

ADVANCED METHODS FOR DYNAMIC AEROELASTIC ANALYSIS
OF ROTORS

A Thesis
Presented to
The Academic Faculty

by

Nicolas D. Reveles

In Partial Fulfillment
of the Requirements for the Degree
Doctor of Philosophy in the
School of Aerospace Engineering

Georgia Institute of Technology
May 2014

Copyright © 2014 by Nicolas D. Reveles

ADVANCED METHODS FOR DYNAMIC AEROELASTIC ANALYSIS
OF ROTORS

Approved by:

Professor Marilyn J. Smith, Advisor
School of Aerospace Engineering
Georgia Institute of Technology

Professor Dewey Hodges
School of Aerospace Engineering
Georgia Institute of Technology

Professor Massimo Ruzzene
School of Aerospace Engineering
Georgia Institute of Technology

Professor Mark Costello
School of Aerospace Engineering
Georgia Institute of Technology

Professor Jonathan Rogers
School of Mechanical Engineering
Georgia Institute of Technology

Date Approved: April 3, 2014

*Dedicated to my grandfather, Dr. Robert Schlaudt,
for encouraging me to pursue my dreams.*

TABLE OF CONTENTS

DEDICATION	iii
LIST OF TABLES	vii
LIST OF FIGURES	viii
LIST OF SYMBOLS OR ABBREVIATIONS	xiii
SUMMARY	xvi
I MOTIVATION	1
1.1 Rotor Aeromechanics Predictions	1
1.2 Rotorcraft Trim Methods for CFD/CSD Analyses	5
1.3 Hybrid CFD/CSD Free-Wake Approaches	7
1.4 Thesis Objectives	9
II AEROELASTIC TRIMMING METHODOLOGY	12
2.1 CFD/CSD Coupling Strategies	12
2.1.1 Loose Coupling	12
2.1.2 Tight Coupling	14
2.1.3 Simultaneous Solution Approach	15
2.2 Baseline Codes	15
2.2.1 OVERFLOW	16
2.2.2 FUN3D	16
2.2.3 DYMORE	16
2.2.4 OVERFLOW/DYMORE	17
2.2.5 FUN3D/DYMORE	17
2.3 Trim of Rotorcraft Vehicles	17
2.3.1 Effect of Aeroelastic Coupling Strategy	17
2.3.2 Inverse Trim Jacobian	18
2.3.3 Trim Algorithm	19
2.4 Development of an Advanced Trim Algorithm	20
2.4.1 Kriging-based Trim Algorithm	20
2.4.2 Mathematical Description of Kriging	21

2.5	Trim Algorithm for Loose and Tight Coupling	25
2.5.1	Inverse Problem	25
2.5.2	Jacobian Formation	26
2.5.3	Control Estimate	29
2.5.4	Improved Algorithm	31
III	HYBRID URANS / FREE-WAKE DESCRIPTION	33
3.1	Hybrid Solver Theory	33
3.2	Baseline and Coupled Codes	36
3.3	Parallel Performance	38
3.4	Wake Effects	39
3.5	Hierarchical Options	41
IV	CFD/CSD TIGHT COUPLING RESULTS	42
4.1	UH-60A Tight Coupling Trim	42
4.1.1	Optimization of Trim Process	43
4.1.2	UH-60A C8534	49
4.1.3	UH-60A C9017	71
4.1.4	Kriging Estimate Applied to Loose Coupling	88
4.2	The HART II Test	91
4.2.1	Background	91
4.2.2	DYMORE Standalone	92
4.2.3	CFD/CSD Loose Coupling	92
4.2.4	Tightly-Coupled CFD/CSD	97
V	HYBRID CFD/CSD	105
5.1	Prescribed Motion Test	105
5.2	Advanced Dynamics Model	111
5.2.1	Loose Coupling Results	113
5.2.2	Tight Coupling Stability Results	116
5.3	UH-60A Dynamic Stall Wind Tunnel Test	117
VI	CONCLUSIONS	129

VII FUTURE WORK	133
7.1 Tight Coupling	133
7.2 Hybrid CFD/Free-Wake	134
REFERENCES	137
VITA	147

LIST OF TABLES

1	UH-60A Test Case Flight Conditions.	42
2	Optimization of convergence using kriging to trim during tight coupling. . .	47
3	Comparison of Trimmed Controls	51
4	List of sensors available and functional during UH-60A C8534.	58
5	Comparison of Converged Controls	71
6	List of sensors available and functional during UH-60A C9017.	83
7	Comparison of Trimmed Controls	98
8	Comparison of statistical metrics for UH-60A C8534 hybrid method predic- tions of pitching moments.	111
9	Comparison of statistical metrics of the reduced models compared to the full-CFD/CSD method for the ADM rotor at $r/R=0.12$ for a 0.35 advance ratio.	116
10	Comparison of controls as predicted by full CFD/CSD and hybrid CFD/CSD with the Spalart-Allmaras turbulence model.	120
11	Comparison of controls as predicted by full CFD/CSD and hybrid CFD/CSD with the $k\omega$ -SST turbulence model.	120
12	Comparison of statistics of the different methods compared with the experi- mental data for the C4540 wind tunnel test case.	126
13	Comparison of statistics of the hybrid-CFD/CSD data compared with the full-CFD/CSD data for the C4540 wind tunnel test case.	126

LIST OF FIGURES

1	A generic diagram of tight coupling iterations and data transfer.	14
2	A diagram of the stable region of the original algorithm.	27
3	Probability density functions at two points, x_a and x_b are shown for a range of predictions.	30
4	The ADM rotor with swept tips was modeled at an advance ratio of $\mu = 0.35$. The isosurfaces visualize the OVERFLOW shed vorticity, while the lines illustrate the free-wake vorticity predicted propagation.	34
5	A diagram showing the flow of data in the coupled hybrid aerodynamics code.	35
6	The boundary of reverse flow for the UH-60A at the high-speed 8534 flight conditions.	37
7	Circulation of the UH60 High-Speed Case (C8534) illustrating the impact of local velocities.	37
8	Tests were conducted on different numbers of processors to validate the parallelization.	38
9	The ADM grids for the hybrid-solver.	39
10	Example of the effects of starting the hybrid code after two revolutions. . .	40
11	Autopilot trimmer for the UH-60A C8534 case.	44
12	Loose coupling trim of the UH-60A C8534 case.	44
13	Trimming of the UH-60A C8534 case in tight coupling from a cold start. . .	45
14	Revolutions needed to reach thrust convergence (1% error) based on loose coupling initialization.	46
15	Comparison of thrust convergence to time the controls are held.	48
16	Thrust convergence of CFD/CSD kriging when using a database provided by CSD alone versus that of using the full CFD/CSD database shows comparable performance.	49
17	Trim of UH-60A C8534 using using the optimal trim procedure for tight coupling results.	50
18	Comparison of hub thrust for the UH-60A C8534 between loose and tight coupling results.	51
19	Comparison of UH-60A case C8534 pitching moment airload predictions at $r/R = 0.865$ for coarse and fine grids.	52
20	Comparisons of loosely and tightly-coupled OVERFLOW/DYMORE for normal forces on UH-60A C8534 inboard and midspan airstations.	53

21	Comparisons of loosely and tightly-coupled OVERFLOW/DYMORE for normal forces on UH-60A C8534 outboard airstations.	54
22	Comparisons of loosely and tightly-coupled OVERFLOW/DYMORE for normal forces on UH-60A C8534 tip.	55
23	Comparisons of loosely and tightly-coupled OVERFLOW/DYMORE for pitching moments on UH-60A C8534 inboard and midspan airstations.	56
24	Comparisons of loosely and tightly-coupled OVERFLOW/DYMORE for pitching moments on UH-60A C8534 outboard airstations.	57
25	Comparisons of loosely and tightly-coupled OVERFLOW/DYMORE for pitching moments on UH-60A C8534 tip.	58
26	Comparisons of loosely and tightly-coupled OVERFLOW/DYMORE for chord forces on UH-60A C8534 inboard and midspan airstations.	59
27	Comparisons of loosely and tightly-coupled OVERFLOW/DYMORE for chord forces on UH-60A C8534 outboard airstations.	60
28	Comparisons of loosely and tightly-coupled OVERFLOW/DYMORE for chord forces on UH-60A C8534 tip.	61
29	Comparison of structural moments for the UH-60A C8534 at $r/R = 0.2$ span.	62
30	Comparison of structural moments for the UH-60A C8534 at $r/R = 0.3$ span.	63
31	Comparison of structural moments for the UH-60A C8534 at $r/R = 0.4$ span.	65
32	Comparison of structural moments for the UH-60A C8534 at $r/R = 0.5$ span.	66
33	Comparison of structural moments for the UH-60A C8534 at $r/R = 0.6$ span.	67
34	Comparison of structural moments for the UH-60A C8534 at $r/R = 0.7$ span.	69
35	Comparison of tip deflections for UH-60A C8534 between loose and tight coupling results.	70
36	Thrust convergence for UH-60A C9017 occurs in about the same amount of time as with loose coupling.	72
37	Comparisons of loosely and tightly-coupled OVERFLOW/DYMORE for normal forces on UH-60A C9017 inboard and midspan airstations.	74
38	Comparisons of loosely and tightly-coupled OVERFLOW/DYMORE for normal forces on UH-60A C9017 outboard airstations.	75
39	Comparisons of loosely and tightly-coupled OVERFLOW/DYMORE for normal forces on UH-60A C9017 tip.	76
40	Comparisons of loosely and tightly-coupled OVERFLOW/DYMORE for pitching moments on UH-60A C9017 inboard and midspan airstations.	77
41	Comparisons of loosely and tightly-coupled OVERFLOW/DYMORE for pitching moments on UH-60A C9017 outboard airstations.	78

42	Comparisons of loosely and tightly-coupled OVERFLOW/DYMORE for pitching moments on UH-60A C9017 tip.	79
43	Comparisons of loosely and tightly-coupled OVERFLOW/DYMORE for chord forces on UH-60A C9017 inboard and midspan airstations.	80
44	Comparisons of loosely and tightly-coupled OVERFLOW/DYMORE for chord forces on UH-60A C9017 outboard airstations.	81
45	Comparisons of loosely and tightly-coupled OVERFLOW/DYMORE for chord forces on UH-60A C9017 tip.	82
46	Vorticity magnitudes showing the differences between loose and tight coupling at $r/R = 0.675$ for $\psi = 0^\circ, 90^\circ, 180^\circ,$ and 270°	83
47	Comparison of structural moments for the UH-60A C9017 at $r/R = 0.2$ span.	84
48	Comparison of structural moments for the UH-60A C9017 at $r/R = 0.3$ span.	85
49	Comparison of structural moments for the UH-60A C9017 at $r/R = 0.5$ span.	86
50	Comparison of structural moments for the UH-60A C9017 at $r/R = 0.6$ span.	87
51	Comparison of structural moments for the UH-60A C9017 at $r/R = 0.7$ span.	89
52	Comparison of structural moments for the UH-60A C9017 at $r/R = 0.9$ span.	90
53	Comparison of loose and tight coupling lead-lag responses via an FFT analysis.	90
54	Comparison with fixed controls for UH-60A C9017 between loose and tight coupling results.	90
55	Comparison of loose coupling convergence with and without a kriging-based solution start.	91
56	Normal Force and Pitching Moments at 87% Radius computed by DYMORE alone for the HART-II baseline case.	93
57	Loose coupling trim results for the HART-II baseline case.	94
58	Loose coupling trim results for the HART-II minimum noise case.	95
59	Loose coupling trim results for the HART-II minimum vibration case.	96
60	Time to trim in tight coupling for the baseline HART-II case.	98
61	Comparison of the HART-II baseline normal force predictions of tight coupling, loose coupling, and comprehensive results.	99
62	Comparison of the HART-II baseline pitching moment predictions (mean removed) of tight coupling, loose coupling, and comprehensive results.	100
63	Comparison of the HART-II baseline oscillatory structural moments between tight coupling, loose coupling, and comprehensive results.	102
64	Comparison of HART-II baseline elastic tip deformations between tight coupling, loose coupling, and comprehensive results.	103

65	Comparison of HART-II baseline frequency content contained in the elastic lag motion predicted by loose and tight coupling simulations.	104
66	UH-60A C8534 blade grids used in the full CFD/CSD analysis.	105
67	UH-60A C8534 cut-through plot of vorticity from the full CFD/CSD solution with the wake convecting down and away from the rotor.	106
68	UH-60A C8534 hub loads of the hybrid simulation before and after the hybrid interface is initialized after two revolutions.	107
69	UH-60A C8534 sectional airloads at 55% radius for the reduced domain grid with and without the hybrid solver are compared with the full CFD/CSD solution.	108
70	UH-60A C8534 in-plane forces viewed looking down at the forward-most section of the rotor disk shows small, fine-scale differences between the two solutions, especially just after 180°.	110
71	Comparison of UH-60A C8534 sectional pitching moment predictions at $r/R = 0.675$ with another hybrid code.	110
72	Comparison of the ADM domains resolved with the CFD/CSD solver and the URANS hybrid CFD/CSD analysis illustrates the reduction in grid size used with the hybrid approach.	112
73	Comparison of the ADM structural moments with the hybrid solver at $r/R=0.12$ for a 0.35 advance ratio with means removed.	114
74	Comparison of statistical relations of the ADM flap bending moments for the CSD and Hybrid-CFD/CSD predictions versus the full-CFD/CSD method for the time period shown in Fig. 73a).	115
75	The chord bending moment signals from all four blades, showing the initial excitement with a vertical line depicting the excitement cut-off and following response for $\mu = 0.35$	117
76	Comparison of the ADM damping predictions for three levels of aerodynamics for the comprehensive, hybrid CFD/CSD, and full CFD/CSD analyses. . .	118
77	Comparison of the UH-60A C4540 original domain and the hybrid analysis domain.	119
78	Comparison of the UH-60A wind tunnel C4540 aerodynamic loads at 92% span as computed with Spalart-Allmaras.	121
79	Comparison of the UH-60A wind tunnel C4540 aerodynamic loads at 92% span as computed with $k\omega$ -SST turbulence model.	122
80	Comparison of the UH-60A wind tunnel C4540 oscillatory structural moments at 40% span for the Spalart-Allmaras turbulence model.	124
81	Comparison of the UH-60A wind tunnel C4540 oscillatory structural moments at 40% span for the $k\omega$ -SST turbulence model.	125

82	Comparison of pitching moments for UH-60A wind tunnel C4540 solved with the $k\omega$ -SST turbulence model for the last four coupling iterations.	127
83	Comparison of with and without ensemble averaging for the UH-60A wind tunnel C4540 aerodynamic pitching moments at 92% span as computed with $k\omega$ -SST.	128

LIST OF SYMBOLS OR ABBREVIATIONS

ADM	Advanced Dynamics Model.
AFDD	Aeroflightdynamics Directorate.
BLUP	Best Linear Unbiased Predictor.
BVI	Blade Vortex Interaction.
CFD	Computational fluid dynamics.
CHARM	Comprehensive Hierarchical Aeromechanics Rotorcraft Model.
FD	Finite Difference.
FFT	Fast Fourier Transform.
HHC	Higher Harmonic Control.
HPC	High Performance Computing.
LES	Large Eddy Simulation.
NASA	National Aeronautics and Space Administration.
PDF	Probability Distribution Function.
RANS	Reynolds-Averaged Navier-Stokes CFD.
TSD	Transonic Small Disturbance.
URANS	Unsteady Reynolds-Averaged Navier-Stokes CFD.

Nomenclature

β	Scalar influence coefficients of $f(x)$
ϵ	Difference in expected and actual value
$\hat{\theta}(L)$	Maximum likelihood controls corresponding to loading condition L
$\hat{L}(\theta)$	Predicted load corresponding to controls θ
$\hat{y}(x)$	Maximum likelihood prediction of value at x
λ	Fit coefficients of BLUP
μ	Rotor advance ratio
Ω	Rotor speed
Ω	Spatial region of interpolation
θ	Rotor controls (collective, lateral cyclic, and longitudinal cyclic)
θ_0	Collective control
θ_j	The j^{th} control of θ
$\theta_{1,c}$	Lateral cyclic control
$\theta_{1,s}$	Longitudinal cyclic control
\vec{v}	A vector
α	Angle of attack
C_l	Sectional lift coefficient
e	Trim error with respect to targets
$E[x]$	Expected value at position x

$f(x)$	A vector of regressors (for n sampling points)
$g(x)$	Arbitrary function
g_s	Gain schedule
J	Trim sensitivity Jacobian
L	Filtered rotor loads to be trimmed (e.g. thrust, roll moment, pitch moment)
L_j	The j^{th} filtered rotor load of L
M_x	Rotor filtered rolling moment
M_y	Rotor filtered pitching moment
N	The number of blades on a rotor
n	An integer
R	Kriging correlation matrix
R_{max}	Maximum trim rate
T	Rotor filtered thrust
T	Time period
t	Time
x	A point in space
x_j	The j^{th} training point
Y	Matrix of data from sampling points
$y(x)$	Actual value at position x
$Z(x)$	A matrix containing the random portion of an arbitrary field
C	An arbitrary constant

SUMMARY

Simulations play an integral role in the understanding and development of rotorcraft aeromechanics. Computational Fluid Dynamics coupled with Computational Structural Dynamics (CFD/CSD) offers an excellent approach to analyzing rotors. These methods have been traditionally “loosely-coupled” where data are exchanged periodically, motion is prescribed for CFD, and the updated loads have a static component for CSD. Loosely-coupled CFD/CSD assumes the solution to be periodic, which may not be true for some simulations. “Tightly-coupled” CFD/CSD, where loads and motion are exchanged at each time step, does not make this periodic assumption and opens up new avenues of simulation to research. A major drawback to tightly-coupled CFD/CSD is an increase in computational cost. Different approaches are explored to reduce this cost as well as examine numerical implications in solutions from tightly and loosely-coupled CFD/CSD.

A trim methodology optimized for tightly-coupled simulations is developed and found to bring trim costs within parity of loosely-coupled CFD/CSD simulations. Aerodynamic loading is found to be nearly similar for fixed controls. However, the lead-lag blade motion is determined to contain a harmonic in the tightly-coupled analysis that is not an integer multiple of the rotor speed. A hybrid CFD/CSD methodology employing the use of a free-wake code to model the far-field effects of the rotor wake is developed to aid in computational cost reduction. Investigation of this approach reveals that computational costs may be reduced while preserving solution accuracy.

This work’s contributions to the community include the development of a trim algorithm appropriate for use in tightly-coupled CFD/CSD simulations along with a detailed examination of the physics predicted by loose and tight coupling for quasi-steady level flight conditions. The influence of the wake in such cases is directly examined using a modular hybrid coupling to a free-wake code that is capable of reduced cost computations.

CHAPTER I

MOTIVATION

Rotorcraft are aircraft whose primary propulsive device consists of rotary wings that rotate about hub(s). The rotary wings are referred to as rotor blades and collectively make up the rotor. Rotorcraft may have any number of rotors, each of which may consist of two or more rotor blades. Vehicles are typically categorized by the placement and purpose of their rotors. For example, traditional helicopters consist of a large main driven rotor designed to provide vertical thrust with a tail rotor whose purpose is to balance the main rotor torque. There exist numerous other configurations such as co-axial helicopters (two driven counter-rotating main rotors), tiltrotors (rotors are mechanically hinged between forward and upwards), autogyros (undriven main rotor), and others.

The study of rotors themselves falls into the broad field of rotor aeromechanics. Rotor aeromechanics deals with aerodynamics, structures, vibration, acoustics, flight dynamics, and performance [1]. The ability to accurately predict rotor aeromechanics is essential to successful rotorcraft designs. Rotors become more aerodynamically efficient (increase power loading) as they grow in size [2]. These rotors typically undergo severe loading in flight due in part to inertial forces such as centrifugal and coriolis forces. This leads designers to attempt to minimize the blade mass while still maintaining enough structural support to lift the vehicle. Thus, rotor blades are usually slender wings which are considerably flexible. As they propel the vehicle, they are subject to strong additional loading from the aerodynamics. This intersection of dynamics, aerodynamics, and elasticity forms the field of dynamic aeroelasticity [3] – a challenging field of study.

1.1 Rotor Aeromechanics Predictions

Computation-based rotorcraft studies have seen remarkable progress in the last four decades. When computers became accessible to engineers in the 1960s, engineers began applying computational techniques to solve rotorcraft problems [4]. Bell Helicopter led the way with

the development of their C81 code. This eventually became the first “comprehensive code” for rotorcraft analyses, from which many others have followed.

A comprehensive rotorcraft code is one which is able to solve all of the relevant rotorcraft problems together, in one single code. Typically, this requires modeling geometric and structural properties, computing structural deformations, approximating airloads, and solving for a set of trim controls. This then allows for performance figures and loading conditions to be examined. By the 1980s structural properties were modeled with finite elements and by the 1990s multi-body dynamics had been included [4]. However, these comprehensive codes still relied on simplified aerodynamic models utilizing combinations of lifting line aerodynamics, dynamic inflow, and sometimes vortex wake models.

Starting in the 1980s, the need for more accurate aerodynamic loading prompted work on coupling comprehensive codes with various fluid dynamics codes. Initially, the focus was on using potential-based codes for improved aerodynamics. In 1984 Tung, Caradonna, and Johnson [5] coupled the comprehensive code CAMRAD to a finite difference solver for the transonic small disturbance equations with a loosely-coupled methodology. The goal of this effort was to improve the predictions on the advancing blade, which experiences a transonic flow that was not well captured by the comprehensive code’s simple aerodynamic models. Bridgeman, Strawn, and Caradonna used an entropy and viscosity corrected potential aerodynamics model in 1988 [6]. These efforts were generally successful in improving normal force loads but suffered from poor prediction of pitching moments. This eventually led to concerted efforts to couple comprehensive codes to CFD codes. Smith [7] and Bauchau and Ahmad [8] investigated the coupling of CFD codes to nonlinear CSD codes in the 1990s.

In 1994, NASA and the U.S. Army jointly instrumented an UH-60A helicopter [9,10] for what has been to date the most comprehensive flight test experiment of a rotorcraft. The helicopter had 221 pressure transducers at nine radial stations, which allowed for accurate characterization of the sectional airloads. Structural loading was also measured in addition to the airloads, along with control forces and the vehicle’s state. The helicopter was flown throughout the flight envelope in both quasi-steady level flight conditions as well as in maneuvers. Attempts to numerically model these flight tests and reproduce experimentally

recorded results led to the creation of the semi-annual airloads workshop at the beginning of the new millennium. Experts in the field of rotorcraft aeromechanics meet at these workshops to show their latest results and work together in an attempt to overcome modeling deficiencies.

In the late 1990s, AFDD tested the Advanced Dynamics Model (ADM) rotor [11, 12]. The four-bladed rotor had two sets of blades: one with straight tips and one with swept tips. Lead-lag stability was examined for a number of advance ratios and shaft tilts on both blades.

In 2001, Higher Harmonic Control (HHC) was investigated in the HART-II program (a follow up of an earlier 1994 HHC test) [13, 14]. The blades were instrumented to record inboard structural loading and outboard airloads. Acoustic measurements were taken and the blade deflections were recorded using photogrammetry. As part of the HART-II International Workshop, two thorough papers reviewed state of the art computations on this rotor. The first by van der Wall et al. [15] reviews comprehensive efforts while the second by Smith et al. [16] reviews CFD/CSD efforts.

The UH-60A rotor was tested in the National Full-Scale Aerodynamics Complex (NFAC) 40 by 80 foot wind tunnel by NASA and the Army [17]. The tests included both thrust sweeps and speed sweeps, including high advance ratios. Many recent studies focus on the slowed rotor at high advance ratios, for example Potsdam et al. [18], Kottapalli [19], and Yeo [20].

Strawn et al. [21] and Datta et al. [22] reviewed the coupling of CFD codes in the 2000s. The coupling of unstructured CFD code to comprehensive codes was examined in the 2000s [23, 24] as a way of reducing the labor required to create a CFD grid geometry. Recent works [25, 26] have shown further improvements to the solution of aeromechanics problems can be made by adaptively refining the CFD meshes. This has allowed flow features to be resolved that otherwise would be too expensive to achieve by uniform grid refinement.

In 2002, Altmikus et al. [27] examined a rotorcraft in loose and tight coupling between an Euler solver and a comprehensive code. Their results revealed that that the airloads were largely the same between the two coupling methodologies for this unstalled high-speed

condition. Their presented results indicate that in regions where differences appear, the normal force prediction contains a 5% deviation ($r/R = 0.9$ and $\psi = 0^\circ$) while pitching moments showed greater differences of approximately 12% of the signal amplitude at $r/R = 0.915$ and $\psi = 180^\circ$. Torsional tip deflections showed the largest overall differences, with a maximum amplitude difference of 0.5° compared to a maximum computed deflection of -1.4° . Phase differences between loose and tight coupling also appear in the tip deflections of as much as 5° . In the same year, another study by Pomin and Wagner [28] tightly coupled a RANS code to comprehensive code. This allowed the low-frequency lead-lag mode to be accurately predicted, but only after 15 revolutions due to large start up transients.

In 2007, Ho et al. [29] examined structural loading on the UH-60A rotor when measured airloads are directly applied to a comprehensive model. They noted significant disparities between predictions and experiment, especially for high-thrust conditions in pitch link loads and torsional moments. They concluded that that the inaccuracies may be due to a specific modeling or analysis error currently unknown to the rotorcraft community.

Many recent works have attempted to shed light on the reason for apparent differences between computations and experiment. For example, in 2014 Kang [30] investigated lead-lag damper models, adding to many previous works, and found that they have a significant impact on the chord bending moment predictions. Burgess and Jain [31] investigated the effects of numerical methods on static and dynamic stall and found they had significant impact.

