

OVERSET ADAPTIVE STRATEGIES FOR COMPLEX ROTATING SYSTEMS

A Thesis
Presented to
The Academic Faculty

by

Rajiv Shenoy

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

Georgia Institute of Technology
May 2014

OVERSET ADAPTIVE STRATEGIES FOR COMPLEX ROTATING SYSTEMS

Approved by:

Professor Marilyn Smith,
Committee Chair
School of Aerospace Engineering
Georgia Institute of Technology

Professor Stephen Ruffin
School of Aerospace Engineering
Georgia Institute of Technology

Professor Daniel Schrage
School of Aerospace Engineering
Georgia Institute of Technology

Professor Ari Glezer
School of Aerospace Engineering
Georgia Institute of Technology

Dr. Michael Park
Computational Aerosciences Branch
NASA Langley Research Center

Date Approved: December 19, 2013

Dedicated to my parents,
Rajendra & Pramila Shenoy.

ACKNOWLEDGEMENTS

I would like to express my deep gratitude to Prof. Marilyn Smith for taking me on as a student and giving me the chance to work on this thesis topic. She encouraged me to stay focused on my thesis goals and gave many useful suggestions for shaping my work here. I also had the good fortune of spending two summers at NASA Langley’s summer internship program in my hometown, which helped me make progress toward my thesis goals. My interaction with many scientists at the Computational Aerosciences Branch and in particular, the FUN3D team, was very fruitful and gave me exposure to several areas of computational fluid dynamics. In particular, I would extend my gratitude to Mike Park, Beth Lee-Rausch, Bob Biedron, Eric Nielsen, and Boris Diskin for all their inputs and ideas. The work that went in to this thesis was built over strong underlying code infrastructure, developed by these scientists. I would also like to thank the rest of the committee, Prof. Ruffin, Prof. Schrage, and Prof. Glezer, for their advisement and inputs toward refining my thesis.

The Office of Naval Research and the Vertical Lift Consortium provided funding for the projects that supported me. The DoD High Performance Computing Modernization Program provided me access to the Navy and Army Engineer Research and Development Center DoD Supercomputing Resource Centers for computational time used for my simulations. These facilities enabled successful and timely completion of my simulations and are gratefully acknowledged.

The Office of Naval Research (ONR) provided financial support for much of this research through ONR Grant N0001409-1-1019. Additional financial support was provided by the Vertical Lift Consortium (VLC) and the National Rotorcraft Technology Center (NRTC), U.S. Army Aviation and Missile Research, Development and

Engineering Center, under Technology Investment Agreement W911W6-06-2-0002, entitled National Rotorcraft Technology Center Research Program. The DoD High Performance Computing Modernization Program provided me access to the Navy and Army Engineer Research and Development Center DoD Supercomputing Resource Centers for computational time used for my simulations. These facilities enabled successful and timely completion of my simulations and are gratefully acknowledged. I would like to thank the ONR Technical monitors, Judah Milgram and John Kinzer, the Navy S/AAA Odessa Murray, and the VLC and NRTC organizations for their help and support. The views and conclusions contained in this document are my own and should not be interpreted as representing the official policies, either expressed or implied, of the U.S. Government.

My parents have been a huge source of support for me and have encouraged me to pursue my Ph.D. degree. I am forever indebted to all their sacrifices and love. My grandmother helped instill belief in my abilities and continues to inspire me in many ways, which I will never forget. My loving sister and her family have also been very supportive over the last several years. The patience, understanding, and affection I received from my girlfriend Namita have been a source of comfort for me. The several trips we made between Illinois and Atlanta went a long way toward keeping me focused and staying motivated as I was wrapping up my thesis.

The several friends I made in Atlanta during my time in graduate school have been important for me. I will never forget Chaitanya, Amruta, Faisal, Prashant, Joel, and Sayan for being there for me and keeping me sane over the last five years. I would like to acknowledge the rest of the lab members I've had a chance to work with at the Nonlinear Computational Aeroelasticity Lab; they've all enhanced my experience of graduate school by maintaining a good work atmosphere in the lab. Nicolas Reveles and Eliot Quon, who are defending their theses shortly, worked along side me over the last five years and my interactions with them motivated me to excel in my own work.

Eric Lynch, James Cook, and Nick Liggett were my seniors who helped familiarize me with different software. My post-processing efforts for this thesis were aided by Marlin Holmes, Philip Cross, Chris Sandwich, and Phillip Richards. Their help has been very instrumental and is acknowledged.

TABLE OF CONTENTS

DEDICATION	iii
ACKNOWLEDGEMENTS	iv
LIST OF TABLES	x
LIST OF FIGURES	xi
SUMMARY	xx
I INTRODUCTION	1
1.1 Rotor-Fuselage and Wind Turbine-Tower Interactions	1
1.1.1 Computational State of the Art	2
1.2 Rotor Hubs	3
1.2.1 Computational State of the Art	5
1.3 Grid Adaptation	6
1.4 Error Estimation Methods	7
1.5 Interpolation for Adaptive Grids	9
1.6 Thesis Objectives	11
II COMPUTATIONAL TOOLS	14
2.1 FUN3D Flow Solver	14
2.1.1 Backward-Differentiation Formula (BDF) Time Integration	15
2.1.2 Turbulence Models	16
2.1.3 Error Estimation Formulations	16
2.1.4 Metric-Based Adaptation	17
2.2 REFINE & BAMG	19
2.3 SUGGAR++ & DiRTlib	19
III TIME-DEPENDENT OVERSET GRID ADAPTATION	22
3.1 Adaptation for Overset Grids	22
3.2 Time-Dependent Error Estimation	23

3.2.1	General versus Periodic Adaptation	25
IV	PARALLEL & ROBUST LOCALIZATION SCHEME	27
4.1	Parallelized Localization Algorithm	28
4.1.1	Nomenclature	28
4.1.2	Barycentric Coordinates and the Neighbor Walk	28
4.1.3	Advancing Front	31
4.1.4	Parallel Paradigm	33
4.2	Robustness Features to Overcome Search Failures	36
4.2.1	Hierarchical Localization	36
4.2.2	Default Search Methods	37
4.3	Generalization to Overset Mixed-Element Unstructured Grids	38
4.4	Cost Analysis	38
4.4.1	Algorithm Complexity	39
4.4.2	Timings	40
V	VERIFICATION OF DESIGN ORDER ACCURACY	44
5.1	Two-Dimensional Inviscid Advecting Vortex	44
5.2	Adaptation Interval Study	47
5.3	Uniform versus Adaptive Refinement	48
5.4	Temporal Resolution	50
5.5	Temporal Evolution of Error	54
VI	ROTOR-FUSELAGE INTERACTION DEMONSTRATION	56
6.1	Georgia Institute of Technology Configuration	56
6.1.1	Approach	58
6.1.2	Methodology Demonstration	60
6.2	Results: Periodic Grid Adaptation	62
6.2.1	Adaptation Strategy Study: Inviscid Computations	62
6.2.2	Turbulence Model Study	68
6.2.3	HRLES Computations	70

6.2.4	Vorticity–Q-criterion Discrepancy	78
6.2.5	Tip Vortex Paths	80
6.2.6	Adaptation Convergence	81
6.3	Results: Adaptation Interval Sensitivity	82
6.3.1	Tip Vortex Paths	85
6.4	Rotor-Fuselage Interaction Aerodynamics	88
VII	GENERIC ROTOR HUB APPLICATION	96
7.1	Rotor Hub Models	96
7.1.1	Baseline Configuration	96
7.1.2	Scissors Configuration	99
7.2	Simulation Approach	100
7.2.1	Temporal Resolution Study	102
7.3	Adaptation Strategy Validation	103
7.4	Static & Dynamic Loads	107
7.4.1	Component Drag Analysis	109
7.4.2	Time-Averaged and Vibratory Loads Analysis	112
7.5	Wake Structure	117
7.5.1	Wake Velocity Profiles	118
7.5.2	Time-Averaged Velocity Trends	120
7.5.3	Velocity Spectral Analysis	123
VIII	CONCLUSIONS	134
8.1	Recommendations for Future Work	137
	REFERENCES	141
	VITA	151

LIST OF TABLES

1	Grid spacings for uniform refinement study.	49
2	Time step sizes for temporal resolution study.	52
3	Vorticity Error Statistics.	55
4	Flow conditions and rotor parameters.	57
5	Summary of tolerances applied for the different adaptation schemes. .	60
6	Inviscid simulations: Summary of grid sizes and ability to predict vortex impingement from the different adaptation schemes.	63
7	HRLES: summary of grid size and ability to predict vortex impingement from the different adaptation schemes	71
8	Summary of adaptation interval study.	83
9	Hub component dimensions at 1/3.5-wind tunnel model.	98
10	Free-stream velocities and associated Reynolds number for the different flow conditions.	98
11	Matrix of flow conditions and analysis methods.	99
12	Effect of adaptation on drag prediction for baseline static simulations (Note: Initial grid size is 11.7 million nodes.)	105
13	Drag coefficient predictions illustrating effect of fuselage on computed hub drag.	107
14	Drag coefficient tabulation for cylindrical components at 0° static azimuth orientations.	112
15	Average drag for rotating cylindrical components at $U_\infty = 8.941$ m/s. .	112
16	Computationally predicted drag coefficients showing the effect of scissors at 240 rpm.	114

LIST OF FIGURES

1	Global index convention illustrated for an example overset grid system.	23
2	Modified transient fixed point iteration method for time-dependent adaptation.	26
3	Barycentric coordinate system on a triangle with associated areas A, B, and C corresponding to points a , b , and c , respectively.	29
4	Determination of node localization and the neighbor step based on barycentric coordinates (two-dimensional representation).	32
5	Neighbor walk algorithm (two-dimensional localization).	33
6	Advancing front iterations depicted in two dimensions.	34
7	Complexity verification of different grid sizes.	40
8	Timing comparison on different grid sizes demonstrating the efficiency of the localization algorithm	42
9	Ratio of time spent during interpolation versus localization	42
10	Localization time estimates for finer grids	43
11	Vorticity ω_y profile of the exact solution.	45
12	Total error comparison of adaptation interval size after 24 characteristic periods (300 time steps). (Note: Second-order slope only accounts for spatial error).	47
13	Total error comparison of adaptation interval size after 24 characteristic periods (600 time steps). (Note: Second-order slope only accounts for spatial error.)	48
14	Total error comparison between uniform and adaptive refinement (2 intervals) after 24 characteristic periods (600 time steps). (Note: Second-order slope only accounts for spatial error.)	49
15	Vorticity mesh contours in each adaptation interval (adaptation cycle 3).	50
16	Vorticity contours in each adaptation interval (adaptation cycle 3).	51
17	Total error comparison of temporal resolution after 24 characteristic periods. Two adaptation intervals are used in these simulations. (Note: Second-order slope only accounts for spatial error).	52

18	Spatial error comparison after 24 characteristic periods. Two adaptation intervals are used in these simulations. Second-order slope is plotted in black.	53
19	Temporal error comparison after 24 characteristic periods. Two adaptation intervals are used in these simulations.	54
20	Time history of vorticity error. Solution interpolation occurs at characteristic times of 6, 12, and 18.	55
21	Model of the GIT RFI configuration.	57
22	Approximate sketches of the vortex-fuselage interaction phenomena. .	59
23	Grid comparison; view is of the blade planform from the top.	61
24	Vorticity-magnitude comparison; view is of the blade planform from the top.	61
25	Inviscid simulations: top centerline time-averaged pressures from one adaptive cycle.	64
26	Inviscid simulations: top centerline time-averaged pressures from two adaptive cycles.	64
27	Inviscid simulations: top centerline instantaneous pressures (first-quarter revolution).	66
28	Inviscid simulations: top centerline instantaneous pressures (second-quarter revolution).	67
29	Contours along the airframe symmetry plane from the $k\omega$ -SST model.	69
30	Contours along the airframe symmetry plane from the HRLES model.	69
31	HRLES: time-averaged pressures on the top centerline from one adaptive cycle.	71
32	HRLES: time-averaged pressures on the top centerline from two adaptive cycles.	72
33	HRLES: top centerline instantaneous pressures (first-quarter revolution).	73
34	HRLES: top centerline instantaneous pressures (second-quarter revolution).	74
35	Vortex behavior (from top: initial grid, adapted to $ \omega $, and <i>vorticity-mixed scheme</i>).	75
36	Vortex behavior (from top: initial grid, adapted to $ \omega $, and <i>vorticity-mixed scheme</i>).	76

37	Vortex behavior (from top: initial grid, adapted to $ \omega $, and <i>vorticity-mixed scheme</i>).	77
38	Inviscid simulations: comparison of the <i>vorticity-mixed</i> and <i>Q-criterion-mixed schemes</i>	79
39	HRLES: comparison of the <i>vorticity-mixed</i> and <i>Q-criterion-mixed schemes</i>	79
40	Comparison of inviscid and viscous modeling of the flowfield for the <i>Q-criterion-mixed scheme</i>	80
41	Vortex-trajectory comparisons from the <i>vorticity-mixed scheme</i>	81
42	Pressure-integral functional convergence for the HRLES <i>vorticity-mixed scheme</i> adaptation sequence.	82
43	Convergence of time-averaged pressures of the 15° adaptation interval using the <i>vorticity-mixed scheme</i> adaptation sequence.	84
44	Comparison of time-averaged pressures varying the adaptation interval size using the <i>vorticity-mixed scheme</i> adaptation sequence.	84
45	Top centerline instantaneous pressures for the <i>vorticity-mixed scheme</i> (first-quarter revolution).	86
46	Top centerline instantaneous pressures for the <i>vorticity-mixed scheme</i> (second-quarter revolution).	87
47	Vortex-trajectory comparisons from the <i>vorticity-mixed scheme</i> for HRLES computations.	88
48	Effect of rotor on fuselage forces.	90
49	Pressure distributions on the fuselage and blades at different azimuths.	91
50	Effect of rotor on fuselage moments.	93
51	Effect of rotor on blade forces and moments. The blade is originally at $\psi = 180^\circ$	95
52	Baseline hub model studied by both computational and experimental campaigns.	97
53	Hub model within John J. Harper wind tunnel test section.	99
54	Hub configuration with and without scissors.	100
55	Planform (looking down) view of the vorticity magnitude for the initial and adapted grid.	101
56	Left (looking forward) side view of the vorticity magnitude for the initial and adapted grid.	102

57	HRLES blending function for $U_\infty = 13.41$ m/s (static configuration). Hot-wire sampling locations are plotted.	103
58	Comparison of normalized power spectral density (PSD) predicted using two different time steps for the $U_\infty = 13.41$ m/s rotating hub at 240 rpm. The wake locations (1-4) correspond to Fig. 57. (Note: Data are presented using the final adapted grid.)	104
59	Drag coefficient convergence with adaptation for the 0° static hub orientation. An error bar of 1% above and below the final drag coefficient is outlined (dashed line).	105
60	Comparison of time-averaged PIV and computationally obtained hub centerline wake velocity deficits at 1 hub diameter downstream ($x = 1D$) along the tunnel axis.	106
61	Correlation between experimental load cell and computational data for the 0° static hub model.	108
62	Surface pressure coefficient distribution on the rotor hub ($U_\infty = 8.941$ m/s). Free-stream flow enters normal to the figure.	109
63	Definition of the different hub components.	110
64	Component drag breakdown for the $U_\infty = 8.941$ m/s simulations. . .	111
65	Comparison of unsteady drag harmonics of the rotating hub with scissors. .	114
66	Comparison of unsteady side force of the rotating hub with scissors. .	115
67	Q-criterion contours of the wake at $Z = 0.0$ m at 240 rpm.	116
68	Wake velocity measurement locations in space. (Red line indicates traverse of each profile).	117
70	Comparison of computational wake velocity profiles for the different configurations at one hub diameter in the wake.	118
69	Time-averaged wake velocity comparison of averaging sampling size. .	119
71	Time-averaged wake velocity comparison at $Z = 0.0$ m at different downstream locations. The rotor hub conditions are at 240 rpm and $U_\infty = 8.941$ m/s.	121
72	Time-averaged wake velocity comparison at $Z = 0.0$ m at different downstream locations. The rotor hub conditions are at 240 rpm and $U_\infty = 13.41$ m/s.	122
73	Time-averaged wake velocity comparison at $Z = 0.0$ m at different downstream locations. The rotor hub conditions are at 240 rpm and $U_\infty = 22.35$ m/s.	123

74	Wake PSD comparison at $z = 0.0$ m and $x = 1D$. The rotor hub conditions are at 240 rpm and $U_\infty = 8.941$ m/s.	125
75	Wake PSD comparison at $z = 0.0$ m and $x = 2D$. The rotor hub conditions are at 240 rpm and $U_\infty = 8.941$ m/s.	126
76	Wake PSD comparison at $z = 0.0$ m and $x = 3D$. The rotor hub conditions are at 240 rpm and $U_\infty = 8.941$ m/s.	127
77	Wake PSD comparison at $z = 0.0$ m and $x = 1D$. The rotor hub conditions are at 240 rpm and $U_\infty = 13.41$ m/s.	128
78	Wake PSD comparison at $z = 0.0$ m and $x = 2D$. The rotor hub conditions are at 240 rpm and $U_\infty = 13.41$ m/s.	129
79	Wake PSD comparison at $z = 0.0$ m and $x = 3D$. The rotor hub conditions are at 240 rpm and $U_\infty = 13.41$ m/s.	130
80	Wake PSD comparison at $z = -0.204D$ and $x = 1D$ at 240 rpm. . . .	131
81	Wake PSD comparison at $z = -0.204D$ and $x = 2D$ at 240 rpm. . . .	132
82	Wake PSD comparison at $z = -0.204D$ and $x = 3D$ at 240 rpm. . . .	133

NOMENCLATURE

Roman

a	Speed of sound, m/s
\mathcal{C}	Complexity coefficient
C	Coarsening factor or vortex strength parameter
c	Rotor blade chord, m
C_ν, C_ϵ	Coefficients used in LES k equations
C_D, C_d	Three-dimensional or two-dimensional drag coefficient $\left(\frac{D}{(1/2)\rho U_\infty^2 S_{ref}}\right)$
C_L	Lift coefficient $\left(\frac{L}{(1/2)\rho U_\infty^2 S_{ref}}\right)$
c_p	Pressure coefficient $\left(\frac{p-p_\infty}{(1/2)\rho U_\infty^2}\right)$
C_T	Thrust coefficient $\left(\frac{T}{\rho\pi R^2(\Omega R)^2}\right)$
C_Y	Side force coefficient
C_{M_x}	Rolling moment coefficient (hub or body) $\left(\frac{M_x}{(1/2)\rho U_\infty^2 S_{ref} c}\right)$
C_{M_y}	Pitching moment coefficient (hub or body)
C_{M_z}	Yawing moment coefficient (hub or body)
$\overline{f_{c_p}}$	Time-averaged fuselage centerline pressure integral
e, e_t, e_s	Error: total, temporal, and spatial
$F_{e,(-)}$	Adaptation formulation across an edge based on a solution feature
F_{tol}	User specified adaptation tolerance

