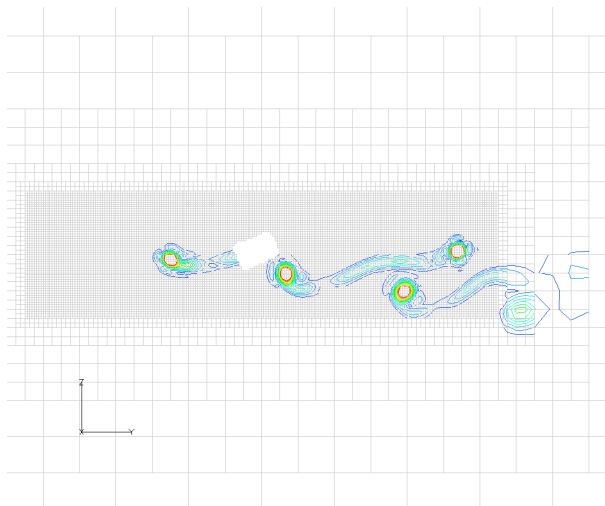


(c) fine L8




(d) fine L8 adapt


Figure 19. High speed (C8534) wake comparisons, Q criteria iso-surfaces overlaid on grid at $z = 0$ (rotor plane axis); (a) coarse grid L6, (b) fine grid L7, (c) fine grid L8, (d) fine grid L8 with adaption, $\mu = 0.368$, $C_T/\sigma = 0.084$.




(a) coarse L6



(b) fine L7




(c) fine L8




(d) fine L8 adapt


Figure 20. High speed (C8534) wake comparisons, vorticity contours overlaid on grid at rotor $r/R=0.62$ spanwise location on retreating side (3rd and 4th quadrant); (a) coarse grid L6, (b) fine grid L7, (c) fine grid L8, (d) fine grid L8 with adaption, $\mu = 0.368$, $C_T/\sigma = 0.084$.




(a) coarse L6




(b) fine L7



(c) fine L8




(d) fine L8 adapt 1 rev




(e) fine L8 adapt 2 revs


Figure 21. High thrust coefficient (C9017) wake comparisons, Q criteria iso-surfaces overlaid on grid at $z = 0$ (rotor plane axis); (a) coarse grid L6, (b) fine grid L7, (c) fine grid L8, (d) fine grid L8 with adaption after 1 rev, (e) fine grid L8 with adaption after 2 revs; $\mu = 0.237$, $C_T/\sigma = 0.126$.




(a) coarse L6




(b) fine L7



(c) fine L8




(d) fine L8 adapt 1 rev




(e) fine L8 adapt 2 revs


Figure 22. High thrust coefficient (C9017) wake comparisons, vorticity contours overlaid on grid at rotor $r/R=0.62$ spanwise location on retreating side (3rd and 4th quadrant); (a) coarse grid L6, (b) fine grid L7, (c) fine grid L8, (d) fine grid L8 with adaptation after 1 rev, (e) fine grid L8 with adaptation after 2 revs; $\mu = 0.237$, $C_T/\sigma = 0.126$.



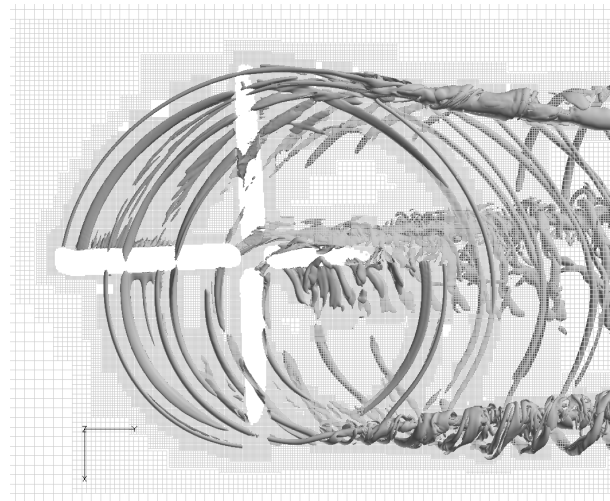
(a) coarse L6




(b) fine L7



(c) fine L8




(d) fine L8 adapt 1 rev




(e) fine L8 adapt 2 revs


Figure 23. Low speed (C8513) wake comparisons, Q criteria iso-surfaces overlaid on grid at $z = 0$ (rotor plane axis); (a) coarse grid L6, (b) fine grid L7, (c) fine grid L8, (d) fine grid L8 with adaptation after 1 rev, (e) fine grid L8 with adaptation after 2 revs; $\mu = 0.150$, $C_T/\sigma = 0.076$.




(a) coarse L6




(b) fine L7



(c) fine L8



(d) fine L8 adapt 1 rev



(e) fine L8 adapt 2 revs

Figure 24. Low speed (C8513) wake comparisons, vorticity contours overlaid on grid at rotor $r/R=0.62$ spanwise location on retreating side (3rd and 4th quadrant); (a) coarse grid L6, (b) fine grid L7, (c) fine grid L8, (d) fine grid L8 with adaptation after 1 rev, (e) fine grid L8 with adaptation after 2 revs; $\mu = 0.150$, $C_T/\sigma = 0.076$.