Maneuvering flight has long been known as a challenging and important phase of flight for computations, which were unable to accurately predict loads [32]. In 2007, Bhagwat et al. [33] succeeded in performing a maneuver using tightly-coupled CFD/CSD techniques and compared their predictions to loose coupling. This was compared to a quasi-steady maneuvering simulation using the averaged periodic motion of a portion of the maneuver. The results showed that the loose coupling produced results similar to those of the tight coupling method, but the peaks were better captured with the tightly-coupled simulation.

It is known that rotor wakes exhibit small aperiodic motion [34], which can have a large effect on local blade loading. Therefore, certain physics-rich environments, such as blade

vortex interactions (BVI) are particularly interesting to examine from a tightly-coupled CFD/CSD standpoint that does not impose periodicity on a solution. As experiments are necessarily somewhat limited in the measurements that can be made, experiments have been performed [35,36] for the purpose of providing data to compare with CFD results.

An examination of the effect of coupling scheme is needed for complex rotor cases, such as those in dynamic stall and with BVI. The analysis should look beyond aerodynamic loading and into the structural loading and response as well.

1.2 Rotorcraft Trim Methods for CFD/CSD Analyses

Trim refers to a vehicle state whereby a set of controls produces an intended response [37]. Trim is often used to refer to a steady-state (or quasi-steady-state) response, although there is no need to limit the definition to those applications. A trimmed condition may correspond to a particular point during a maneuver, for example. Often the targeted vehicle response is a set of resultant loads acting on the vehicle, although other possibilities, such as flight path or a particular tip-path plane for rotorcraft are appropriate.

Generally trimming a vehicle does not imply determining all of a vehicle’s controls, but only the most relevant subset. For example, in fixed-wing applications, it is not usually required to solve for elevator, aileron, and horizontal stabilizer position to predict performance. In rotorcraft applications, however, knowledge of the collective and two cyclic control positions is required to produce accurate performance predictions since these controls affect the tip-path plane of the main rotor. As in fixed-wing applications, rotorcraft trim usually neglects certain controls such as the tail rotor collective and throttle (which instead is assumed to produce a given rotational rate on the rotor).

Trim is a challenge due to the nonlinear nature of rotors. Failure to capture a rotor’s trim controls in a simulation will produce a trim state at a different operating point than intended, rendering the simulation of little value. For rotorcraft trimmed to hubloads (aerodynamic blade loading integrated in the rotor hub frame), there exists an additional problem as hubloads are generally quasi-steady and so are inappropriate for directly targeting in trim. Therefore, oscillations need to be removed (or “filtered”) to converge hubloads to a fixed

target. While there are a number of methods to accomplish this, there are two prominent methods in current simulations. A Fourier decomposition of the hubload waveform can be applied so the overall mean remains. Another method is to apply a moving average of length $\frac{2\pi}{\Omega N}$, where Ω is the main rotor frequency and N is the number of blades. This has the effect of averaging the last partial revolution ($1/N$ of a revolution) and, therefore, is quick to respond to changes. Given a continuous and periodic waveform, both methods will produce the same filtered result. It is only when transient changes due to the updating of controls that the two methods may result in slightly different filtered results where one leads or lags behind the other. If the waveform is allowed to return to a quasi-steady periodic motion, the two approaches will again arrive at the same result.

Loosely-coupled simulations are especially advantageous due to their ability to decouple the trim process from the CFD computations. In 2004, Potsdam, Yeo, and Johnson [38] succeeded in improving predictions of a UH-60A “blackhawk” rotor in forward flight by coupling the CFD code OVERFLOW to the CSD code CAMRAD II with a loosely-coupled methodology. This represents the first loosely-coupled CFD/CSD simulations.

Early applications of tight coupling, such as that performed by Bauchau and Ahmad [8] and Sheffer et al. [39] avoided trim altogether and instead used measured flight controls at the expense of discrepancies between prediction and experiment. Others, such as Pomin and Wagner [28], used estimated controls from a comprehensive code which were held fixed in tight coupling.

Trim is largely reserved for cases appropriate for loose coupling. There are few examples of trim performed in tight coupling, largely due to the inseparable cost of CFD computations. Lee et al. [40] were the first to trim a rotor in tight coupling. They coupled a full potential code to a comprehensive code and achieved airload results that were comparable to loose coupling. The study by Altmikus et al. [27] did trim in tight coupling at approximately 2.5 times the computational cost for that of loose coupling. Differences in control angles of approximately 0.3 degrees were found in the trimmed results.

Today most researchers avoid trim in tight coupling by either first trimming in loose coupling or prescribing controls. The first CFD/CSD stability analyses [41, 42] all used

the former method, where trim was completed in loose coupling and restarted with tight coupling.

State-of-the-art maneuvers are done using either prescribed controls in tight coupling or the motion is averaged and performed in quasi-steady loose coupling. In the recent study by Sheng et al. [43], trim was started in loose coupling followed by prescribed control motion in a tightly coupled maneuver. In her recent work, Marpu [44] examined maneuvers using a quasi-steady approximation at a given time in the maneuver, effectively averaging each portion. In the tightly-coupled maneuvering case studied by Bhagwat et al. [33], motion was prescribed from flight test data and compared to the quasi-steady loose coupling approach. The costs of the tight coupling were already significantly higher than those of loose coupling, and so a traditional trim algorithm would have been infeasible for the maneuver. If a viable trim option existed, future maneuvers could “fly” the vehicle within a simulation while actively trimming.

There currently exists a need for a trim methodology appropriate for tightly-coupled simulations. The ability to trim in tight coupling is especially important for maneuvers where the current state of the art instead relies on prescribed controls. However, because the trim process cannot be separated from the CFD computations in tight coupling it has been generally avoided in even quasi-steady level flight.

1.3 Hybrid CFD/CSD Free-Wake Approaches

Rotor blades leave lasting tip vortices that do not convect away as they do in fixed-wing vehicles due to the passing of following blades. The structure of the wake was visualized with smoke in hover by Gray [45] in 1955 and in forward flight was visualized using bubbles in a water tunnel by Lehman [46] in 1968.

Elucidation of rotor wake structures led to the development of semi-empirical numerical tools to model the wake. In these methods, the wake is segmented into vortex line elements that may be integrated by application of the Biot-Savart law. Piziali was among the first to develop such an approach [47] using a prescribed wake structure.

“Free-wake” computations, where the wake structure is permitted to interact with itself

via induced velocities, soon followed in the work of Scully [48]. To avoid a singularity and more accurately model the physics, the elements have a core size at the center of which velocity is zero. Scully's investigations into the effects varying semi-empirical parameters revealed that the vortex core size has little effect on the solution.

In 1986, Egolf and Sparks [49] coupled a full potential solver to a free-wake code. The full potential equations were solved using finite differencing techniques on a computational grid for a near-body blade while the free-wake was used to model the far-field. The inviscid and irrotational nature of the full potential equations meant that stall and transonic effects could not be modeled with this method.

In 1997, Berkman et al. [50] coupled a Navier-Stokes solver with a free-wake applied to the tip vortex. Yang et al. [51, 52] improved upon the methodology with the addition of blade deformation and trim capabilities. Their approach was evaluated on rotors in forward and low-speed descent flight and found comparable to full Navier-Stokes simulations.

Engineers are often faced with a difficulty in solving a given problem to satisfy engineering accuracy within a practical amount of time and at a reasonable cost. With respect to computational aeromechanics, this has led to a number of efforts [52–54] at creating lower cost hybrid or mixed solver codes.

In 2006, Sitaraman and Baeder [55] compared wake capturing verses wake coupling methods. For wake coupling, they evaluated TURNS/PWAM, a hybrid CFD/free-wake code. They noted large sensitivity of wake coupling methods in the prediction of stall events that wake capturing methods do not suffer from. Computational cost was six times higher for the wake capturing methods than the wake coupling methods, suggesting that wake coupling methods are appropriate for design situations where computational cost would otherwise be too large.

In 2008, Agusta Westland developed a hybrid panel code free-wake methodology for reduced computational cost analyses [56]. Using a multiblock approach they were able to reduce computational cost from $\mathcal{O}(N^2)$ operations to $\mathcal{O}(N \log(N))$, for N panels.

In 2013, Smith et al. [16] published a thorough review of CFD/CSD predictions for

the HART-II rotor. Their analyses included two hybrid codes: TURNS/PWAM and GT-Hybrid. It was noted that the hybrid codes predicted as much as a 10° phase offset of the BVI events that full CFD/CSD methods did not. Hybrid methodologies were also prone to significant over-prediction of pitching moments BVI events were present.

While many recent efforts have relied on free-wakes to model far-field effects, recent efforts [57,58] have examined using the velocity-vorticity transport equations, a more consistent approach but at the penalty of significantly higher computational costs. In 2013, Whitehouse and Boschitsch [59] used VorTran-M to demonstrate improvements in stability as well as accurate solutions with at least an order of magnitude reduction in cell count as compared to similar methods. While their methodology is serial, a parallel version is under development.

Hybrid methodologies offer an additional benefit to lowering cost, provided they are modular. A modular hybrid code that has the ability to exchange computational components of the analysis can be used to study a problem at various hierarchical levels [57]. This can then be used to identify causal physics through selective use of hybrid components, an especially useful feature when faced with complex problems.

1.4 Thesis Objectives

There are several gaps in current dynamic rotorcraft aeroelastic analyses. While rotorcraft are known to encounter aperiodicity in some conditions, due to wake structure or in maneuvers and stability analyses, state-of-the-art simulations largely use a loosely-coupled approach that assumes periodicity. Prior investigations of quasi-steady level flight have examined unstalled conditions and found few differences in predictions, however more analyses are needed to confirm this. The lack of a viable trim option for tightly-coupled simulations has prevented current maneuver analyses from actively trimming. Development of a trim methodology for tightly-coupled predictions would permit greater confidence in results.

It is of enormous utility that researchers find ways of reducing the computational costs of their work. Today's state-of-the-art tools may be too expensive for certain engineering applications where perhaps hundreds or thousands of simulations must be performed (such

as is common in design problems or complete exploration of the flight envelope). Therefore engineers require tools of various levels of fidelity that allows them to select the appropriate model and simulation type based on their engineering judgment.

This thesis seeks to advance the state of the art by:

- Investigate the effects of tightly-coupled versus loosely-coupled aeromechanics simulations. Examinations should not be limited to the aerodynamic loading, which prior research has found to generally be consistent with loosely-coupled simulations. Structural moments and deformations should be examined. To broadly assess the differences, three conditions should be evaluated: high-speed forward flight, dynamic stall, and blade vortex interaction.
- Develop a trim methodology appropriate for tightly-coupled simulations. State of the art simulations are beginning to investigate maneuvering cases, however no practical trim methodology exists for such simulations. A developed trim methodology should first be tested in quasi-steady level flight conditions to assess its performance and limitations.
- Develop a carefree hybrid CFD/CSD free-wake solver that is built upon pre-existing and validated solvers. This would allow future updates to the solver modules to immediately be available in the future; this is necessary to ensure that the methodology stays relevant to future engineers and scientists.
- Ensure that hierarchical modularity exists within a single code so it becomes possible to identify causal physics and separate problems stemming from a purely numerical nature.
- Examine the effects of a hybrid CFD/free-wake approach on a prescribed motion test to ensure that the correct physics are reproduced. Comparisons should be made to the full solver solution on traditional meshes as well as the cropped, hybrid meshes. This permits direct hierarchical examination of the hybrid methodology.

- Understand any accuracy tradeoff between computational speed-up gained from a hybrid methodology and the full CFD/CSD simulations. This would be accomplished by performing a stability analysis and comparing results with full CFD/CSD simulations. This is necessary to ensure that the solver is appropriate for these aeromechanics problems.
- Investigate a dynamic stall condition using the hybrid approach. As dynamic stall conditions are computationally challenging, it is important to understand the performance of these methodologies to know if the aeromechanic predictions are valid.

This thesis seeks to answer questions regarding the following:

- **Periodic/Aperiodic:** Does aperiodicity exist in physics-rich quasi-steady level flight conditions with blade vortex interactions and dynamic stall? Is it present in airloads, structural loads, or deflections?
- **Cost:** Can the computational cost of trimming in tight coupling be reduced from the previously demonstrated 2.5 times increase in cost over loose coupling to a cost more comparable to that of loose coupling trim? If so, can similar methods reduce the cost of loose coupling?
- **Accuracy:** Is the developed hybrid approach accurate for prescribed motion? Can modularity be used to explain any inaccuracies? Are hybrid techniques appropriate for complex cases with dynamic stall and aperiodic conditions such as stability analyses?

CHAPTER II

AEROELASTIC TRIMMING METHODOLOGY

2.1 CFD/CSD Coupling Strategies

The solution of a dynamic aeroelastic problem requires an accurate representation of the governing equations of the problem's dynamics, elasticity, and aerodynamics. Modern comprehensive codes do this by simplifying the aerodynamics, frequently using either lifting line approximations, potential wakes, or table lookups. It is important to recognize that this simplification affects the entire solution and not simply the aerodynamics since the aeromechanics are physically coupled. Improvements of the aerodynamics are a logical start, and computational fluid dynamics (CFD) are often employed to capture nonlinear aerodynamic effects.

There are three ways to computationally couple a comprehensive code and a CFD code. These are known as loose coupling, tight coupling, and what will be termed here as a "simultaneous solution" approach. In a brief overview, the differences in these coupling strategies stem from how they exchange and keep track of the data of each module. In loose coupling approaches, data are exchanged over a time window while in tight coupling approaches, data are exchanged at each timestep. Finally, in a simultaneous solution approach, one code would solve all the governing equations (dynamics, elasticity, and aerodynamics) simultaneously.

2.1.1 Loose Coupling

The loose coupling algorithm is an iterative approach that links together two separate but coupled sets of equations, and these are not limited to CFD/CSD simulations. Loose coupling of CFD computations to a comprehensive code is accomplished by passing an entire revolution of data at each data interchange. This does not require that either code be run a full revolution since each blade is usually assumed to be identical, and thus a revolution of data can be formed by aggregating data from all blades. In this way, each code only

needs to be run a minimum of $1/N$ revolutions to fully populate the data required for an iteration of data exchange.

The original formulation of loose coupling, by Tung, Caradonna, Boxwell, and Johnson in 1984 [5], did not couple modern CFD code to a comprehensive code, but instead sought to couple a transonic small disturbance (TSD) code into a comprehensive code. They used a finite difference approach (denoted in Eqn. 1 by the abbreviation “FD”) applied to the TSD equations.

$$C_l(\alpha) = C_{l,FD}(\alpha_{old}) + C_{l,table}(\alpha) - C_{l,table}(\alpha_{old}) \quad (1)$$

The term $C_{l,FD}(\alpha_{old}) - C_{l,table}(\alpha_{old})$ is static for a current angle of attack, α , and is commonly known as a “delta-load”.

While the loose coupling approach was the first to be widely adapted for solving the aeroelastic problem, it is certainly not the most obvious. The use of a naïve coupling scheme that lacks the delta-loads and instead applies only the previously computed airloads:

$$C_l(\alpha) = C_{l,FD}(\alpha_{old}) \quad (2)$$

A significant problem with this scheme is that it does not respond to the change in the new angle of attack due to the static loading. Thus, if such a method were to converge, it is at a much slower rate.

There are three principal advantages to using a loose coupling approach. The first deals with its ability to decouple the trim process, which will be discussed in detail in Section 2.3.1. The second advantage is that it requires no modifications to the URANS code (provided that the airloads are already calculated), making it a natural choice for early attempts at solving the aeroelastic problem. The third advantage to such an algorithm is that it permits different solvers to be joined together without significant modification.

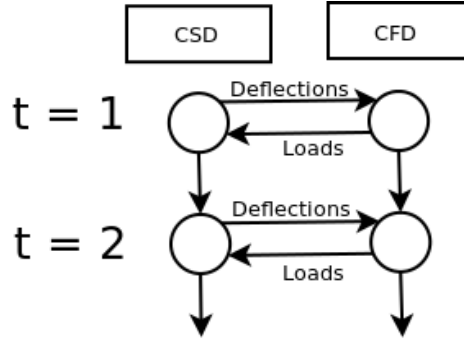


Figure 1: A generic diagram of tight coupling iterations and data transfer.

2.1.2 Tight Coupling

The tightly-coupled paradigm is a direct approach to solving a coupled problem without requiring the solvers to be rewritten. Data are exchanged frequently between separate but coupled equations. Two codes that are tightly-coupled must be run concurrently for practical reasons. While theoretically the codes could be sequentially started and stopped each timestep, although the start up costs make this impractical. Therefore, the scientist or engineer performing the coupled simulation should have a good grasp of all the coupled fields. For efficiency, two codes execute concurrently have several options for data transfers. The slowest is a read/write of data to the disk. A better option is to directly pass data through memory making use of direct library calls. A third option is to pass data through memory through a scripting interface, such as Python [60]. The latter two are popular and have found wide use in the technical community [61–63].

Figure 1 illustrates an example of a generic tightly-coupled CFD/CSD flowchart. First the comprehensive code takes a timestep, computing the deflections for the CFD code. Next, the CFD code applies these deflections and computes aerodynamic loads, which are returned to the comprehensive module. This process continues until the desired number of iterations have been reached.

Note that one code will always take the next timestep before the other, and that such a system is necessarily first order accurate. Therefore, as the timestep is reduced towards zero, the solution approaches the “exact” solution for a given model. One possible means of increasing the order of the coupling is to exchange data multiple times each timestep.

Another option is to reduce numerical coupling errors by taking small timesteps.

2.1.3 Simultaneous Solution Approach

“Simultaneous solution” will be used here to refer to the technique where the aeromechanics governing equations are solved simultaneously. A simultaneous solution approach clearly offers an advantage over the loosely-coupled and tightly-coupled methodologies previously discussed as it contains no time lag and has no assumptions of periodicity [64].

Alternatively, this approach comes at a price. The governing equations of fluids and structures have disparate time and length scales resulting in large computational expense to resolve. Simultaneous solution requires a dedicated code to solve the governing equations. While this is not a problem itself, it does mean that if advances are made in solving one of the sets of governing equations it may require significant resources to implement the changes in a single solver. It certainly does not allow modularity since codes dedicated to solving a subset of the governing equations can no longer be leveraged.

In the case of aeroelastic problems, there are disparate time and length scales that make simultaneous solution impractical. To theoretically resolve both the fluid dynamics and the structural dynamics requires resolving the small Kolmogorov time and length scales [65] of the fluid dynamics and fundamental periods and characteristic lengths for the structural dynamics. While today researchers may not be resolving the smallest scales of the Navier-Stokes equations in rotorcraft applications, the solution of URANS equation require very small time and length scales in the boundary layer computational cells compared to the structural time and length scales.

2.2 Baseline Codes

This section is dedicated to an overview of the baseline methodologies applied in this work: two URANS computational codes (OVERFLOW and FUN3D), a structural dynamics solver (DYMORE), and a free-wake solver (CHARM).

2.2.1 OVERFLOW

OVERFLOW [66] is a structured URANS CFD solver developed by NASA over the past three decades [67]. It employs overset grids [68] for simultaneous modeling of different reference frames as well as to reduce the complexity of modeling complex geometries. The code solves the unsteady Reynold’s averaged Navier-Stokes (URANS) equations in a parallel framework. OVERFLOW contains a plethora of URANS turbulence models, although there are several hybrid RANS-LES (large eddy simulation) turbulence approaches available.

OVERFLOW has been extensively evaluated in rotorcraft aeromechanics [69, 70], fixed-wing vehicles [71], and spacecraft configurations [72], as well as in numerous two-dimensional airfoil studies [73, 74].

2.2.2 FUN3D

FUN3D [75] is a unstructured CFD solver developed by NASA. FUN3D resolves the unsteady Reynold’s Averaged Navier-Stokes equations in a parallel framework. While the unstructured mesh topology permits ease in the modeling of complex geometries, FUN3D also contains an overset option for handling moving reference frames of motion. Like OVERFLOW, this is achieved by using a stationary “background grid” overlaid by body components that are allowed to move. Overset capabilities are achieved using third party libraries, DirTlib [76, 77] and SUGGAR/SUGGAR++ [78, 79].

FUN3D has been well validated for rotorcraft applications [24, 80–82]. Its unstructured framework permits its application in a variety of computational fields such as sling loads [83], radiation [84], and launch vehicles [85].

2.2.3 DYMORE

DYMORE [86] is a multibody dynamics code developed at the Georgia Institute of Technology. It employs finite elements to achieve exact geometric beam models in rotorcraft, as well multibody elements such as linkages and hinges. Time integration is accomplished via either the Radau time integration scheme or the generalized- α time integration scheme. Although

the code is often used as a rotorcraft comprehensive code [42,87–89], it has additional capabilities beyond rotorcraft applications. DYMORE contains no inherent knowledge of rotor geometries; instead, DYMORE models consist of a collection of topological entities that are connected to a selection of joints, hinges, springs, masses, shells, membranes, and beams, with appropriate boundary conditions applied so that they interact together as they would in a physical system.

2.2.4 OVERFLOW/DYMORE

For this work, OVERFLOW and DYMORE have been coupled together. Both loose and tight coupling options are available. In tight coupling, data may be passed via a Python interface or directly through library calls.

Airstations are automatically computed internal to the CFD code as the midpoint between spanwise grid stations along the quarter-chord line. These should be carefully matched internal to DYMORE to avoid interpolation.

2.2.5 FUN3D/DYMORE

FUN3D and DYMORE have also been coupled together. Both loose and tight coupling options are available. Once again, in tight coupling data may be passed through a Python interface or directly through library calls.

Due to the unstructured nature of the CFD code, the user selects airstations at radial locations that are used to integrate sectional forces and moments. The same locations should be selected within DYMORE to avoid having to interpolate.

2.3 Trim of Rotorcraft Vehicles

2.3.1 Effect of Aeroelastic Coupling Strategy

When performing a trim analysis during an aeroelastic simulation, the choice in the aeroelastic coupling methodology significantly influences the process of trim. In a loosely-coupled approach, the comprehensive code is run separately from the CFD simulation, and trimming within the comprehensive code is advantageous as the costs are orders of magnitude less per revolution than running the CFD solver. Therefore, a stable but slow trimmer is

an acceptable choice.

In a tightly-coupled CFD/CSD framework or a simultaneous solution, the trim process can not be decoupled, and a more cost-effective algorithm is required to maintain a computational efficiency. Inefficiencies within the algorithms typically applied in loose coupling trim approaches (e.g. the autopilot trimmer) must be replaced with faster methods.

2.3.2 Inverse Trim Jacobian

A trim Jacobian describes a vehicle's sensitivity of a response that is to be adjusted by the controls. It is defined as:

$$\underline{J} = \begin{bmatrix} \frac{\partial T}{\partial \theta_0} & \frac{\partial T}{\partial \theta_{1,c}} & \frac{\partial T}{\partial \theta_{1,s}} \\ \frac{\partial M_x}{\partial \theta_0} & \frac{\partial M_x}{\partial \theta_{1,c}} & \frac{\partial M_x}{\partial \theta_{1,s}} \\ \frac{\partial M_y}{\partial \theta_0} & \frac{\partial M_y}{\partial \theta_{1,c}} & \frac{\partial M_y}{\partial \theta_{1,s}} \end{bmatrix}. \quad (3)$$

The trim Jacobian relates the hub loads (thrust, T , roll moment, M_x , and pitching moment, M_y) to the vehicle controls (collective and cyclics). If a trim Jacobian is known, then it is a relatively simple task to update the controls to approach the trim target. If the airloads are linearly related to the airloads, then this would occur in a single iteration, but as this is almost never the case, it becomes an iterative process.

An algorithm to update the controls follows from the Newton-Raphson algorithm:

$$\underline{\theta}_{n+1} = \underline{\theta}_n - \underline{J}^{-1} \underline{g}(\theta_n). \quad (4)$$

Recognizing that the controls are the independent variable and therefore the function to drive to zero is the difference in loads between the target and actual values is:

$$\Delta L \equiv g(\theta_0, \theta_{1,c}, \theta_{1,s}) = \begin{bmatrix} T_{target} \\ M_{x,target} \\ M_{y,target} \end{bmatrix} - \begin{bmatrix} T_{actual} \\ M_{x,actual} \\ M_{y,actual} \end{bmatrix}. \quad (5)$$

Therefore the control adjustment may be given as:

$$\Delta\theta = J^{-1}\Delta L. \quad (6)$$

2.3.3 Trim Algorithm

The trim algorithm selected for these runs was a slightly modified version of the quasi-steady trim algorithm implemented internal to DYMORE [86]. The control law:

$$\theta_{final} = \theta_{initial} + \frac{t_{final} - T_{initiate}}{T_{adjust}}\Delta\theta \quad (7)$$

was designed to steer a system to match a set of target values without requiring one to specify what the system physically represents.

A generic control law is completely unaware of system dynamics and may be subject to instabilities. To remedy this, there are three additional parameters which may be optimized for a particular system. These are a set of gains, a gain schedule, and rate limiter.

The set of gains acts on each individual control separately and can be defined so that a rotor collective is able to move more quickly than the cyclic controls, which are frequently required to make smaller scale adjustments. The gain schedule is particularly useful as large initial changes in input controls that can be reduced as the control signals near the target values. This is accomplished by defining a “minimum error”, beneath which the gain rapidly tends towards zero:

$$g_s = \tanh(4e/e_{min}), \quad (8)$$

where error is defined as:

$$e = \frac{\|L_{target} - L_{actual}\|}{\|L_{target}\|}. \quad (9)$$

Finally, without proper controls in place, such a control law would attempt to move the controls to reach the solution in the span of the adjustment phase (T_{adjust}). For complex non-linear systems with short adjustment phases, this approach is not feasible or accurate, and

it may also introduce multiple transients that could drive the system unstable. Therefore a third parameter is used to limit the maximum rate of the controller:

$$\|\theta_{final} - \theta_{initial}\| < R_{max}. \quad (10)$$

This equation limits the change in controls at each time step, causing controls to move smoothly one iteration to the next. This is also a more physical approach, as pilots are unable to cause sudden step changes in their controls.