H	Hessian matrix
h	Grid spacing, grid units, m^2/s^2
h_{eff}	Effective grid spacing
\hat{I}	Adaptation intensity
i	Adaptation time interval
j	Fixed-point iteration
k	Turbulence kinetic energy,
k^{SGS}	Subgrid scale turbulence kinetic energy
l_e	Edge length, grid units
M	Anisotropic adaptation metric
n	Time step index within adaptation interval
n_1, n_2	Node (1 or 2) of given edge
N_I	Number of adaptation intervals
N_s	Number of samples within adaptation interval
N_t	Number of time steps within adaptation interval
N_{FP}	Number of fixed point adaptation iterations
p	Static pressure, N/m^2
Pr	Prandtl number
Q_{crit}	Q-criterion $(\frac{1}{2}(\ \Omega\ ^2 - \ S\ ^2))$, s^{-2}
R	Rotor radius, m, or vortex core radius

Re_c	Reynolds number based on blade chord and tip speed $\left(\frac{V_{tip}c}{\nu}\right)$
S	Mean strain-rate, s^{-1}
S_{ref}	Reference area for force and moment nondimensionalization, m^2
T	Rotor thrust, N
t	Time, s
u, v, w	Instantaneous local velocity components in (x,y,z), m/s
U_∞	Free stream velocity, m/s
\mathcal{V}	Dual volume
V_h	Effective hot-wire measured velocity, m/s
V_{tip}	Rotor tip speed, m/s
X	Grid metric eigenvectors
x, y, z	Streamwise, spanwise and vertical directions, m
y^+	Nondimensional normal distance of the first cell from the wall $\left(\sqrt{\frac{\tau_w}{\rho}} \frac{y}{\nu}\right)$

Greek

β_{1c}	Longitudinal first flapping harmonics, deg.
β_{1s}	Lateral first flapping harmonics, deg.
β_i	Barycentric coordinates for tetrahedra
Δ	Characteristic length for grid scale
δ_{ij}	Kronecker delta
δ_j	Coefficients of BDF time integration

Λ	Grid metric eigenvalues
μ	Rotor advance ratio or coefficient of molecular viscosity, N·s/m ²
μ_T	Turbulence or eddy viscosity
ν	Kinematic viscosity, m ² /s
ρ	Density, kg/m ³
τ	Specific viscous stress tensor, m ² /s ²
τ_w	Specific viscous stress at wall
ϕ	Solution variable
ψ	Blade azimuth angle, deg.
∇^j	Backward-differencing operator
Ω	Rotation rate tensor, s ⁻¹ , or rotational velocity of rotor, rad/s
ω	Specific turbulence dissipation rate or vorticity magnitude, s ⁻¹

SUMMARY

The resolution of the complex physics of rotating configurations is critical for any engineering analysis that requires multiple frames of reference. Two well-known applications are in the rotorcraft and wind energy industries. Rotor wake impingement from rotor-fuselage and wind turbine-tower interactions impact structural and acoustic characteristics. Additionally, parasite drag resulting from rotorcraft hubs may result in severe limitations on forward flight vehicle performance. Complex turbulent wakes from rotors and hubs impinging on downstream empennage can create adverse aeroelastic behavior and can affect handling qualities. Numerical simulations of these flows require state-of-the-art Navier Stokes methods using dynamic overset grids. However, many current methods typically used in industry result in wakes that dissipate essential features. In order to address these concerns, two advancements are introduced in this thesis.

Feature-based grid adaptation on dynamic overset grids has been developed and demonstrated with an unstructured Navier Stokes solver. The unique feature of the adaptation technique is that it is applied globally on the overset grid system except within the boundary layer. In concert with grid adaptation, an efficient parallelized search algorithm for solution interpolation over massively distributed systems has been created. This results in cost-effective interpolation that retains the numerical order of accuracy and has been verified in both space and time. The improvements have been demonstrated for rotor-fuselage interaction and a generic rotating hub. Detailed analysis of convergence of the methodology and sensitivity of the results to relevant parameters have also been included.

CHAPTER I

INTRODUCTION

Rotating flows are encountered in many common engineering applications. In the regime of aerospace and mechanical engineering applications, these flows can be classified as either internal or external. Internal flows are encountered in turbofan and turbojet engines at the compressor and turbine stages. The most common rotating systems pertaining to external flows are encountered in rotorcraft and wind turbines. Although the purposes of helicopter and wind turbine rotors are different, the aerodynamic and aeroelastic issues they encounter are quite similar due to the similar nature of their wakes, which in many cases are complex, turbulent, and a challenge to predict using high fidelity computational tools.

1.1 Rotor-Fuselage and Wind Turbine-Tower Interactions

Rotor-Fuselage Interaction (RFI) is a major area of concern to the rotorcraft industry. The interdependent interactions of the rotor blades and the helicopter fuselage can have a significant influence on each other's steady and unsteady pressures. In addition to structural effects, RFI plays a major role on the helicopter's performance, handling qualities, and acoustic response. While vehicle performance is governed by steady loading, the fluctuating surface pressures caused by rotor wake impingement on the fuselage and empennage structure can lead to component fatigue or failure, and dictate acoustic characteristics [1]. From an aerodynamic efficiency perspective, these strong vortex interactions pose problems because they occur at most flight conditions. Payload capabilities also suffer as a result of these complex, unsteady phenomena [2]. RFI has been identified as a fundamental problem whose evaluation using unsteady

Reynolds-averaged Navier-Stokes (URANS) and Large Eddy Simulation (LES) computational fluid dynamics (CFD) is a major long-term goal for researchers [3].

Wind turbines similarly operate in aerodynamic conditions that are challenging to numerically simulate. If the wind turbine is not oriented normal to the wind, the blade loads vary cyclically due to rotation. The turbine wake impingement on the tower and nacelle makes it essential that these structures are able to withstand such loads. Furthermore, in a downwind configuration the tower’s turbulent wake influences airloads on the rotor blades. This highlights the need to accurately model the effect of rotating blades in the presence of stationary geometries. Lynch [4] recommended that more accurate wake methods are necessary to capture and preserve the physics essential to this region where such interactions occur. Furthermore, Abras [5] endorsed the development of an unsteady grid adaptation method for the accurate resolution of rotorcraft wakes.

1.1.1 Computational State of the Art

Several different computational campaigns have studied rotor-fuselage interactions as detailed by O’Brien [6]. For Euler and Navier-Stokes simulations, the rotor modeling has over the past two decades seen the development of the actuator-based method [7–10] followed by overset grid methods [11–13] to capture the motion in multiple frames. Another method that has been used by several researchers is the sliding mesh approach [14,15], which avoids the extra overhead of the overset method, but is limited in that it cannot account for complex grid motion, e.g. intermeshing rotors.

The most popular configuration for demonstrating improved rotor-fuselage interactions in the literature is the Georgia Institute of Technology RFI configuration [16,17]. Many prior computational campaigns [13,15,18–23] have used this configuration for validation. Most researchers have studied this configuration using the Euler equations, while failing to capture important vortex-fuselage impingement

phenomena. O'Brien [6] and Steijl et al. [21] underscored that viscous effects show improved correlations with experimental data because of their ability to capture flow separation and other effects. Additionally, anisotropic feature-based grid adaptation was recommended in conjunction with the overset method [6] for improved predictions with minimal computational grid increases.

1.2 Rotor Hubs

Parasite drag on rotorcraft can become a crucial factor in forward flight especially during high speed flight [24, 25] and can limit the range, maximum speed, and payload of the vehicle. Reduction in parasite drag can improve vehicle stability and control [26] and significantly decrease vibrational and blade loads to reduce vehicle weight and extend the rotor blade life [27]. Hub assemblies for single main rotor helicopters can contribute nearly 25%-30% of vehicle parasite drag, while hub assemblies for coaxial rotors, such as those on the XH-59 [28], can contribute as much as 50% of the parasite drag. The complex and highly turbulent rotor and hub wake is convected to the empennage and results in aeroelastic behavior such as tail buffet and aerodynamic behavior known as tail "wag". These phenomena result in fatigue and reduced handling qualities [25]. Therefore, reduction of the drag in the design of hub systems is critical to the success of high-speed rotorcraft design. In order to achieve these goals, the drag sources associated with complex hub designs must be thoroughly investigated, well understood, and reliably predicted using analysis tools.

For fixed wing vehicles, there is a large literature base on flight to wind tunnel drag correlation [29,30] preceding the establishment of cryogenic wind tunnels capable of separately or simultaneously analyzing flight Reynolds and Mach numbers with sub-scale models. However, rotary wing designers must deal with aerodynamically complicated appendages and interference of rotating components, where the issue of drag is much more complex and significant.

A typical hub is a plethora of interacting bluff bodies, which are associated with poor performance characteristics and higher drag. The primary drag contributions in these bluff bodies arise from flow separation (pressure drag) rather than viscous effects (friction drag). Bluff body wakes affect the performance of both commercial and military air vehicles [31], in particular impacting tail component fatigue and handling qualities. Accurate hub drag predictions are inhibited by many different complex flow interactions. Some hub wake characteristics include periodic forcing and vortex interactions. In many cases, these interactions originate from fine structures such as tubes, wires and linkages. The identification of the influence of the full assembly and component Reynolds numbers can not be minimized; over the Reynolds number range of interest, distinct bifurcations produce significant and measurable differences in the wake flow field [32].

Much of the experimental research of hub drag is directed toward improving drag characteristics of current hub designs by the addition of fairings. Fairing designs have been explored by Sikorsky and others to reduce flow separation and interference drag between the hub and fuselage [28]. To date, frontal swept area of the hub design has been the leading parameter tied to hub drag [33], therefore the fairing of an existing hub design does not address the issue directly. This is especially true for articulated hubs, where empty space is required for the control hinges, resulting in channel flows that interact with cylindrical components. While empirically corrected analytic estimates have been developed to predict hub drag based on frontal area, there is no consistent trend when accounting for interference effect and frontal swept area [34]. Considerations for hub displacement from the fuselage have been made, weighing the effects of increased frontal area to decreased interference [33].

1.2.1 Computational State of the Art

Recently, there have been several fundamental experimental studies of hub drag, accompanied by state-of-the-art computational fluid dynamics (CFD) predictions of these complex flows [28, 35–39]. These studies have all focused on models that are a fraction ($1/5$ - $1/4$ scale) of the full-scale rotor hub. Thus, the scaling of these complex hub systems, including rotational and interference effects, must also be understood.

Wake et al. [28] investigated $1/4$ -scale faired hubs for the X2 TechnologyTM Demonstrator aircraft using an unstructured computational solver. These simpler faired elliptical geometries can prove to be challenging because of the difficulty of predicting separation (and potentially transition). The focus of their effort was to investigate the impact of aerodynamic fairings on drag for the dual hub configuration. Using a grid refinement study, they were able to obtain agreement with experiment within 15% for their tetrahedral grids, within 3% for their hexahedral grids, and they matched within 8% the drag estimates obtained by two experimental studies for various configuration changes. The configuration that they analyzed was static and did not include components such as root stubs or hardware in the analysis. A follow-on study in 2011 by Sikorsky [35] using another unstructured method confirmed the overall findings of their initial undertaking. Similar issues have been recently encountered during another experimental-computational collaboration [39].

Bridgeman and Lancaster [36] have studied a $1/5$ -scale Bell rotor hub and fuselage both experimentally and computationally. Using an extensive grid independence study, they found that total drag predictions within 5%-10% of experimental values using an unstructured solver could be achieved for the non-rotating hub-fuselage configuration. Details such as hardware, pitch links, and root stubs were included in the computational model; a breakdown of the individual contributions of these components was not part of the focus of this work. A subsequent effort [37] indicated that comparable results could be achieved by any of the unstructured solvers that

were evaluated. Hill and Louis [38] extended the methodology to rotating hubs.

Reich et al. [39] studied a notional rotor hub at 1/3 and 2/3 scale Reynolds number both computationally and experimentally at a single advance ratio of 0.2 in the Penn State University water tunnel. Measurements from experiment included total hub drag and wake diagnostics. Computations with a commercial code (STAR-CCM+) were also performed. They found that the most prominent wake structures were from the two-per-revolution (scissors) and four-per-revolution (main hub arms). Both of these vortical structures persisted far downstream of the hub, with the four-per-rev structure dissipating faster than the two-per-rev structure. The average drag predicted by the computations was under predicted by about 8% and the computational results did not include turbulent wake spectra correlations with experiment.

1.3 Grid Adaptation

Feature-based grid adaptation for unsteady problems has been applied on non-overset grids using various methodologies. Accurate predictions of hovering rotors in a single rotating adaptive mesh have been performed by several researchers [40–43]. However, these scenarios cannot be immediately applied to the prediction of rotors in forward flight where adaptation is needed in both the background inertial reference frame and the near-body rotating frame. Further, the interaction of rotors with non-moving bodies such as fuselages, wind-tunnel struts, and other configurational components also require moving-grid capability to simulate multiple motion frames. As an alternative to the overset configuration, Nam et al. [14] have demonstrated an unstructured sliding mesh approach where a rotating grid communicates with a stationary background grid. Here, articulation of the rotor blades was made possible using grid deformation based on a spring analogy, along with feature-based grid adaptation. Another non-overset based approach is the application of unstructured grid movement and deformation to enable moving body adaptation, which has been

utilized by Cavallo et al. [44] and Berglind et al. [45].

Past research efforts in overset adaptation have in many instances relied on an off-body (background) Cartesian grid-based adaptive capability. Meakin [46, 47] presented a grid component grouping algorithm with overset structured grids using a method of adaptive spatial partitioning and refinement and applied it to background Cartesian grids. Variations of this technique have been subsequently demonstrated by Henshaw and Schwendeman [48] and Kannan and Wang [49]. Canonne et al. [50] used an overset structured cylindrical grid topology to simulate rotor motion in hover where the background grid is adapted. Hybrid-solver developmental efforts have focused on rotor methodologies where two separate solvers are applied in the near body and background regions, respectively. Duque et al. [51] have employed a structured near-body and an unstructured wake grid approach to evaluate rotors using isotropic adaptation, which was limited only to the wake grid. Their work recommended the use of anisotropic adaptation to accurately capture inherently anisotropic phenomena such as tip-vortices. Sankaran et al. [52] and Wissink et al. [53–55] have successfully implemented automatic mesh refinement in the Helios solver to resolve the wake based on flow field features. This high-order Cartesian solver is coupled to a body-fitted (non-adaptive) unstructured solver to resolve complex geometries. Chaderjian et al. [56] similarly have utilized off-body adaptive grids to capture complex wake features for both forward flight and hover applications.

1.4 Error Estimation Methods

Feature-based adaptation relies on the solution features to provide local error estimates within the flow. Refinement and coarsening are performed based on a solution variable satisfying a certain criteria. Regions of high local errors indicate refinement, while regions of low errors may be coarsened to save computational effort.

Most methods in literature above for rotating body wake simulations used vorticity-magnitude based indicators to create the adaptation error-estimation criterion. However, Kamkar [57] found that non-dimensional formulations of wake features may be more suitable since they require little intuition for tuning tolerances for different flow applications and regimes.

The effect of time-dependency on the flow field is an essential aspect in applying grid adaptation to simulate time-dependent flows with dynamic motion. For steady state simulations, the classical grid adaptation algorithm is suitable, where the grid and solution are adapted iteratively until they converge at a fixed point. On the other hand, time-dependent phenomena do not possess the property of a steady state fixed point because the solution changes with time. Therefore, for unsteady flows the grid either has to be adapted frequently or a projected area of refinement needs to be prescribed around critical regions [58]. Canonne et al. [50] have shown that adapting the solution at a given frequency (based on flow time) have proven effective, with increasing frequency yielding higher accuracy. Investigations involving off-body Cartesian-based adaptive mesh refinement [48, 49, 52–55] have extended this rationale to adapt the solution at a frequency comparable to that of the solver time step, allowing for a coupled adaptive flow solver. However, this capability is not computationally efficient for mixed-element unstructured methods since they do not have the advantage of octree data structures to provide for the adaptive mechanics. Additionally, for anisotropic grids, frequent adaptation could cause interpolation errors to saturate the solution and prevent it from remaining time-accurate [59].

An alternative approach is thus required which addresses the time-dependency issue without the computational overhead of frequent adaptation. Kang and Kwon [41] present an adaptation technique that detects local maxima of a vortex core every five degrees and using a 3-D parabolic blended curve to represent the vortex core path. Nam et al. [14] describe a “quasi-unsteady” adaptive procedure for rotors in forward

flight based on a time period or interval dependent on the blade passing frequency. Cells satisfying an adaptation indicator are marked at each time step within the interval and adaptation is performed for those cells at the conclusion of each interval. Jung and Kwon [22] extended this technique for unstructured overset grids using Euler calculations.

Similar to the above schemes for unstructured grids is the transient-fixed point scheme described by Alauzet et al. [60], which is suitable for metric based adaptation methods [59,61]. However, this method, thus far, has only been applied toward Runge-Kutta based explicit time integration schemes, which only requires performing the error-estimation on the current time step and at intermediate steps. This scheme is an extension of the classical adaptation scheme in which phenomena may be adapted over a time interval comprising several time steps. The adapted grid then should provide improved accuracy for unsteady flow features over that same time interval. After grid adaptation and a more accurate simulation are obtained for a given time interval, this process is repeated over a subsequent time interval until the entire simulation is complete. Such a methodology is suitable for mixed-element unstructured grids since it eliminates the cost of adapting too frequently while providing highly accurate grids to resolve the unsteady flow phenomena. An extension of this adaptation methodology for anisotropic grids involving dynamic bodies for an explicit time integration method is delineated by Alauzet and Oliver [62].

1.5 Interpolation for Adaptive Grids

Due to recent advances in the computational capability for computational fluid dynamics (CFD) simulations, typical problems require parallel computing power involving a large number of processors. Therefore, it becomes imperative for developers to optimize codes for efficiency of any process during a simulation and ensure parallel scalability. Data interpolation over overlapping domains, such as in overset grids

or adaptive grids, require sufficiently fast search algorithm to find donors and their corresponding interpolation weights. One such class search algorithms requires identification of a node's of a host element on the donor grid. This process of searching for a node's host element is called *localization*.

Grid adaptation may be utilized in a time-dependent simulation to retain the desired time accuracy. Any subsequent grid modifications require a solution transfer of the current time step and those of backplanes for backward-differentiation formula (BDF) implicit time integration schemes from the background grid to the new grid. In order to ensure an accurate solution transfer, trilinear interpolation is a requisite, which ensures second-order accuracy in space. A further increase in accuracy may be achieved using a conservative interpolation scheme [63]. A major impediment of an accurate interpolation scheme is the ability to efficiently localize nodes of the new grid to the enclosing elements of the background grid. In addition to solution transfers for grid adaptation, applications of such localization algorithms include other CFD related concerns such as grid sequencing [64], overset grid interpolation [65], and particle tracing algorithms [66]. Noack [65] applied octree based data structures to perform the stencil search for overset interpolation. An improved method using kd-tree based data structures was demonstrated for actuator blade source interpolation by Lynch et al. [67]. Roget and Sitaraman [68] employ an alternating digital tree based data structure for overset interpolation stencil searches. Additionally, they explore subdividing the domain into blocks to help speed up the search procedure. Other efforts that use tree-based data structures to perform searches include Plimpton et al. [69] and Lee et al. [70].

1.6 Thesis Objectives

Many advances have been made in computational methods for improved simulations using different grid adaptation approaches. The overset grid method has been advantageous to obtain accurate predictions for dynamic moving bodies. Further, a non-Cartesian based fully unstructured method is amenable for simpler grid generation for body-fitted grids of complex geometries. The current state-of-the-art methods for Navier-Stokes-based overset grid adaptation have generally restricted the adaptation region to off-body grids, with near-body grids requiring some prior tuning to capture features. One exception where near-body grids are adapted is in the OVERFLOW overset structured grid methodology [71]. This method, however, is feasible mainly for simplistic geometries because body-fitting overset structured grids for more complex geometries becomes a bottleneck in the analysis of complex configurations. Additionally, among metric-based adaptation methods, previous work has been restricted to explicit time integration schemes.

This thesis focuses on necessary extensions and improvements to an unstructured computational solver to permit overset time-dependent grid adaptation for complex wake flows using BDF-based implicit time integration. A feature-based grid adaptation methodology is extended for dynamic overset grids so that accurate solutions may be obtained without *a priori* knowledge of the flow field. The applications investigated here include an RFI configuration and a rotating hub model for which sufficient experimental data are available for correlating the predicted physics. This thesis seeks to advance the state of the art for flow simulations of complex rotating systems by:

- Developing a **feature-based adaptation method for unstructured mixed-element** topologies specifically for **time-dependent, dynamic body simulations**. The unique feature presented here that is not in prior art for unstructured methods, is that all component grids (near-body and off-body grids) are subject to grid adaptation within the overset system. Also, another important key attribute is that this method addresses the added complexity of adaptation with implicit time integration schemes.
- Implementing a **cost effective parallel localization scheme** (search algorithm) that permits solution interpolation with **locally second-order spatial accuracy**. An improved scheme would perform the search in $\mathcal{O}(n)$ operations and it should be complete at least fast as many other solver tasks. The assumption here is that if the interpolation is sufficiently accurate, the interpolation errors are smaller than the spatial and temporal errors of the numerical scheme, which implies that the adaptive grids should remain time accurate to the desired order.

This thesis seeks to answer the following questions pertaining to method verification as well as validation and application:

- **Cost:** Will the new parallel localization scheme result in cost effective solution transfers? Does the observed complexity of the scheme match the proposed complexity in both number of searches and overall time, relative to other major solver tasks?
- **Accuracy:** Does the grid adaptation and interpolation method retain design order accuracy in both time and space?
- **Applicability:** How effective is the implementation of periodic or interval-based adaptation in resolving wakes of complex rotating systems without *a*

priori knowledge about the physics? Does this new method improve accuracy and converge in a relatively few number of adaptive iterations?

- **Sensitivity:** What are major sensitivities of this methodology? For which features does the method obtain good results? Is there a clear advantage of modeling the flow as either inviscid or viscous? Is there a sensitivity to the adaptation interval size?

CHAPTER II

COMPUTATIONAL TOOLS

The new computational advances presented in this thesis are demonstrated using the FUN3D unstructured-overset solver [72, 73]. The software had grid adaptation capability previously only for non-overset, steady-state simulations [74–77]. Therefore, the extension of this capability for time-dependent overset simulations was requisite. Additionally, these extensions for overset grid adaptation can then be leveraged for future development of adjoint-based grid adaptation for unsteady overset simulations. The FUN3D framework is used for demonstration of the new capabilities presented in this thesis. All improvements demonstrated in this framework should be extendible to any tetrahedral or mixed-element unstructured solver with overset capability.

2.1 FUN3D Flow Solver

FUN3D, an unstructured Navier Stokes solver developed at NASA, was selected as the flow solver with which to demonstrate the new adaptation strategy. FUN3D utilizes an implicit, node-centered finite volume scheme to solve the governing equations on unstructured, mixed-topology grids [73]. The flow variables are stored at each node of the primal (original) grid. Each node is assigned a control volume by computing the dual grid, which is composed of a polyhedral shape that depends on the number and type of elements that share that node. Both compressible and incompressible [78] Mach regimes can be resolved by the flow solver. Time accuracy is achieved using a second-order backward-differentiation formula (BDF2-opt) [79], described in Section 2.1.1. Roe’s flux-difference splitting scheme [80] is used compute the inviscid fluxes, while an equivalent central difference approximation is utilized to resolve the viscous fluxes. A red-black relaxation scheme with a point implicit procedure solves

the resulting linear system of equations [81].

2.1.1 Backward-Differentiation Formula (BDF) Time Integration

FUN3D employs the BDF family of implicit schemes for time integration. These schemes are robust for nonlinear problems and allow for large time-steps, while remaining numerically stable [79]. Unlike Runge-Kutta based explicit methods, which rely on intermediate stages between the current and the subsequent time step, the BDF schemes require information at the current time step and at a given number of previous time steps, called *backplanes*.

Hairer et al. [82] describe the BDF discretization of the differential equation $u_t = f(u, t)$ as

$$\sum_{j=0}^k \delta_j \nabla^j u_{n+1} = \Delta t f_{n+1}. \quad (1)$$

The backward differencing operator is denoted by ∇^j and δ_j are coefficients for each term. The number of terms on the left hand-side of the equation depends on the order of the scheme. For example, the BDF2 scheme is a second-order scheme requiring three levels of storage (two backplanes) and the BDF3 is a third-order scheme requires four levels of storage (three backplanes). The BDF3 scheme, however, is less stable than the BDF2 scheme.

The BDF2-opt scheme, which is applied in this thesis, is designed as a stable blend of the BDF2 and BDF3 schemes [79]:

$$BDF2OPT = (0.5 \times BDF3) + (0.5 \times BDF2). \quad (2)$$

The scheme, therefore, combines the optimal accuracy and stability properties of both schemes. It is formally second-order accurate and requires four levels of solution storage. All computational advances described in this work require consistency with BDF time integration by incorporating the complexity of backplane solutions.

2.1.2 Turbulence Models

FUN3D has several turbulence methods for unsteady Reynolds-averaged Navier Stokes (URANS) and hybrid RANS-Large Eddy Simulation (HRLES) computations. In this work, the Menter $k\omega$ -SST [83] and the Georgia Tech hybrid RANS-LES (GT-HRLES) [84] models were applied. Implementations of both these models allow for an option of the production term to be modeled by either vorticity and strain rate. Closure for the shear stress tensor is obtained from the Boussinesq approximation:

$$\bar{\rho}\tau_{ij} = 2\mu_T \left[S_{ij} - \frac{1}{3} \frac{\partial u_k}{\partial x_k} \delta_{ij} \right] - \frac{2}{3} \bar{\rho} k \delta_{ij}. \quad (3)$$

The GT-HRLES model combines the $k\omega$ -SST model with LES by calculating turbulence kinetic energy (TKE) at every grid node in the domain. The RANS calculation of TKE is complemented with an LES calculation using a subgrid-scale (SGS) kinetic energy k^{SGS} equation:

$$\begin{aligned} \frac{\partial}{\partial t} (\bar{\rho} k^{SGS}) + \frac{\partial}{\partial x_j} (\bar{\rho} \tilde{u}_j k^{SGS}) &= \bar{\rho} \tau_{ij}^{SGS} \frac{\partial \tilde{u}_i}{\partial x_j} - \\ C_\epsilon \bar{\rho} \frac{(k^{SGS})^{3/2}}{\Delta} + \frac{\partial}{\partial x_j} \left[\left(\frac{\tilde{\mu}}{Pr} + \frac{\mu_T^{SGS}}{Pr_T} \right) \frac{\partial k^{SGS}}{\partial x_j} \right]. \end{aligned} \quad (4)$$

Here, C_ϵ is a constant 0.916 and Δ is the characteristic length given by $\Delta = \mathcal{V}^{1/3}$ where \mathcal{V} is the dual volume. The SGS eddy viscosity is next computed from

$$\mu_T^{SGS} = \bar{\rho} C_\nu \Delta \sqrt{k^{SGS}} \quad (5)$$

where C_ν is a constant value 0.0667.

A blending function is then used to weight the calculation between the RANS and LES dominated regions of flow [84]. The blending function permits a RANS-LES weighted calculation of TKE and eddy viscosity (μ_T) [85].

2.1.3 Error Estimation Formulations

This thesis explores three adaptation indicators: vorticity, pressure difference, and the Q-criterion. The vorticity adaptation formulation, $F_{e,|\omega|}$, is similar to the one

applied by Duque et al. [51], which scales vorticity (ω) with an edge length, l_e . For a given edge connecting nodes n_1 and n_2 , the vorticity formulation is computed based on the averaged vorticity magnitude across the edge,

$$F_{e,|\omega|} = l_e \frac{|\omega|_{n_1} + |\omega|_{n_2}}{2}. \quad (6)$$

The pressure difference formula, $F_{e,\Delta p}$, was defined as the magnitude of the pressure difference Δp over an edge scaled by the edge length (l_e) as

$$F_{e,\Delta p} = l_e |p_{n_1} - p_{n_2}|. \quad (7)$$

The formulation of the Q-criterion indicator is based on Kamkar's non-dimensional method [57] and uses the rotation rate (Ω) and strain rate (S) tensors. Here the maximum value across the edge is applied,

$$F_{e,Q-crit.} = \max_{n_1, n_2} \left(\frac{1}{2} \left(\frac{\|\Omega\|^2}{\|S\|^2} - 1 \right) \right). \quad (8)$$

The vorticity-based and Q-criterion-based methods invoke the vorticity-magnitude Hessian to determine anisotropy, while the pressure difference adaptation utilizes the Mach number Hessian, described by Bibb et al. [75]

2.1.4 Metric-Based Adaptation

FUN3D's anisotropic tetrahedral adaptation capability [74, 86] formed the basis for the new adaptation strategy, described later in Chapter 3. This functionality has been applied to investigate several applications such as sonic-boom propagation [76, 77], viscous transonic drag prediction [77], and re-entry vehicle configurations [75]. Feature-based adaptation requires the identification of a feature or indicator, as well as a formulation for an error estimate to define the grid modification. Using any these formulations described in Section 2.1.3, the normalized local adaptation intensity, \hat{I} , is derived for each node as the maximum of the edge formulation, F_e , over all incident edges of a given node,

$$\hat{I} = \max_{edges} \left(\frac{F_e}{F_{tol}} \right), \quad (9)$$

where F_{tol} is a user-specified tolerance. The new isotropic mesh spacing is calculated using an estimate of the spacing on the original mesh h^0 , a coarsening factor C (typically around 115%), and an adaptation intensity as

$$h_1 = h^0 \min \left(C, \left(\frac{1}{\hat{I}} \right)^{0.2} \right). \quad (10)$$

The power of 0.2 is an under-relaxation parameter that controls the aggressiveness of the refinement process. This parameter relates the convergence rate of error to the grid spacing for *adjoint-based adaptation*. The value of 0.2 (or 1/5) has been found sufficient since the convergence rates for *adjoint-based adaptation* are about $O(h^4)$ – $O(h^5)$ [87]. Although there is no formal connection to the local error estimates, this value has been typically chosen as a sufficient under-relaxation parameter for feature-based adaptation [75].

Consequently, an anisotropic adaptation metric may be derived using a scalar quantity for the isotropic spacing and a Hessian to stretch the resulting mesh. The Hessian of a quantity $(-)$ can be described as

$$H = \begin{bmatrix} \frac{\partial^2(-)}{\partial x^2} & \frac{\partial^2(-)}{\partial x \partial y} & \frac{\partial^2(-)}{\partial x \partial z} \\ \frac{\partial^2(-)}{\partial x \partial y} & \frac{\partial^2(-)}{\partial y^2} & \frac{\partial^2(-)}{\partial y \partial z} \\ \frac{\partial^2(-)}{\partial x \partial z} & \frac{\partial^2(-)}{\partial y \partial z} & \frac{\partial^2(-)}{\partial z^2} \end{bmatrix}. \quad (11)$$

The Hessian is computed and decomposed into eigenvalues and eigenvectors as $H = X \Lambda X^T$. The anisotropic adaptation metric is formed by taking the absolute value of the eigenvalues of this matrix $M = X |\Lambda| X^T$. Here X are the eigenvectors and Λ are the diagonal matrix of eigenvalues. The spacing requirement h_1 is incorporated into the anisotropic framework by scaling the eigenvalues so that the largest eigenvalue (tightest spacing) corresponds to h_1 :

$$\Lambda = \begin{bmatrix} \Lambda_1^2 & & \\ & \Lambda_2^2 & \\ & & \Lambda_3^2 \end{bmatrix} = \begin{bmatrix} \left(\frac{1}{h_1} \right) & & \\ & \left(\frac{1}{h_2} \right) & \\ & & \left(\frac{1}{h_3} \right) \end{bmatrix}. \quad (12)$$

Here, the eigenvalues and eigenvectors are sorted such that $|\Lambda_1| > |\Lambda_2| > |\Lambda_3|$. Further details of the computation of the Hessian-based metric and its significance to the adaptation process are described by Park [86].

2.2 *REFINE & BAMG*

FUN3D’s metric-based adaptation as detailed in Section 2.1.4 is used for both feature and adjoint based-adaptation. The *refine* library, which is available via the FUN3D software suite, enables three-dimensional metric-based grid mechanics. Details of the parallelized adaptation mechanics, which include grid operations such as node insertion and removal by splitting or collapsing edges, edge and face swapping, and node smoothing are given in Park [86].

Bidimensional Anisotropic Mesh Generator (BAMG) [88] is a grid mechanics software for two-dimensional domains. It is used in this thesis for adaptation mechanics for the two-dimensional inviscid advecting vortex problem discussed in Section 5.1.

2.3 *SUGGAR++ & DiRTlib*

SUGGAR++ [89] and DiRTlib [90] provide overset capability with FUN3D and have been successfully used for compressible and incompressible rotorcraft applications [6, 91]. In such simulations, the background grid, which consist of the fuselage and other wind-tunnel static geometries up to the far-field boundaries, are assembled with near-body grids for each of the moving rotor blades. Some overset grid terminology are introduced here:

- **Component grid:** A constituent grid of an overset grid system that may be stationary or moving. Moving component grids require a transformation matrix.
- **Composite grid:** The assembled grid of different component grids after hole cutting.

- **Hole cutting:** Different component grids cut each other during the assembly process, regions that lie behind existing geometry become blanked in the solution process. This is referred to as hole cutting.
- **Domain connectivity:** The assignment of nodes in the composite grid as either field points, holes, fringes, donors, and orphans.
- **Transformation matrix:** Moving component grids require this matrix in order for the solver to perform the necessary grid motion and obtain updated grid coordinates. Hole cutting is performed on these updated coordinates.

SUGGAR++ enables a solver to obtain domain connectivity information. SUGGAR++ identifies the following different types of nodes in the composite grid:

- **In points:** Nodes where the governing equations are solved.
- **Holes:** Nodes that are outside the domain of interest and are, therefore, *blanked* out in the solution process. These nodes reside typically within a physical body belonging to a separate component grid.
- **Fringes:** Nodes adjacent to the holes are identified as fringes. These nodes form the inter-grid boundary definition on the given component grid. Nodes on outer boundaries of a component grid that are completely overlapped by another component grid are also marked as fringes.
- **Donors:** Nodes that are the source for solution interpolation to fringes between component grids are donors. The solution at these donor nodes along with their interpolation weights produce the interpolated value at fringes.
- **Orphans:** Nodes where no suitable donor is available are orphans. This usually arises when there is insufficient overlap between component grids. A donor suitability function permits the identification criteria for orphans.

DiRTlib assigns interpolation weights for fringe points from donor points. Both parallel and serial execution of DiRTlib is available via message passing libraries.

CHAPTER III

TIME-DEPENDENT OVERSET GRID ADAPTATION

3.1 Adaptation for Overset Grids

The overset grid adaptation capability does not restrict adaptation to any component grid. This enables each grid to evolve independently based on the flow features within each grid. The current adaptation capability for viscous flows is restricted only to nodes beyond the boundary layer for this demonstration of overset applications. FUN3D invokes *refine* (Section 2.2) in order to perform grid mechanics to satisfy the requested metric. The extension of this method to include overset adaptation requires communication with DiRTlib [90], the grid connectivity module, to assign a component mesh identification number for each node in the composite mesh. The code performs adaptation over the entire composite grid system by tracking the component mesh ID for all added nodes.

Since overset assembly of the component meshes is handled by a library outside of the FUN3D framework (SUGGAR++) [89], a generalized global index convention was requisite so that subsequent assembly of the adapted grid with its domain connectivity information would be compatible with the solution information. This process is required to perform valid solution transfers between the unadapted and adapted grid systems. The convention requires both the flow solver and adaptation code to assign composite grid global indexes by arranging nodes in contiguous fashion by mesh ID over the list of component meshes. As described in Fig. 1 (a), nodes added due to adaptation are initially assigned new global indexes by appending them to the current global index list. Node removal results in unused global indexes, which is handled by a reverse, global-index shifting procedure. In order to satisfy the condition of

Mesh ID	Original Node Indexes	Updated Node Indexes	New Node Indexes
1	1 2 3	1 2	14 15
2	4 5 6	3 4	16
3	7 8 9	5 6 7	17
4	10 11 12	8 9	18 19 20
5	13 14 15 16 17	10 11 12 13	21

(a) Node removal (struck through) and insertion

Mesh ID	Updated Node Shifted Indexes	New Node Shifted Indexes
1	1 2	3 4
2	5 6	7
3	8 9 10	11
4	12 13	14 15 16
5	17 18 19 20	21

(b) Resorted global indexes

Figure 1: Global index convention illustrated for an example overset grid system.

contiguous mesh IDs, a new procedure was introduced to re-sort the global indexes of the adapted grid system (Fig. 1 (b)). After adaptation, the component meshes are then saved, and the resultant domain connectivity information is obtained by invoking SUGGAR++ for subsequent grid assembly.

3.2 Time-Dependent Error Estimation

Time-dependent adaptation is obtained using the transient fixed-point method developed by Alauzet et al. [60] and analyzed by Najafiyazdi et al. [59]. The anisotropic metric is computed for every grid node at a given time step and is progressively accumulated over a selected time interval such that a single metric at each node is retained to form the time-dependent grid metric. The time-dependent metric is formed by collecting N_s solution samples within a time interval i as given by,

$$M_{i,max} = \bigcap_{n=1}^{N_s} M_{i,n}, \quad (13)$$

where \bigcap is the intersection operator. This equation implies that for all the samples obtained within the interval, the metric corresponding to the maximum magnitude

at each node is stored. This means that the resulting metric will correspond to the highest errors recorded at each node within the interval. Using this time-intersected metric, a new adapted mesh may be obtained suitable for a given time interval characterized by the flow phenomena obtained within the interval. Previous authors [59,60] have employed Runge-Kutta based explicit time integration and, therefore, have not addressed the complexity of backplanes. Therefore, in this method, the backplane solutions are included in the intersection process to extend time accuracy to the error estimation process. This process may be repeated over the same interval depending on the choice of feature formulations. The general adaptation procedure that uses N_{FP} adaptation formulations requires the following steps for simulation:

- Repeat the following procedure for each N_{FP} adaptation iterations:
 - Execute the flow solver for N_t time steps, collecting N_s solution samples and applying Eq. 13 to combine or intersect the metric in time. Note, since a BDF time integration is used, the intersection includes the backplane metrics of the first time step.
 - Perform grid adaptation using the resulting time-dependent metric.
 - For overset grid problems, assemble the new composite grid from the component meshes and create domain connectivity information, as described in Section 3.1.
 - For moving grid problems with BDF schemes, the coordinates at the backplanes need to be updated to obtain relevant grid motion information. The backplane transformation matrices for each moving geometry are used to update grid motion at each backplane time.
 - Revert to the solution at the beginning of the interval and interpolate this solution as well as its backplanes onto the adapted grid.
- Execute the flow solver for N_t time steps over the adapted grid.