The original DYMORE source code did not operate in this manner. Originally, the trim rate (the left hand side of Eqn. 10) was set equal to the norm of the current control values instead of the difference in the controls between timesteps. Changing the right hand side Eqn. 10 provided the ability to limit the rate regardless of the magnitude.

2.4 Development of an Advanced Trim Algorithm

If trim is desired in a tightly-coupled elastic framework, the trim procedure cannot be separated from the more computationally costly aerodynamic calculations. Therefore it is imperative that the trim algorithm be as cost effective as possible. Traditional autopilot approaches can require a significant and unacceptable amount of time to trim. When controls are updated smoothly (as done in flight) the Jacobian cannot be updated.

2.4.1 Kriging-based Trim Algorithm

A trim algorithm requires knowledge of the dependent variable sensitivities to trim targets, which are used to formulate the inverse trim Jacobian. Supplied with initial controls, various schedules of control adjustments are possible. In CFD/CSD trimming, the trim sensitivity is formed by explicitly perturbing each trim control angle and measure the response, so that the derivative to be computed by finite differencing techniques. This is costly, as transients introduced at each control perturbation must be allowed to damp out and reach a steady value. As the initial controls are not known *a priori*, instead reasonable values are used as a starting point from which the trim can begin. These initial controls may be not be near the desired trim control angles, adding further time to trim.

An alternative approach is to apply a metamodel to describe the relationship between trim controls and their responses. This approach permits the interpolation of response values so that the trim controls corresponding to the trim targets can be estimated. An inverse trim Jacobian can be built from the metamodel through a mathematical examination of the data, requiring no explicit perturbation of the trim controls. This approach removes the need to wait for the model to damp out transient effects.

A metamodel permits a trim approach that offers the following advantages over a traditional autopilot approach:

1. Solution estimation.
2. Jacobian estimate at different operating points.
3. Requires no perturbation of controls.

2.4.2 Mathematical Description of Kriging

Kriging is an interpolation technique that originated in geostatistics [90] but today is recognized to be useful over a broad range of fields. In the late 1980s, engineers applied kriging to reduce computational time in expensive numerical simulations where differential equations were solved by fast approximations of the solution for a given set of inputs [91]. Kriging has been successfully applied in the field of aerodynamics [88, 92–95].

Kriging assumes that a solution field is comprised of a random field with a superimposed deterministic trend [91]. Typically the random field is modeled with either a semivariogram or a covariance model:

$$\begin{bmatrix} \hat{y}(x) \end{bmatrix} = \begin{bmatrix} \beta \end{bmatrix}^T \begin{bmatrix} f(x) \end{bmatrix} + \begin{bmatrix} Z(x) \end{bmatrix} \quad (11)$$

where \hat{y} is the best estimate at a given point, x . The deterministic trend is represented by scalars, β , which are applied to a vector of regressors, $f(x)$. The random component is described by the random function, $Z(x)$.

There exist several types of kriging associated with varying underlying model assumptions. The three most common are ordinary kriging, simple kriging, and universal kriging.

Ordinary kriging assumes that the deterministic trend is an unknown constant, while simple kriging extends this assumption and assumes that the mean value is zero or otherwise known. Universal kriging fits a linear trend to the deterministic component of the data.

In this effort, as is common in computational uses of kriging, ordinary kriging is applied, which assumes that the field of data are flat/constant (e.g. a plane in 3D) so that $\{f(x)\} = \{1\}$, with “random” perturbations ($Z(x)$). Ordinary kriging was selected for its ability to capture complex results with few training points while maintaining a simple form [95]. The value of each constant, β , is unknown in ordinary kriging, and so the estimator may be described as:

$$\begin{bmatrix} \hat{y}(x) \end{bmatrix} = \begin{bmatrix} \beta \end{bmatrix} + \begin{bmatrix} Z(x) \end{bmatrix}$$

This may be expressed as a linear predictor:

$$\hat{y}(x) = \lambda^T(x)Y \tag{12}$$

The predictor, $\lambda^T(x)Y$, is known as the Best Linear Unbiased Predictor (BLUP) and is found by satisfying the following two conditions:

1. The predictor should be completely unbiased.
2. The predictor should have minimal variance.

Each of these requirements will be discussed in detail, but first recall the meaning of a useful and necessary statistical tool, the *expected value*. The expected value may be thought of (in this context) as the value one would obtain by sampling an infinite number of equally spaced points across a domain of interest and then taking their average:

$$E[x] \equiv \text{expected value}$$

Thus, the expected value of a function, $g(x)$, where x is an k -dimensional independent variable, over a region, Ω , would be:

$$E[g(x)] = \frac{\iint_{\Omega} g(x) \partial x}{\iint_{\Omega} \partial x}, x \in \Omega$$

Clearly this is the equivalent of the average of $g(x)$ over the spatial region Ω .

2.4.2.1 Unbiased Predictor

The requirement of having an unbiased predictor can be restated as a requirement that the difference in the expected value and the actual value is zero over the domain. Let:

$$\epsilon \equiv \hat{y}(x) - y(x).$$

Thus, from the definition of unbiased, the expected error is zero:

$$E[\epsilon] = 0.$$

Substituting in the previous definition for the predictor gives

$$E[\lambda^T(x)Y - y(x)] = 0.$$

Recognizing that the expected value in ordinary kriging is the unknown mean, β , the equation becomes:

$$\beta \lambda^T(x)\{1\} - \beta = 0.$$

Simplifying this expression gives the condition for an unbiased predictor:

$$\lambda^T(x)\{1\} = 1.$$

Simplifying this expression gives the condition for an unbiased predictor: Therefore, in order for the predictor to be unbiased, the weights must sum to one.

2.4.2.2 Minimization of Predictor Variance

The second requirement of the kriging predictor is that the variance of the estimator needs to be minimized. It is not sufficient for the model to only be unbiased, but the model should attempt to fit the data in a way that produces the lowest possible error estimate. For example, *if* the second requirement were *not* present, a trivial solution could be given as:

$$\hat{y}(x) = \beta.$$

However, it is obvious that such a model is of limited practical use. Therefore it is desired to improve the estimate by minimizing variance.

To begin, examine the variance of ϵ :

$$\text{Var}(\epsilon) = \text{Var}(\hat{y}(x) - y(x)).$$

Substituting Eqn 12 yields:

$$\text{Var}(\epsilon) = \text{Var}(\lambda^T(x)Y - y(x)).$$

Algebraic manipulation results in the following:

$$\begin{aligned} \text{Var}(\epsilon) &= \text{Var}\left(\left[\lambda^T - 1\right] \cdot \begin{bmatrix} Y \\ y(x) \end{bmatrix}\right) \\ \text{Var}(\epsilon) &= \left[\lambda^T - 1\right] \left(\begin{bmatrix} R \end{bmatrix}\right) \begin{bmatrix} \lambda \\ -1 \end{bmatrix}, \end{aligned}$$

where R is the correlation matrix, defined as:

$$R = \begin{bmatrix} R(x_1, x_1) & R(x_1, x_2) & \cdots & R(x_1, x_n) \\ R(x_2, x_1) & R(x_2, x_2) & \cdots & R(x_2, x_n) \\ \vdots & \vdots & \ddots & \vdots \\ R(x_n, x_1) & R(x_n, x_2) & \cdots & R(x_n, x_n) \end{bmatrix}.$$

Solving for both the minimization of variance with the unbiased constraint [96] yields the most likely predictor for an unknown point, x :

$$\hat{y}(x) = \beta + r^T(x) \left[R^{-1} \right] (Y - \beta), \quad (13)$$

where β is found by the method of least-squares:

$$\beta = (\{1\}^T R \{1\})^{-1} \{1\}^T R^{-1} Y. \quad (14)$$

Finally, the variance at a point can be solved for:

$$Var(\hat{y}(x)) = \sigma^2 \left(1 - \begin{bmatrix} \{1\}^T & r^T(x) \end{bmatrix} \begin{bmatrix} 0 & \{1\}^T \\ \{1\} & R \end{bmatrix} \begin{bmatrix} 1 \\ r(x) \end{bmatrix} \right). \quad (15)$$

The process variance of the region, σ^2 , may be found as:

$$\sigma^2 = \frac{1}{n} \left[Y - \{1\}\beta \right]^T R^{-1} \left[Y - \{1\}\beta \right]. \quad (16)$$

With these formulae known, it is now possible to model a set of observation (training) points with kriging.

2.5 Trim Algorithm for Loose and Tight Coupling

The original formulation of the kriging-based trim algorithm in DYMORE was implemented by Zake [97]. This implementation was used to trim the UH-60A “blackhawk” vehicle in different flight regimes by Reveles et al. [88]. However, the trim algorithm had difficulty maintaining stability with some flight conditions and was not robust.

2.5.1 Inverse Problem

The original implementation considered the loads as an independent variable while the controls were dependent on the loads, and was therefore an inverse problem. There existed concern about the possibility of a plurality of load solutions being mapped to different controls. The algorithm was changed to allow the control angles to be independent to the

mapped loads. Since a fixed set of controls should only give rise to a single set of quasi-steady loads, this is also a more physical model.

The original implementation applied an inaccurate representation of the inverse trim Jacobian, because it attempted to solve the inverse problem. Recall the form of the trim Jacobian (Eqn. 3). The original implementation attempted to approximate its inverse as:

$$\underline{\underline{J}}^{-1} \approx \begin{bmatrix} \frac{\partial \theta_0}{\partial T} & \frac{\partial \theta_0}{\partial M_x} & \frac{\partial \theta_0}{\partial M_y} \\ \frac{\partial \theta_{1,c}}{\partial T} & \frac{\partial \theta_{1,c}}{\partial M_x} & \frac{\partial \theta_{1,c}}{\partial M_y} \\ \frac{\partial \theta_{1,s}}{\partial T} & \frac{\partial \theta_{1,s}}{\partial M_x} & \frac{\partial \theta_{1,s}}{\partial M_y} \end{bmatrix} \quad (17)$$

This approximation, though not accurate, appears attractive as the controls are expressed as a function of the loads. An accurate representation of the inverse trim Jacobian would have required a search for a given difference in controls if the inverse problem were left in place. This is illustrated for the first term, by extension of Cramer's Rule:

$$\underline{\underline{J}}_{1,1}^{-1} = \frac{\left(\frac{\partial M_x}{\partial \theta_{1,c}} \frac{\partial M_y}{\partial \theta_{1,s}} - \frac{\partial M_x}{\partial \theta_{1,s}} \frac{\partial M_y}{\partial \theta_{1,c}} \right)}{\det(J)}, \quad (18)$$

which is not the same as the approximation. With the problem recast directly (i.e. loads as a function of control angles), the trim Jacobian may be directly computed and its inverse taken.

2.5.2 Jacobian Formation

The original algorithm employed a finite difference scheme to compute the trim Jacobian's sensitivities with a fixed positive 5% change in controls. This proved to be cumbersome to use as it required that the kriging model predict data at exactly a 5% point, where potentially the variance could be unacceptably high. In practice, this limited it to only situations where the estimated control point was essentially the same as the actual solution point. An automated approach that could reliably predict an accurate Jacobian was desired.

This problem is illustrated in Figure 2. The region surrounding each training point is known to a high degree of confidence. However, as one moves away from a training point into the interstitial space between points, the uncertainty grows. As an estimate, this

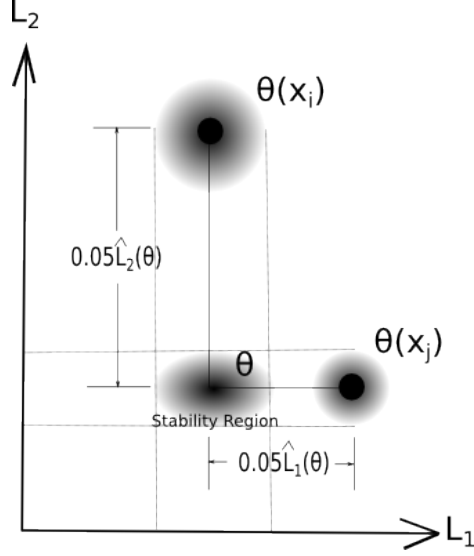


Figure 2: A diagram of the stable region of the original algorithm.

uncertainty is not necessarily an issue. The concern stems from the simple finite difference equation used to estimate the sensitivity:

$$\frac{\partial L_i}{\partial \theta_j} \approx \frac{0.05L_i}{\hat{\theta}_j(1.05\theta) - \theta_j}. \quad (19)$$

It is very important to obtain the diagonal (i.e. $i = j$) term signs correctly to steer the solution to trim. Eqn. 19 has a number of weaknesses that make this a challenge. First, the finite difference is based on the difference between predicted and the current values of loads (L_i). This naturally limits the speed of the trim algorithm, as the faster the controls are updated, the more transients appear possibly resulting in sign errors. As the controls (θ_j) are predicted from the loads (L_i) at a fixed 5% increment, Eqn. 19 requires that the predicted points be near enough to training points so that the prediction error does not result in a sign error. Figure 2 elucidates the narrow region of stability (determined by the region that does not cause error in sign) caused by such a scheme. Further, to maintain the stability, the controls must not stray from this region so that the trimmed control solution must exist within the stability region.

As more training points are added, it is possible that more stable regions may appear, especially as clustered points give rise to larger regions of high confidence. However, this is complicated by the inverse formulation where the loads are treated as independent variables

requiring the engineer pre-select loads rather than controls as training points. This is generally very difficult to do as it is essentially identical to the original trim problem that the engineer seeks to solve.

Due to the inflexibility of Eqn. 19, several alternatives were explored. One idea was to estimate the sensitivity directly from the kriging metamodel, however such an approach suffers from noise [98].

The search for a more accurate trim Jacobian led to improvements that minimize the stability issues. As the training points are almost never in the control directions, the formulation of the trim sensitivities requires solving another problem, albeit a much easier one in the case of few controls. The directional derivatives must be decomposed into the contributions in the control directions. Four points of zero variance (training point locations) are required, and they must span the \mathbb{R}^3 control-space basis. The method for choosing them is:

1. Sort the controls by those closest to the current controls,
2. Select the first two closest controls for data points,
3. Seek the next closest control which is not colinear with the first two as the third data point,
4. Seek the next closest control which is not coplanar with the other three data points for the fourth point.

With the four data points now selected, vectors, \vec{v} , can be drawn connecting them:

$$\vec{v}_j = (\theta_0(x_{j+1}) - \theta_0(x_1))\hat{i} + (\theta_{1,c}(x_{j+1}) - \theta_{1,c}(x_1))\hat{j} + (\theta_{1,s}(x_{j+1}) - \theta_{1,s}(x_1))\hat{k} \quad (20)$$

The sensitivity of a trim target with respect to the direction of a vector is given by:

$$D_{\vec{v}_j} L_i = \frac{L_i(x_{j+1}) - L_i(x_1)}{\|\vec{v}_j\|}. \quad (21)$$

After numerically evaluating Eqn. 21, a system of equations can be solved to decompose the sensitivities into components located in the control directions by examining the definition of the directional derivative:

$$D_{\vec{v}_j} L_i = \frac{\partial L_i}{\partial \theta_0} \left(\frac{\vec{v}_j \cdot \hat{i}}{\|\vec{v}_j\|} \right) + \frac{\partial L_i}{\partial \theta_{1,c}} \left(\frac{\vec{v}_j \cdot \hat{j}}{\|\vec{v}_j\|} \right) + \frac{\partial L_i}{\partial \theta_{1,s}} \left(\frac{\vec{v}_j \cdot \hat{k}}{\|\vec{v}_j\|} \right). \quad (22)$$

Thus, the system of equations for the case of three unknown controls may be explicitly written as:

$$\begin{bmatrix} \frac{\vec{v}_1 \cdot \hat{i}}{\|\vec{v}_1\|} & \frac{\vec{v}_1 \cdot \hat{j}}{\|\vec{v}_1\|} & \frac{\vec{v}_1 \cdot \hat{k}}{\|\vec{v}_1\|} \\ \frac{\vec{v}_2 \cdot \hat{i}}{\|\vec{v}_2\|} & \frac{\vec{v}_2 \cdot \hat{j}}{\|\vec{v}_2\|} & \frac{\vec{v}_2 \cdot \hat{k}}{\|\vec{v}_2\|} \\ \frac{\vec{v}_3 \cdot \hat{i}}{\|\vec{v}_3\|} & \frac{\vec{v}_3 \cdot \hat{j}}{\|\vec{v}_3\|} & \frac{\vec{v}_3 \cdot \hat{k}}{\|\vec{v}_3\|} \end{bmatrix} \begin{bmatrix} \frac{\partial L_i}{\partial \theta_0} \\ \frac{\partial L_i}{\partial \theta_{1,c}} \\ \frac{\partial L_i}{\partial \theta_{1,s}} \end{bmatrix} = \begin{bmatrix} D_{\vec{v}_1} L_i \\ D_{\vec{v}_2} L_i \\ D_{\vec{v}_3} L_i \end{bmatrix}. \quad (23)$$

While for three trim controls the cost of solving this system is virtually zero compared to the cost of obtaining the aeroelastic solution, it is important to realize that in general this method scales poorly. Since the cost of solving a linear system is $\mathcal{O}(n^3)$ for n controls, the total cost scales as $\mathcal{O}(n^4)$ since it must be solved for each of the n trim targets. Therefore, while this methodology has very low cost for only a handful of controls it may not be appropriate for cases where there are a very large number of trim controls.

2.5.3 Control Estimate

The original algorithm located the controls which had a maximum likelihood of corresponding to the trim target. The variance can be applied to determine which controls maximized the likelihood of mapping to the trim target.

Perhaps the most valuable asset of kriging is that it provides two estimates. First, a given position's most likely solution is estimated. This is no different than deterministic interpolation algorithms. Kriging also provides a second estimate of the amount of variance at the most likely solution. Together, a set of controls can be sought that globally has the maximum likelihood of solving the trim problem rather than finding a set of controls whose individual maximum likelihood solves the trim problem.




Figure 3: Probability density functions at two points, x_a and x_b are shown for a range of predictions.

Figure 3 illustrates this concept. Consider the case where a solution $y(x) = C$ is sought, where C is an arbitrary constant. This is an identical problem to the solution estimate, $L(\theta) = L_{target}$. In the figure, the probability distribution function (PDF) of two different points, x_a and x_b are plotted across the range of the dependent variable y .

If, as an example, the constant C is set to zero, then the position x is sought that corresponds to $y(x) = 0$. Given the PDFs of Figure 3, there exists a greater probability of $y(x_b) = 0$ than $y(x_a) = 0$. However, deterministic interpolation algorithms provide only the most likely value of a given point (i.e. $\hat{y}(x_a)$ and $\hat{y}(x_b)$). As $\hat{y}(x_a) = 0$, the deterministic interpolation algorithm finds x_a as the solution despite x_b having a greater likelihood of solving the problem.

In the control angle problem, the variance that kriging estimates can be used to find the most likely controls that map to the target loads. To accomplish this, the errors have been assumed to form a Gaussian distribution, permitting the likelihood of a value to be described as:

$$p(y) = \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{1}{2}\left(\frac{y-\mu}{\sigma}\right)^2}, \quad (24)$$

where μ is the maximum value.

By applying a simple grid-based search over the region of interest, the point with the absolute maximum likelihood of satisfying the solution may be found. This is especially important as ordinary kriging predicts controls that result in loads that are near the spatial averages of the region in areas where few data exist. As the spatial averages are potentially near the solution point, the ability to separate predicted solutions based on their relative likelihood offers significant advantages over other methods.

2.5.4 Improved Algorithm

In summary, the original algorithm suffered several shortcomings. It was formulated as an inverse problem where controls were a function of loads, had stability problems in trim, and used an incorrect approximation of the inverse trim Jacobian. The improved algorithm was recast with the controls as the independent variable, therefore removing the inverse problem. Secondly, the 5% constant offset was removed, and the point of lowest variance was selected to ensure that the highest quality predictions were obtained. Finally, the loads were no longer required to correspond to a delta in each control individually, but instead an arbitrary direction in control space can be selected.

Taken together, these changes result in a practical trim algorithm for use in tightly-coupled CFD/CSD simulations. The second and third improvements are especially useful because, given an arbitrary direction with unbounded magnitude, the lowest variance points will always be the training points themselves. This permits a simple construction of the inverse trim Jacobian that is not a function of the current loads, which may themselves have error with respect to the quasi-steady values at the current controls.

While far more robust than the original, the described improved trim algorithm could fail to determine correct trim sensitivities along the Jacobian diagonal under two circumstances. As this will result in unstable trim, it is necessary to understand where this can occur. The first condition causing sign identification failure arises when the observation points (i.e. training database) does not box in the estimated control angles. In these situations, the trim sensitivity is estimated from the closest points to the solution, but the region may

not adequately represent the local sensitivity potentially causing a sign error. The second situation occurs when the observation points are spread over too great an area, with too few nearby the solution. In these situations, the finite difference is unable to adequately resolve the local sensitivities, possibly causing a sign error. To avoid these problems, at least four observation points that span the basis need to be in the vicinity of the desired solution.

CHAPTER III

HYBRID URANS / FREE-WAKE DESCRIPTION

The traditionally high cost of computing aerodynamic CFD/CSD solutions has led many to attempt to create hybrid solvers [99]. The development of a hybrid solver with the ability to identify causal physics in a single code via modularity is outlined. By coupling a URANS CFD solver with a free-wake code and a comprehensive code, different hierarchical components of the code can be activated. The following hierarchy can be constructed:

- Comprehensive (CSD) code alone
- Hybrid CFD / free-wake with comprehensive (CSD) code
- Full CFD with comprehensive (CSD) code

3.1 Hybrid Solver Theory

Rotating systems such as rotorcraft produce a shed wake that influences the aerodynamic flowfield of the vehicle [2]. A hybrid aeroelastic solver that reduces computational requirements of an unsteady RANS CFD solver while maintaining solution accuracy has been designed for rotating systems such as rotor blades. This approach permits the computational domain to be significantly reduced in size, saving computer memory and costs. The URANS equations resolve the near-field single or overset structured grid system about the rotor blade or rotor disk. A free-wake calculation based on the solution of the potential equations models the effect of blade vortices in the far-field.

A major goal in the application of a hybrid-CFD approach is to reduce the cost of computing an aerodynamic solution as compared to a full-CFD approach, while retaining the majority of the physical accuracy of the CFD methodology. This is accomplished by reducing extent of the grid domain that needs to be resolved, thus reducing computational costs. A low-cost free-wake model is applied to update the far-field boundary on the reduced domain to include the influence of a wake outside of the CFD grid domain. The free-wake

far-field solver accurately models complex wakes, including blade-vortex interactions, that are developed as a rotor generates lift, illustrated in Figure 4.




Figure 4: The ADM rotor with swept tips was modeled at an advance ratio of $\mu = 0.35$. The isosurfaces visualize the OVERFLOW shed vorticity, while the lines illustrate the free-wake vorticity predicted propagation.

The coupled URANS/free-wake methodology exploits URANS CFD’s highly accurate solutions in the near-field, while using the free-wake model in the far-field. The coupling is accomplished using the process outlined in Figure 5. The CFD solver first computes the aerodynamic loads on a rotor blade and passes them to the free-wake, which converts them into circulation values. The circulation advances the free-wake, which then computes the induced velocities along the outer boundary of the reduced CFD domain. These induced velocities update the boundary conditions internal to OVERFLOW, which act to effectively extend the modeled domain size.

Early in the development of the hybrid solver, it was discovered that sufficiently small time steps could result in incorrect values of circulation. Due to the low-velocity region surrounding areas of reverse flow (for example, Figure 6) on which the Kutta-Joukowski computation of bound circulation strength is numerically inaccurate. These computational errors were observed at some radial locations (r) within the region of reverse flow and were the result of the presence of a zero local velocity (V_L) in the denominator of the equation for the bound circulation, Γ :

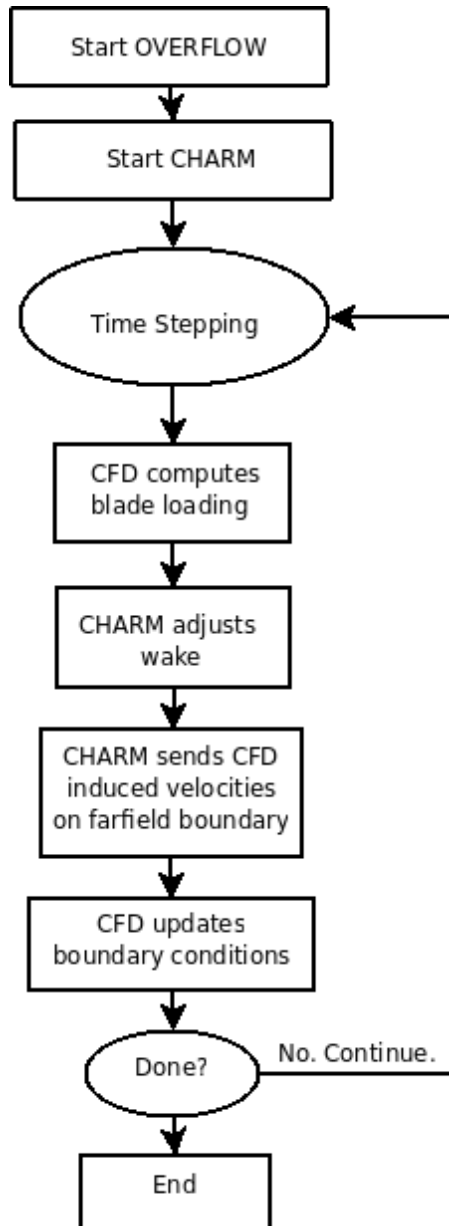


Figure 5: A diagram showing the flow of data in the coupled hybrid aerodynamics code.

$$\Gamma = \frac{L'}{\rho_\infty V_L}, \quad (25)$$

where L' is the lift per unit span at the radial location, and ρ_∞ is the free stream density. The local velocity is defined as the velocity normal to the leading edge of a rotor blade at the radial location where the circulation is computed. The local velocity is the superposition of the rotation (Ωr) and the vehicle flight speed (V_∞), so that $V_L = \Omega r + V_\infty$. For the rotor, there may be occurrences when the local velocity can be zero ($\Omega r = V_\infty$) in the reverse flow region (e.g., Figure 6). If the local velocity is zero, then a numerical singularity in the computation of the circulation will occur. This singularity can be avoided if the circulation is expressed in terms of the chord (c) and the sectional lift coefficient (c_ℓ), so that the bound circulation will be zero only when the local velocity is also zero:

$$\begin{aligned} L' &= \frac{1}{2} \rho_\infty V_L^2 c c_\ell \\ \Gamma &= \frac{V_L c c_\ell}{2} \end{aligned} \quad (26)$$

Numerous discontinuities were also observed in the computed bound circulation (Figure 7) when local velocities approach zero near the reverse flow region. A velocity cutoff or tolerance was added to the circulation, so that the numerical computation of the bound circulation more closely approximates its physical behavior (Figure 7). The value of the cutoff is determined on a case-by-case basis by selecting the lowest value that permit the free-wake to execute without encountering non-physical circulation strengths. In the illustrated 8534 case, this was determined to be 20 feet per second.