This adaptation procedure may be repeated over N_I intervals over the entire simulation time of interest. The total number of time steps in the full simulation would be $N_t \times N_I$. In this thesis, the solution is sampled at each time step throughout the time interval ($N_s = N_t$) to obtain the time-dependent metric. The above algorithm is described in the flowchart in Figure 2. The outer loop occurs over different adaptation intervals, while the inner loop is occurs over the number of fixed-point iterations. Here, the time t_i^0 denotes the starting time of each interval i , and solution interpolation occurs for the solution at t_i^0 and its backplanes.

3.2.1 General versus Periodic Adaptation

Time-dependent grid adaptation may be performed over different interval sizes. The choice of the time interval is dependent on the application of the problem [60,61]. For example, if highly unsteady or transient behavior is of interest then a small interval is desired. However, the drawback is in the cost of the operations of the adaptation iterations, which include the grid modifications, overset grid assembly (if required), and the solution interpolation.

For certain applications, having adaptation occur over large intervals may be advantageous in order to obtain accurate solutions for a moderate cost. The number of adaptation intervals may be as low as one if this is the desired interval size of interest. This is particularly useful for rotating flows on rigid geometries where grid motion and large scale features of interest are, in general, periodic. An example would be rigid-body rotor simulations, where a periodic interval could be identified as the time corresponding $1/n_{blades}$ revolutions or its multiples. A more practical reason for having a periodic adaptation procedure is the case where an accurate solution interpolation method is not available. This would involve executing the flow solver again using the adapted grid from a freshly initialized flow field, rather than restarting from the solution saved at the beginning of the interval.




Figure 2: Modified transient fixed point iteration method for time-dependent adaptation.

CHAPTER IV

PARALLEL & ROBUST LOCALIZATION SCHEME

The previous adaptive grid interpolation capability of FUN3D is performed by *re-fine* to provide an improved initial flow solution for the restart of steady-state problems [75]. The accuracy of the solution interpolation was *zeroth-order*, which does not fulfil the requisite accuracy for a time-accurate simulation. An accurate interpolation scheme requires a sufficient interpolation stencil. The search process of each node's interpolation stencil in an adaptive grid may become a bottleneck in terms of computational cost. The exhaustive naïve search process for stencils can approach prohibitive limits for high end grids resulting in $\mathcal{O}(n^2)$ searches. Load balancing of the adapted grid will complicate the parallel search algorithm. Therefore, a novel approach is presented involving a sweeping advancing front scheme that robustly expedites localization and results in an overall cost of $\mathcal{O}(n)$ searches, where n is the node count of the grid. An advancing front scheme using neighbor walks to speed up searches is employed. Such a scheme has shown promise on serial applications to minimize the number of searches [63,92]. The key development here, is the extension of this scheme over massively distributed systems.

The localization method is parallelized and handles general overset unstructured grids over massively distributed systems. Robustness enhancements are described so that the method may be used for realistic complex geometries for dynamic grid adaptation. The complexity and cost of the method is discussed. Furthermore, since time dependent simulations may involve multiple solution transfers between grids, the accuracy of solution transfers is assessed for solver verification of design order accuracy.

4.1 *Parallelized Localization Algorithm*

4.1.1 Nomenclature

A few terms are introduced in this chapter and are defined here for clarity:

- **Background grid:** The underlying grid where the solution is defined.
- **New grid:** The collection of points (grid nodes) which require a solution interpolation from the background grid. These points are actually element agnostic, that is, they do not require the entire grid data structure for the localization algorithm to function. These points may either be grid nodes of an adapted grid, an overlapping grid in an overset system, or any general points which require some solution interpolation.
- **Localization:** The search process of all nodes in the new grid to corresponding host elements in the background grid.
- **Neighbor walk:** The march on different elements towards the background grid element that encloses a given new grid node.
- **Advancing front:** The active set of new grid nodes that require localization.

4.1.2 Barycentric Coordinates and the Neighbor Walk

Nodes of the new grid need to be localized to the enclosing tetrahedral element of the background grid. A barycentric coordinate system for tetrahedra is introduced in order to determine whether a node is inside (or local) to a background element [63]. If a node P is local to a given element, then its barycentric coordinates, β_i for $i = 1, \dots, 4$, with respect to that element are non-negative. Barycentric coordinates are useful for providing interpolation weights for simplices such as line segments, triangles, tetrahedra, etc. An illustration of barycentric coordinates in two dimensions (triangle) is given in Fig. 3. The barycentric coordinates for node P are $(\beta_a, \beta_b, \beta_c)$ as defined




Figure 3: Barycentric coordinate system on a triangle with associated areas A , B , and C corresponding to points **a**, **b**, and **c**, respectively.

by:

$$\beta_a = \frac{A}{A + B + C} \quad (14)$$

$$\beta_b = \frac{B}{A + B + C} \quad (15)$$

$$\beta_c = \frac{C}{A + B + C}, \quad (16)$$

where A , B , and C are the areas of sub-triangles formed by dividing triangle **abc** at node P . These barycentric coordinates are normalized such that $\beta_a + \beta_b + \beta_c = 1$

The same analogy can be extended to a tetrahedron, which is a three dimensional simplex. For classical trilinear interpolation, these coordinates also are the interpolation weights:

$$\phi_{interp}(P) = \sum_{i=1}^4 \beta_i(P) \phi_i, \quad (17)$$

where P is the point that requires interpolated solution and ϕ_i is the solution (or any variable) on the nodes of a tetrahedral element in the background grid. An example of interpolation in the two-dimensional analogy is considered if node P is close to the

centroid of triangle **abc**. In this case, the interpolation weights will be nearly equal from points **a**, **b**, and **c**. If node P is close to point **a** then the coordinates will reflect that and the interpolation weight of point **a** dominates the other points.

This barycentric coordinate-based trilinear interpolation is requisite for a solution transfer of error order 2 [63]. Furthermore, an estimate of the interpolation error can be obtained given the new grids' coordinates $\mathbf{X}_{new} = (x, y, z)_{new}$ and the barycentric interpolation weights. The deviation $\delta_{\mathbf{X}}$ of the interpolated grid coordinates of the new grid is

$$\delta_{\mathbf{X}} = \|\mathbf{X}_{interp} - \mathbf{X}_{new}\|, \quad (18)$$

where \mathbf{X}_{interp} is found using Eq. 17 with $\phi = \mathbf{X}_{old}$. In theory, since grid coordinates are linear functions in space, this deviation should be machine zero. However, due to round-off errors and the accuracy of barycentric interpolation on highly skewed elements, it may result in small non-zero values. Therefore, as an interpolation check, the maximum deviation is computed for all solution transfers and is monitored.

The neighbor walk is performed by making steps to adjacent elements until the all the barycentric coordinates are non-negative. An example of a neighbor walk for one node is shown Fig. 5. A step to an adjacent element can be performed depending on the number of negative barycentric coordinates:

- If only one barycentric coordinate is negative, say β_i , indicating that node i has a negative coordinate, the step should be made to the neighboring element adjacent to the face opposite node i .
- If more than one barycentric coordinate is negative, then a choice of steps is available to make the step. A random selection is performed in this case, which also serves as a robustness feature, ensuring that should the neighbor walk reach the same element due to a closed loop, a cyclic process can be terminated by

this random selection process.

The determination of the localization of a node and (if required) the neighbor step is described in Fig. 4. The background grid's element adjacency information is first populated in order to perform neighbor walks. This enables a given element to record the elements adjacent to it. The search is effectively confined to one-dimension, which makes it more efficient than a multi-dimensional search algorithm. The neighbor walk, however, is susceptible to search failure in concave domains, which is described in Section 4.2. If for any reason the neighbor walk search fails, a default method may be used to complete the search.

4.1.3 Advancing Front

In order to optimize the localization for the entire new grid, given the grid connectivity information, an advancing front approach is utilized. This approach enables a fast search process by keeping the search space small [92]. This approach requires the population of the new grid's node neighbors; no information about the new grid's elements is needed. The nodes to be localized sweep as an advancing front [92] across the new grid until all localizations are complete. The idea is analogous to the advancing front algorithm of unstructured grid generation. A marking system is introduced where in each node of the new grid is marked as either *active*, *untouched*, or *complete*. A front needs to be initiated by marking at least one node as *active*. The nodes used for front initiation may be situated anywhere in the domain. A *guess element* is assigned to each activated node. Once the node is localized, it is marked *complete*, and all of its *untouched* node neighbors are marked *active*. The neighbors are also inherit a *guess element*, the enclosing element of the node that has been localized.

Figure 6 describes the advancing front sweeping algorithm in two dimensions. For example, consider a random node in the middle of the domain that is assigned to be






Figure 4: Determination of node localization and the neighbor step based on barycentric coordinates (two-dimensional representation).




Figure 5: Neighbor walk algorithm (two-dimensional localization).

the first node to be localized. The front at this point comprises this single node. After it gets localized using either the neighbor walk or a default search method, it gets marked *complete* and its neighbors are activated, as seen in Fig. 6(a). These newly activated nodes inherit the *guess element* of that first node. Following the localization of those nodes on the current front, all their node neighbors will be activated and the front will be advanced as depicted in Fig. 6(b).

4.1.4 Parallel Paradigm

In the parallel environment, each processor handles a partition of the background and new grids. A separate load balancing of the new grid in many instances require the neighbor walk and advancing front algorithms to be able to make off-processor queries in order to complete the localization. For these off-processor queries, the *client* is the processor that manages the new grid's *active* node and the *server* is the processor or partition of the *guess element*.

In order to parallelize the algorithms, a new vector *guess partition* is introduced to




Figure 6: Advancing front iterations depicted in two dimensions.

tag the partition number of the node's *guess element* until the localization is complete. The *guess element* and *guess partition* vectors span all local nodes and ghost nodes for a given partition. The values of *guess partition* can be any number between one and the number of processors over which the grid is distributed (inclusive of the current partition). If a given neighbor walk requires crossing a partition boundary, the *guess partition* is updated with a partition number of a ghost node on the failed element face. Since no information is available locally at this point to provide a *guess element* on that partition, the global index of that ghost node (which is available) is stored, and the query is tagged as an non-local neighbor walk. Upon completion of the localization, three vectors need to be populated to make subsequent interpolation straightforward: the *enclosing element*, *enclosing part*, and the *weights*.

The parallel paradigm is divided into two processes, namely, *parallel advancing front* and *parallel activate front*. These two processes use collective communication calls between all servers and all clients. Therefore, even if the server and the client processor are the same, information is passed via parallel collective communication. These calls are repeated until all nodes (globally) are marked *complete*. All processors

get synchronized at the end of every such iteration. The *parallel advancing front* procedure works on all *active* nodes and performs the following operations:

- Sends the coordinates of all nodes marked *active* from the client partition to its respective *guess partition* (server) to perform the neighbor walk on that partition.
- If a query is tagged as a non-local neighbor walk, the global node index is converted to a local index using a lookup operation. Next, the background grid's *node-to-element* information is used to provide a *guess element* to resume the neighbor walk. One random element that surrounds that node is selected. The randomness of the element picked in the off-partition query causes the neighbor walk to not always agree with the element that would have been selected had the partition boundary not existed. Therefore these non-local neighbor walks result in an equal or higher number of total searches.
- Once all neighbor walks are exhausted, that is, they either are localized or encounter another partition boundary, then the *guess element*, *guess part*, and its *weights* are sent back to the client partition to get updated.
- If a node is marked *complete*, the *enclosing element*, *enclosing partition*, and its associated *weights* are stored.

The *parallel activate front* procedure operates on the new grid and serves to advance the front beyond the current partition by the following operations:

- Updating the *guess element* and *guess partition* information on ghost nodes. This is performed by the same collective communication calls as earlier, but only on the ghost nodes.
- If a ghost node is marked *complete*, then all local *untouched* nodes that are its neighbors are marked *active* and neighbor's *guess element* and *guess partition*

are assigned to the node.

4.2 Robustness Features to Overcome Search Failures

Search failures may occur either because of unavoidable cyclic searches [63] or because these one-dimensional neighbor walks can encounter geometries in a concave domain. Concave domains are formed due to complex geometrical protrusions. Because the neighbor walk is effectively one-dimensional in space, a successful neighbor walk requires a nearly direct path between the *guess element* and the node via adjacent elements. Therefore, geometrical protrusions in a complex domain may create search failures for neighbor walks. These search failures may be overcome using the robustness features described in this section.

4.2.1 Hierarchical Localization

The second improvement addresses a problem encountered when concave geometries are present. Since the neighbor walk relies on a geometrically linear search from the initial *guess element* to the final enclosing element, it relies on grid contiguity for each neighbor walk. This breaks down when dealing with concave domains which result in the linear neighbor walks encountering geometry boundaries. This is encountered in most realistic grids for aerospace applications. One way to overcome these geometry failures is to queue these searches and re-seed them with guesses from neighbors that have successfully localized. A more robust method is to initiate the front differently by activating nodes in hierarchical fashion starting from geometry boundaries and progressing into the volume. This prevents any neighbor walks from encountering geometry boundaries during the search. Nodes are activated and localized in the following sequence:

- Corner nodes, which coincide with three or more boundaries. An exhaustive search can be performed in order to ensure that these critical points in the domain are successfully localized.

- Edge nodes, which coincide with two boundaries.
- Face nodes, which coincide with one boundary and volume nodes, which comprise the rest of the nodes in the domain.

This hierarchical approach aids in both robustness as well as efficiency. Efficiency is increased in each stage of the hierarchy since the searches become increasingly shorter.

4.2.2 Default Search Methods

In certain situations, the criterion for localization fails and results in cyclic neighbor walks [63]. Cyclic neighbor walks occur when the same element is visited multiple times. The computation of barycentric coordinates becomes inaccurate for highly anisotropic tetrahedra with roundoff error. The inaccuracy results in very small negative values in the barycentric coordinates, when the node has already found its true host element.

The kd-tree node search method, implemented in FUN3D for actuator-blade simulations [67], is used when the above localization algorithm fails. The method identifies the nearest background grid node to the new grid node that needs to be localized. By identifying this nearest node, a search across all elements associated with that node enables identification of the host element. The complexity of this method is $\mathcal{O}(N_{nodes}^{new} \times \log(N_{nodes}^{new}))$, where N_{nodes}^{new} is the (remaining) new grid nodes that need to be localized. This is more expensive than that of the present localization algorithm, which is effectively $\mathcal{O}(N_{nodes}^{new})$, but it is much cheaper than that of the exhaustive search method. If, for any reason, there are a few nodes that are still not localized to their host element, then an exhaustive (naïve) search across all elements on all partitions may be performed.

4.3 *Generalization to Overset Mixed-Element Unstructured Grids*

Overset grids are handled by localizing each new grid node of a given component to an element belonging to the same component in the background grid. This occurs in sequence for each component grid in the overset system. Mixed element grids, which are commonly encountered in Navier-Stokes problems, need to be compatible with the localization method. Since the existing algorithm relies on tetrahedral background elements in order to perform the neighbor walk and create interpolation weights based on barycentric coordinates, the background grid is converted into a pure tetrahedral grid. This conversion is only performed in memory, and the original mixed element grid is not overwritten. Only the background grid's *element-to-node* data structure, which maps each element to its corresponding nodes, needs to be converted to its tetrahedral counterpart. Non-tetrahedral elements such as prisms, pyramids, and hexahedra are converted into tetrahedra by the method described by Dompierre et al. [93].

4.4 *Cost Analysis*

A simple box grid was used to test the algorithms to ascertain the complexity and provide timing results varying both the size of the grid and the number of processors. Adaptation was not attempted on the grid in order to keep the number of nodes the same between the background and new grids. Grid modifications in the form of edge swaps and grid smoothing were performed that keep the number of nodes constant [86]. Additionally, the new grid was load balanced separately from the background grid so that the parallel algorithm could be tested. A solution vector of conserved variables with only the current plane of information (such as in a steady-state problem) comprised the information to be interpolated from the background grid to the new grid.

Four isotropic grids of 1,000 nodes, 125,000 nodes, 1,000,000 nodes, and 27,000,000 nodes were utilized in the study. The typical simulation in FUN3D is on the order of the finest grid used here, so that the cost of a realistic problem may be obtained from these results. The number of processors ranges from 2^j for $j = 3, \dots, 9$ (from 8 up to 512 processors). Since smaller grids are not decomposed over a large number of partitions and large grids are not decomposed over a small number of partitions, only a partial test matrix of all the grids and decompositions is studied.

The cost analysis was performed on an SGI ICE Altix 8200 cluster, comprising of dual socket quad core processors with an interconnection speed of approximately 16 Gbit/s. Tests were repeated using 5 trials to ascertain that the randomness of the neighbor walk did not produce large variation and to verify that consistency was not sacrificed due to CPU loading.

4.4.1 Algorithm Complexity

Alauzet and Mehrenberger [63] state that the sequential algorithm complexity (total number of searches) should be $\mathcal{O}(\mathcal{C} \times N_{nodes}^{new})$. Here \mathcal{C} is the average number of background grid elements visited and N_{nodes}^{new} refers to the number of nodes on the new grid. A correct implementation of the parallelized scheme should obtain the sequential complexity regardless of the number of processors.

The total number of searches is the number of total elements visited by all new grid nodes during localization. When this is plotted against the number of processors (Fig. 7(a)), the trend is nearly constant for all the grids. Variations are partly due to randomness of the neighbor walks and result as high as 10% for the 1,000 node grid and within 1% for the 27,000,000 node grid. Another interesting result to note is that the average number of visited elements \mathcal{C} for all grids (except for the 1,000 node grid) is consistently between four and five (see Fig. 7(b)). This is noted to be optimal since the advancing front algorithm progressively decreases the search region

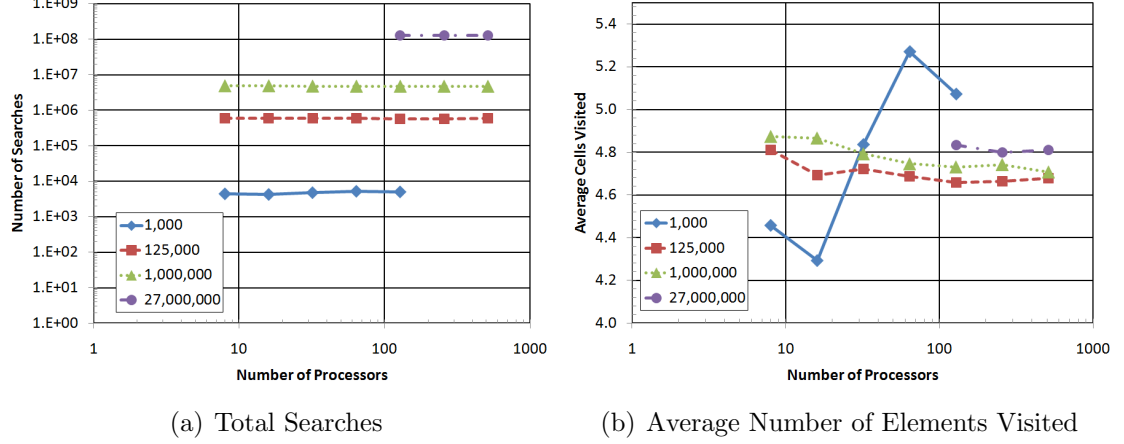


Figure 7: Complexity verification of different grid sizes.

and, therefore, reduces the number of visited elements. As predicted by Alauzet and Mehrenberger [63] for the sequential case, \mathcal{C} is in the order of the number of elements of the background grid that are overlapped by a node of the new grid, is now verified for the parallel algorithm regardless of the number of processors.

The discrepancy found in the trend of average number of visited elements for the 1,000 node grid is related to the number of partitions and the shapes of the partitions. Since the number of nodes per partition range from 125 to as low as 8, the number off-processor queries become more common since neighbor walks encounter more partition boundaries. As stated earlier, these non-local neighbor walks select random elements surrounding a partition boundary node in the vicinity of the actual element that should be picked. This causes the search cost to increase slightly. The non-monotonic behavior with the number of processors is attributed to the decomposition method, which is not tuned for optimizing the neighbor walks.

4.4.2 Timings

Localization was timed by including the preprocessing tasks of populating the background grid's element adjacencies and the new grid's node neighbors in addition to the advancing front operations. The interpolation was timed separately in order to

determine the benefits of the localization to the unit interpolation request. Only the one plane of the solution is interpolated, but projections may be made for multiple backplanes that are necessary for time-accurate solution transfers.

Timing results in the form of wall time for localization and interpolation are given in Fig. 8. Because the algorithm complexity is essentially constant regardless of the decomposition, the time was expected to decrease monotonically as the number of processors increases. At around 100 ms, an increase in the time required is observed for all grids, and this behavior is attributed to the interconnection speed between processors. For the 1,000,000 node grid, increasing the number of processors beyond 128 yields smaller time reduction. Increasing the processors beyond 32 for the 125,000 node grid shows nearly constant wall time. A very similar behavior of diminishing return at around 64 processors was also reported in interpolation timing studies conducted by Plimpton et al. [69].

Fitting a power function on the 27,000,000 node grid yields a power of -0.89 for the localization and a power of -0.82 for the interpolation. These fits were ascertained when they were scaled for the 1,000,000 node grid, yielding comparable values for decompositions up to 64 processors. If parallel communication is ideal, then the power value should be -1 [69]. Therefore communication overhead is the primary cause of this degradation in the optimal speed.

Another important result that can be derived from Fig. 9 is that for grids above 1,000,000 nodes, the time spent in interpolation is around 10% the time spent localizing the grid. If three backplanes in addition to the current plane are to be interpolated for a given time-accurate scheme, the time required to interpolate the solutions will require about 40% of the time spent localizing the grid. For smaller grids, this result is not expected due to the time increase that occurs in the region of diminishing return.

The localization power function fit obtained earlier, was used to obtain projections




Figure 8: Timing comparison on different grid sizes demonstrating the efficiency of the localization algorithm




Figure 9: Ratio of time spent during interpolation versus localization

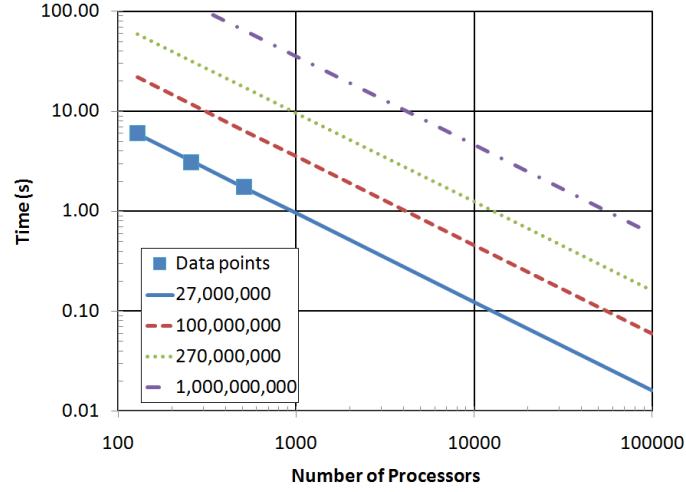


Figure 10: Localization time estimates for finer grids

of localization wall time for an upper-end FUN3D simulation. Scaling a grid of one billion nodes (Fig. 10) on 512 processors is estimated to require approximately one minute for localization.

CHAPTER V

VERIFICATION OF DESIGN ORDER ACCURACY

5.1 *Two-Dimensional Inviscid Advecting Vortex*

Accuracy assessments have been performed to ensure that the solver's spatial and temporal design order accuracy is maintained using this interpolation method. The objective was to ascertain that both the accumulation of interpolation error and spatial and temporal discretization errors conform to the design order in space and time. The numerical scheme was chosen is second-order in space and second-order in time (BDF2-opt) [79]. The verification case is a two-dimensional inviscid advecting vortex which has an exact solution at any given time. This case was studied extensively by Shearer and Scott [94]. The vortex aligned along the y -axis. Defining the vortex center location as (x_v, z_v) , the vortex velocity and pressure in time and space is given as:

$$u = 1 - \frac{C(z - z_v)}{R^2} \exp(-r^2/2) \quad (19)$$

$$w = \frac{C(x - x_v)}{R^2} \exp(-r^2/2) \quad (20)$$

$$p = p_\infty - \frac{\rho C^2}{2R^2} \exp(-r^2/2) \quad (21)$$

$$r^2 = \frac{(x - x_v - t)^2 + (z - z_v)^2}{R^2}, \quad (22)$$

where u and w are the velocity components in the x and z directions, respectively, p is the static pressure, and R is the core radius. Time t appears in the exponential portion of velocity and pressure. The quantity C is defined such that $C/(U_\infty R)$ is the nondimensional vortex strength parameter. In this simulation, this parameter is set to 0.02 and the freestream Mach number M_∞ is set to 0.1. The definition of vorticity




Figure 11: Vorticity ω_y profile of the exact solution.

based on the velocities is

$$\omega_y = \frac{\partial u}{\partial z} - \frac{\partial w}{\partial x} \quad (23)$$

$$= -\frac{C(z - z_v)}{R^2} \exp(-r^2/2) \left[\frac{(x - x_v)^2 + (z - z_v)^2}{R^2} \right]. \quad (24)$$

The exact vertical variation in vorticity through the center of the vortex at any given longitudinal location is plotted in Fig. 11. The variation of vorticity spans from large negative values in the core, transitioning to positive values outside of the core (until two core radii), and, finally, an asymptotic decrease to zero away from the vortex region.

The incompressible Navier-Stokes equations were used to model the advecting vortex. The vortex moves until $t = 24$ and the root mean square (RMS) of the swirl velocity error (w momentum equation) is monitored. The monitored error is volume-weighted and normalized so that it becomes a suitable metric for the non-uniform adapted grids. The formula for this error metric is

$$e_{RMS} = \sqrt{\frac{1}{\mathcal{V}_{tot}} \sum_{i=1}^{N_{nodes}} (e_i^2) \mathcal{V}_i}, \quad (25)$$

where e_i is the error of the swirl velocity, \mathcal{V}_i is the node's dual volume, and \mathcal{V}_{tot} is the total domain volume. This error is the total spatial and temporal error, and, should approach zero as h and Δt approach zero ($O(h^2, \Delta t^2)$).

Grid adaptation using the vorticity formulation, as described in Section 2.1.4, with $F_{tol} = 0.003$, is applied so that at least 10 nodes to resolve the vortex core. This method produces a time-intersected metric, and grid mechanics were performed by the two-dimensional meshing software Bidimensional Anisotropic Mesh Generator (BAMG), described in Section 2.2. As many as three cycles of grid adaptation (four levels including original grid) are performed using a constant prescribed complexity for each cycle. The drop in error was nearly one order of magnitude by the third adaptation cycle, and this was considered sufficient to be able to perform the accuracy assessments. This prescribed complexity is a natural way of controlling the grid growth for metric-based adaptation [62]. Effective grid spacing $h_{eff} = N_{nodes}^{-1/2}$ is tabulated as a measure of the spatial refinement for each cycle of adaptation in addition to the number of grid points N_{nodes} used.

While the error of the swirl velocity has been monitored for comparison, the error estimation process based on the vorticity formulation does not imply that errors in swirl velocity will converge to its exact solution. There is no direct association between local errors of vorticity and the error in the swirl velocity. However, since the vorticity formulation is a classical feature-based error estimation method for vortex dominated flows [57], it is a suitable method for this case. Further, monitoring the swirl velocity error permits an analysis of this error estimation method to an unrelated global quantity of interest, which is typically utilized as a convergence metric in complex three-dimensional configurations.




Figure 12: Total error comparison of adaptation interval size after 24 characteristic periods (300 time steps). (Note: Second-order slope only accounts for spatial error).

5.2 Adaptation Interval Study

The number of adaptation intervals was varied to evaluate if interpolation errors degrade solution accuracy. This study was performed for the temporal resolution requiring 300 and 600 time steps. The interval length size was divided such that it was simulated using one interval (requiring no solution transfer), two intervals (one solution transfer), and four intervals (three solution transfers). Figures 12 and 13 plot the errors versus effective spacing and grid size for 300 and 600 time steps, respectively. Clearly, as the grid is more refined, the error reduction is more substantial as the number of intervals increase. This is partly attributable to prescription of a constant complexity at each adaptation cycle. More importantly, the accumulated error follows the expected second-order slope with adaptation for both time steps analyzed.

By applying Eq. 18, all interpolation deviations are ascertained to be less than 10^{-14} , which is close to machine zero (double precision). Therefore, these studies show no indication of interpolation errors within the total error.




Figure 13: Total error comparison of adaptation interval size after 24 characteristic periods (600 time steps). (Note: Second-order slope only accounts for spatial error.)

5.3 Uniform versus Adaptive Refinement

Adaptive grid refinement should accelerate error convergence with grid size or spacing for a consistent method, where interpolation does not dominate. By studying different uniformly refined grids, the expected second-order accuracy should be obtained as long as it is in the asymptotic range of convergence. Uniform refinement is performed for four levels as shown in Table 1. The effect of uniform refinement on the error and its comparison against adaptive refinement (two intervals) is plotted in Figure 14 using 600 time steps. The slope of uniform refinement converges to nearly second-order by the second level refinement. However, using the adaptive method, slopes higher than second-order are obtained immediately illustrating accelerated convergence. Since the temporal error is not isolated from the spatial error, both curves show degradation of the slope at smaller spacings.

Table 1: Grid spacings for uniform refinement study.

h	N_{nodes}
0.060	1884
0.030	7577
0.015	28769
0.0075	115824

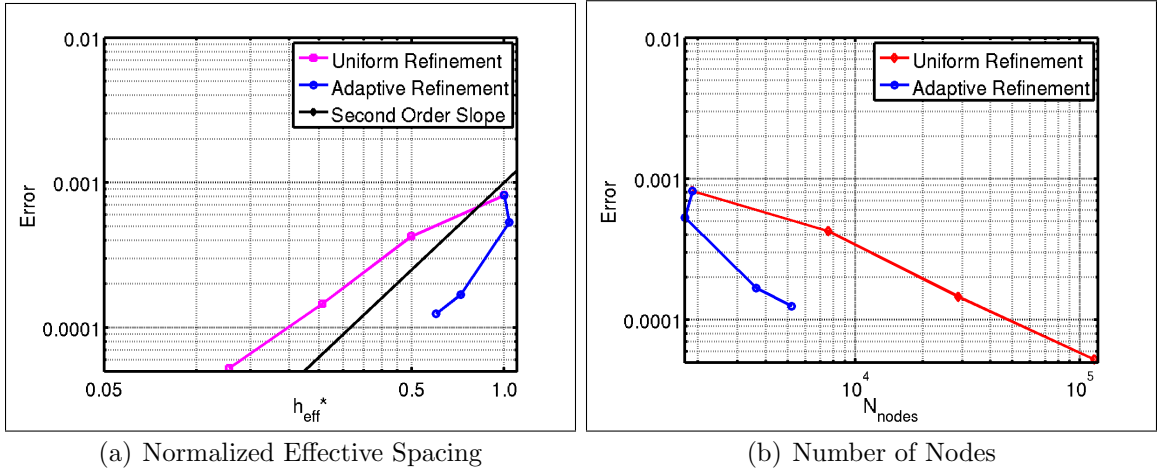


Figure 14: Total error comparison between uniform and adaptive refinement (2 intervals) after 24 characteristic periods (600 time steps). (Note: Second-order slope only accounts for spatial error.)

5.4 Temporal Resolution

An additional time-step refinement study was performed to assess the effect of temporal resolution on the accumulated error. This study was performed using two adaptation intervals which involves one solution transfer. The progression of the vortex as it advects is shown over four adaptation intervals in Figs. 15 and 16. The contours (Fig. 16) show a smooth distribution of vorticity levels in the flow field, which the mesh contours (Fig. 15) do not clearly indicate. In these figures, the grid is in its final adaptation cycle and clearly shows that the grid in each adaptation interval is tuned to preserve the vortex core.




Figure 15: Vorticity mesh contours in each adaptation interval (adaptation cycle 3).

Using different time step sizes, as tabulated in Table 2, the effect of the temporal resolution is observed in Fig. 17, where two adaptation intervals are performed. In all these figures, the effective grid spacing is presented as a normalized quantity h_{eff}^* . The spacing is normalized by the spacing original uniform (isotropic) unadapted grid




Figure 16: Vorticity contours in each adaptation interval (adaptation cycle 3).

$h_{eff} = 0.06$ (1884 nodes). Therefore, original grid is plotted at $h_{eff}^* = 1$, and the three subsequent adaptation cycles are the next three data points each of the lines. For each time-step size, the first adaptation cycle results in a net coarsening since h_{eff}^* increases (N_{nodes} decreases). Each subsequent cycle obtains a net refinement since h_{eff}^* decreases. Temporal effects clearly increase as more adaptation cycles are performed. By comparing the slope with the theoretical second-order slope, the spatial second-order accuracy is ascertained starting at temporal resolutions of 300 time steps or greater. Further, the accumulated error becomes less sensitive to time step choice at this point. The degradation of the error at small time steps is because the temporal error is not isolated from the total error in Fig. 17.

One approach to isolate the temporal error from the spatial error for the solution

Table 2: Time step sizes for temporal resolution study.

Δt	N_{steps}
0.16	150
0.08	300
0.04	600
0.02	1200

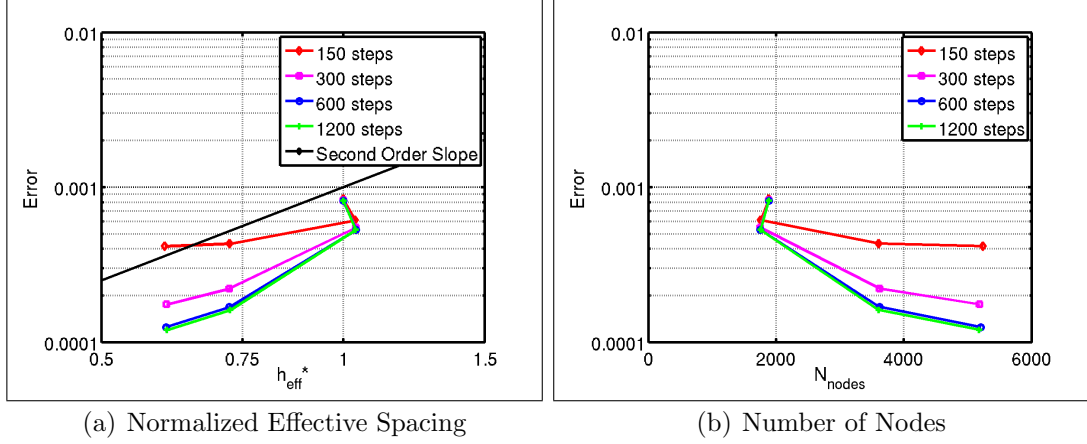


Figure 17: Total error comparison of temporal resolution after 24 characteristic periods. Two adaptation intervals are used in these simulations. (Note: Second-order slope only accounts for spatial error).

$\phi(h, \Delta t)$ is by using the following decomposition, as described by Hay et al. [95]:

$$\begin{aligned}
 e(\phi) &= \phi_{ex} - \phi(h, \Delta t) \\
 &= e_t(\phi) + e_s(\phi) \\
 &= [\phi(h, *) - \phi(h, \Delta t)] + [\phi_{ex} - \phi(h, *)],
 \end{aligned} \tag{26}$$

where $e_t(\phi)$ is the temporal error, $e_s(u)$ is the spatial error, and u_{ex} is the exact solution. The spatial error can be driven to zero by defining $u_{h,*}$ as the solution as the time step approaches zero:

$$\phi(h, *) = \lim_{\Delta t \rightarrow 0} \phi(h, \Delta t). \tag{27}$$

Therefore, the temporal error can be found by subtracting the spatial error from the

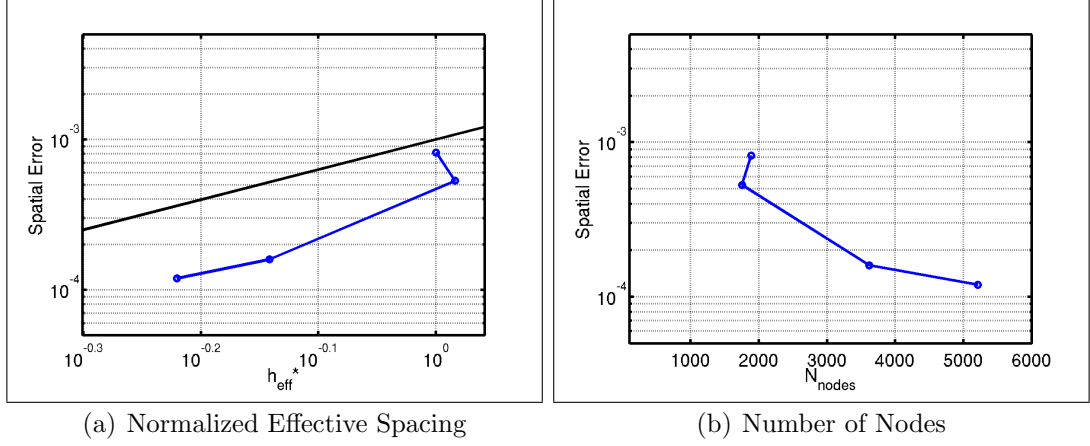


Figure 18: Spatial error comparison after 24 characteristic periods. Two adaptation intervals are used in these simulations. Second-order slope is plotted in black.

total error: $e_t(\phi) = e(\phi) - e_s(\phi)$. The spatial error can be obtained by taking a very small time step and subtracting its total error from a given solution's total error. Another method is to use a higher-order temporal scheme as the reference solution when computing the difference. An alternate method to isolate the spatial error is to perform a Richardson extrapolation of the errors obtained at different time steps and to extrapolate the error as the time step approaches zero [96], which has been applied here.