3.2 Baseline and Coupled Codes

Comprehensive Hierarchical Aeromechanics Rotorcraft Model (CHARM) [100,101] is a free-wake code developed by Continuum Dynamics Incorporated (CDI). CHARM directly calculates vortex positions and contains its own panel code so that it is capable of analyzing




Figure 6: The boundary of reverse flow for the UH-60A at the high-speed 8534 flight conditions.




Figure 7: Circulation of the UH60 High-Speed Case (C8534) illustrating the impact of local velocities.

rotor/wake/airframe interaction problems. Its primary uses include computation of performance characteristics, blade loading, blade deformations and dynamics, wake geometry, far-field aeroacoustics, and interactional aerodynamics.

This hybrid solver incorporates a number of features that sets it apart from other hybrid solvers. NASA’s URANS computational fluid dynamics research solver OVERFLOW has been coupled with the industry standard free-wake code CHARM by Continuum Dynamics, Inc. (CDI). Aeroelastic simulations are achieved through coupling with the computational structural dynamics codes (CSD), DYMORE4 or RCAS [102], which are also widely used in industry.

3.3 Parallel Performance

The CFD code OVERFLOW requires data be sent to various processors within its parallel framework. With the hybrid approach, wake-induced velocities must be sent to all compute nodes that contain farfield boundaries to be updated.




Figure 8: Tests were conducted on different numbers of processors to validate the parallelization.

To accomplish this, the interface coupling code was parallelized and then tested on different numbers of processors (Figure 8). Although grid partitioning is necessarily not the same

for different numbers of processors, excellent agreement of the results was demonstrated regardless. Processor independence is necessary for a general parallel code, and validation of parallel accuracy is crucial.

3.4 *Wake Effects*

The influence of the far wake on the prediction of the performance is illustrated in Figure 10. The parameters usually evaluated for trim (thrust, roll moment and pitching moment) are examined with and without the **CHARM** wake. These effects were evaluated by utilizing the same cropped off-body grids with and without **CHARM** (shown in Figure 9). A simulation with the ADM rotor at an advance ratio, $\mu = 0.20$, is started with **OVERFLOW** alone using the near-body grids. The blade grids have 60 spanwise airstations and represent 77% of the 1.65M grid node count. The loads are allowed to stabilize for two revolutions, and then the far wake was included (by initializing the **CHARM** calculations).




Figure 9: The ADM grids for the hybrid-solver.

As the wake convects, it induces a downwash on the rotor, lowering the total thrust. As the downwash is greater on the advancing (right) side of the rotor disk where the local

velocities are greater compared to the retreating (left) side of the rotor disk, there is a left rolling moment. The interaction of the aft portion of the rotor disk with the wake introduces a larger downwash than the forward portion of the rotor disk, resulting in a net positive pitching moment, as observed in Figure 10.




Figure 10: Example of the effects of starting the hybrid code after two revolutions.

3.5 Hierarchical Options

Several hierarchical options are accessible via namelist options. This permits a variety of run options all within one code that can either be selected to accomplish a specific task (e.g. low-cost stability analysis) or to evaluate causal physics (e.g. by limiting the wake).

1. Full CFD/CSD via tight coupling
2. Full CFD/CSD via loose coupling
3. Hybrid CFD/CSD via tight coupling
4. Hybrid CFD/CSD via loose coupling
5. Comprehensive CSD code alone

The CSD code DYMORE has been used in this work. It is also possible to use any pre-existing interface within OVERFLOW for other comprehensive codes, such as RCAS [102] or CAMRAD-II [103]. The code is designed to be *carefree*, and so as new options are added, more interface options may become available.

CHAPTER IV

CFD/CSD TIGHT COUPLING RESULTS

4.1 UH-60A Tight Coupling Trim

The UH-60A flight test program [10] has generated high-quality experimental data including structural loads and integrated pressure tap measurements with which computations may be compared. Quasi-steady level flight cases (shown in Table 1) such as a high-speed forward flight condition (C8534), and a high thrust condition with dynamic stall (C9017) are excellent test cases with which to evaluate the tightly-coupled CFD/CSD trim algorithm. Case C8534 is a high-speed forward flight case operating in quasi-steady level flight. Due to the high-speed, the wake convects away from the vehicle quickly, reducing complex wake interactions. As a frequently studied case [29, 38, 104–106], it is ideal as a testbench for loose and tight coupling comparisons.

The CFD code OVERFLOW was selected for these computations. A coarse grid containing 5.1 million nodes was used in these analyses. Each of the four blade grids contains 125 points in the chordwise direction and 81 spanwise airstations. The y^+ value for the grids is approximately one across the blades. Left-hand side terms in OVERFLOW were discretized with an ARC3D Beam-Warming block tridiagonal scheme while right-hand side terms were discretized using Euler central differencing. Smoothing was applied to second-order and fourth-order terms of enthalpy with coefficients of 2.0 and 0.04, respectively.

The structural dynamic code DYMORE has 81 airstations co-located with the CFD airstations to remove interpolation error and maintain spatial conservation. Each blade has

Table 1: UH-60A Test Case Flight Conditions.

Case	Rotor Speed (RPM)	Density (slug/ft ³)	Temperature (F)	Airspeed (ft/s)	Pitch (deg)	Sideslip (deg)
c8534	258.1	0.0020823	71.814	266.5	-4.31	1.27
c9017	255.8	0.0013242	24.761	170.2	2.85	-1.59

been discretized into 10 third-order finite elements. The loose coupling algorithm applied incremental timesteps of 1° within the comprehensive code to trim the solution via an autopilot controller over 300 to 400 revolutions. Loads were filtered in DYMORE using a Fourier decomposition within the trim algorithms. All CFD integrated hubloads were examined and plotted with a moving average technique.

For the tight coupling process, the CFD and CSD modules advanced in time with the same azimuthal increment (0.05°), and data are transferred between the modules at each time step. The initial control settings are initialized by the CSD simulation, applying a training database with the kriging model. These are held fixed until tight coupling trim begins, at which point the controls are smoothly updated each timestep to continuously drive the system to a trimmed state.

4.1.1 Optimization of Trim Process

The tight coupling trim process was evaluated with the UH60A high-speed forward flight case C8534. The objective was to minimize the cost associated with trimming within parity of the loose coupling algorithms. Clearly this is a formidable task, as the loose coupling trim procedure (run in a cost-effective comprehensive code) requires hundreds of revolutions in a coupling iteration for each partial revolution of CFD computations, as seen in Figure 11. Figure 12, the loose coupling trims iteratively and trims in just five coupling iterations, corresponding to just two revolutions of CFD.

A naïve trim approach where the CFD simulation has been initialized from free streams values only (i.e. a so-called “cold start”) is illustrated in Figure 13. While the results indicate that the tight coupling simulation required considerably longer to trim than the loose coupling simulation did, it did converge to a trim condition, indicating a successful application of the trim algorithm.

Careful examination of the initial trim results (Figure 13) led to the understanding that the severe loads encountered in the beginning of the trim were caused by start-up transients in the CFD solver’s flow field. These transients are then exchanged between the CSD and CFD modules via the tightly-coupled algorithm. Therefore this “ringing” effect adversely

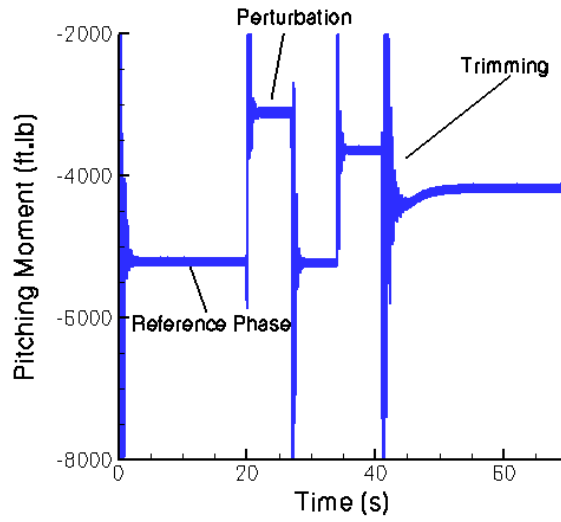


Figure 11: Autopilot trimmer for the UH-60A C8534 case.

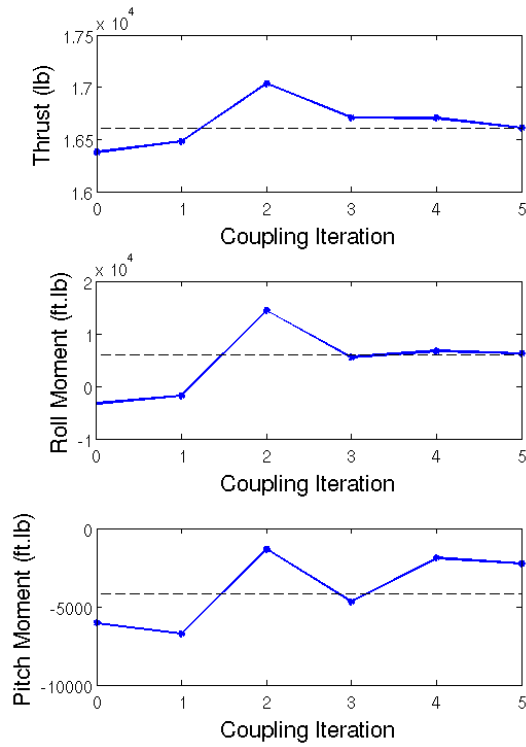


Figure 12: Loose coupling trim of the UH-60A C8534 case.

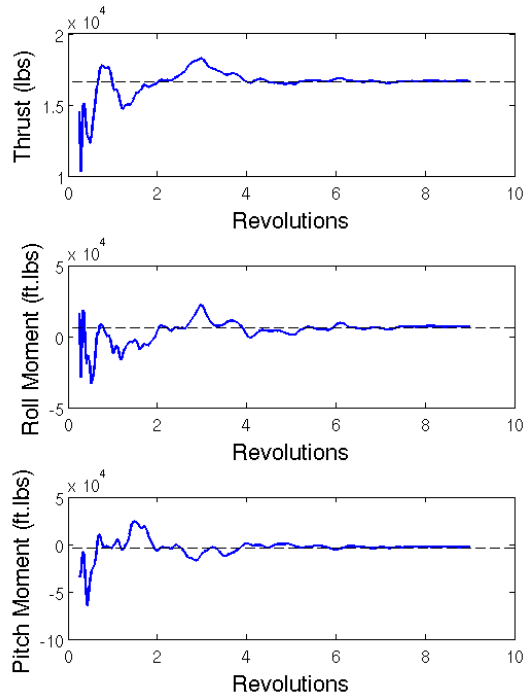


Figure 13: Trimming of the UH-60A C8534 case in tight coupling from a cold start.

affects the trim algorithm and could potentially compromise the stability of the trim process for some flight conditions.

An alternative start-up procedure using prescribed motion was investigated to see if the time to trim could be further reduced. If, as was hypothesized, the transients from the aerodynamic free stream initialization are responsible for inducing the large perturbations in the system, then the alternative approach would greatly reduce the computational cost of the trim time. Removal of these transients would also enhance the stability of the trim algorithm. Therefore, an important question arises, which is the most efficient method to start up the tight coupling.

The C8534 UH-60A flight case was restarted from various solutions of prescribed motion between the zeroth and first revolution. The motion came from the comprehensive code’s solution at predicted controls that was then run with constant controls in the CFD solver to allow transients to pass. Five cases were evaluated: a “cold start”, and prescribed motion of *slfrac*18, *slfrac*14, *slfrac*12, and 1 revolution, the details of which are provided in




Figure 14: Revolutions needed to reach thrust convergence (1% error) based on loose coupling initialization.

Table 2: Optimization of convergence using kriging to trim during tight coupling.

Total Time Required (revs)	Loose Coupling (revs)	Tight Coupling (revs)	Kriging Trim On (revs)	CPU (hours)
6+	0	6.0	5.0	2814
3.6	1/8	3.5	2.5	1728
2.4	1/4	2.2	1.2	1166
2.7	1/2	2.2	1.2	1306
5.1	1	4.1	3.1	2432

Table 2. The thrust was used as the parameter to determine the efficiency of the transfer between methods. Figure 14 illustrates the behavior of the simulation at each of these initializations. Convergence for the thrust parameter is defined as the first instance in which the error drops below 1% and remains for a full revolution. “Cold starting” the tight coupling solution is not efficient, as it requires a number of revolutions before the error remains below the 1% threshold. Immediately upon restarting from some fractional portion of a loose coupling revolution (which applied the initial kriging estimate for the controls), the errors drop below 5% towards the 1% demarcation of the trim target in approximately half a revolution. Thus, some amount of a prescribed motion restart is recommended prior to starting a tightly-coupled simulation. When evaluated and visually inspected as in Figure 14, the optimal loose coupling start appears to be $\frac{1}{4}$ to $\frac{1}{2}$ revolution for the four-bladed rotor.

The controls were held constant, or “locked”, for one revolution (which has not yet been optimized) after switching from loose to tight coupling computations. This approach was found to be necessary as an additional transient during the transfer from loose to tight coupling appears. If the kriging-based trimmer is started immediately upon the transfer to tight coupling, then large transients, such as the example in Figure 13, will appear at the start of the simulation. At the end of this period, however, an additional training point can be added to the kriging database for application in the current as well as future simulations.

An alternative to this approach was to vary the time for which the controls were held fixed. If the controls are released too early, the simulation will attempt to trim with transients still present, while too long of a hold increases the time to trim by postponing the

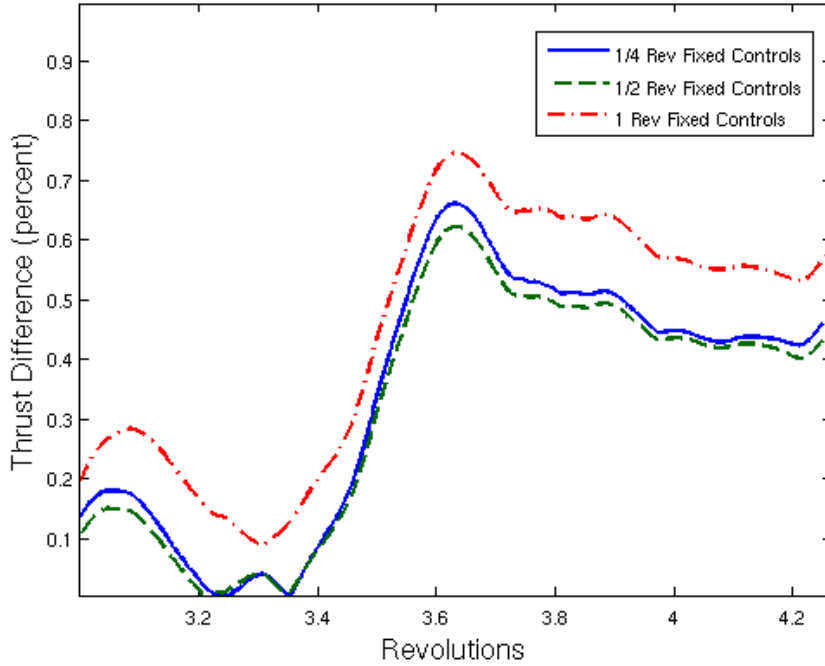


Figure 15: Comparison of thrust convergence to time the controls are held.

start. Figure 15, reveals that a half-revolution or even a quarter-revolution of fixed controls does offer a small advantage in trim time. When such short period of locked controls is applied, the solution may have too many transients to include the result as a training point, and so engineering discretion is required.

Due to the large cost of developing CFD/CSD training points, an investigation was undertaken to determine if efficient trimming could be achieved via the kriging-based trim algorithm using far less costly CSD training points. The CSD databases evaluated for the C8534 database had very good correspondence (less than 5% error) with the controls predicted by OVERFLOW loose coupling. C8534 tightly-coupled simulations were performed using both a twelve point CSD and the nine point CFD/CSD database (Figure 16). As previously noted that initialization from a quarter revolution of loose coupling speeds up convergence, both cases used this initialization to optimize the simulation. The results indicate that, at least for some conditions, comparable convergence behavior can be obtained at significantly reduced training cost from a database with comprehensive code training points.




Figure 16: Thrust convergence of CFD/CSD kriging when using a database provided by CSD alone versus that of using the full CFD/CSD database shows comparable performance.

4.1.2 UH-60A C8534

It was hypothesized that due to the “aircraft-like” wake structure in the high-speed case, aperiodicity in the solution would be minimal, and loose and tight coupling procedures should produce the same physics and trim to the same result. To test this theory, both loosely-coupled and tightly-coupled CFD/CSD simulations were evaluated when the hub loads converged to their target values.

For this initial test of the tightly-coupled trim algorithm, a “box” of controls was selected around the solution using 1.6° spacing. The CFD computations were evaluated using a hybrid RANS-LES turbulence model [107]. Each control point was run in loosely-coupled CFD/CSD approach to compute the corresponding hub loads. The control angle solution to the trim problem was estimated with Kriging via its map between loads and controls. The tight coupling controls were started with the kriging-estimated controls of the target hub loads. A Jacobian was computed and updated every 15° using a 2% first-order approximation of the trim sensitivity derivatives between the current state and the kriging




Figure 17: Trim of UH-60A C8534 using using the optimal trim procedure for tight coupling results.

Table 3: Comparison of Trimmed Controls

Simulation	Collective	Lateral Cyclic	Longitudinal Cyclic
Tight Coupling	13.21°	2.28°	-8.09°
Loose Coupling	13.21°	2.46°	-8.04°




Figure 18: Comparison of hub thrust for the UH-60A C8534 between loose and tight coupling results.

model.

When trimmed using optimal trim procedures, the time to trim is approximately three revolutions (Figure 17) and is comparable to loose coupling simulations that require two revolutions to reach trim. Earlier studies (Figure 13) indicate that the trim solution is very stable. Examination of converged controls (Table 3) reveals that the resulting control angles are very similar (within 0.2 degrees) with the largest differences seen in the lateral cyclic.

The influence of the grid was next evaluated on a refined grid containing 80.9 million nodes (15 times higher node count). The blade grids had twice as many spanwise stations and nearly twice as many chordwise cells. As observed in Figure 19, the airloads for this case change only minimally for the refined grid. The improvements resulting from refining the grid consist of a phase difference of approximately 4° and an increase in the peak-to-peak magnitude of 7.8%.

In Figure 20, loose and tight coupling simulations produce essentially identical results in terms of normal loads. The largest difference observed is at the midspan section of $r/R = 0.55$ at $\psi = 110^\circ$, and is barely discernible. The predicted results capture the





Figure 19: Comparison of UH-60A case C8534 pitching moment airload predictions at $r/R = 0.865$ for coarse and fine grids.


overall trend of the experimental data. The most inboard section ($r/R = 0.225$) misses an event at approximately 0° as the blade passes near the vehicle fuselage, which was not modeled in these simulations. The outboard (Figure 21) and tip (Figure 22) regions continue to produce excellent agreement between loose and tight coupling results. The predictions generally capture the depth of the transonic dip on the advancing side, with a phase lag of approximately 30° .

Pitching moments are next examined. These predictions are more challenging than normal forces, but are important to achieve overall accurate results. For example, it has been shown [104] that improving pitching moment calculations causes sectional airloads to improve in phase with respect to the experimentally measured values. The inboard and midspan sections (Figure 23) confirm that loose and tight coupling pitching moment predictions are nearly identical. These results capture the waveform of the experimental measurements. Moving outboard (Figure 24) and to the tip (Figure 25), the loose and tight coupling results remain nearly identical while the phase lag peaks at $r/R = 0.965$ with the coarse grid.

Chord force is challenging to predict due to its strong dependence on in-plane drag forces. Biedron and Lee-Rausch identified [108] that one particular source of error with respect to the chord forces is the integration itself. Due to the limited number pressure taps, accurate chord forces are difficult to measure experimentally. If the CFD/CSD solutions are integrated using the same procedure as the experiment, the correlation between experiment




(a) $r/R = 0.225$




(b) $r/R = 0.675$

Figure 20: Comparisons of loosely and tightly-coupled OVERFLOW/DYMORE for normal forces on UH-60A C8534 inboard and midspan airstations.




(a) $r/R = 0.865$



(b) $r/R = 0.965$

Figure 21: Comparisons of loosely and tightly-coupled OVERFLOW/DYMORE for normal forces on UH-60A C8534 outboard airstations.



(a) $r/R = 0.99$


Figure 22: Comparisons of loosely and tightly-coupled OVERFLOW/DYMORE for normal forces on UH-60A C8534 tip.

and computations significantly improve. This effect was most prominent near the transonic dip on the advancing side.

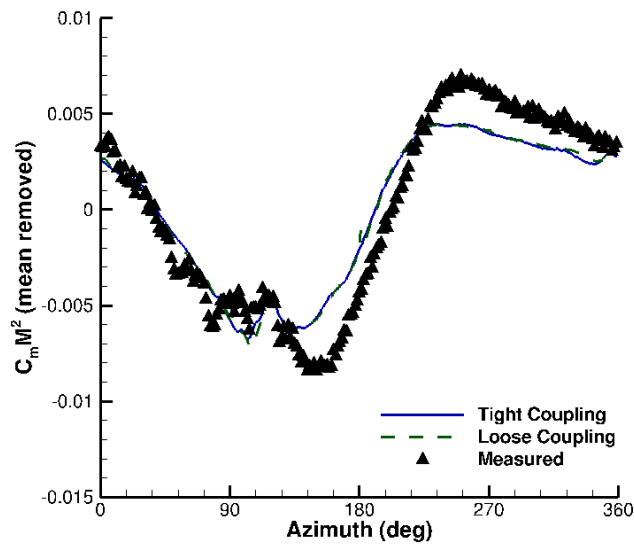
Again examining chord force from root to tip, the loose and tight coupling results contain few discernible differences near the inboard and midspan sections (Figure 26). An event around $\psi = 10^\circ - 20^\circ$ on the station closest to the root ($r/R = 0.225$) is believed to be an effect of the perturbation in the normal force caused by the rotor-fuselage interaction and therefore cannot be predicted without the fuselage in the grid.

Loose and tight coupling predictions still indicate very close agreement of chord force at the outboard (Figure 27) and tip (Figure 28) sections. The advancing side chord forces generally exaggerate the transonic dips with respect to experiment, but as previously discussed this is caused by limited experimental pressure taps. Integration of the CFD results about the same pressure tap points significantly improves the correlation.

Structural loads are examined from root to tip a representative selection of sensors that provided reliable data in this case. Table 4 lists which measurements were working for this case. Due to the harsh environment of the flight test, some sensors were damaged at




(a) $r/R = 0.225$




(b) $r/R = 0.675$

Figure 23: Comparisons of loosely and tightly-coupled OVERFLOW/DYMORE for pitching moments on UH-60A C8534 inboard and midspan airstations.




(a) $r/R = 0.865$



(b) $r/R = 0.965$

Figure 24: Comparisons of loosely and tightly-coupled OVERFLOW/DYMORE for pitching moments on UH-60A C8534 outboard airstations.



(a) $r/R = 0.99$


Figure 25: Comparisons of loosely and tightly-coupled OVERFLOW/DYMORE for pitching moments on UH-60A C8534 tip.

stations that were otherwise instrumented.


While structural responses often have much in common with one another [109], they have persisted as a challenging problem to predict computationally. At the most inboard station, $r/R = 0.2$, (Figure 29) the flap bending moment produces essentially identical results between loose and tight coupling approaches, and the general shape of the waveform is captured numerically with an approximate 15° phase lag with respect to the experimentally measured values. The loose and tight coupling approaches produce chord bending moment

Table 4: List of sensors available and functional during UH-60A C8534.

r/R	Torsion Moment	Flap Bending Moment	Chord Bending Moment
0.2	No	Yes	Yes
0.3	Yes	Yes	Yes
0.4	No	Yes	Yes
0.5	Damaged	Yes	Yes
0.6	No	Yes	Yes
0.7	Yes	Yes	Yes
0.8	No	Damaged	Damaged
0.9	Yes	Yes	No




(a) $r/R = 0.225$




(b) $r/R = 0.675$

Figure 26: Comparisons of loosely and tightly-coupled OVERFLOW/DYMORE for chord forces on UH-60A C8534 inboard and midspan airstations.




(a) $r/R = 0.865$



(b) $r/R = 0.965$

Figure 27: Comparisons of loosely and tightly-coupled OVERFLOW/DYMORE for chord forces on UH-60A C8534 outboard airstations.




(a) $r/R = 0.99$

Figure 28: Comparisons of loosely and tightly-coupled OVERFLOW/DYMORE for chord forces on UH-60A C8534 tip.

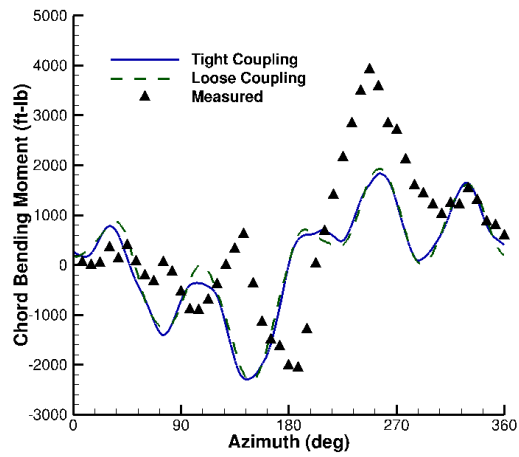
predictions with small differences, where loose coupling results have overshoots with respect to the tight coupling results, especially on the advancing side at approximately $90^\circ - 100^\circ$. A phase lag of approximately 35° exists between the predictions and experiment, while the predicted peak moment at around 270° is half of the measured value.

At $r/R = 0.3$ (Figure 30), all three measurements are available for comparison. The flap bending moment results indicate the existence of minute, but discernible differences between the tight and loose coupling solutions. While the trend of the waveform is generally well captured, a phase lag of approximately $15-25^\circ$ is present between the computations and experiment. The minimum on the retreating side at 270° is missing from the numerical predictions. The chord bending moment shows significant differences with the experiment (generally over predicting). However, the predictions do match the experimental frequency content. The loose and tight coupling torsion moment predictions are essentially identical. Both capture the general trend of the measured waveform, but show an offset of approximately 400 lbs to the experimental values.

Both flapwise bending moments and chordwise bending moments were measured at the




(a) Flap Bending Moment




(b) Chord Bending Moment


Figure 29: Comparison of structural moments for the UH-60A C8534 at $r/R = 0.2$ span.



(a) Flap Bending Moment



(b) Chord Bending Moment




(c) Torsion Moment

Figure 30: Comparison of structural moments for the UH-60A C8534 at $r/R = 0.3$ span.


midspan stations of $r/R = 0.4$ (Figure 31) and $r/R = 0.5$ (Figure 32). Key findings in these data are largely similar at these two locations. The flap bending moment indicate loose and tight coupling predictions are nearly identical with only a small discernible difference at $\psi = 240^\circ$. The predictions contain 30° and 35° phase lags with respect to the experiment at $r/R = 0.4$ and $r/R = 0.5$, respectively. The chord bending moment predictions contain greater disparity to the experimental measurements, except on the retreating side which generally captures the local waveform. The tight coupling chord bending moment predictions are nearly identical to the loose coupling results, although the loose coupling predictions contain higher amplitudes on the advancing side than the tight coupling results. The aft portion of the rotor disk from approximately $240^\circ - 45^\circ$ is best predicted with the correct frequency content and with comparatively little to no phase lag. At the bow of the rotor disk from approximately $45^\circ - 240^\circ$, however, the simulated predictions behave inversely to experiment in amplitude although the frequency content is correct. In this region, the measured local maximum at 150° is a predicted local minimum while the measured global minimum at 180° is predicted as a local maximum.

The outboard midspan section of $r/R = 0.6$ (Figure 33) the flap bending moment contains minor but discernible differences between loose and tight coupling predictions. The phase lag has increased to between $30^\circ - 45^\circ$, but otherwise the trend is well captured. The chord bending moment predictions also demonstrate these minor differences between loose and tight coupling results. The predictions of both approaches have considerable differences with respect to the experimental measurements, capturing the frequency content but reduced peak-to-peak amplitudes relative to experiment. As with the predictions at $r/R = 0.4$ and 0.5 , the chord bending moment predictions are out of phase between $45^\circ - 240^\circ$.

At the outboard station location of $r/R = 0.7$ (Figure 34), all moment measurements are available. The flap bending moment predictions maintain a phase lag of as much as 45° with respect to the experiment. The predictions themselves again show minor discernible differences between loose and tight coupling approaches. The simulated results do not capture the waveform as well as the inboard and midspan sections did due to the differences




(a) Flap Bending Moment




(b) Chord Bending Moment

Figure 31: Comparison of structural moments for the UH-60A C8534 at $r/R = 0.4$ span.




(a) Flap Bending Moment




(b) Chord Bending Moment

Figure 32: Comparison of structural moments for the UH-60A C8534 at $r/R = 0.5$ span.



(a) Flap Bending Moment



(b) Chord Bending Moment


Figure 33: Comparison of structural moments for the UH-60A C8534 at $r/R = 0.6$ span.

in the tip normal force loading (Figure 20). The chordwise bending moment predictions still contain discernible differences between the predicted loose and tight coupling results. While the frequency content of the measured data are clearly visible in the predictions, the waveform shape has poor correlation. The torsion moment predictions are essentially identical between the loose and tight coupling approaches. Similar to the flap predictions, a 45° phase lag has developed with respect to the experiment although the waveform is otherwise captured.


As can be seen in Figure 35, the tip deflections in pitch are essentially identical between loose and tight coupling predictions. Examining the flap deflection, small differences between the two methodologies can be discerned. These differences produce a maximum relative error of 27% at 182 degrees, although this corresponds to an over prediction of the flap angle in the loose coupling approach by approximately $\frac{4}{5}$ of an inch. Most interestingly, the lead-lag deflection of the blades is not periodic with rotor speed in the tightly-coupled simulation, but is in the loosely-coupled simulation. This is possible because the tightly-coupled simulation is free from any assumptions of periodicity while the loosely-coupled simulations implicitly assume that all solution frequencies are comprised of harmonics of rotor revolutions. As the rotor lead-lag frequency is well known [2, 110, 111] to be approximately $0.2\Omega - 0.3\Omega$ for articulated rotors (e.g. the UH-60A), a converged loose coupling simulation could not contain such a frequency. This is verified by a simple expression for the approximation of the lead-lag frequency in Eqn. 27. Inserting the UH-60A lead-lag hinge yields an uncoupled frequency of 0.27Ω .

$$\nu_\zeta = \sqrt{\frac{3}{2} \left(\frac{eR}{R - eR} \right)} \quad (27)$$


As the lead-lag motion couples nonlinearly with the flapping motion due to Coriolis effects [110], this explains the small differences between loose and tight coupling previously seen in the tip flap deflections. This coupling produces no effect on the feathering axis [111], and explains why the pitch closely matches between the loose and tight coupling approaches.



(a) Flap Bending Moment




(b) Chord Bending Moment




(c) Torsion Moment


Figure 34: Comparison of structural moments for the UH-60A C8534 at $r/R = 0.7$ span.



(a)



(b)



(c)

Figure 35: Comparison of tip deflections for UH-60A C8534 between loose and tight coupling results.

Table 5: Comparison of Converged Controls

Simulation	Collective	Lateral Cyclic	Longitudinal Cyclic
Tight Coupling	11.23°	1.87°	-8.12°
Loose Coupling	9.95°	2.86°	-6.61°

4.1.3 UH-60A C9017

The UH-60A case 9017 poses a more interesting case for loose and tight coupling evaluations, as dynamic stall occurs on the retreating side of the rotor disk [112]. Since dynamic stall is very sensitive to simulation parameters and this case is stall limited, the impact of the tight coupling approach on the predictions has been a topic of speculation at the UH60-A airloads workshop. For this case, the coarse grid was again run to minimize computational time, applying the same overall CFD and CSD solution options from the C8534 study. Here the use of the Menter $k\omega$ -SST turbulence model [113] was applied. The original implementation of the kriging-based trimmer did not work for this case, and the improved version was used instead.

The loose coupling approach was started with one full revolution and subsequent iterations occurred every $\frac{1}{4}$ revolution until convergence. The tight coupling approach was started after $\frac{1}{4}$ revolution of prescribed motion, controls were then held fixed for $\frac{1}{2}$ revolution and updated every iteration thereafter throughout the simulation. Convergence (Figure 36) occurred within approximately one revolution of each approach.

Table 5 reveals that the control angles between the loose and tight coupling simulations resulted in greater than a 1° difference in collective control. This is significant and arises due to the stall limit causing low thrust sensitivity with collective control angle at this high-thrust condition.

Inboard and midspan stations are first examined (Figure 37). At the most inboard stations of $r/R = 0.225$ and 0.4 the tight coupling, loose coupling, and comprehensive analyses all produce similar results, although the comprehensive results lack higher harmonics as is to be expected. Significant offsets are seen with respect to the experimental data. Towards the midspan of the blade ($r/R = 0.675$), the offset decreases and the CFD/CSD





Figure 36: Thrust convergence for UH-60A C9017 occurs in about the same amount of time as with loose coupling.

simulations more accurately approximate the measured data. At these stations, amplitude differences exist between loose and tight coupling results, although the phase is nearly identical. The higher amplitude of the tight coupling predictions significantly improves the results over the loose coupling approach at $r/R = 0.675$ at $\psi = 180^\circ$ while the loose coupling approach under predicts the maximum sectional lift coefficient. Alternatively, at approximately $\psi = 305^\circ$, the tight coupling approach contains a spurious oscillation and so the trimmed loose and tight coupling results are mixed.


The outboard results (Figure 38) confirm that the tight coupling solution is more prone to higher amplitude oscillations than the resulting loose coupling solution. The greater amplitude is seen to generally improve predictions around the aft portion of the rotor disk around $\psi = 180^\circ$. The comprehensive code is unable to predict the transonic dip located on the advancing side of the blade. The larger amplitudes of the tight coupling predictions more accurately approximate the experimental data with respect to the loose coupling results, especially in the third quadrant of the rotor at the blade tip (Figure 39). The comprehensive code's results indicate that the lifting line results are unable to accurately predict tip loads, being insensitive to azimuth.

Pitching moments are next examined, starting with the inboard and midspan stations (Figure 40). At $r/R = 0.4$, the CFD/CSD simulations significantly under predict the the moment at the rear of the rotor disk while the comprehensive analysis is missing higher harmonics. At the midspan station of $r/R = 0.55$ the loose coupling shows a lack of sensitivity to the events on the retreating side while the tight coupling shows a phase lag of approximately of 20° with respect to the dynamic stall event in the fourth quadrant. At the $r/R = 0.675$ location the loose coupling continues to indicate low sensitivity to the dynamic stall events while the tight coupling predicts the event in the third quadrant while spuriously predicting an event in the fourth quadrant.

The pitching moment on the outboard stations (Figure 41) indicates the tight coupling solution is generally more prone to moment stall than the loose coupling solution. The comprehensive code is once again unable to predict these stall events. The tight coupling predictions produce a complex combination of three events at $r/R = 0.865$, where the




(a) $r/R = 0.225$




(b) $r/R = 0.675$

Figure 37: Comparisons of loosely and tightly-coupled OVERFLOW/DYMORE for normal forces on UH-60A C9017 inboard and midspan airstations.




(a) $r/R = 0.865$



(b) $r/R = 0.965$

Figure 38: Comparisons of loosely and tightly-coupled OVERFLOW/DYMORE for normal forces on UH-60A C9017 outboard airstations.




(a) $r/R = 0.99$


Figure 39: Comparisons of loosely and tightly-coupled OVERFLOW/DYMORE for normal forces on UH-60A C9017 tip.

middle event is spurious. The first event, at $\psi = 270^\circ$ is captured by both solutions while the last event ($\psi = 340^\circ$) is only captured in the tight coupling solution. At $r/R = 0.965$, only one event at $\psi = 270^\circ$ is predicted by tight coupling results that loose coupling approach misses. The experimental data contains a second moment stall at approximately 330° that both solutions miss. The comprehensive codes maintains a constant azimuth solution in the pitching moment at the blade tip (Figure 42), similar to the predicted normal forces. The loose and tight coupling solutions both capture the waveform of the data, however the tight coupling results have slightly greater under prediction with respect to the experiment at $\psi = 220^\circ$.

Chord forces are next examined at the inboard and midspan stations (Figure 43). Due to its limited aerodynamics model, the comprehensive code is not able to predict the higher harmonics that the CFD/CSD simulations are able to capture due to the use of lifting line aerodynamics. At $r/R = 0.675$ the loose coupling results contains phase lag at $\psi = 180^\circ$ such that the tightly-coupled solution better captures the flight test results. At the outboard stations (Figure 44), the comprehensive code continues to significantly over predicts the




(a) $r/R = 0.225$




(b) $r/R = 0.675$

Figure 40: Comparisons of loosely and tightly-coupled OVERFLOW/DYMORE for pitching moments on UH-60A C9017 inboard and midspan airstations.




(a) $r/R = 0.865$



(b) $r/R = 0.965$

Figure 41: Comparisons of loosely and tightly-coupled OVERFLOW/DYMORE for pitching moments on UH-60A C9017 outboard airstations.




(a) $r/R = 0.99$


Figure 42: Comparisons of loosely and tightly-coupled OVERFLOW/DYMORE for pitching moments on UH-60A C9017 tip.

measured data, especially in the third and fourth quadrant where the dynamic stall occurs. The general underprediction of chord force on the advancing side is attributed to limited integration points in the experiment, as was previously discussed. The tip chord forces (Figure 45), show the comprehensive code is able to capture the general trend (unlike pitching moments at the tip) but, as expected, is unable to predict higher harmonics. The loose and tight coupling solutions follow the same trend with minor phase and amplitude differences and neither shows a clear better prediction of the experimental data.

Similar to the 8534 case, the loose and tight coupling predictions for the 9017 case at the most inboard radial sections show comparable predictions. Correlation with experiment is mixed, but it is consistent with other predictions using grids of this size [38]. The outboard stations contain significant differences in the predictions appear, in particular for the retreating side of the rotor. Phase and amplitude differences result, with loose coupling results providing more accurate correlation to experiment in some instances, while tight coupling results are more accurate at other locations. There is significant unsteadiness in the flow during the fourth quadrant, which manifests itself as a stall phenomena in the first




(a) $r/R = 0.225$




(b) $r/R = 0.675$

Figure 43: Comparisons of loosely and tightly-coupled OVERFLOW/DYMORE for chord forces on UH-60A C9017 inboard and midspan airstations.




(a) $r/R = 0.865$



(b) $r/R = 0.965$

Figure 44: Comparisons of loosely and tightly-coupled OVERFLOW/DYMORE for chord forces on UH-60A C9017 outboard airstations.



(a) $r/R = 0.99$

Figure 45: Comparisons of loosely and tightly-coupled OVERFLOW/DYMORE for chord forces on UH-60A C9017 tip.

quadrant as well. These differences are discernible through vorticity magnitude slices (e.g. Figure 46).

Next, the structural loads are examined starting from the blade root and moving towards the blade tip. Not all measurements are available at each spanwise measurement station (see Table 6 for a complete list) and a representative sample is examined. For the section closest to the root (Figure 47, two measurements are available: flap and chordwise bending moment. The flap bending moment reveals that loose and tight coupling results have similar frequency content, however the tight coupling results contain larger amplitude oscillations than the loose coupling predictions. This is most apparent on the retreating side where the integrated effects of dynamic stall are over predicted by both coupling schemes. The chord bending moment reveals that the integrated effects of dynamic stall are captured by the CFD/CSD simulations. Unfortunately, due to strain gauge malfunction [29], clipping is clearly seen in the chordwise bending moment experimental data, and a comparison in these regions cannot be made.

At $r/R = 0.3$ (Figure 48), flap bending moments and torsion moments are available for






Figure 46: Vorticity magnitudes showing the differences between loose and tight coupling at $r/R = 0.675$ for $\psi = 0^\circ, 90^\circ, 180^\circ,$ and 270° .

Table 6: List of sensors available and functional during UH-60A C9017.

r/R	Torsion Moment	Flap Bending Moment	Chord Bending Moment
0.2	No	Yes	Yes
0.3	Yes	Yes	Damaged
0.4	No	Damaged	Damaged
0.5	Yes	Damaged	Damaged
0.6	No	Yes	Yes
0.7	Yes	Yes	Yes
0.8	No	Damaged	Damaged
0.9	Yes	Damaged	No




(a) Flap Bending Moment




(b) Chord Bending Moment

Figure 47: Comparison of structural moments for the UH-60A C9017 at $r/R = 0.2$ span.




(a) Flap Bending Moment



(b) Torsion Moment

Figure 48: Comparison of structural moments for the UH-60A C9017 at $r/R = 0.3$ span.




(a) Torsion Moment

Figure 49: Comparison of structural moments for the UH-60A C9017 at $r/R = 0.5$ span.


comparison. Again, flap bending moment contains an over prediction of the flap response due to dynamic stall. The tightly-coupled simulation again has similar harmonic content to the loosely-coupled simulation, but higher amplitude oscillations. At $r/R = 0.5$ (Figure 49), only torsion moments are available. While tight and loose coupling results continue to follow the same trend, the tight coupling simulation continues to produce higher amplitudes, especially at $\psi = 30^\circ$ and 240° . At 30° , the relative oscillation of the loose coupling prediction is closer to experiment, however neither predict the correct slope. At 240° , the tight coupling more accurately captures the local waveform.

Clipping is again seen in the experimental data at a span of $r/R = 0.6$, and comparisons cannot be made. For oscillatory flap and chord bending moments, both CFD/CSD loose and tight coupling capture the general waveforms, however tight coupling continues to overshoot parts of the experiment that loose coupling does not (such as the flap bending moment at $\psi = 140^\circ$ and 200°).

The outboard station of $r/R = 0.7$ (Figure 51) has all three structural measurements (flap bending moment, chord bending moment, and torsion moment) available for comparison. The flap bending moment illustrates loose and tight coupling predictions are generally close with the large overshoot in tight coupling at $\psi = 250^\circ$ still present. The chord bending moment is dominated by clipping and few conclusions can be drawn from it. The region



(a) Flap Bending Moment



(b) Chord Bending Moment

Figure 50: Comparison of structural moments for the UH-60A C9017 at $r/R = 0.6$ span.

between approximately 220° - 260° is not clipped, and the tight coupling solution is closer in phase to the experiment than the loose coupling solution. The torsion moments all produce the same trends with minor phase differences between tight and loose coupling predictions. No data are available for $r/R = 0.8$ due to the failure of strain gauges, however torsion moments are available at the blade tip region (Figure 52). Both results of loose and tight coupling characterize the general waveform, except at 90° where neither produces the peak torsion moment, but instead predicts a local trough.


An FFT analysis of the lead-lag response confirms that, as in C8534, the lead-lag natural frequency is present in the tight coupling prediction (Figure 53). Considerably less data (3.5 revolutions) was available for the tight coupling FFT analysis than the loose coupling FFT analysis, which was able to use 100 revolutions from the comprehensive trim. This resulted in a much coarser frequency spectrum for the tight coupling approach, although it is evident that natural frequency of the lead-lag mode is present.

As the tightly-coupled solution did not develop significant aperiodicity, the cause of the differences in the predictions was next examined. By holding the controls fixed at a constant value in both loose and tight coupling, the possibility that the differences arise from the solution procedure can be examined. As seen in Figure 54, both loose and tight coupling produce nearly identical results when the controls are held fixed. This reveals that the differences in the solutions can be largely attributed to the vehicle control angles that each trim algorithm predicted.

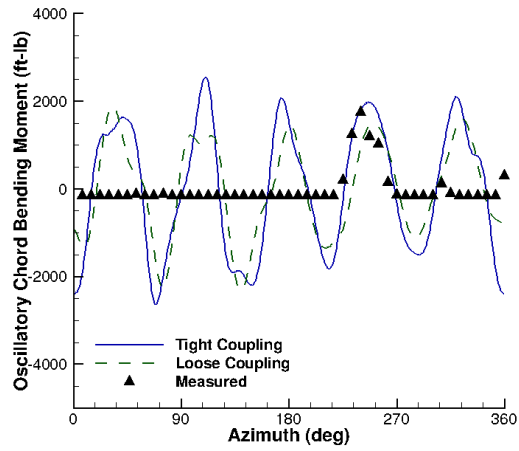
4.1.4 Kriging Estimate Applied to Loose Coupling

If a kriging trim database exists for a configuration and an engineer wanted to run loose coupling procedure instead of a tight coupling procedure, the loose coupling initial guess may be used to accelerate convergence. This idea was evaluated using the C8534 high-speed forward flight case as a test bench.


The comprehensive code (DYMORE) was initialized with its controls locked to the the same kriging-based control estimate from the tight coupling trim case. Aerodynamics were computed solely by DYMORE's simplified aerodynamics model (i.e. without delta-loads).



(a) Flap Bending Moment




(b) Chord Bending Moment



(c) Torsion Moment

Figure 51: Comparison of structural moments for the UH-60A C9017 at $r/R = 0.7$ span.



(a) Torsion Moment

Figure 52: Comparison of structural moments for the UH-60A C9017 at $r/R = 0.9$ span.




Figure 53: Comparison of loose and tight coupling lead-lag responses via an FFT analysis.

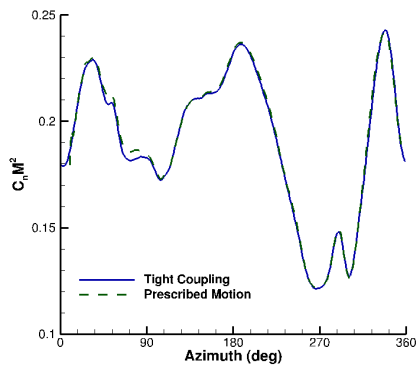


Figure 54: Comparison with fixed controls for UH-60A C9017 between loose and tight coupling results.




Figure 55: Comparison of loose coupling convergence with and without a kriging-based solution start.

The corresponding motion was then used to start the CFD simulation, and delta-loads were then applied during the autopilot trim sequence until convergence was achieved. The simulation results indicate that approximately two fewer coupling iterations were required to reach trim with this approach compared to loose coupling approach that is trimmed with the comprehensive aerodynamics model during the first iteration.

The application of the kriging control estimates to start the CFD simulation (CFD/CSD iteration 0) resulted in trimmed loads that were less accurate than initial estimates from a traditional autopilot. However, the use of the kriging during the loosely-coupled simulation resulted in more rapid convergence to the target loads than the original autopilot, as depicted in Figure 55. The larger error observed at iteration 0, which is the first CFD simulation using the initial estimates of the CSD trim without the feedback from the delta airloads, is rapidly adjusted in subsequent coupling iterations.

4.2 The HART II Test

4.2.1 Background

Higher Harmonic Control (HHC) was investigated in 1994 with the HART I program [13,14]. HART I was successful at determining that HHC could reduce rotorcraft noise which was demonstrated to be partially caused by wake related phenomenon, and merited further study. This spawned the HART II tests in October 2001, which had considerably more wake data than the HART I tests did [13,14].

Three cases were considered from the HART-II test – a baseline case and two cases with higher-harmonic control (HHC). The baseline case was used as a validation case for the HART-II rotor inputs. The two HHC cases were used to demonstrate the flexibility of the

coupling method, namely that it is not limited to traditional controls.

Higher Harmonic Control was accomplished using hydraulic actuators that were tuned to 4/rev motion. A controller was then used to maintain equal amplitude between the actuators and maintain the proper phase relation between them allowing 3, 4, or 5/rev actuation in the rotating frame, although only 3/rev was used [14].

An azimuthal sweep of the HHC phase at 30° increments was performed with the 3/rev higher harmonics, and revealed two cases of interest [14]. The first case was a minimum noise case that occurred when the phase was set to 300°. In the second configuration, the HHC produced minimum vibrations at a phasing of 180°.

4.2.2 DYMORE Standalone


The DYMORE4 input deck for the HART-II case was previously validated by Lynch [114] and was modified to encompass cases with higher harmonic control. DYMORE4 was run with very simple aerodynamics. A table lookup was employed for airfoil data, and no wake or stall models were applied. As is common with a comprehensive code, a tiploss factor (valued at 0.97) [2] was applied.

In Figure 56, DYMORE alone with its reduced order aerodynamics was able to capture the basic airload trends. However, the blade-vortex interaction (BVI) observed in the baseline case (both advancing and retreating side) was entirely missed with the simplified aerodynamic model, as is expected without the inclusion of a wake model [115].


4.2.3 CFD/CSD Loose Coupling

A set of unstructured overset grids consisting of the rotor blades and fuselage was resolved with FUN3D to compute aeroelastic responses for the HART II model. Each rotor blade had 1.2 million nodes and the combined fuselage and off-body grid contained 9.0 million. The total node count for the grids was 13.6 million nodes.


The loose coupling simulation was initialized with one revolution of prescribed motion, allowing transients to wash out of the system. This was followed by half revolution coupling iterations until convergence was achieved. All three cases were evaluated: baseline (Figure 57), minimum noise (Figure 58), and minimum vibration (Figure 59). These results indicate




(a) HART BL Cn




(b) HART BL Cm




(c) HART MN Cn



(d) HART MN Cm



(e) HART MV Cn




(f) HART MV Cm


Figure 56: Normal Force and Pitching Moments at 87% Radius computed by DYMORE alone for the HART-II baseline case.

that the largest discrepancies between experiment and predictions occur in the baseline case, where BVI is prominent.


The simulations were performed on one of a High-Performance Computing (HPC) cluster with SUSE Linux and AMD Opteron 64-bit processors. The cost was 7.33 hours per half-revolution on 128 processors.



(a) Thrust




(b) Rolling Moment




(c) Pitching Moment


Figure 57: Loose coupling trim results for the HART-II baseline case.