Since this study varies both the space (adaptation cycles) and time discretizations, the spatial and temporal errors can be decomposed correctly here and comparisons can be made against the design order slope. Figure 18 plots the spatial error that was projected using Richardson extrapolation from results of the temporal resolution study (Fig. 17). The error drops at a higher rate than second-order for the first two adaptation cycles but the error drop rate decreases at the third cycle.

The temporal error reduction for the different adaptation cycles is plotted in Fig. 19. Overall, the slopes agree with the second-order slope and at times approach third-order behavior. The temporal error is higher at the coarser time steps as the spatial refinement increases. As the time-step is refined, this pattern does not persist,

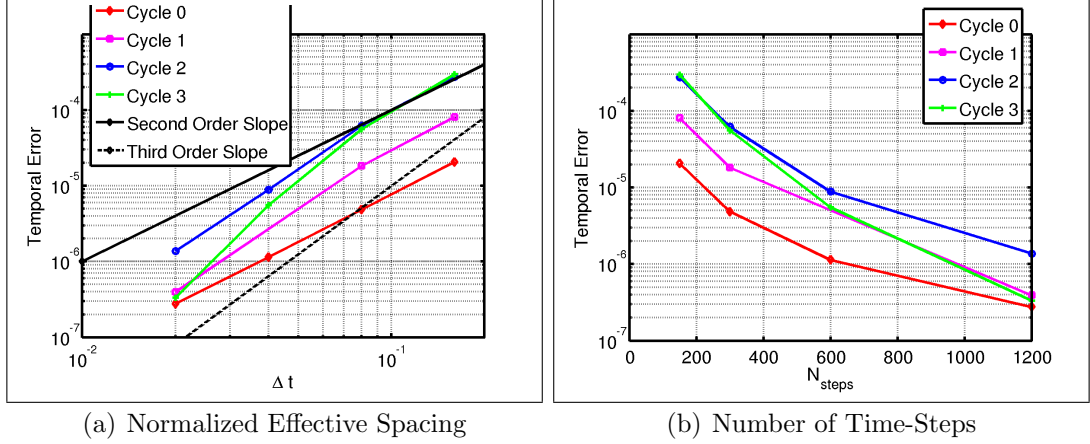


Figure 19: Temporal error comparison after 24 characteristic periods. Two adaptation intervals are used in these simulations.

indicating a possible interdependence of time-step and spatial refinement.

5.5 Temporal Evolution of Error

When performing solution interpolation between the different time intervals, small errors may be introduced that were not discussed in the error decomposition described in Eq. 26. The interpolation error that is associated with a particular quantity of interest may be quantified by plotting its temporal history. Since vorticity-based adaptation is performed, monitoring the vorticity field over each adaptation interval will quantify the error evolution.

Figure 20 describes the RMS error (Eq. 25) of vorticity as a function of the characteristic time of the vortex. Since four equal intervals are used in the simulation, three solution interpolations occur at characteristic times of 6, 12, and 18. The error exhibits a 20% discontinuous transient at each adaptation indicating that there is a component of interpolation error involved. The error increase in these discontinuities is approximately 20% of the total error at the end of the simulation. During each interval, the error varies up to 25%, neglecting the initial transience. The transient error damps out and is within the overall variation (25%). Table 3 describes the

statistics of the vorticity errors plotted in Fig. 20. The mean of the error in each interval does not grow with time indicating that transient interpolation errors are damped. With exception to the first interval which undergoes a initial transience from zero error the standard deviation of the errors are within 10% of the mean of the error in each interval.




Figure 20: Time history of vorticity error. Solution interpolation occurs at characteristic times of 6, 12, and 18.

Table 3: Vorticity Error Statistics.

Interval	Maximum	Minimum	Mean	Standard Deviation
1	0.0202	0.0019	0.0168	0.0037
2	0.0214	0.0165	0.0189	0.0013
3	0.0231	0.0164	0.0182	0.0015
4	0.0226	0.0153	0.0170	0.0015

CHAPTER VI

ROTOR-FUSELAGE INTERACTION DEMONSTRATION

A rotor-fuselage interaction (RFI) configuration is studied for validation analysis of the new capabilities introduced in this work. This test case has been extensively evaluated at the Harper Wind Tunnel at the Georgia Institute of Technology [16]. This model has been evaluated by numerous prior computational efforts [6, 13, 15, 18, 19, 21–23, 97] with a variety of computational approaches, including overset grid adaptation. A common limitation of all these analyses was the inability to resolve the physics of vortex-fuselage impingement, which is an important component to the analysis of such complex interactions, in part due to the choice to model an inviscid fluid [6, 21]. In addition, O’Brien [6] recommended that feature-based adaptation would provide improved results while minimizing computational cost. This chapter addresses this shortcoming with the computational enhancements introduced in this thesis.

6.1 Georgia Institute of Technology Configuration

The Georgia Institute of Technology RFI configuration, as depicted in Fig. 21, comprises a cylindrical fuselage and a hemispherical nose to permit easier identification of RFI. The rotor blades have a rectangular planform with a NACA-0015 airfoil section. The rotor blades were designed to be very stiff [16] which allow for computational analyses to neglect elastic blade deformations, which was assumed by all prior efforts [6, 13, 15, 21–23]. One advance ratio, $\mu = 0.10$, was selected for investigation since comprehensive experimental data are available for this condition and because most prior efforts have used the same test condition for demonstration. The relevant

blade angles and thrust coefficient are reported in Table 4. Data from this effort include instantaneous and time-averaged pressures along the fuselage, as well as vortex behavior via laser light sheets. The fuselage station is non-dimensionalized (x/R) by the rotor radius ($R = 0.4572$ m) for ease in presentation.




Figure 21: Model of the GIT RFI configuration.

Table 4: Flow conditions and rotor parameters.

Variable	Value
μ	0.10
V_{tip} (m/s)	100.5
c (m)	0.086
Re_c	577,870
β_{1s} ($^\circ$)	-2.02
β_{1c} ($^\circ$)	-1.94
C_T	0.009045


O’Brien [6] used this as a validation case for his series of actuator to overset rotor models implemented into FUN3D. Numerous details of the fuselage pressure coefficients have not been captured by these methods, in spite of the simplistic model geometry. O’Brien [6] noted that some time-averaged features just aft of the rotor were captured when the entire model (rotor strut and hub geometry) were included. Since the model investigated here does not include the hub pin geometry, the pressures aft of the rotor region ($x/R > 1$) are not expected to exhibit improved correlation. However, the major goal of this chapter is to accurately resolve the vortex-fuselage

impingement observed in the original experiments near the nose ($x/R \approx 0.45$) [16] and ascertain the reasons for the lack of resolution in prior efforts.


The vortex-fuselage impingement process is sketched in Fig. 22. The evolution of two tip vortices shed from the two blades are depicted. Brand [16] observed that the tip vortex from the prior blade ($\psi = 60^\circ$ in Fig. 22 (a) and $\psi = 90^\circ$ in Fig. 22 (b)) dissociates into two circulatory regions due to the interaction with the following blade. These vortical features create primary and secondary vortex-fuselage interactions. The primary interaction is the remnant of the original tip vortex after it interacts with the following blade, while the secondary interaction arises from the blade's vortex sheet [16]. Since the vortex sheet is characterized by opposite circulation to the primary tip vortex, they become counter-rotating circulatory regions. As the flow progresses, as shown in the sketch, Brand reported that the primary vortex interaction weakens and remains relatively stationary at $x/R \approx 0.45$, while the secondary interaction is rapidly swept downstream. He explained this phenomena by image vortices representing the fuselage surface causing different streamwise reactions: an upstream component for the primary interaction and a downstream component which accelerates the secondary interaction [16]. This process repeats periodically every half a revolution (180°) as each subsequent tip vortex encounters the same evolution.

6.1.1 Approach

Periodic adaptation was initially performed in Section 6.2, in which the metric was sampled for a dynamic motion of 180° blade sweep or 180 steps after the solution became periodic (after two revolutions). This may be performed over multiple adaptation iterations. After adapted grid was obtained, the flow field was re-initialized at $t = 0$ and recomputed assuming the periodic nature of the large scale features would be accurately captured by the resulting grid.



(a) $\psi = 240^\circ$ ($\psi = 60^\circ$)



(b) $\psi = 270^\circ$ ($\psi = 90^\circ$)

Figure 22: Approximate sketches of the vortex-fuselage interaction phenomena, adapted from Brand [16].

Using periodic adaptation, the methodology was initially evaluated using inviscid simulations so that the best strategies could be applied for the more expensive fully turbulent simulations, which has been recommended for improved correlations [6, 21] due to the complex viscous interactions between the tip vortices and geometry (blades and fuselage). Solution advancement was performed with a time step equivalent to 1° azimuthal sweep, which was found to be sufficient in previous studies [6, 21, 23]. During each time-step, a maximum of 40 sub-iterations were used in conjunction with the temporal error control option [79] to ensure two orders of magnitude reduction in residual. To obtain valid comparisons with experiment, the computed thrust coefficient was ascertained to be within 1% of the experimental value listed in Table 4. The

tolerance levels for the different feature-based schemes are given in Table 5. These tolerance levels are selected via solution interrogation based on the recommendation of resolving the vortex core of interest by seven to ten nodes [57].

Subsequently, a turbulence model study was performed to assess the efficacy of the $k\omega$ -SST model against GT-HRLES method with respect to prediction of the physics. Finally, fully turbulent computations were carried out to assess the validity of the most beneficial adaptation strategies.

Following these studies, a study of sensitivity to the adaptation interval size was also performed against that of periodic adaptation in Section 6.3. Again, the best adaptation strategies resulting from the periodic simulations are analyzed.

Table 5: Summary of tolerances applied for the different adaptation schemes.

Method	Tolerance (F_{tol})
$ \omega $	0.001
Δp	0.003
Nondimensional Q-criterion	0.01

6.1.2 Methodology Demonstration

The ability of this new methodology to capture complex unsteady features in the flow field is confirmed in Figs. 23 and 24, showing a view from the top of the blade. Figure 23 (a) illustrates the initial grid, and the distinction between the background and blade grids can be easily observed. The tip vortex of the previous blade appears diffused by the time it enters the current blade grid (bottom of Fig. 24 (a)). A new grid with a combined adaptation to $|\omega|$ and Δp (over two cycles) captures the previous tip vortex as it interacts with the blade as shown in Figs. 23 and 24 (b). Furthermore, the overset adaptation capability clearly allows these features to pass seamlessly between the blade grid and the fuselage grid.

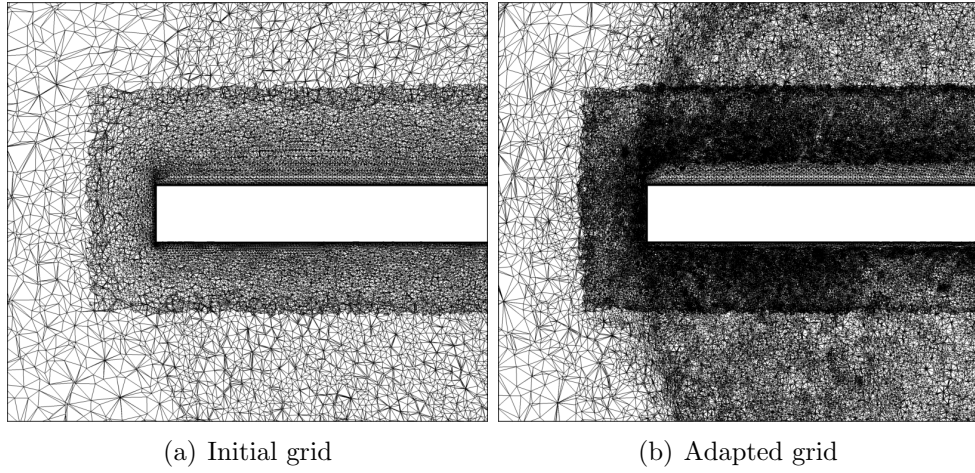


Figure 23: Grid comparison; view is of the blade planform from the top.

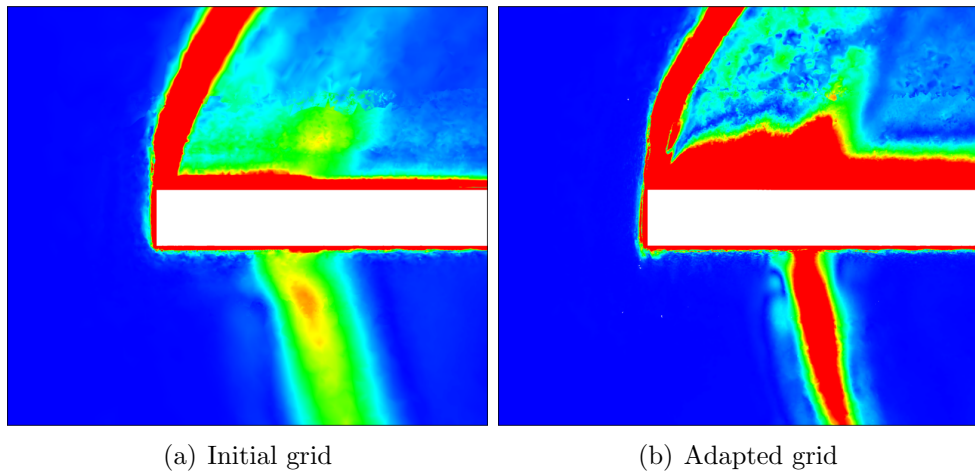


Figure 24: Vorticity-magnitude comparison; view is of the blade planform from the top.

6.2 *Results: Periodic Grid Adaptation*

The first set of results utilized periodic adaptation to perform extensive analysis of the best adaptation strategies, the effects of modeling, and the convergence of the method. Using the best results in this section, the interval sensitivity analysis is conducted in Section 6.3.

6.2.1 *Adaptation Strategy Study: Inviscid Computations*

Several feature-based adaptation schemes were evaluated for the GIT RFI configuration to demonstrate the validity of the method, as well as to determine the appropriate flow field adaptation feature. Table 6 details the grid sizes resulting from the initial grid to the different feature-based schemes that were studied. Additionally, the ability of the each method to predict improved vortex-fuselage impingement from the initial grid is qualitatively provided. A detailed quantitative analysis is discussed by examining the instantaneous pressures. The initial grid contained pre-refinement in the expected rotor wake region. The first three adaptation schemes ($|\omega|$, Δp , and Q-criterion) apply grid modifications on the initial grid. The following three schemes are double adaptation schemes, that is, they use the solution from one of the previous adaptation schemes and adapt that grid a second time using the same or a different feature. The grid initially adapted to $|\omega|$ was again adapted to $|\omega|$, referred to as the $|\omega|$ (*iter. 2*) *scheme*. In an alternate fashion, the grid originally adapted to $|\omega|$ was next adapted to Δp , referred to as the *vorticity-mixed scheme*. The final scheme uses the grid initially adapted to Q-criterion and adapts it to Δp , henceforth referred to as the *Q-criterion-mixed scheme*. Double adaptation to Δp was not included because of its inability to sufficiently capture the tip vortex intensity. Also, since double adaptation to $|\omega|$ was found to be ineffective, and because both $|\omega|$ and Q-criterion target high vorticity regions such as tip vortices, double adaptation to Q-criterion was also not considered.

Table 6: Inviscid simulations: Summary of grid sizes and ability to predict vortex impingement from the different adaptation schemes.

Scheme Description	Total Nodes (Millions)	Improvement of Impingement Accuracy
Initial Grid	2.03	–
Adapted to $ \omega $	6.68	Poor
Adapted to Δp	2.92	Poor
Adapted to Q-crit.	9.95	Poor
Adapted to $ \omega $ (iter. 2)	12.8	Poor
Adapted to $ \omega $ & Δp	8.86	Good
Adapted to Q-crit. & Δp	23.3	Poor

The Δp adaptation resulted in a much more efficient adaptation than the $|\omega|$ and Q-criterion-based schemes based solely on the grid size (Table 6). Only minor differences between the $|\omega|$ and Δp schemes were observed for the time-averaged fuselage pressure coefficient (\bar{c}_p) peaks at the $x/R = 0.5$ and 2.0 locations (Fig. 25) after one cycle. With the exception of the Q-criterion scheme, the result of the first adaptation was to refine the initial vortex interaction at the nose ($x/R = 0.1$), so that the magnitude and pressure rise are more accurately captured. The Q-criterion scheme does not differ greatly from the initial grid at $0.2 \leq x/R \leq 0.5$ but shows significant differences from all the other schemes at and behind the hub region $x/R \geq 1$. The effects of the vortex shed from the second blade, observed in the pressures along $0.2 \leq x/R \leq 0.5$, are not captured with either the baseline or the single adaptation. The vortex-fuselage interaction at $x/R = 2.3$ is captured by all three meshes as a weak pressure pulse, and the magnitude does not change with grid adaptation, although a minor ($x/R < 0.05$) shift forward is observed upon grid adaptation.

The time-averaged pressures from the second adaptation are plotted in Fig. 26. While the $|\omega|$ (*iter. 2*) scheme results in a much larger mesh size, its influence on the fuselage pressure characteristics is minimal. On the other hand, the *vorticity-mixed scheme* yields significant improvement, and employs about 75% of the mesh required




Figure 25: Inviscid simulations: top centerline time-averaged pressures from one adaptive cycle.





Figure 26: Inviscid simulations: top centerline time-averaged pressures from two adaptive cycles.


by double adaptation to $|\omega|$. This scheme captures the vortex-fuselage interaction described by Brand [16], denoted by the sharp rise in pressure for $0.4 \leq x/R \leq 0.5$. Adapting the grid using the *Q-criterion-mixed scheme* obtains improvement in the average pressures in the nose region, but it does not improve the pressures where the vortex-fuselage interaction exists.

Further examining the instantaneous flow field (Figs. 27 and 28), the characterizations noted in the time-average pressures are re-enforced in these data. For the first quarter revolution where the rotor blades depart from the fuselage (Fig. 27), the $|\omega|$ (*iter. 2*) *scheme* and the *Q-criterion-mixed scheme* improve the instantaneous pressure prediction only slightly and include a significant lag ($\Delta x/R \approx 0.1$) in the location of the primary vortex interaction at $0.4 \leq x/R \leq 0.5$. The *vorticity-mixed scheme* still indicates lower magnitude at $\psi = 30^\circ$ and 60° due to an apparent phase lag of this vortex interaction, particularly at $\psi = 90^\circ$, however the results are significantly improved over the single adaptation. It is also important to note that the magnitude and character of the pressure rise resulting from the primary interaction, which is completely missed by the baseline and other adaptation schemes, are overall well-captured. Minor differences are observed in the centerline pressure exclusive of this vortex interaction.