(a) Thrust

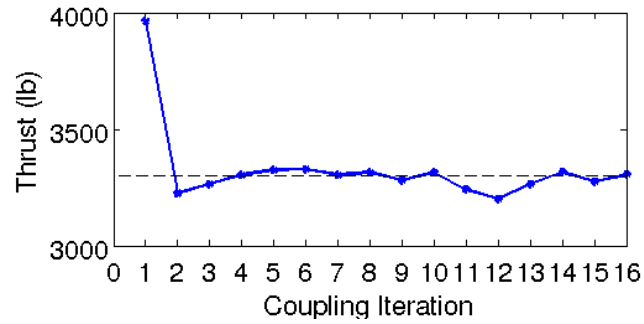


(b) Rolling Moment




(c) Pitching Moment


Figure 58: Loose coupling trim results for the HART-II minimum noise case.



(a) Thrust



(b) Rolling Moment



(c) Pitching Moment

Figure 59: Loose coupling trim results for the HART-II minimum vibration case.

4.2.4 Tightly-Coupled CFD/CSD

The baseline condition was selected to evaluate the tightly-coupled CFD/CSD technique due to the numerous blade-vortex events present. Due to the high cost of the unstructured CFD, the training points required careful selection. The initial control angles for the CSD simulation were selected to verify that the tight coupling computations within FUN3D and so this point was included in the model. The solution predicted by CSD via the autopilot was evaluated via a tightly-coupled simulation for inclusion in the observation points as part of the kriging meta-model.

Since the autopilot’s trim Jacobian roughly approximates the physical trim Jacobian, the autopilot’s computed inverse trim matrix was applied to help guide the selection of training points:

$$\vec{C}_n = \underline{\underline{J}}^{-1}(\vec{L}_{desired} - \vec{L}_0), \quad (28)$$

where $\vec{L}_{desired}$ represents the desired loads, and \vec{L}_0 are the initial loads from the tight coupling at the CSD control point. This permits a targeted and selective way of establishing the training point locations. That is, the controls selected for training are those which are expected to produce a desired loading condition. Since it is crucial to surround the target loads with observation points, loading conditions within a certain percent of target can be chosen.

For the HART-II baseline case, the desired training parameters were +/-5% of the thrust target, +/-20% of the rolling moment target, and +/-20% of the pitching moment target. A training point was also obtained at the estimated solution. Far more important than achieving the correct loads, the desired locations for these training points were selected such that they would bound the target (above and below) with confidence.

Together the training points were applied with kriging to estimate the actual solution point as a starting condition for the tightly-coupled trimming procedure. Control angles remained constant for one revolution to allow the majority of initialization transients to leave the system before control adjustments began. Trim occurred within 15 revolutions

Table 7: Comparison of Trimmed Controls

Simulation	Collective	Lateral Cyclic	Longitudinal Cyclic
Tight Coupling	3.81°	1.38°	-0.98°
Loose Coupling	3.99°	1.49°	-0.98°

(Figure 60) and the final loads were within 50 pounds of the thrust target, and within 20 ft-lbs of the roll and pitch moments. The resulting loose and tight coupling simulations trimmed control angles (Table 7) were within 0.2° of one another.




Figure 60: Time to trim in tight coupling for the baseline HART-II case.

Airloads were examined at 87% of the rotor span where experimental pressure tap measurements were obtained. Normal forces, shown in Figure 61, indicate that switching from a comprehensive code (CSD) to CFD/CSD in either loose or tight coupling offers tremendous improvements. The comprehensive code captures the general trend only without capturing




Figure 61: Comparison of the HART-II baseline normal force predictions of tight coupling, loose coupling, and comprehensive results.

high frequency content. The difference between the tight and loose coupling predictions is primarily a slight over prediction in the normal loads of the tightly-coupled CFD/CSD between approximate azimuth angles of 240°-300°. Conversely, between azimuth angles of 80°-130°, the loosely-coupled solution indicates a slight over prediction in amplitude. Outside of the advancing and retreating regions previously discussed, the normal loads are otherwise comparable.

The pitching moment comparisons (Figure 62) exhibit similar trends as with normal force predictions. The loose coupling simulation under predicts the advancing side moment between approximately 80° - 100° and the tight coupling simulation over predicts the second quadrant (135°-180°). Taken together, the differences in loose and tight coupling predictions of the normal forces and pitching moments are consistent with the slight differences in the predicted control angles at convergence.





Figure 62: Comparison of the HART-II baseline pitching moment predictions (mean removed) of tight coupling, loose coupling, and comprehensive results.


Structural moments (Figure 63) are generally comparable in loose and tight coupling, with the largest difference occurring for the oscillatory lag moment. The tightly-coupled simulation contains some aperiodicity in the structural loading. This is primarily due to transients in the tightly-coupled solution which do not have time to dissipate as they do in a loosely-coupled approach where the structural dynamics may be run for hundreds or even thousands of revolutions at little cost.

Figure 64 depicts the elastic deflections at the blade tip. Blade elastic flap is the most accurately predicted for all simulations: tightly-coupled, loosely-coupled, and comprehensive. All predict the waveform and fall primarily within the bounds of experimental measurement. On the retreating side of the rotor disk, where all simulations collectively under predict the elastic deflection, the waveform is still present, but a phase lag develops. The predictions for elastic lag deflections mirror those seen for the elastic lag moment (Figure 63b). All predictions exhibit an offset from the experimental data. Elastic twist shows a strong correlation between tight and loose coupling process, both of which accurately characterize the waveform with an offset. The comprehensive code fails to predict the high-frequency content of the twist due to its limited aerodynamic model.


As with the UH-60A C8534 case, the elastic lag is observed to contain the lead-lag natural frequency. Application of an FFT to the elastic lag motion (Figure 65) proves this. The tight coupling simulation indicates two peaks: one at the driven frequency and another at 0.79Ω . The latter frequency is consistent with the natural frequency predicted by comprehensive codes as a part of the HART-II workshop [115].



(a) Flap Bending Moment

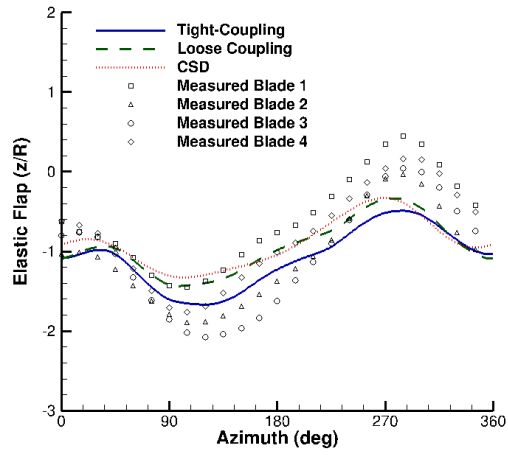


(b) Chord Bending Moment




(c) Torsion Moment

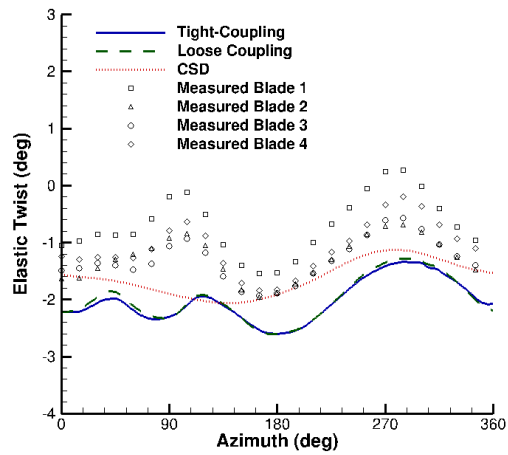
Figure 63: Comparison of the HART-II baseline oscillatory structural moments between tight coupling, loose coupling, and comprehensive results.



(a) Elastic Flap



(b) Elastic Lag



(c) Elastic Twist

Figure 64: Comparison of HART-II baseline elastic tip deformations between tight coupling, loose coupling, and comprehensive results.

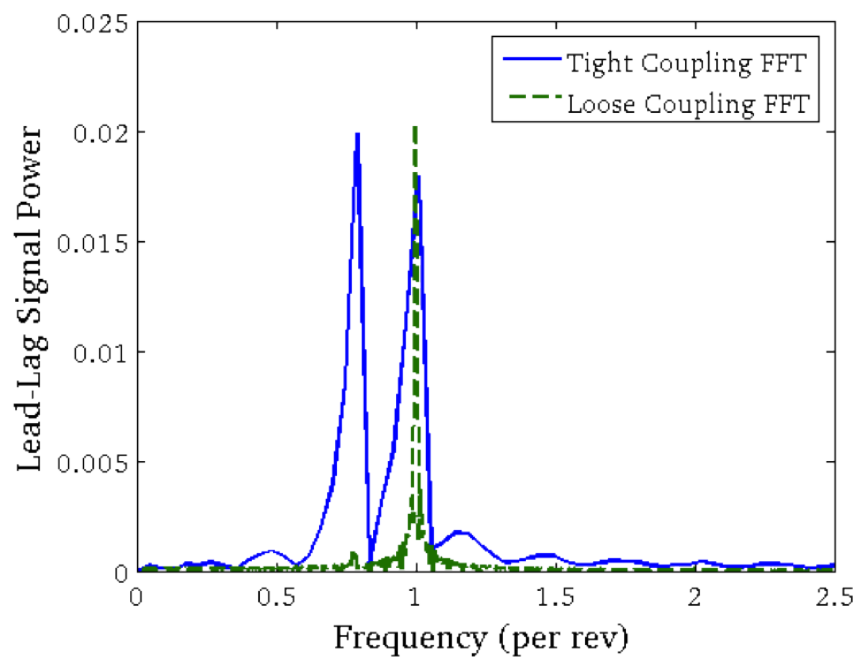


Figure 65: Comparison of HART-II baseline frequency content contained in the elastic lag motion predicted by loose and tight coupling simulations.

CHAPTER V

HYBRID CFD/CSD

This chapter delves into the results obtained using the hybrid CFD/CSD code previously discussed. The goal of this development is two-fold. First, a reduction in computational costs is sought that preserves computational accuracy. Second, a hierarchical modularity is desired to enable the understanding of causal physics.

5.1 Prescribed Motion Test

The first verification that the hybrid solver predicts correct physics utilized prescribed motion to examine the effects of the far wake via the free-wake solver on the solution. The UH-60A high-speed forward flight case C8534 was chosen for this analysis and the motion used came from the full CFD/CSD trimmed solution.

To test the communication of vorticity via the Lagrangian free-wake solver, the near body grids were coarsened for the hybrid solver. The original grids contained 4.1 million nodes with 1.9 million in the near-body grids and 2.2 million in off-body grids. The blade grids (Figure 66) had 81 airstations, which was reduced to 52 for the hybrid solver for a coarse, non-optimized demonstration case. The background grid was clipped from 5.0 rotor radii to 1.06 rotor radii and the cell size of the off-body grid was doubled. The final




Figure 66: UH-60A C8534 blade grids used in the full CFD/CSD analysis.




Figure 67: UH-60A C8534 cut-through plot of vorticity from the full CFD/CSD solution with the wake convecting down and away from the rotor.

URANS/free-wake grid node count contained 2.1 million nodes compared to the 4.1 million used in the URANS solver.

The URANS solver applied identical run options for both the URANS and the hybrid URANS/freewake simulations. The timestep was one-twentieth of a degree with no Newton sub iterations. The left-hand-side temporal terms were discretized with the ARC3D diagonalized Beam-Warming scalar penta-diagonal scheme. The right-hand-side spatial terms were discretized with an Euler central difference scheme. The Spalart-Allmaras turbulence model [116] was used with the full three-dimensional viscous terms.

Two analyses were compared with the CFD/CSD results. The CFD/CSD analysis was first examined with the cropped and coarsened grids, but with the hybrid solver turned off. This had the effect of preventing a far wake from developing. Since the wake convects away from and behind the rotor in forward flight (Figure 67), this simulation permits the examination of the wake's influence on the rotor performance. After two revolutions, the hybrid solver was initiated, enabling a comparison of the rotor performance with and without a wake.

The predicted hub loads this simulation are illustrated in Figure 68. The initial unsteadiness at the initiation of the run arise from transients followed by the development of a quasi-steady state. The transition to the hybrid solver after two revolutions is characterized by a sudden change in hub loads as the effects of the wake in forward flight impact the rotor inflow. A new quasi-steady state then forms after the transients from the free-wake diminish.




Figure 68: UH-60A C8534 hub loads of the hybrid simulation before and after the hybrid interface is initialized after two revolutions.





Figure 69: UH-60A C8534 sectional airloads at 55% radius for the reduced domain grid with and without the hybrid solver are compared with the full CFD/CSD solution.

The predicted sectional airloads (Figure 69) from the hybrid solver and the CFD/CSD solution are compared. The normal force illustrates the lack of a far wake in the cropped domain causes the load to overshoot, in particular towards the rear of the rotor disk, where the wake effects are most significant. The influence of the far wake in the normal force is observed as the loads approach those of the full domain CFD/CSD solution. The pitching moment is largely unaffected by the presence of the far wake, although close inspection of the hybrid solution indicates low-amplitude high-frequency oscillations, especially as the rotor blade approaches the rear of the rotor disk. These are due to unsteady wake effects. Examination of the chord force reveals that, similar to the normal force behavior, the cropped domain results in an over estimation of the chord force. The addition of the free-wake largely corrects this, and the prediction overlaps the full domain results. On the forward portion of the rotor disk, the cropped domain with and without the free-wake under estimate chord force. This discrepancy reveals that the chord force is not a function of the wake, but rather the coarsening of the grid. This latter point highlights why the modularity of this tool is so useful, as here it allows the engineer to easily separate grid independence from wake effects.


The region of under predicted in-plane chord forces on the forward section of the rotor disk is examined in Figure 70. The code modularity attributes the under prediction of the chord force to the reduced grid density along the span, which was decreased from 82 to 52 points. This manifests itself primarily as a phase lag most noticeably at 180° azimuth. The results are otherwise strikingly similar with only fine-scale differences visible.

Pitching moment predictions of the current work were compared to another hybrid URANS/free-wake code (GT-Hybrid). Figure 71 illustrates a typical airstation at $r/R = 0.675$. The OVERFLOW/CHARM blade grids contain fewer points along the blade section, although with the off-body grids does contain more nodes.

To obtain a quantitative comparison of the pitching moment predictions, a statistical analysis can be performed. This method was suggested by Bousman [117] and refined by Smith [118]. This method has also been used to assess other hybrid code accuracy [119,120]. In this analysis, the reference signal (e.g. the experiment) is selected and plotted on the



(a) CFD/CSD



(b) Hybrid-CFD/CSD

Figure 70: UH-60A C8534 in-plane forces viewed looking down at the forward-most section of the rotor disk shows small, fine-scale differences between the two solutions, especially just after 180° .




Figure 71: Comparison of UH-60A C8534 sectional pitching moment predictions at $r/R = 0.675$ with another hybrid code.

abscissa and a comparison signal is plotted on the ordinate. If the signals are identical, the data will scale with a slope of 1.0 and yield a perfect coefficient of determination value of $r^2 = 1.0$. As the signals differ, the slopes and coefficients of determination will vary from these ideal values.

Using the experimental data as the reference, the two hybrid methodologies can be quantitatively assessed. The results are tabulated in Table 8 and indicate that the OVERFLOW/CHARM method contains less scatter with respect to the experiment as well as a better overall trend. One possible explanation for this may be the use of overset (chimera) grids applied in this work.

Table 8: Comparison of statistical metrics for UH-60A C8534 hybrid method predictions of pitching moments.

Code	r^2	Slope
OVERFLOW/CHARM	0.91	0.71
GT-Hybrid	0.80	0.56

5.2 *Advanced Dynamics Model*

The Advanced Dynamics Model (ADM) rotor [11, 12] was experimentally evaluated for lead-lag stability parameters by the Aeroflightdynamics Directorate (AFDD) in the 1990s for a variety of flight conditions with both straight and swept blades. Several efforts [41, 42, 89] have focused on predicting the damping via simulation. These simulations have shown that accurate prediction of stability damping is possible with CFD/CSD techniques. However, this is associated with significant computational costs. The objective of the hybrid approach is to determine if the CFD/free-wake is comparable to the more costly CFD/CSD simulations.

The grids for these computations were provided by AFDD. The original CFD/CSD grids contain approximately 3.9 million points in 43 meshes. These grids have been shown by Yeo et al. [41] to compare well in terms of airloads, structural response, and stability to a more refined 29.7 million node grid. Each blade is comprised of three meshes: a root cap, the main blade (with or without a swept tip), and the tip cap. The main blade contains 133 points in the chordwise direction and 108 spanwise points (or 107 airstations). The




Figure 72: Comparison of the ADM domains resolved with the CFD/CSD solver and the URANS hybrid CFD/CSD analysis illustrates the reduction in grid size used with the hybrid approach.

remaining grids are off-body grids that account for 46.8% of the total grid points. The comprehensive solver was setup with airstations coincident to the CFD grid airstations to remove interpolation error.

The grids have been modified for application with the hybrid solver. The farfield background grid extent was reduced from 5.0 rotor radii to 2.2 rotor radii. The off-body grid cell size was increased by a factor of two. These modifications reduced the original off-body grid count from 1.7 million nodes to 0.4 million, as illustrated in Figure 72. The near-body blade grids were also coarsened in the radial direction from 107 spanwise cells to 60. This further reduced the near-body grids cell count by 77% from 1.7 million nodes to 1.27 million nodes (or 317K nodes per blade). These grid modifications permitted a computational cost reduction from approximately 58.9 cpu-hours/rev to 24.8 cpu-hours/rev on a linux cluster consisting of dual quad-core Intel processors with 15 GB/node. This represents a reduction in runtime of 57.9%. These grids were not optimized for the hybrid approach, modifications such as grid adaption [25] in the near-body grids will result in further cost savings.

The URANS module applied quarter-degree timesteps and six Newton subiterations per timestep. The temporal terms were discretized with the ARC3D diagonalized Beam-Warming scalar penta-diagonal scheme and spatial terms were discretized using Euler central differences. The Spalart Allmaras turbulence model was invoked to compare with prior simulations computed by Morillo et al. [89].


The free-wake code applied a maximum of 14 vortex trailers in each of the circulation zones with 1° timesteps. Three circulation zones representing inboard and positive circulation, outboard and negative circulation, and a varying strength vortex pair for inboard positive circulation in the farfield were assigned, as recommended in the CHARM manual [100]. The vortex trailer core sizes were limited for the inboard and outboard zones to a maximum of 100% of the blade radius. The farfield vortex trailer core sizes were set to 10% of the blade radius.

5.2.1 Loose Coupling Results


A comprehensive solution was obtained with 1° timesteps using loose coupling simulation. In the tight coupling simulations the timestep was reduced to match the CFD code's quarter degree timesteps. The structural model did not include structural damping to avoid estimating an unknown quantity. Since the intent is to compare the different aerodynamic predictions between the URANS and hybrid approaches, the lack of structural damping aided in the hierarchical comparison since all damping in the simulations are due to aerodynamic predictions.

The hybrid solver was then executed in loose coupling mode to trim the rotor. Once trim was achieved, the rotor structural loads between the full CFD/CSD and the hybrid CFD/CSD predictions were compared at an advance ratio of $\mu = 0.30$. In Figure 73, the various levels of solution hierarchy are evaluated at an inboard radial span location of $r/R = 0.12$.


The comprehensive (CSD) code is unable to capture the higher frequency oscillations in the flap bending moment that the CFD-based solver predicted. The CSD result has a mean value 29% higher than the CFD/CSD. The CFD/CSD predictions have similar harmonic



(a) Flap bending moment




(b) Chord Bending moment




(c) Torsion moment

Figure 73: Comparison of the ADM structural moments with the hybrid solver at $r/R=0.12$ for a 0.35 advance ratio with means removed.



(a) CSD vs Full-CFD/CSD



(b) Hybrid-CFD/CSD vs Full-CFD/CSD

Figure 74: Comparison of statistical relations of the ADM flap bending moments for the CSD and Hybrid-CFD/CSD predictions versus the full-CFD/CSD method for the time period shown in Fig. 73a).

content and mean values as the hybrid CFD/CSD. The latter predicts 7.2% higher peak-to-peak amplitudes than the full CFD/CSD simulations. The chord bending moment is well captured by all three solvers, although the CSD prediction has a significantly different mean, with a 25% offset from the full CFD/CSD prediction. The torsion moment results indicate that the CSD code misses the harmonic content of the full CFD/CSD results, while the hybrid code captures the frequency but includes 0.15 ft-lb amplitude overshoots.

As before, a statistical analysis of the results can be performed to quantitatively assess the performance of the various approaches. Using the full CFD/CSD results as the reference, the comprehensive and the hybrid CFD/CSD results can be quantitatively assessed. An example plot as applied to the flap bending moment is illustrated in Figure 74.

Examination of the resulting statistical fit (Table 9) reveals that the hybrid code results generally offer significant improvements to the comprehensive code. With respect to the chord bending moment, the CSD result has slightly less scatter, but the hybrid result indicates a more accurate slope. The torsion moment slope is more accurately predicted by the CSD code, although the hybrid code had significantly less scatter with respect to the full CFD/CSD solution, indicating improved phase angle predictions (as observed temporally in Figure 73).

Table 9: Comparison of statistical metrics of the reduced models compared to the full-CFD/CSD method for the ADM rotor at $r/R=0.12$ for a 0.35 advance ratio.

Moment	CSD		Hybrid-CFD/CSD	
	Slope	r^2	Slope	r^2
Flapwise	0.47	0.25	0.95	0.87
Chordwise	0.82	0.98	0.88	0.93
Torsion	0.79	0.59	1.31	0.86

5.2.2 Tight Coupling Stability Results

The stability analysis was performed using an identical procedure to that outlined in Morillo et al [89]. The solution was first trimmed with a loose coupling approach, then switched to tight coupling simulation. Once the system was verified to be at a quasi-steady state, a perturbation to the structural model was applied via a tip force of 10.0 lb in the lead-lag direction. The tip force was applied over one entire revolution and then removed so that the response could be recorded. At least eight additional revolutions of the response were performed to collect the data needed for analysis. The first revolution after the tip force was applied was discarded from the analysis due to transient response. The final revolutions were then evaluated with a Prony analysis to extract the modal parameters.

To maintain consistency with the analysis by Morillo et al [89], the response signal in the Prony analysis was the chord bending moment on each of the four blades. The damping estimates for each blade were then averaged. While DYMORE’s advanced analysis capabilities allow for multiple signals to be used together to improve the modal decomposition process, this feature was not applied to ensure a valid comparison.

Figure 75 illustrates the four-bladed response for the high-speed advance ratio of $\mu =$




Figure 75: The chord bending moment signals from all four blades, showing the initial excitement with a vertical line depicting the excitement cut-off and following response for $\mu = 0.35$.

0.35. The first revolution (approximately 0.04 seconds) is a continuous quasi-steady signal. As soon as the perturbation has been applied the signal amplitude experiences a rapid increase in amplitude. Subsequently, a slow decay in the signal amplitude occurs indicating that this condition experiences low aerodynamic damping.

The modal estimation procedure was performed at several operating points with the hybrid CFD/CSD solvers and then compared to full CFD/CSD solution (Figure 76). As the structural model contained no structural damping, the objective was to match the trend of the structural damping and not the actual values. The results indicate that the comprehensive code predicts a flat trend and is unable to follow experiment. The hybrid URANS/free-wake solver follows the trend established by the CFD/CSD simulation, although advance ratios $\mu = 0.2$ and 0.25 were more closely predicted by the comprehensive code.

5.3 UH-60A Dynamic Stall Wind Tunnel Test

NASA recently tested a UH-60A rotor [17] in the National Full-Scale Aerodynamics Complex (NFAC) 40 by 80 foot wind tunnel. The objective was to evaluate a challenging flight condition, so a high-thrust condition with dynamic stall was selected (case 4540). Here




Figure 76: Comparison of the ADM damping predictions for three levels of aerodynamics for the comprehensive, hybrid CFD/CSD, and full CFD/CSD analyses.


the advance ratio is 0.299 and the thrust coefficient is near its maximum of 0.010. A dynamic stall event is observed to occur on the retreating outboard side of the rotor disk (e.g. $slfracrR = 0.92$ and $\psi = 300^\circ$).

The hybrid approach was compared directly to the standard full CFD/CSD methodology. To eliminate near-body grid density effects, the near-body blade grids were identical in both cases. Each blade contained 99 airstations, and the total near-body node count for all four blades was 12.0 million.

The off-body grids were cropped from five rotor radii to just 1.2 rotor radii, and the resulting domain is shown in Figure 77. The domain cropping had the effect of reducing the off-body node count from 10.8 million to 8.9 million. Unlike the prior cases that have been evaluated, the finest level of off-body grid spacing was kept constant.

The URANS solver applied quarter-degree timesteps with ten Newton subiterations. Two turbulence models were examined: Spalart-Allmaras and the $k\omega$ -SST turbulence model. Both used only thin-layer viscous terms in the normal direction rather than the full three-dimensional viscous terms. The temporal terms were discretized with the ARC3D diagonalized Beam-Warming scalar penta-diagonal scheme, while the spatial terms were discretized

OVERFLOW/CHARM



OVERFLOW




Figure 77: Comparison of the UH-60A C4540 original domain and the hybrid analysis domain.

with the Roe upwind scheme.

Both the full CFD/CSD simulation and the hybrid simulation were independently trimmed with loose coupling. The resulting trimmed control angles from the hybrid analysis were within half a degree of the full domain analysis with both turbulence models, tabulated in Tables 10 and 11.

Table 10: Comparison of controls as predicted by full CFD/CSD and hybrid CFD/CSD with the Spalart-Allmaras turbulence model.