As observed in Fig. 28, the second quarter of the rotor revolution continues this overall trend, but with increasing differences with experiment as the two rotor blades approach the fuselage to their original positions. Again, the primary vortex interaction at $0.4 \leq x/R \leq 0.5$ continues to be more accurately captured with the *vorticity-mixed scheme*. The secondary vortex interaction in the range $0.7 \leq x/R \leq 0.9$ during this quarter is not captured by either the $|\omega|$ (*iter. 2*) *scheme* or the *Q-criterion-mixed scheme*. The *vorticity-mixed scheme* appears to capture this feature, albeit lagging by $\Delta x/R \approx 0.1$, and as a weaker interaction at $\psi = 120^\circ$ and 150° . The significant high pressure region near the nose and forward portion of the fuselage



(a) $\psi = 30^\circ$




(b) $\psi = 60^\circ$




(c) $\psi = 90^\circ$


Figure 27: Inviscid simulations: top centerline instantaneous pressures (first-quarter revolution).



(a) $\psi = 120^\circ$



(b) $\psi = 150^\circ$



(c) $\psi = 180^\circ$

Figure 28: Inviscid simulations: top centerline instantaneous pressures (second-quarter revolution).

($0.2 \leq x/R \leq 0.5$) at $\psi = 180^\circ$ when the blades are directly over the fuselage is not predicted using any method. Brand's experiment reported the unsteady pressure data via high frequency microphones [16]. Because large amounts of ensemble-averaging of the data are performed, he chose to record averaged data at 6° intervals, while these computational results are not averaged. He conducted an analysis of the sensitivity of the interval size between 6° and 1.5° , resulting in small changes at many azimuthal locations in the period at $x/R = 0.45$ [16]. However, since most of the discrepancy in Fig. 28 (c) occurs at $x/R < 0.45$, and the largest changes of the pressures occur in this region close to $\psi = 180^\circ$, the ensemble-averaging of the experimental pressures may be responsible for the differences in the distribution.

It is evident that the *vorticity-mixed scheme* clearly most accurately predicts instantaneous pressures among the various schemes evaluated. The inability of the *Q-criterion-mixed scheme* to mimic the *vorticity-mixed scheme* in these inviscid simulations is discussed later.

6.2.2 Turbulence Model Study

The capability of this adaptation methodology to capture viscous flow phenomena was considered. Two turbulence methods, a RANS model ($k\omega$ -SST) and the GT-HRLES model, were studied. Both models were applied using rotation correction [98]. The production term was based on the strain-rate rather than vorticity since this method showed better correlations for rotating flows [99]. Three free-stream turbulence intensity $Tu = \frac{\sqrt{u'^2}}{U_\infty}$ is set to 0.2% which is comparable to the maximum intensity reported for the Harper tunnel (0.5%) [16]. Using a highly pre-refined composite grid (15.4 million nodes), wherein the background grid was refined in the wake region between the blade and the fuselage, the effect of tip vortex dissipation is studied by applying both these turbulence models. The eddy viscosity and vorticity magnitude predicted using the $k\omega$ -SST model and HRLES models are shown in Figs. 29 and 30,

respectively. These figures depict the prediction at $\psi = 120^\circ$, but the same result is obtained at different blade azimuths.




Figure 29: Contours along the airframe symmetry plane from the $k\omega$ -SST model.




Figure 30: Contours along the airframe symmetry plane from the HRLES model.

Two observations can be made with respect to Figs. 29 and 30. First, the eddy viscosity prediction from the $k\omega$ -SST simulation is significantly higher and widespread in the region coinciding with the forward tip vortex. The vortex also appears to be visibly diffused or spread out in comparison to the HRLES simulation. Furthermore,

the vortex core region of the HRLES simulation predicts significantly lower eddy viscosity, implying that this vortex core exhibits expected laminar behavior. The other observation is that the high eddy viscosity prediction from the $k\omega$ -SST simulation that dominates the rotor wake region clearly diffuses the vortex propagating toward the fuselage following interaction with the oncoming blade. Additionally, rotor wake vorticity contours of the HRLES simulation show two distinct high vorticity regions that are characteristic of the expected vortex-fuselage interaction. This study demonstrates the superiority of the HRLES model in the preservation of the tip vortex and capturing a more complex interaction with the subsequent blade passage.

6.2.3 HRLES Computations

The initial grid used for the HRLES computations had the same pre-refinement in the rotor-wake region as in the inviscid simulations. The primary difference is the addition of a boundary layer grid with a $y^+ = 1$ and at least 35 boundary layer cells. Mixed-elements are used to generate this boundary layer grid, allowing for accurate simulations with the HRLES turbulence model [84]. The grid sizes resulting from the initial grid to the different feature-based schemes are listed in Table 7. Similar to Table 6, a qualitative description of each method’s ability improve the impingement prediction is given. Among the single adaptation schemes, adaptation to Δp was not considered because of its inability to preserve regions of high vorticity, which are essential in order to capture the magnitude of the fuselage surface pressures. The double adaptation schemes studied were the *vorticity-mixed scheme* and the *Q-criterion-mixed scheme*. The $|\omega|$ (*iter. 2*) *scheme* was not performed because the inviscid simulations showed very little improvement from the single adaptation to $|\omega|$.

Time-averaged fuselage pressures comparing the $|\omega|$ and Q-criterion single adaptation schemes (Fig. 31) show very small differences, unlike the results from the inviscid

Table 7: HRLES: summary of grid size and ability to predict vortex impingement from the different adaptation schemes

Scheme Description	Total Nodes (Millions)	Improvement of Impingement Accuracy
Initial Grid	5.52	—
Adapted to $ \omega $	13.0	Poor
Adapted to Q-crit.	16.8	Minimal
Adapted to $ \omega $ & Δp	24.2	Good
Adapted to Q-crit. & Δp	30.7	Good

simulation (Fig. 25). Both these schemes improve the magnitude of the pressures in the forward part of the fuselage ($x/R \leq 0.3$), with exception to the nose region. Additionally, the primary vortex interaction resulting from both these schemes indicate the presence of a small pressure pulse at $x/R \approx 0.5$.




Figure 31: HRLES: time-averaged pressures on the top centerline from one adaptive cycle.

The *vorticity-mixed scheme* and *Q-criterion-mixed scheme* improve the accuracy of the time-averaged pressures (Fig. 32), particularly in predicting the intensity of the primary vortex interaction $x/R \approx 0.5$. Both schemes have clear similarities, but the *vorticity-mixed scheme* predicts a stronger pressure pulse than the *Q-criterion-mixed scheme*. Again, since the hub pin geometry is not modeled [6], the HRLES simulations are not able to accurately predict the pressures in the aft portion of the rotor ($x/R > 1.5$).






Figure 32: HRLES: time-averaged pressures on the top centerline from two adaptive cycles.

The instantaneous fuselage pressures, plotted in Figs. 33 and 34, show good correlation with experimental data at intermediate azimuths ($\psi = 90^\circ - 150^\circ$) with both the *vorticity-mixed scheme* and *Q-criterion-mixed scheme*. Both schemes agree very well with each other and the only difference observed is that the *vorticity-mixed scheme* predicts a higher magnitude of the primary vortex interaction. With respect to experiment, both schemes result in a small lead of $\Delta x/R = 0.01$ in the spatial location for the primary vortex interaction. The magnitude of the interaction is under predicted for the first quarter revolution and is over predicted for the second quarter revolution. The secondary vortex interaction initially has a marginal lead in the spatial location of $\Delta x/R = 0.05$, but during the second quarter revolution, this interaction lags behind the experiment by the same amount. The faster convection rate of the secondary vortex interaction is not captured. As with the inviscid simulations, the high pressure regions in the forward portion of the fuselage ($0.2 \leq x/R \leq 0.5$) at $\psi = 180^\circ$ do not correlate well.


To further understand the significance of the grid refinement in each scheme, the vortex behavior is examined in Figs. 35 – 37. The adaptation sequence from the initial grid to the *vorticity-mixed scheme* are plotted from top to bottom at each selected azimuthal location. Since the adaptation sequence leading to *Q-criterion-mixed scheme* show the similar results as with the *vorticity-mixed scheme*, they are



(a) $\psi = 30^\circ$




(b) $\psi = 60^\circ$




(c) $\psi = 90^\circ$

Figure 33: HRLES: top centerline instantaneous pressures (first-quarter revolution).


not presented here. It is clear from scanning from top to bottom that the forward vortex core is more crisply predicted after the first adaptation (middle plot) and refined further upon the second adaptation (bottom plot). In addition, the initially weaker vortical features in the rotor wake, diffused over large areas, become further defined due to adaptation.



(a) $\psi = 120^\circ$



(b) $\psi = 150^\circ$



(c) $\psi = 180^\circ$

Figure 34: HRLES: top centerline instantaneous pressures (second-quarter revolution).

Differences between the various adaptation schemes can also be discerned from the magnitude and shape of the vorticity contours in these figures. For example, the shape of the vortex at $0.4 \leq x/R \leq 0.5$ where the primary interaction occurs is very different. Specifically, tracing the *vorticity-mixed scheme* across Figs. 35– 37, it is




Figure 35: Vortex behavior (from top: initial grid, adapted to $|\omega|$, and *vorticity-mixed scheme*).

possible to track the path of the tip vortex as it leaves the blade, interacts with the previous blade's wake sheet, and finally collides with and encompasses the fuselage centerline. Brand [16] reported that the tip vortex from the prior blade interacts with




Figure 36: Vortex behavior (from top: initial grid, adapted to $|\omega|$, and *vorticity-mixed scheme*).

the following blade at $x/R = 0.3$ at $\psi = 188^\circ$ (or $\psi = 8^\circ$), which is comparable to the grid adaptation results in Fig. 37 (b). The weakening vorticity of the primary interaction, observed at approximately $x/R = 0.45$ in Figs. 35– 37, correlates to the




Figure 37: Vortex behavior (from top: initial grid, adapted to $|\omega|$, and *vorticity-mixed scheme*).

experimental visualization, as sketched in Fig. 22. The vortex sheet roll-up, which was experimentally observed to traverse in the range $0.4 \leq x/R \leq 0.5$ during this azimuthal time period can also be observed traveling downstream at $\psi = 210^\circ$ and

with a distinct rotation by $\psi = 240^\circ$, located in the same fuselage locations. The development of the secondary vortex-fuselage interaction and its subsequent rapid downstream convection (Fig. 22 (b)) is also observed in the *vorticity-mixed scheme* and *Q-criterion-mixed scheme* (not shown).

6.2.4 Vorticity–Q-criterion Discrepancy

The reason why there is a lack of agreement between the *vorticity-mixed* and *Q-criterion-mixed schemes* in the inviscid simulations but not in the viscous simulations has been investigated. The Q-criterion indicator formulation targets regions where the rotation rate $\|\Omega\|$ dominates the strain rate $\|S\|$, since $F_{tol} = 0.01$. In regions where $\|S\|$ exceeds $\|\Omega\|$, the Q-criterion (dimensional or non-dimensional) values are negative. Therefore such regions are not selected for refinement. The vorticity magnitude method, on the other hand, does not discriminate regions where strain rates dominate and its range is always non-negative.

Figure 38 illustrates the differences in the flow fields resulting from these schemes for the inviscid simulations. The vortex sheet region clearly displays uniformly high values of vorticity, but the values of Q-criterion in those regions are negative or very close to zero. Regions where both vorticity-magnitude and Q-criterion are high include the vortex core region and few localized regions in the vortex sheet. The fully turbulent flow fields resulting from these adaptation schemes is compared in Fig. 39. High values of vorticity throughout the vortex sheet are observed in Fig. 39 (a). In comparison to Fig. 38 (b), the vortex sheet in Fig. 39 (b), shows several more spots of positive Q-criterion values. Therefore, these regions in the vortex sheet are selected for refinement in the turbulent simulation and are excluded from refinement in the inviscid simulation. Since the adaptation metric is accumulated over 180° azimuthal sweep, the highly refined vortex sheet benefits the preservation of the tip vortex as it passes through these regions. Another difference between inviscid and turbulent

simulations is that the latter exhibits an interaction with the rotor wake and fuselage boundary layer about the juncture of the hemispherical nose and cylindrical fuselage. This region indicates both high values of vorticity magnitude and positive Q -criterion. This behavior arises mainly due to the no-slip boundary condition of the turbulent simulations.

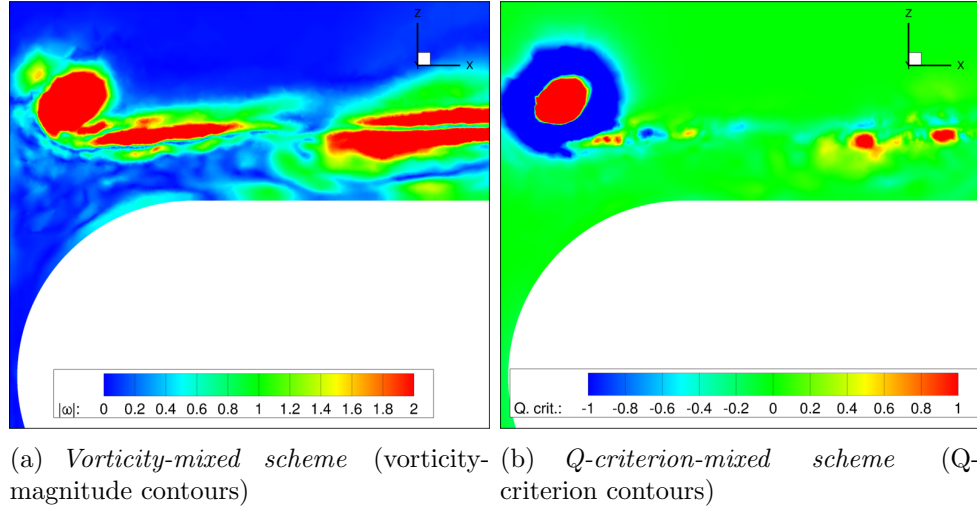


Figure 38: Inviscid simulations: comparison of the *vorticity-mixed* and *Q-criterion-mixed* schemes.




Figure 39: HRLES: comparison of the *vorticity-mixed* and *Q-criterion-mixed* schemes.

The genesis of positive Q -criterion values in the turbulent vortex sheet is explained

by examining the flow field near the blade trailing edge in Fig. 40. The trailing edge region exhibits high Q -criterion values over a significant portion of the blade span for the turbulent simulation, absent from the inviscid simulation. This is attributed to boundary layer-trailing edge vortex shedding due to the blade’s no-slip boundary condition.




Figure 40: Comparison of inviscid and viscous modeling of the flowfield for the Q -criterion-mixed scheme

6.2.5 Tip Vortex Paths

Experimental visualizations document the tip vortex locations as soon as they become visible at $\psi = 188^\circ$ and are plotted at 30° intervals until impingement with the fuselage in Fig. 41 (a). Vortex locations from the *vorticity-mixed scheme* resulting from both the inviscid and HRLES simulations, are shown in comparison to experiment. Figure 41 (b) plots the streamwise spatial location lead (or lag) with respect to experiment. Both simulations show the same lead at the first vortex location. However, the inviscid simulation shows a vortex lag for the rest of the azimuthal locations, as high as $\Delta x/R = 0.032$. These results are corroborated by the instantaneous fuselage pressures in Figs. 27 and 28. The HRLES simulation correlates much better with the spatial location, generally leading the experiment, with the maximum vortex lead of

$\Delta x/R = 0.013$. This spatial lead is also observed via the surface pressures in Figs. 33 and 34. The uncertainty in the streamwise location reported with experiment was 15 mm or $\Delta x/R = 0.033$. Both simulations with the *vorticity-mixed scheme* are within the experimental error, but Fig. 41 (b) shows significantly better spatial correlation from the HRLES simulation of the vortex location with experiment.

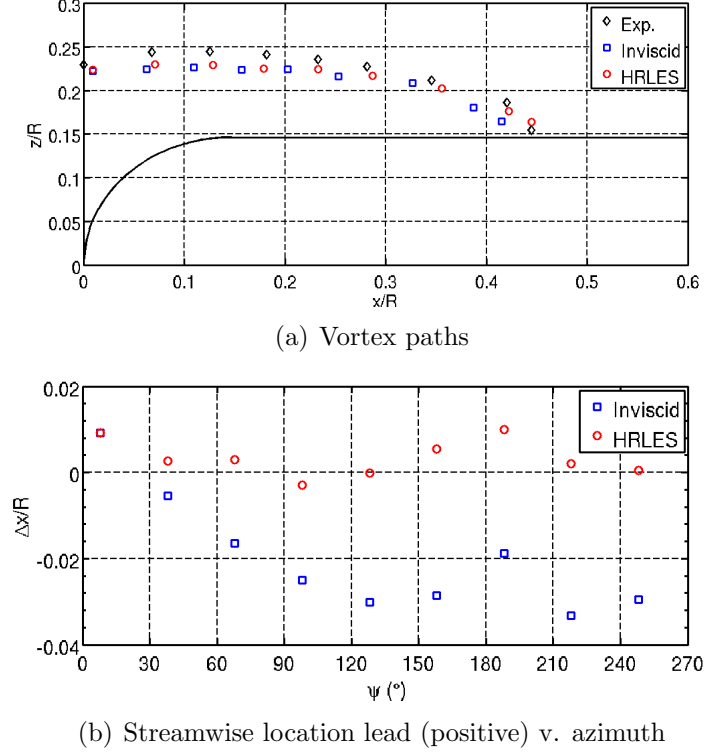


Figure 41: Vortex-trajectory comparisons from the *vorticity-mixed scheme*.

6.2.6 Adaptation Convergence

The convergence of this feature-based adaptation process has been assessed by evaluating an integrated quantity of interest obtained over a series of adaptation cycles. Because the time-averaged pressure distribution has been previously applied as a suitability criterion of an adaptation scheme, the time-averaged centerline pressure integral is used here to identify convergence. The functional quantity is:

$$f_{\bar{c}_p} = \int_0^{3R} \bar{c}_p dx, \quad (28)$$

where the fuselage length is $3R$.

The convergence of the *vorticity-mixed scheme* has been assessed by performing an additional adaptation sequence, i.e. additional adaptations that include both $|\omega|$ and Δp . The pressure integral is plotted in Fig. 42 for the *vorticity-mixed* HRLES simulation. Convergence is established for the four adaptation cycles over two iterations of the *vorticity-mixed scheme*. The change in the functional between cycles 3 and 4 is observed to be within 0.05% and the overall change over the second adaptation iteration is within 1.25%.




Figure 42: Pressure-integral functional convergence for the HRLES *vorticity-mixed scheme* adaptation sequence.

6.3 Results: Adaptation Interval Sensitivity

General adaptation intervals were performed in order to assess the sensitivity of the interval size. The adaptation procedure is delineated in Section 3.2 and is similar to the periodic adaptation procedure. The major difference here is that solution transfers were performed between grids tuned to differing intervals. Solution interpolation was performed using trilinear interpolation based on the localization method described in Chapter 4. This study applied the HRLES turbulence model with the *vorticity-mixed*

scheme.