Control (deg)	Full-CFD/CSD	Hybrid-CFD/CSD
Collective	7.52	7.23
Lateral Cyclic	-0.02	-0.49
Longitudinal Cyclic	-4.64	-4.45

Table 11: Comparison of controls as predicted by full CFD/CSD and hybrid CFD/CSD with the $k\omega$ -SST turbulence model.


Control (deg)	Full-CFD/CSD	Hybrid-CFD/CSD
Collective	7.70	7.34
Lateral Cyclic	-0.01	-0.46
Longitudinal Cyclic	-4.87	-4.53

The hybrid CFD/CSD airloads were compared with the CFD/CSD airloads and experiment (Figures 78 and 79). The normal force predictions for both the hybrid and the full domain CFD/CSD analyses capture the dynamic stall event that occurs at approximately 300° but then predict a second, non-physical dynamic stall event at approximately 330° . This is attributed to the inadequacies of the turbulence models (especially the thin-layer approximation). While the results are not optimal for correlation with the wind tunnel results, the goal has been met. The hybrid model accurately reproduces the results of the full domain CFD/CSD simulation. The sectional pitching moment of the outboard station contains a significant amount of high-frequency content in the hybrid waveform. This is attributed to the unsteady wake structure, and is a common feature in many hybrid codes.

Oscillatory structural loads were examined for both turbulence models (Figures 80 and 81). As is typical of CFD/CSD methods, the flap bending moments are well captured, as they are driven by the blade normal loading which itself is well captured. The hybrid and full domain methodologies are in close agreement with each other for both turbulence




(a) Normal force




(b) Pitching moment

Figure 78: Comparison of the UH-60A wind tunnel C4540 aerodynamic loads at 92% span as computed with Spalart-Allmaras.



(a) Normal force



(b) Pitching moment


Figure 79: Comparison of the UH-60A wind tunnel C4540 aerodynamic loads at 92% span as computed with $k\omega$ -SST turbulence model.

models. The results for the chord bending moments miss the dip at approximately 240° , but predictions from both hybrid and full domain methodologies are a close match with each other using either turbulence model. The torsion moment predictions do not fully characterize the waveform of the experiment, in particular towards the rear of the rotor disk. The hybrid and full domain CFD/CSD solutions both follow the same mean trend, but the hybrid approach introduces small amplitude, high-frequency oscillations toward the rear of the rotor disk. These oscillations are associated with the high-frequency content in the free-wake module previously discussed for the pitching moments.


As with the ADM case, a statistical analysis was performed to assess the accuracy of the two approaches with respect to the experiment and to one another. Table 12 tabulates statistical results for the data in Figures 80 and 78. The results indicate that the hybrid CFD/CSD predictions are comparable statistically to the full domain CFD/CSD results. The pitching moment coefficient predictions are also comparable, despite the high-frequency content in the signal. The only quantity where the hybrid approach did not result in an equal or better fit to the experimentally measured values was in the flapwise bending moment, where the results were comparable (within 2% of one another).

Table 13 repeats this statistical analysis, comparing hybrid CFD/CSD predictions to the full domain CFD/CSD results. The statistical fit slopes are in close agreement, averaging less than 5% difference. Chordwise bending moment predictions contain the greatest discrepancy in slope with 10% difference. This generally indicates that the hybrid results closely follow the mean trend of the full domain signal. The coefficients of determination reveal that there is very low scatter with respect to the full domain signal except for the pitching moment coefficient comparisons, due to the high-frequency free-wake content in the hybrid pitching moment results.


An examination of literature [121] from hybrid CFD/freewake solutions reveals that these methods are prone to high-frequency content in the pitching moment predictions that do not appear in traditional CFD/CSD solvers. Using the hierarchical modularity, the high-frequency content was found to be directly attributed to the free-wake unsteadiness. This acts as a temporary solution to a problem that should be further pursued with future



(a) Flap bending moment




(b) Chord bending moment




(c) Torsion moment


Figure 80: Comparison of the UH-60A wind tunnel C4540 oscillatory structural moments at 40% span for the Spalart-Allmaras turbulence model.



(a) Flap bending moment



(b) Chord bending moment



(c) Torsion moment

Figure 81: Comparison of the UH-60A wind tunnel C4540 oscillatory structural moments at 40% span for the $k\omega$ -SST turbulence model.

Table 12: Comparison of statistics of the different methods compared with the experimental data for the C4540 wind tunnel test case.

Analysis	Full-CFD/CSD		Hybrid-CFD/CSD	
	Slope	r^2	Slope	r^2
$C_n M^2$	0.88	0.86	0.92	0.87
$C_m M^2$	0.46	0.36	0.51	0.41
Flapwise	0.88	0.94	0.90	0.92
Chordwise	0.39	0.67	0.43	0.67
Torsion	0.47	0.41	0.49	0.42

Table 13: Comparison of statistics of the hybrid-CFD/CSD data compared with the full-CFD/CSD data for the C4540 wind tunnel test case.

Analysis	Hybrid-CFD/CSD	
	Slope	r^2
$C_n M^2$	1.03	0.98
$C_m M^2$	0.95	0.82
Flapwise	1.03	0.99
Chordwise	1.10	0.99
Torsion	1.03	0.98

free-wake developments.

As the wake is aperiodic, it is possible to ensemble average multiple solutions together to effectively filter out the unsteady effects of the wake. As seen in Figure 83 this technique aids in further recovery of the full domain CFD/CSD results.


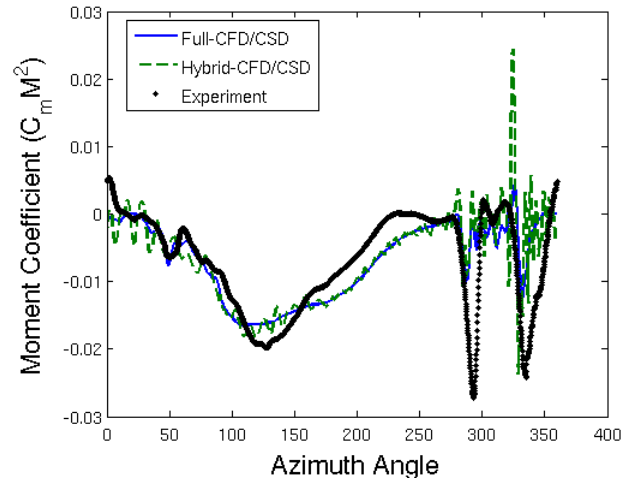



Figure 82: Comparison of pitching moments for UH-60A wind tunnel C4540 solved with the $k\omega$ -SST turbulence model for the last four coupling iterations.



(a) Original



(b) Ensemble Averaged

Figure 83: Comparison of with and without ensemble averaging for the UH-60A wind tunnel C4540 aerodynamic pitching moments at 92% span as computed with $k\omega$ -SST.

CHAPTER VI

CONCLUSIONS

This thesis has explored advanced aeroelastic methods for rotor blades. Chief contributions are listed here.

- A practical trim algorithm suitable for tightly-coupled CFD/CSD simulations has been developed and evaluated on a variety of quasi-steady level flight conditions that include transonic advancing blades, dynamic stall, and blade vortex interaction. This is an especially useful development for future tightly-coupled maneuvering simulations, which are currently not actively trimmed.
- A numerical experiment was conducted where the differences in loose and tight coupling predictions for quasi-steady level flight conditions were examined. While airloads were found to be comparable if the trim control angles are identical, the lead-lag deflections were observed to contain a natural frequency component, causing small differences in deflections.
- A modular hybrid URANS/free-wake aeromechanics methodology that maintains solution integrity in the near body was developed and investigated in high-speed forward flight, dynamic stall, and stability analyses. The methodology offers significant computational cost savings of greater than 50% on non-optimized grids. The preservation of solution fidelity was confirmed using statistical analyses. Pitching moments were seen to offer improvements over existing methodologies.

This work significantly improved on a previous implementation of a kriging-based tight coupling trim algorithm. Specific improvements regarding the tight coupling trim algorithm over the original include:

- Instabilities were removed by no longer forcing a fixed finite difference applied to a 5% change the current hub loads. This imposed strict requirements of the observation

point locations for stability. These requirements were difficult to satisfy due to the inverse problem. Instead, the new algorithm computes trim sensitivities from nearby observation points via a directional decomposition between points.

- The original implementation attempted to solve an inverse problem, where the controls were treated as a function of the loads. The new algorithm fixed this by recasting the kriging metamodel so that loads were a function of controls.
- The improved algorithm estimates the most likely control angle by applying a global search that includes the predicted variance in its evaluation. This leads to an estimate that is globally the most likely solution. The original algorithm was unable to do this with the inverse loads/controls problem.
- The improved algorithm computes the inverse trim Jacobian directly and accurately from the Jacobian, leading to faster trim. The original implementation attempted to apply an approximation necessary with the inverse problem.

Specific finds from the tight coupling results include:

- A tightly-coupled CFD/CSD trim algorithm has been developed and optimized to keep trim costs in parity with a loose coupling algorithm. The application of a kriging metamodel to accurately estimate control angles combined with a stable algorithm resulted in trim times that were significantly reduced from that of a traditional autopilot approach. The new trim algorithm is more robust than the original, which failed for some configurations.
- Trimmed controls were comparable to previously reported results [27] with an Euler solver in unstalled conditions. Under these flight conditions there is a strong sensitivity of thrust to collective input reducing the complexity of trim. For unstalled conditions, control angle differences were within 0.2° for loose and tight coupling predictions. In conditions with dynamic stall there is a nonlinear relation between stall-limited thrust and collective, and the resulting collective controls differed by nearly 1.3° .

- Integrated structural moments predictions from loose and tight coupling simulations exhibit larger disparities than the aerodynamic load predictions. This is due to a combination of a lower number of revolutions run within the structural model in tight coupling simulation, as well as differences in the lead-lag motion.
- The use of a kriging meta-model was examined in a loose coupling algorithm by predicting the initial control angles for the first coupling iteration. This was found to speed up convergence by approximately 40%.
- A refined grid with over 15 times the grid node count has been shown to improve phase predictions by 4° and peak-to-peak amplitudes of aerodynamic loading by 7.8% in the loose coupling results.

From the hybrid CFD/CSD analyses:

- A parallel carefree hybrid CFD/CSD free-wake methodology has been developed and validated for multiple processors. The hybrid methodology has been demonstrated to maintain solution accuracy at less than 50% of the computational cost.
- The hybrid CFD/CSD methodology is especially useful in situations where stability analyses are required. It was demonstrated that cost can be reduced by greater than a factor of two while preserving the same lead-lag damping trends that the full CFD/CSD predicted.
- A statistical analysis of hybrid results indicates that aerodynamic predictions are comparable to a full CFD/CSD simulation. When using the same near-body blade grids, predicted structural moment trends were within 10% of full CFD/CSD computations with low scatter ($r^2 = 0.98$ or better) while the controls were within 0.5° of the full CFD/CSD predictions.
- Aperiodic high-frequency noise in the aerodynamic pitching moments has been found to come from some unsteady wake structures predicted within the free-wake model integrated into the methodology. It is possible to reduce this high-frequency noise by

applying an ensemble average over the converged iterations, which is shown to reduce scatter.

CHAPTER VII

FUTURE WORK

This chapter focuses on recommendations for future work.

7.1 Tight Coupling

There are several recommendations for tightly-coupled CFD/CSD analyses. The effects of a higher-order coupling scheme should be examined. This would ensure that high-order accurate algorithms are not reduced to first-order accuracy at the coupling interface. These algorithms could potentially enable larger time step sizes, reducing computational costs. A one- or two-degree-of-freedom airfoil (e.g. with pitch and plunge motion bound by springs) would make an ideal testbench to assess the improvements available from such methods before application to rotorcraft vehicles. Stall flutter, a self-excited phenomenon whereby dynamic stall couples with structural forces [122, 123], is particularly well suited for examination with this approach, as it applies to two-dimensional airfoils, and so a full three-dimensional model is not required.

Investigation of a full maneuver with active trim would greatly benefit the rotorcraft community. Current studies of maneuvers are limited to prescribed blade and hub motion from flight-test data and sometimes the motion is averaged to be able to simulate a loosely-coupled maneuver. As maneuvers typically define the high stress loads on a vehicle that are not well captured by current methods, improvements in accuracy could eventually lead to better vehicle designs through reduced structural margins.

The cause of rotor wake aperiodicity seen in some experiments, especially hover, is uniquely suited for tightly-coupled analyses as loosely-coupled analyses are inappropriate due to their assumption of periodicity. Computations with high grid density in the rotor wake, which can be made computationally affordable using adaptive grid refinement rather than global refinement, can be applied to accurately track the location of the wake.

The inverse trim Jacobian employed by the trim algorithm in tight coupling could be

used with advanced trim algorithms. Such algorithms can be optimized for stability or performance [124]. Rotorcraft simulations have utilized these advanced controls to demonstrate carefree maneuvers throughout the flight envelope [125] and could be adapted for CFD/CSD simulations.

7.2 Hybrid CFD/Free-Wake

Further improvements to the hybrid methodology may be possible by utilizing grid adaption techniques [25] to create more optimal grids to predict performance with hybrid CFD/free-wake solvers. In this manner, a coarse grid can be refined to accurately capture the near-body physics, leaving other areas unrefined. Advancements in transition and turbulence modeling [126] also promise to more accurately resolve the underlying governing equations.

Grid density can be significantly reduced by removal of the overset background grid. Each blade would be directly influenced by the rotor wake without passing the data to the overset process. The large reduction in number of grid nodes would further decrease the cost of the hybrid analysis. Challenges to this method are potential numerical reflections off the near-body blade grids due to the close proximity of the blade's far-field boundary.

The hybrid methodology should be evaluated in a hover condition, where the wake dynamics may be unsteady [127]. Hover offers a challenging condition to further examine the hybrid versus full CFD/CSD methods. If the hybrid approach maintains an accurate solution in these cases, a descent case into a ring-vortex state would further detail the limits of hybrid methodologies.

Acknowledgements

First and foremost, the author would like to acknowledge his advisor, Professor Marilyn J. Smith for all of her help and guidance. While it goes without saying that without her this thesis would have not been possible, it is no understatement to say that her guidance enabled the author to develop the skills and knowledge necessary to pursue graduate studies themselves. The author is also very grateful to the committee members for taking the time to serve on his Ph.D. committee.

The efforts presented here were partially funded by several projects:

- The Vertical Lift Consortium (VLC), formerly the Center for Rotorcraft Innovation and the National Rotorcraft Technology Center (NRTC) under Technology Investment Agreement WBS 2011-B-11-03.1-P3 entitled National Rotorcraft Technology Center Research Program.
- NRA Grant NNX07AP43A, “Innovative Strategies for Rotary-Wing Coupled Aeroelastic Simulations” (Technical Monitor, Elizabeth Lee-Rausch).

The author would like to acknowledge that this research and development was accomplished with the support and guidance of the NRTC and VLC. The views and conclusions contained in this document are those of the author and should not be interpreted as representing the official policies, either expressed or implied, of the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation thereon. Computing resources were provided in part by the NASA High-End Computing (HEC) Program, NASA Advanced Supercomputing (NAS) Division at Ames Research Center, as well as the DOD High Performance Computing (HPC) program.

The author was fortunate to spend a summer working at Bell Helicopter where he was able to apply his research to some very interesting aeromechanics problems. Jorge Morillo was instrumental in that and, along with Loren Ahaus, has been very supportive of the research performed here at Georgia Tech.

The author's laboratory colleagues all deserve a round of appreciation. Rajiv Shenoy, who helped with the intricacies of FUN3D and from whom the author shamelessly stole his PBS filenaming structure (humorously referred to as "rajiv.run"), stands out as both a personal friend and a good colleague. Eliot Quon, for frequent discussions of hybrid techniques. Eric Lynch, for being an incredible role model. Dan Prosser, for discussions regarding rotor aeromechanics. The entire lab – for being a wonderful group to work with. Finally, the author is extremely grateful for the support of his friends and family.

REFERENCES

- [1] JOHNSON, W., *Rotorcraft Aeromechanics*. Cambridge University Press, 2013.
- [2] LEISHMAN, J. G., *Principles of Helicopter Aerodynamics*. New York, NY: Cambridge University Press, 2006.
- [3] HODGES, D. H. and PIERCE, G. A., *Introduction to Structural Dynamics and Aeroelasticity*. Cambridge University Press, 2002.
- [4] JOHNSON, W., “A History of Rotorcraft Comprehensive Analyses,” in *69th Annual Forum of the American Helicopter Society*, (Phoenix, Arizona), May 2013.
- [5] TUNG, C., CARADONNA, F. X., BOXWELL, D. A., and JOHNSON, W. R., “The Prediction of Transonic Flows on Advancing Rotors,” in *40th Annual Forum of the American Helicopter Society*, (Reston, VA), May 1984.
- [6] BRIDGEMAN, J. O., STRAWN, R. C., and CARADONNA, F. X., “An Entropy and Viscosity Corrected Potential Method for Rotor Performance Prediction,” in *Proceedings of the American Helicopter Society 44th Annual Forum*, (Washington, D.C.), June 16–18 1988.
- [7] SMITH, M. J., *A Fourth-Order Euler/Navier-Stokes Prediction Method for the Aerodynamics and Aeroelasticity of Hovering Rotor Blades*. PhD thesis, Georgia Institute of Technology, Atlanta, Georgia, 1994.
- [8] BAUCHAU, O. A. and AHMAD, J. U., “Advanced CFD and CSD Methods for Multidisciplinary Applications in Rotorcraft Problems,” in *6th AIAA, NASA, and ISSMO Symposium on Multidisciplinary Analysis and Optimization*, no. AIAA-1996-4151, (Bellevue, WA), American Institute of Aeronautics and Astronautics, September 1996.
- [9] KUFELD, R.M. AND BALOUGH, D.L. AND CROSS, J.L. AND STUDEBAKER, K.F., AND JENNISON, C.D. AND BOUSMAN, W.G., “Flight Testing of the UH-60A Airloads Aircraft,” in *Proceedings of the American Helicopter Society 50th Annual Forum*, (Washington, DC), May 11–13 1994.
- [10] BOUSMAN, W. G. and KUFELD, R. M., “UH-60A Airloads Catalog,” Tech. Rep. TM-2005-212827, NASA, 2005.
- [11] MAIER, T., SHARP, D., and LIM, J., “Fundamental Investigation of Hingless Rotor Aeroelastic Stability, Test Data and Correlation,” in *Proceedings of the American Helicopter Society 51st Annual Forum*, (Fort Worth, TX), May 9-11 1995.
- [12] MAIER, T., SHARP, D., and ABREGO, A. I., “Aeroelastic Stability for Straight and Swept-Tip Rotor Blades in Hover and Forward Flight,” in *Proceedings of the American Helicopter Society 51st Annual Forum*, (Montreal, Canada), May 25-27 1999.

- [13] VAN DER WALL, B. G., BURLEY, C. L., YU, Y., RICHARD, H., PENGEL, K., and BEAUMIER, P., “The HART II Test - Measurement of Helicopter Rotor Wakes,” *Aerospace Science and Technology*, vol. 8, no. 4, pp. 273 – 284, 2004.
- [14] VAN DER WALL, B., JUNKER, B., BURLEY, C., BROOKS, T., Y., Y., TUNG, C., RAFFEL, M., RICHARD, H., WAGNER, W., MERCKER, E., PENGEL, K., HOLTHUSEN, H., BEAUMIER, P., and BELRIEUX, Y., “The HART II Test in LFF of the DNW - a Major Step Towards Rotor Wake Understanding,” in *Proceedings of the 28th European Rotorcraft Forum, Bristol, England*, September 2002.
- [15] VAN DER WALL, B. G., LIM, J. W., SMITH, M. J., JUNG, S., BAILLY, J., BAEDER, J., and BOYD, D.D., J., “The HART II International Workshop: An Assessment of the State of the Art in Comprehensive Code Prediction,” *CEAS Aeronautical Journal*, vol. 4, pp. 223–252, September 2013. DOI 10.1007/s13272-013-0077-9.
- [16] SMITH, M. J., LIM, J. W., VAN DER WALL, B. G., BAEDER, J., BIEDRON, R. T., BOYD, D.D., J., JAYARAMAN, B., JUNG, S., and MIN, B.-Y., “The HART II International Workshop: An Assessment of the State-of-the-Art in CFD/CSD Prediction,” *CEAS The Aeronautical Journal*, vol. 4, pp. 345–372, December 2013. DOI 10.1007/s13272-013-0078-8.
- [17] NORMAN, T. R., SHINODA, P., PETERSON, R., and DATTA, A., “Full-Scale Wind Tunnel Test of The UH-60A Airloads Rotor,” in *Proceedings of the American Helicopter Society*, (Virginia Beach, VA), May 3-5 2011.
- [18] POTSDAM, M., DATTA, A., and JAYARAMAN, B., “Computational Investigation and Fundamental Understanding of a Slowed UH-60A Rotor at High Advance Ratios,” in *Proceedings of the American Helicopter Society 64th Annual Forum*, (Fort Worth, Texas), American Helicopter Society, May 1–3 2012.
- [19] KOTTAPALLI, S., “Performance and Loads Correlation of a UH-60A Slowed Rotor at High Advance Ratios,” in *American Helicopter Society Vertical Lift Aircraft Design Conference*, (San Francisco), January 2012.
- [20] YEO, H., “Investigation of UH-60A Rotor Performance and Loads at High Advance Ratios,” *Journal of Aircraft*, vol. 50, pp. 576 – 589, March–April 2013.
- [21] STRAWN, R. C., CARADONNA, F. X., and DUQUE, E. P. N., “30 Years of Rotorcraft Computational Fluid Dynamics Research and Development,” *Journal of the American Helicopter Society*, vol. 51, pp. 5–21, January 2006.
- [22] DATTA, A., NIXON, M., and CHOPRA, I., “Review of Rotor Loads Prediction With the Emergence of Rotorcraft CFD,” in *Proceedings of the 31st European Rotorcraft Forum*, (Florence, Italy), September 2005.
- [23] ABRAS, J. N. AND LYNCH, C. E. AND SMITH, M. J., “Advances in Rotorcraft Simulations with Unstructured CFD,” in *Proceedings of the American Helicopter Society 63rd Annual Forum*, (Virginia Beach, Virginia), May 1–3 2007.
- [24] BIEDRON, R. T. AND LEE-RAUSCH, E. M., “Rotor Airloads Prediction Using Unstructured Meshes and Loose CFD/CSD Coupling,” in *26th AIAA Applied Aerodynamics Conference*, no. AIAA-2008-7341, (Honolulu, HI), Aug. 2008.

- [25] SHENOY, R., SMITH, M. J., and PARK, M. A., “Unstructured Overset Mesh Adaptation with Turbulence Modeling for Unsteady Aerodynamic Interactions,” *AIAA Journal of Aircraft*, on line. doi: 10.2514/1.C032195.
- [26] WISSINK, A. M., KAMKAR, S., PULLIAM, T. H., SITARAMAN, J., and SANKARAN, V., “Cartesian Adaptive Mesh Refinement for Rotorcraft Wake Resolution,” in *AIAA 28th Applied Aerodynamics Conference*, no. AIAA-2010-4554, (Chicago, Illinois), June 28–July 1, 2010.
- [27] ALTMIKUS, A., WAGNER, S., BEAUMIER, P., and SERVERA, G., “A Comparison: Weak Versus Strong Modular Coupling for Trimmed Aeroelastic Rotor Simulations,” in *American Helicopter Society 58th Annual Forum*, (Montreal, Canada), June 11–13 2002.
- [28] POMIN, H. and WAGNER, S., “Aeroelastic Analysis of Helicopter Rotor Blades on Deformable Chimera Grids,” *Journal of Aircraft*, vol. 41, no. 3, pp. 577–584, 2004.
- [29] HO, J., YEO, H., and R.A., O., “Investigation of Rotor Blade Structural Dynamics and Modeling Based on Measured Airloads,” in *63rd Annual Forum of the American Helicopter Society*, (Virginia Beach, VA), May 2007.
- [30] KANG, H. and HASBUN, M., “Hydraulic Damper Model for Rotorcraft Comprehensive Analysis,” in *5th Decennial American Helicopter Society Aeromechanics’ Specialists’ Conference*, (San Francisco, California), January 22–24 2014.
- [31] BURGESS, N. K. and JAIN, R., “Effects of Numerical Methods on Static and Dynamic Stall Simulations,” in *Proceedings of the American Helicopter Society International 5th Decennial Specialists’ Conference on Aeromechanics*, (San Francisco, CA), January 22–24 2014.
- [32] KUFELD, R. M. and BOUSMAN, W. G., “High Load Conditions Measured on a UH-60A in Maneuvering Flight,” *Journal of the American Helicopter Society*, vol. 43, no. 3, pp. 202–211, 1998.
- [33] BHAGWAT, M., ORMISTON, R., SABERI, H., and XIN, H., “Application of CFD/CSD Coupling for Analysis of Rotorcraft Airloads and Blade Loads in Maneuvering Flight,” in *Proceedings of the American Helicopter Society 63rd Annual Forum*, no. 3, p. 2048, American Helicopter Society, Inc, May 1–3 2007.
- [34] RAMASAMY, M., PAETZEL, R., and BHAGWAT, M., “Aperiodicity Correction for Rotor Tip Vortex Measurements,” in *Proceedings of the American Helicopter Society 67th Annual Forum*, vol. 67, American Helicopter Society, 2011.
- [35] CARADONNA, F., LAUTENSCHLAGER, J., and SILVA, M., “An Experimental Study of Rotor-Vortex Interactions,” in *AIAA 26th Aerospace Sciences Meeting*, no. AIAA-88-0045, (Reno, Nevada), January 11–14 1988.
- [36] CARADONNA, F., STRAWN, R., and BRIDGEMAN, J., “An Experimental And Computational Study Of Rotor-Vortex Interactions,” in *Proceedings of the 14th European Rotorcraft Forum*, (Milan, Italy), 1988.
- [37] PETERS, D. A. and BARWEY, D., “A General Theory of Rotorcraft Trim,” *Mathematical Problems in Engineering*, vol. 2, no. 1, pp. 1–34, 1996.

- [38] POTSDAM, M., YEO, H., and JOHNSON, W., “Rotor Airloads Prediction Using Loose Aerodynamic/Structural Coupling,” in *Proceedings of the 60th Annual Forum of the American Helicopter Society*, (Baltimore, MD), June 2004.
- [39] SHEFFER, S.G. AND ALONSO, J.J. AND MARTINELLI, L. AND JAMESON, A., “Time-Accurate Simulation of Helicopter Rotor Flows Including Aeroelastic Effects,” in *35th AIAA Aerospace Science Meeting and Exhibit*, (Reno, Nevada), AIAA-97-0399, January 6–9 1997.
- [40] LEE, C. S., SABERI, H., and ORMISTON, R. A., “Aerodynamic and Numerical Issues for Coupling CFD into Comprehensive Rotor Analysis,” in *Proceedings of the 53rd Annual Forum of the American Helicopter Society*, (Virginia Beach, VA), April 1997.
- [41] YEO, H., POTSDAM, M., and ORMISTON, R., “Application of CFD/CSD to Rotor Aeroelastic Stability in Forward Flight,” in *Proceedings of the American Helicopter Society 66th Annual Forum*, (Phoenix Arizona), May 2010.
- [42] ZAKI, A. AND REVELES, N. AND SMITH, M. J. AND BAUCHAU, O. A., “Using Tightly Coupled CFD/CSD Simulation for Rotorcraft Stability Analysis,” in *Proceedings of the American Helicopter Society 66th Annual Forum*, (Phoenix, Arizona), May 11–13 2010.
- [43] SHENG, C., ICKES, J., WANG, J., and ZHAO, Q., “CFD/CSD Coupled Simulations for a Helicopter Rotor in Forward and Maneuvering Flight,” in *31st AIAA Applied Aerodynamics Conference*, no. AIAA-2013-2792, (San Diego, California), June 24–27 2013.
- [44] MARPU, R. P., *Physics Based Prediction of Aeromechanical Loads for UH-60A Rotor*. PhD thesis, Georgia Institute of Technology, 2013.
- [45] GRAY, R., “On The Motion Of The Helical Vortex Shed From A Single-Bladed Hovering Helicopter Rotor And Its Application To The Calculation Of The Spanwise Aerodynamic Loading,” Tech. Rep. Princeton University Aeronautical Engineering Department Report 313, 1955.
- [46] LEHMAN, A. F., “Model Studies of Helicopter Rotor Patterns,” Tech. Rep. US-AAVLABS Technical Report 68-17, 1968.
- [47] PIZIALI, R., “Method for the Solution of the Aeroelastic Response Problem for Rotating Wings,” *Journal of Sound and Vibration*, vol. 4, no. 3, pp. 445 – 489, 1966.
- [48] SCULLY, M., “A Method of Computing Helicopter Rotor Wake Distortion,” *Massachusetts Inst. of Technology, ASRL TR*, pp. 138–1.
- [49] EGOLF, T. and SPARKS, S., “A Full Potential Rotor Analysis with Wake Influence Using an Inner-Outer Domain Technique,” in *Proceedings of the American Helicopter Society 42nd Annual Forum, Washington, DC*, pp. 997–1011, 1986.
- [50] BERKMAN, M. E., SANKAR, L. N., BEREZIN, C. R., and TOROK, M. S., “A Navier-Stokes/Full Potential/Free Wake Method for Advancing Multi-Bladed Rotors,” in *Proceedings of the American Helicopter Society 53rd Annual Forum*, (Virginia Beach, Virginia), April 29 – May 1 1997.

- [51] YANG, Z., SANKAR, L. N., SMITH, M. J., and BAUCHAU, O. A., “Recent Improvements to a Hybrid Method for Rotors in Forward Flight,” in *AIAA 38th Aerospace Sciences Meeting and Exhibit*, no. AIAA-2000-0260, (Reno, Nevada), January 10–13 2000.
- [52] YANG, Z., SANKAR, L. N., SMITH, M. J., and BAUCHAU, O. A., “Recent Improvements to a Hybrid Method for Rotors in Forward Flight,” *AIAA Journal of Aircraft*, vol. 39, pp. 804–812, October–November 2002.
- [53] MOULTON, M. A., WENREN, Y., and CARADONNA, F. X., “Free-Wake Hover Flow Prediction with a Hybrid Potential / Navier-Stokes Solver,” in *Proceedings of the 55th Annual Forum*, (Montreal, Canada), American Helicopter Society, May 1999.
- [54] MIN, B. and SANKAR, L., “Hybrid Navier-Stokes/Free-Wake Method for Modeling Blade-Vortex Interactions,” *Journal of Aircraft*, vol. 47, no. 3, 2010.
- [55] SITARAMAN, J. and BAEDER, J., “Evaluation of the Wake Prediction Methodologies used in CFD Based Rotor Airload Computations,” in *AIAA 24th Applied Aerodynamics Conference*, (San Francisco, California), June 5–8 2006.
- [56] D’ANDREA, A., “Development of a Multi-Processor Unstructured Panel Code Coupled with a CVC Free Wake Model for Advanced Analyses of Rotorcraft and Tiltrotors,” in *Proceedings of the American Helicopter Society 64th Annual Forum*, Montréal, Canada, 2008.
- [57] QUON, E. W., SMITH, M. J., WHITEHOUSE, G. W., and WACHSPRESS, D. A., “Hierarchical Variable Fidelity Methods for Rotorcraft Aerodynamic Design and Analysis,” in *Proceedings of the American Helicopter Society 67th Annual Forum*, (Virginia Beach, Virginia), May 3–5 2011.
- [58] WHITEHOUSE, G. R., BOSCHITSCH, A. H., QUACKENBUSH, T. R., WACHSPRESS, D. A., and BROWN, R. E., “Novel Eulerian Vorticity Transport Wake Module for Rotorcraft Flow Analysis,” in *Proceedings of the American Helicopter Society 63rd Annual Forum*, (Virginia Beach, Virginia), May 1–3, 2007.
- [59] WHITEHOUSE, G. R. and BOSCHITSCH, A. H., “Towards the Next Generation of Grid-Based Vorticity-Velocity Solvers for General Rotorcraft Flow Analysis,” in *Proceedings of the American Helicopter Society 69th Annual Forum*, (Phoenix, Arizona), May 21 –23 2013.
- [60] LUTZ, M. and ASCHER, D., *Learning Python*. Sebastopol, California: O’Reilly and Associates, Inc., 2003.
- [61] WISSINK, A.M. AND SITARAMAN, J. AND SANKARAN, V. AND MAVRIPLIS, D.J. AND PULLIAM, T.H., “Multi-Code Python-Based Infrastructure for Overset CFD with Adaptive Cartesian Grids,” in *46th AIAA Aerospace Sciences Meeting and Exhibit*, (Reno, NV), AIAA-2008-0927, January 7–10 2008.
- [62] SANKARAN, V., SITARAMAN, J., WISSINK, A., DATTA, D., JAYARAMAN, B., POTSDAM, M., MAVRIPLIS, D., YANG, Z., O’BRIEN, D., SABERI, H., CHENG, R., HARIHARAN, N., and STRAWN, R., “Application of Helios Computational Platform to Rotorcraft Flowfields,” in *AIAA 48th Aerospace Sciences Meeting and Exhibit*, no. AIAA-2010-1230, (Orlando, Florida), January 4–7, 2010.

- [63] SITARAMAN, J., POTSDAM, M., WISSINK, A., JAYARAMAN, B., DATTA, A., MAVRIPLIS, D., and SABERI, H., “Rotor Loads Prediction Using Helios: A Multi-solver Framework for Rotorcraft Aeromechanics Analysis,” *Journal of Aircraft*, vol. 50, no. 2, pp. 478–492, 2013.
- [64] BENDIKSEN, O. O., “A New Approach to Computational Aeroelasticity,” in *32nd Structures, Structural Dynamics, and Materials Conference*, no. AIAA-91-0939-CP, (Baltimore, MD), American Institute of Aeronautics and Astronautics, 1991.
- [65] KOLMOGOROV, A., “The Local Structure of Turbulence in Incompressible Viscous Fluid for Very Large Reynolds Numbers,” *Dokl. Akad. Nauk SSSR*, vol. 30, no. 4, 1941.
- [66] BUNING, P., JESPERSEN, D., PULLIAM, T., CHAN, W., SLOTNICK, J., KRIST, S., and RENZE, K., *OVERFLOW Users Manual, Version 1.8*. NASA Langley Research Center, 1998.
- [67] JESPERSEN, D., PULLIAM, T., and BUNING, P., “Recent Enhancements to OVERFLOW,” in *AIAA 35th Aerospace Sciences Meeting and Exhibit*, no. AIAA-97-0644, (Reno, Nevada), January 6–9, 1997.
- [68] STEGER, J., DOUGHERTY, F., and BENEK, J., “A Chimera Grid Scheme,” *American Society of Mechanical Engineers, New York*, pp. 59–69, 1993.
- [69] POTSDAM, M. A. and STRAWN, R. C., “CFD Simulations of Tiltrotor Configurations in Hover,” in *Proceedings of the 58th Annual Forum*, (Montreal, Canada), American Helicopter Society, June 2002.
- [70] NARDUCCI, R., “OVERFLOW Simulation of Rotors in Hover: The Boeing Company,” in *52th AIAA Aerospace Sciences Meeting*, no. AIAA-2014-0208, (National Harbor, Maryland), American Institute of Aeronautics and Astronautics, January 2014.
- [71] BUNING, P., GOMEZ, R., and SCALLION, W., “CFD Approaches for Simulation of Wing-Body Stage Separation,” in *AIAA 22nd Applied Aerodynamics Conference and Exhibit*, no. AIAA-04-4838, (Providence, Rhode Island), August 16–19 2004.
- [72] BUNING, P., PARKS, S., CHAN, W., and RENZE, K., “Application of the Chimera Overlapped Grid Scheme to Simulation of Space Shuttle Ascent Flows,” *Proceedings of the Fourth International Symposium on Computational Fluid Dynamics*, vol. 1, pp. 132–137, 1991.
- [73] SMITH, M. J., LIGGETT, N., and KOUKOL, B. C. G., “The Aerodynamics of Airfoils at High and Reverse Angles of Attack,” *AIAA Journal of Aircraft*, vol. 48, no. 6, pp. 2012–2023, 2011.
- [74] LIGGETT, N. and SMITH, M. J., “Temporal Convergence Criteria for Time-Accurate Viscous Simulations of Separated Flows,” *Computers & Fluids*, vol. 66, pp. 140–156, 2012. doi: 10.1016/j.compfluid.2012.06.010.
- [75] “<http://fun3d.larc.nasa.gov>.” Website, Last Accessed, February 2014.

- [76] NOACK, R., “DiRTlib: A Library to Add Overset Capability to Flow Solvers,” in *6th Symposium on Overset Grids and Solution Technology*, (Fort Walton Beach, FL), October 2002.
- [77] NOACK, R., “DiRTlib: A Library to Add an Overset Capability to Your Flow Solver,” in *17th Computational Fluid Dynamics Conference*, no. AIAA-2005-5116, (Toronto, Ontario), American Institute of Aeronautics and Astronautics, June 2005.
- [78] NOACK, R. W., “SUGGAR: a General Capability for Moving Body Overset Grid Assembly,” in *AIAA 17th Computational Fluid Dynamics Conference*, no. AIAA-2005-5117, (Toronto, Ontario), June 2005.
- [79] NOACK, R. W., BOGER, D. A., KUNZ, R. F., and CARRICA, P. M., “Suggar++: An Improved General Overset Grid Assembly Capability,” in *AIAA 19th Computational Fluid Dynamics Conference*, no. AIAA-2009-3992, (Austin, Texas), June 22-25 2009.
- [80] BIEDRON, R. and THOMAS, J., “Recent Enhancements to the FUN3D Flow Solver for Moving-Mesh Applications,” in *47th AIAA Aerospace Sciences Meeting Including the New Horizons Forum and Aerospace Exposition*, no. AIAA-2009-1360, (Orlando, Florida), American Institute of Aeronautics and Astronautics, January 2009.
- [81] ABRAS, J. and HARIHARAN, N., “CFD Solver Comparison of Low Mach Flow over the ROBIN Fuselage,” in *52th AIAA Aerospace Sciences Meeting*, no. AIAA 2014-0752, (National Harbor, Maryland), American Institute of Aeronautics and Astronautics, January 2014.
- [82] MASSEY, S.J. AND KRESHOCK, A.R., “Coupled CFD/CSD Analysis of Rotor Blade Structural Loads with Experimental Validation,” in *31st AIAA Applied Aerodynamics Conference*, no. AIAA-2013-3158, (San Diego, CA), June 24–27 2013.
- [83] PROSSER, D. and SMITH, M. J., “Navier-Stokes-Based Dynamic Simulations of Sling Loads,” in *54th AIAA/AHS/ASME Structural Dynamics and Materials Conference*, no. AIAA-2013-1922, (Boston, MA), April 2013. DOI: 10.2514/6.2013-1922.
- [84] WOOD, W., “Radiation Coupling with the FUN3D Unstructured-Grid CFD Code,” in *43rd AIAA Thermophysics Conference*, no. AIAA-2012-2741, (New Orleans, Louisiana), American Institute of Aeronautics and Astronautics, June 25–28 2012.
- [85] BARTELS, R., VATSA, V., CARLSON, J., PARK, M., and MINECK, R., “FUN3D Grid Refinement and Adaption Studies for the Ares Launch Vehicle,” in *28th AIAA Applied Aerodynamic Conference*, no. AIAA-2010-4372, (Chicago, Illinois), June 18–July 1 2010.
- [86] BAUCHAU, O.A., “Dymore User’s Manual,” tech. rep., Georgia Institute of Technology, June 2010.
- [87] BAUCHAU, O. A., BOTTASSO, C. L., AND NIKISHKOV, Y. G., “Modeling Rotorcraft Dynamics with Finite Element Multibody Procedures,” *Mathematical and Computing Modeling*, vol. 33, pp. 1113–1137, 2001.
- [88] REVELES, N. AND SMITH, M. J. AND ZAKI, A. AND BAUCHAU, O.A, “A Kriging-Based Trim Algorithm For Rotor Aeroelasticity,” in *Proceedings of the 37th European Rotorcraft Forum*, (Gallarata, Italy), September 12–15 2011.

- [89] MORILLO, J., AHAUS, L., REVELES, N., and SMITH, M. J., “Lead-Lag Stability Prediction: Why CFD is Needed,” in *Proceedings of the 68th American Helicopter Society Forum and Technology Display*, (Fort Worth, Texas), May 1–3 2012.
- [90] KRIGE, D. G., “Geostatistics and the Definition of Uncertainty,” *Transactions - Institution of Mining and Metallurgy*, vol. 93, pp. 41–47, 1984.
- [91] GIBSON, D. R., *Modeling Buliding and Design Augmentation for Improved Response Surface Estimation*. Georgia Institute of Technology: Material Handling Research Center, 1994.
- [92] LIU, J., HAN, Z., and SONG, W., “Efficient Kriging-Based Aerodynamic Design of Transonic Airfoils: Some Key Ideas,” in *AIAA Aerospace Sciences Meeting*, no. AIAA-2012-0967, (Nashville, Tennessee), January 2012.
- [93] DUCHAINE, F., MOREL, T., and GICQUEL, L., “Computational-Fluid-Dynamics-Based Kriging Optimization Tool for Aeronautical Combustion Chambers,” *AIAA Journal*, vol. 47, pp. 631–645, March 2009.
- [94] SIMPSON, T., MAUERY, T., KORTE, J., and MISTREE, F., “Kriging Models for Global Approximation in Simulation-Based Multidisciplinary Design Optimization,” *AIAA Journal*, vol. 39, December 2001.
- [95] MARTIN, J., “Robust Kriging Models,” in *AIAA Aerospace Sciences Meeting*, no. AIAA-2010-2854, (Orlando, Florida), April 2010.
- [96] MARTIN, J. and SIMPSON, T., “Use of Kriging Models to Approximate Deterministic Computer Models,” *AIAA Journal*, vol. 43, no. 4, pp. 853–863, 2005.
- [97] ZAKI, A. A., *Using Tightly-Coupled CFD/CSD Simulation for Rotorcraft Stability Analysis*. PhD thesis, Georgia Institute of Technology, Atlanta, Georgia, 2012.
- [98] VAZQUEZ, E. and WALTER, E., “Estimating Derivatives and Integrals with Kriging,” in *44th IEEE Conference on Decision and Control, and European Control Conference*, (Seville, Spain), IEEE, December 12–15 2005.
- [99] RAJMOHAN, N., *Application of Hybrid Methodology to Rotors in Steady and Maneuvering Flight*. PhD thesis, Georgia Institute of Technology, Atlanta, Georgia, 2010.
- [100] WACHSPRESS, D.A., “CHARM Wake-Panel Module User’s Manual (Version 4.0),” tech. rep., Continuum Dynamics, Inc., July 2009.
- [101] WACHSPRESS, D. A., QUACKENBUSH, T. R., and BOSCHITSCH, A. H., “First-Principles Free-Vortex Wake Analysis for Helicopters and Tiltrotors,” in *Proceedings of the American Helicopter Society 59th Annual Forum*, (Phoenix, Arizona), May 6–8, 2003.
- [102] SABERI, H., KHOSHLAHJEH, M., ORMISTON, R., and RUTKOWSKI, M., “Overview of RCAS and Application to Advanced Rotorcraft Problems,” in *Proceedings of the American Helicopter Society 4th Decennial Specialists’ Conference on Aeromechanics*, (San Francisco, California), January 2004.

- [103] JOHNSON, W., “Rotorcraft Dynamics Models for a Comprehensive Analysis,” in *Proceedings of the 54th American Helicopter Society Annual Forum*, (Washington D.C.), May 1998.
- [104] DURAISAMY, K., SITARAMAN, J., and BAEDER, J., “High Resolution Wake Capturing Methodology for Improved Rotor Aerodynamic Computations,” in *Proceedings of the American Helicopter Society 61st Annual Forum*, (Grapevine, Texas), 2005.
- [105] ABRAS, J. N. AND SMITH, M. J. AND BAUCHAU, O. A., “Rotor Airload Predictions Using Coupled Finite Element and Free Wake Methods,” in *Proceedings of the American Helicopter Society 62nd Annual Forum*, (Phoenix, Arizona), May 9–11 2006.
- [106] ANUSONTI-INTHRA, P., “Development of Rotorcraft Wake Capturing Methodology using Fully Coupled CFD and Particle Vortex Transport Method,” in *Proceedings of the American Helicopter Society 62nd Annual Forum*, Phoenix, AZ, 2006.
- [107] LYNCH, C. E. and SMITH, M. J., “Hybrid RANS-LES Turbulence Models on Unstructured Grids,” in *Proceedings of the 38th AIAA Fluid Dynamics Conference and Exhibit*, no. AIAA-2008-3854, (Seattle, WA), June 23–26 2008.
- [108] BIEDRON, R. and LEE-RAUSCH, E., “An Examination of Unsteady Airloads on a UH-60A Rotor: Computation versus Measurement,” in *Proceedings of the American Helicopter Society 68th Annual Forum*, (Fort Worth, TX), May 1–3 2012.
- [109] BOUSMAN, W., “The Response of Helicopter Rotors to Vibratory Airloads,” *Journal of the American Helicopter Society*, vol. 35, pp. 53–62, October 1990.
- [110] JOHNSON, W., *Helicopter Theory*. New York: Dover Publications, Inc., 1994.
- [111] BRAMWELL, A., DONE, G., and BALMFORD, D., *Helicopter Dynamics – 2nd Edition*. Woburn, MA: Butterworth Heinemann, 2001.
- [112] BOUSMAN, W. G., “A Qualitative Examination of Dynamic Stall from Flight Test Data,” *Journal of the American Helicopter Society*, vol. 43, no. 4, pp. 279 – 295, 1998.
- [113] MENTER, F. R., “Two-Equation Eddy-Viscosity Turbulence Models for Engineering Applications,” *American Institute of Aeronautics and Astronautics Journal*, vol. 32, no. 8, pp. 1598–1605, 1994.
- [114] LYNCH, C. E., *Advanced CFD Methods for Wind Turbine Analysis*. PhD thesis, Georgia Institute of Technology, Atlanta, Georgia, 2011.
- [115] VAN DER WALL, B. G., LIM, J. W., SMITH, M. J., JUNG, S., BAILLY, J., BAEDER, J., and BOYD, D.D., J., “An Assessment of Comprehensive Code Prediction State-of-the-Art Using the HART II International Workshop Data,” in *Proceedings of the 68th American Helicopter Society Forum and Technology Display*, (Fort Worth, Texas), May 1–3 2012.
- [116] SPALART, P. R. and ALLMARAS, S. R., “A One-Equation Turbulence Model for Aerodynamic Flows,” *LaRecherche Aerospatiale*, vol. 1, no. 5, pp. 5–21, 1994.
- [117] BOUSMAN, W. and NORMAN, T., “Assessment of Predictive Capability of Aeromechanics Methods,” *Journal of the American Helicopter Society*, vol. 55, January 2010.

- [118] SMITH, M. J., LIM, J. W., VAN DER WALL, B. G., BAEDER, J., BIEDRON, R. T., BOYD, D.D., J., JAYARAMAN, B., JUNG, S., and MIN, B.-Y., “The HART II International Workshop: An Assessment of the State of the Art in CFD/CSD Prediction,” in *CEAS Aeronautical Journal*.
- [119] MAKINEN, A. M., WAKE, B., and OPOKU, D., “Quantitative Evaluation of Rotor Load Prediction Results Correlated to Flight Test Data,” in *Proceedings of the American Helicopter Society 67th Annual Forum*, (Virginia Beach, VA), May 3–5 2011.
- [120] MARPU, R. P. AND SANKAR, L. N. AND NORMAN, T. R. AND EGOLF, T. A. AND MAKINEN, S., “Analysis of the UH-60A Rotor Loads Using Wind Tunnel Data,” in *51st AIAA Aerospace Sciences Meeting*, no. AIAA-2013-0635, (Dallas/Ft. Worth, TX), January 7-10 2013.
- [121] SMITH, M. J., LIM, J. W., VAN DER WALL, B. G., BAEDER, J., BIEDRON, R. T., BOYD, D.D., J., JAYARAMAN, B., JUNG, S., and MIN, B.-Y., “An Assessment of CFD/CSD Prediction State-of-the-Art Using the HART II International Workshop Data,” in *Proceedings of the 68th American Helicopter Society Forum and Technology Display*, (Fort Worth, Texas), May 1–3 2012.
- [122] YABIL, S., SMITH, M. J., and DIMITRIADIS, G., “Unsteady Navier-Stokes Simulation of Low-Reynolds Stall Flutter,” in *AIAA 50th Aerospace Sciences Meeting and Exhibit*, no. AIAA-2012-0037, (Nashville, Tennessee), January 9–12 2012.
- [123] WATRIN, D., PERRY, T., DIMITRIADIS, G., and SMITH, M. J., “Computational Considerations for the Prediction of Stall Flutter,” in *AIAA Structural Dynamics, Structures and Materials Conference*, no. AIAA-2012-1706, (Honolulu, Hawaii), April 2012.
- [124] GIRON-SIERRA, J. M., PRIETO, J., DELGADO, I., GMV, I. N., and CANTOS, P. T., “Overview of Advanced Control Trends for Aerospace Missions,” in *54th International Astronautical Congress of the International Astronautical Federation, the International Academy of Astronautics, and the International Institute of Space Law*, (Bremen, Germany), September 29 – October 3 2003.
- [125] SAHANI, N. A. and HORN, J., “Command Limiting for Full-Envelope Guidance and Control of Rotorcraft,” in *AIAA Guidance, Navigation, and Control Conference and Exhibit*, no. AIAA-2005-6348, (San Francisco, California), August 15–18 2005.
- [126] COOK, J., SMITH, M. J., SANCHEZ-ROCHA, M., SMITH, M. J., SHENOY, R., and MENON, S., “Improved Prediction of Complex Rotorcraft Aerodynamics,” in *69th American Helicopter Society Annual Forum*, (Phoenix, AZ), May 21–23, 2013.
- [127] QUACKENBUSH, T. and WACHSPRESS, D., “CDI Review And Assessment Of Selected Issues In Hovering Rotor Tip Vortex Dynamics,” in *Proceedings of the Army Research Office Rotorcraft Wake Prediction Basic Research Workshop*, Georgia Institute of Technology, Atlanta, GA, March 2009.

VITA

Nicolas Reveles was born in California in 1985. He quickly developed an intense interest in aircraft and rockets that led him to design, build, and fly high powered rockets in the Texas desert. In 2004, he entered the aerospace engineering program at the University of Texas in Arlington. After becoming enamored with the field, he transferred to the Georgia Institute of Technology.

During his time as an undergraduate at Georgia Tech, Nicolas served as the president of the Ramblin' Rocket Club and was an active participating member of the Yellow Jacket Flying Club. With a growing interest in computation-based simulation, he interned at the Nonlinear Computational Aeroelasticity Lab (NCAL). Upon completing his Bachelor's degree with highest honors, he entered graduate school in the same laboratory.

While completing his graduate studies, Nicolas continued to active in the aerospace community, serving in various leadership roles in the local chapter of the American Helicopter Society. He proudly received two Vertical Lift Consortium scholarships during this period. He also interned at Bell Helicopter's Rotor Dynamics group for one summer. In 2014, Nicolas completed a Doctorate of Philosophy in aerospace engineering at the Georgia Institute of Technology.