Since convergence of the *vorticity-mixed scheme* on the time-averaged pressures within 1.25% was ascertained by the second adaptation iteration in Section 6.2.6, only one adaptive iteration of this scheme was performed in this study. The same tolerance levels as in Section 6.2, listed in Table 5, were applied. The periodic adaptation method is similar to adapting with an interval of 180° except that the adaptation was not repeated over different intervals. The new interval choices are 5° and 15° , as listed in Table 8. For an adaptation strategy to be consistent, the periodic adaptation results should either be retained or improved in comparison to the experiment with these finer intervals. Solution degradation due to interpolation is monitored and tabulated in Table 8. Here, the maximum deviation includes the both the $|\omega|$ and *vorticity-mixed schemes*. Finally, the statistics of the grid sizes are also tabulated. The initial grid at the beginning of each simulation comprised 5.5 million nodes. For interval based adaptation, since the wake grows into the domain, the grid is expected to continually grow due to net refinement. Since the periodic grid’s results were obtained first, the grid growth of the smaller interval simulations was controlled such that the grid at the end of the revolution (corresponding to the same phenomenon captured over the periodic interval) had a comparable number of nodes to the periodic case. Control of the grid size is achieved by scaling the metric to the desired complexity (number of nodes) [62]. Valid comparisons of the results between the different interval sizes may be made because the final grid sizes compare within 2% of each other.

Table 8: Summary of adaptation interval study.

Interval Size	No. of Intervals	Max. $\delta_{\mathbf{x}}$	Min/Max. No. Nodes
5°	72	5.366×10^{-15}	5.4/24.4 million
15°	24	3.689×10^{-15}	5.6/24.2 million
Periodic (180°)	0	N/A	24.2 million

After following the adaptation procedure over all the intervals over one revolution, the grids obtained were evaluated on a subsequent revolution. Figure 43 plots the centerline time-averaged pressures resulting from the 15° interval simulation for the revolution where the metric was sampled (first rev.) and the following revolution (second rev.). The excellent agreement between the pressures for $0 < x/R < 1.0$ shows that the flow field is converged. More importantly, these converged pressures show that interpolation errors over an extra 24 solution transfers do not degrade the solution.




Figure 43: Convergence of time-averaged pressures of the 15° adaptation interval using the *vorticity-mixed scheme* adaptation sequence.




Figure 44: Comparison of time-averaged pressures varying the adaptation interval size using the *vorticity-mixed scheme* adaptation sequence.


The time-averaged centerline pressures for the 5° and 15° interval simulations are plotted against that resulting from the periodic interval in Fig. 44. A small

improvement in the pressures around the nose and fore-body region ($0 < x/R < 0.3$) is observed. There is, however, minimal improvement from the periodic interval of the pressure distribution at $0.35 < x/R < 0.45$. The 15° and 45° interval pressures compare excellently with respect to each other in the region $0 < x/R < 1.0$, except at $x/R \approx 0.45$, the area of significant vortex-fuselage impingement, where the 15° interval exhibits a slightly greater impingement pulse. In this region, the gradients in the pressure distribution are captured much better and the comparison with the experiment is much better. The agreement with respect to each other for the 5° and 15° interval simulations implies that these pressures converge with decreasing the adaptation interval size. All three pressure distributions show excellent correlation with experiment for $1.0 < x/R < 1.5$.


The instantaneous pressure distributions are compared in Figs. 45 and 46. At all azimuthal locations here is agreement in the pressures between the different interval computations, especially between the 5° and 15° interval computations. The improvements from the periodic adaptation are noticed from both non-periodic intervals, where the periodic interval incorrectly predicts secondary impingement at $\psi = 30^\circ$ and shows lags in the secondary interaction phenomena from the non-periodic and experimental data for $\psi = 120^\circ$ to $\psi = 180^\circ$ (Fig. 46). Additionally, the magnitude of the primary vortex interaction is also over predicted during the second quarter-revolution. Overall, excellent correlation between the location of the primary vortex interaction is observed here, indicating that the method is consistent with the prediction of the primary vortex-fuselage impingement.

6.3.1 Tip Vortex Paths


The tip vortex trajectories as well as their lead/lag with respect to the experiment are illustrated in Fig. 47. In terms of vertical location of the vortex z/R , the non-periodic intervals diverge slightly from the experimental location. With respect to



(a) $\psi = 30^\circ$




(b) $\psi = 60^\circ$




(c) $\psi = 90^\circ$

Figure 45: Top centerline instantaneous pressures for the *vorticity-mixed scheme* (first-quarter revolution).


the streamwise location's lead/lag with respect to the experimental location (Fig. 47 (b)), there is no clear trend in convergence of the vortex location at the different azimuthal locations. A consistent lead against the experimental prediction is observed



(a) $\psi = 120^\circ$



(b) $\psi = 150^\circ$



(c) $\psi = 180^\circ$

Figure 46: Top centerline instantaneous pressures for the *vorticity-mixed scheme* (second-quarter revolution).

for the non-periodic intervals, that is, the lead decreases until around $\psi \approx 100^\circ$ and increases again until $\psi \approx 200^\circ$. With a few exceptions such as when the vortex is at $\psi = 218^\circ$ where the lead is divergent, consistent predictions between all interval

sizes are observed between different azimuthal locations within $\Delta x/R = 0.005$. All predictions of the tip vortex here are within $\Delta x/R = 0.01$ (or 1% rotor radius), which is well within the experimental uncertainty $\Delta x/R = 0.033$.




Figure 47: Vortex-trajectory comparisons from the *vorticity-mixed scheme* for HRLES computations.


6.4 Rotor-Fuselage Interaction Aerodynamics

The effects of the rotor and the fuselage interactional aerodynamics are investigated here. Prior to this work, only O'Brien [6] had previously conducted an analysis of the fuselage aerodynamics using an actuator blade model. However, his work did not analyze the blade aerodynamics due to interaction with the fuselage. The analysis


in this section will be compared against O'Brien's results in order to highlight the differences of using the overset adaptive grid approach. Applying the adapted grid resulting from the periodic interval adaptation from Section 6.2, an analysis over two revolutions has been conducted here. Aerodynamic loads on the fuselage are characterized as the lift, drag, and side force coefficients (C_L, C_D, C_Y , respectively) and roll, pitch, and yaw moment coefficients ($C_{M_x}, C_{M_y}, C_{M_z}$, respectively). The lift coefficient is oriented in the positive z direction (up) while the side coefficient is oriented positive to the right of the fuselage.

The fuselage forces are examined in both the time and frequency domains in Fig. 48. The rotor blade imparts a two-per-rev (180°) periodic behavior on the fuselage forces and moments. At each half-revolution multiple (13, 13.5, \dots , 15), the two blades pass over the fuselage, 180° apart from one another. Figure 48 (a)'s predictions of the fuselage lift and drag coefficients indicate a sharp response within approximately $\pm 45^\circ$ of blade passage. Without the influence of the blades, the nearly constant fuselage drag is the result of free-stream pressure effects, while lift coefficient is nearly zero since the fuselage is an axisymmetric non-lifting body. Small deviations in lift are observed for the azimuth range $60^\circ \leq \psi \leq 120^\circ$, which arise from the vortex-fuselage impingement process (see Figs. 45- 46 and Fig. 49 (a)). The pressure distribution on the fuselage and blades are shown in Fig. 49 at different azimuths describing the effect of blade passage. Due to the high pressures on the top of the fuselage primarily in the forward fuselage region (Fig. 49 (e)), the drag coefficient more than doubles and the lift coefficient decreases from zero to nearly -15. The side force coefficient on the fuselage has a net left (negative) response and exhibits a negative impulse at each blade passage. Figure 48 (b) plots the fuselage unsteady force harmonics in terms of the dominant frequencies. As expected, the two-per-rev and associated harmonics are present. The side force shows higher frequency content

than the lift and drag at higher harmonics such as six-per-rev and twelve-per-rev content (Fig. 48 (b)). This phenomenon is attributed to the close blade passage on the forward part of the fuselage due to the forward shaft tilt that results in a significantly high pressure region with large pressure variations in time [16] due to blade passage.



(a) Fuselage force history



(b) Unsteady fuselage force harmonics

Figure 48: Effect of rotor on fuselage forces.

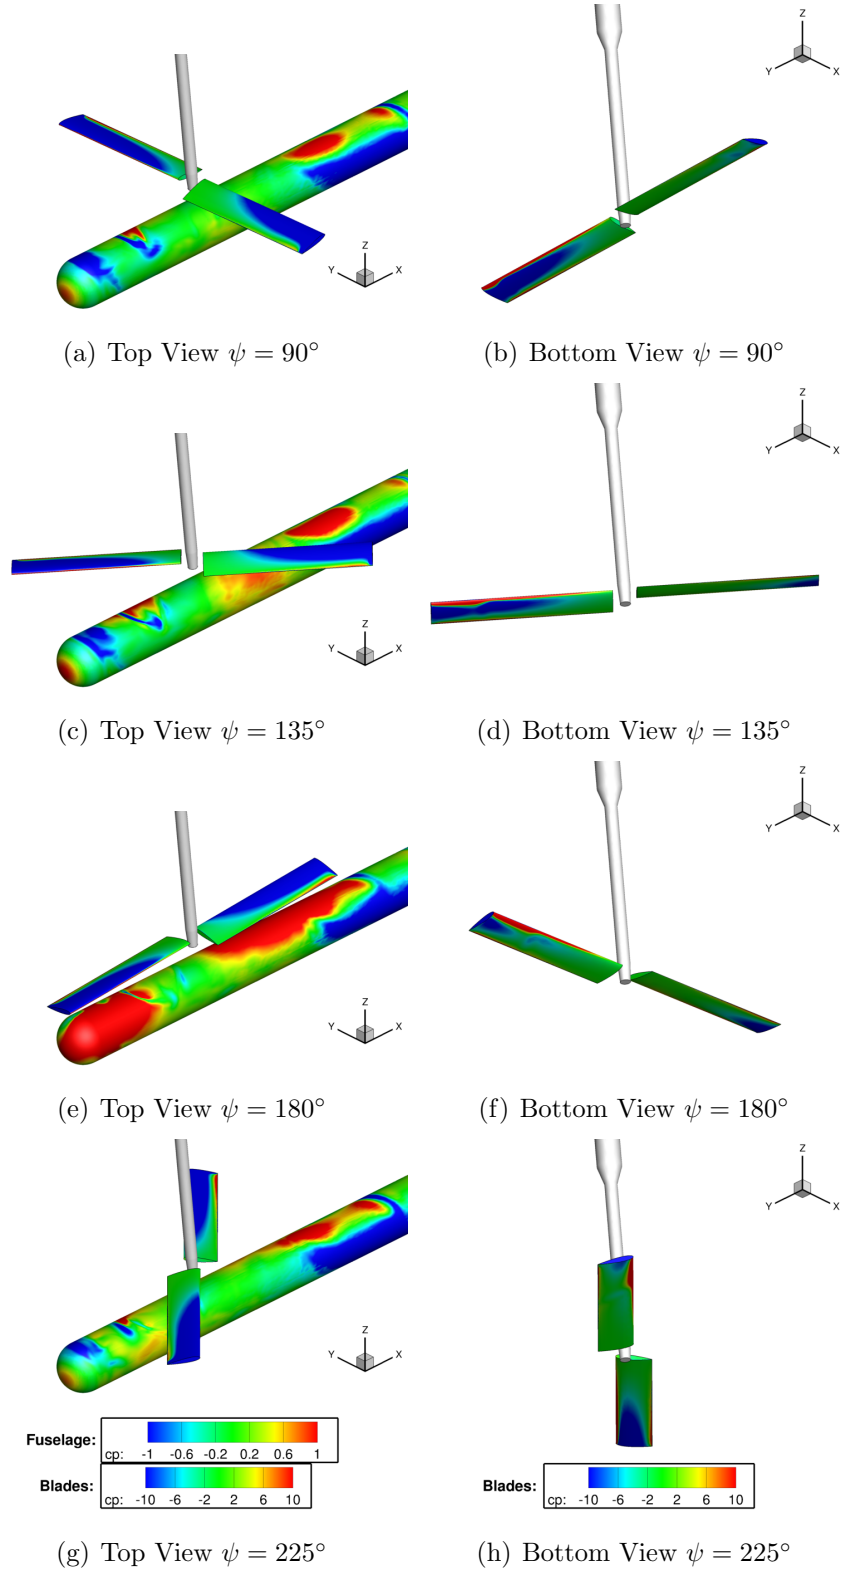



Figure 49: Pressure distributions on the fuselage and blades at different azimuths.


The rotor’s effect on the fuselage moments are studied in Fig. 50. Analogous to the lift coefficient response, the pitching moment of the fuselage also shows a decrease at each blade passage, as observed in Fig. 50 (a). The pitching moment also exhibits a vibratory response to the vortex-impingement. The fuselage yaw moment exhibits a significant rolling moment change from a nose left to right yaw due to blade passage. Between subsequent blade passages, a relatively smooth and monotonic transition from the right to left yaw is experienced by the fuselage. The rolling moment of the fuselage also shows a response due to blade passage and is correlated to the response of the side force, where clear six-per-rev and twelve-per-rev behavior is observed, plotted in Fig. 50 (b). Due to the dominating presence of two-per-rev harmonics, this behavior in the pitching and yawing moments frequencies are not apparent (Fig. 50 (b)), but it is observed in the time domain upon blade passage. If this configuration were representative of a helicopter in forward flight, these strong blade passage effects would cause a significant vibratory response on the fuselage, adversely influencing the handling qualities. Figures 48 (b) and 50 (b) indicate that the strongest high frequency vibratory effects are observed for the side force, rolling moment, and yawing moment.

Similar physics was predicted by O’Brien’s [6] analysis of fuselage forces and moments. The lift decreases, drag decreases, and the side force exhibits similar oscillations at each occurrence of blade passage. However, the magnitude in lift coefficient he predicted was not as great as that predicted here (lower by approximately 20%). Additionally, while he predicted the same behavior in the yaw moment and pitching moment history, the amplitude of the rolling moment he predicted was substantially smaller than that of the current analysis. While the same solver (FUN3D) was utilized, his rotor model and turbulence model were different, and he did not apply grid adaptation in his work.

Unlike the effect of the blades on the fuselage, the effect of the fuselage on the



(a) Fuselage moment history



(b) Unsteady fuselage moment harmonics


Figure 50: Effect of rotor on fuselage moments.

blades load has not been analyzed for this configuration by O'Brien [6] and other previous investigations. The blade aerodynamic loads are integrated quantities centered at the rotor hub in the blade reference frame. The blade loads are the lift, drag, and


side force coefficients (C_L, C_D, C_Y , respectively) and roll, pitch, and yaw moment coefficients at the hub ($C_{M_x}, C_{M_y}, C_{M_z}$, respectively). Figure 51 plots the loads history the second blade, originating at $\psi = 180^\circ$. At this azimuth, the blade is aligned with the front fuselage centerline (Fig. 49 (e) and (f)) at revolutions 13, 14, and 15 (time steps 4680, 5040, and 5400, respectively). The lift coefficient increases as the blade encounters either the forward or rear part of the fuselage due to the high pressure region between the blade and the fuselage depicted in Fig. 49 (e) and (f). However, the lift coefficient increase is smaller (by 10%) as the blade aligns with forward fuselage. This is attributed to the close proximity of the blade with the forward fuselage causing a high pressure region over a relatively smaller volume in comparison to the rear fuselage blade passage. Figure 49 (e) and (f) show that the high pressure region is distributed over a larger area of the fuselage when the blade is aligned with the fuselage rear. When comparing the advancing and retreating sides (away from fuselage passage), the lift coefficient is slightly higher on the advancing side (about 5%) due to the difference in the local dynamic pressure, which is not surprising since this is a low speed forward flight condition. The drag force history mimics the expected drag of typical rotorcraft blades; it is highest when the blade is on the advancing side and lowest when the blade on the retreating side of the revolution. The drag on the blade as it passes over the fuselage has a negligible effect as the drag signal is smoother than that of lift. The blade side force experiences smaller magnitudes than the lift and drag forces but as it passes over the fuselage (forward or rear) there is a net decrease in the side force indicating that the force is biased toward the hub.

The blade rolling moment shape corresponds to the shape of the lift history, which is not surprising since the lift directly acts on the lower side of the blade contributing most to this hub moment. On the advancing blade, the magnitude of the roll moment increases, while on the retreating side, it decreases. The blade pitching moment is smooth and negligible in comparison to the other loads the blade encounters. The yaw

moment of the blade is also smooth and mimics the shape of the drag (in magnitude) history because the blade drag contributes to the hub yaw moment.



(a) Blade force history



(b) Blade moment history

Figure 51: Effect of rotor on blade forces and moments. The blade is originally at $\psi = 180^\circ$.

CHAPTER VII

GENERIC ROTOR HUB APPLICATION

7.1 Rotor Hub Models


Two rotor hub configurations are investigated for loads and wake analysis. The computational modeling of this complex geometries was enabled by the use of the overset body-fitted unstructured approach described in this work. Both hub models were studied experimentally in the John J. Harper low speed wind tunnel located within the Daniel Guggenheim School of Aerospace Engineering of Georgia Institute of Technology.

7.1.1 Baseline Configuration


Experiments on a generic four-bladed hub model were conducted initially. The model (Fig. 52) is approximately one-quarter scale of a ten-ton helicopter. The complete model includes a number of geometric components found in a typical rotor hub: hub plates, blade shanks, swashplate, pitch links, drive shaft and requisite hardware (nuts, bolts, etc). Some components found in full-scale rotors, such as hydraulic lines and control wire bundles, were not included in the model. While there were different models with plugged and unplugged rotor shanks, only the plugged shank model was tested with rotation. A more complete detailed description of the experiments can be found in Refs. [100, 101].

The experimental evaluation was initiated with static hub tests. Various azimuthal orientations at zero angle of attack were tested. Force data were obtained in 15° increments over a quarter revolution (since this is a four-bladed hub) for a range of available tunnel speeds. The plugged shank configuration was evaluated over the similar wind tunnel speed ranges and for rotation rates up to 240 rpm and was the

only configuration modeled computationally. Table 9 describe the dimensions of the different hub components. The reference area is 0.04057 m^2 , based on the frontal area of the hub taken at the 45° azimuthal position.



(a) 0° azimuth view



(b) 45° azimuth view

Figure 52: Baseline hub model studied by both computational and experimental campaigns.

Table 9: Hub component dimensions at 1/3.5-wind tunnel model.

	Diameter (m)	Length (m)	Width (m)	Height (m)
Blocks	-	0.0683	0.0694	0.0489
Driveshaft	0.0267	-	-	0.3194
Hub mount	0.1524	-	-	0.00635
Hub plates	-	0.3124	0.3124	0.00953
Pitch Links	0.0127	-	-	0.1761
Shanks	0.0349	-	-	0.0889
Swashplate	0.3556	-	-	0.00635

Table 10: Free-stream velocities and associated Reynolds number for the different flow conditions.

U_∞ (m/s)	Re_D
8.941	0.29×10^6
13.41	0.43×10^6
22.35	0.72×10^6

The wind tunnel free-stream density was 1.1901 kg/m^3 with a reference kinematic viscosity of $1.496 \times 10^{-5} \text{ m}^2/\text{s}$. The different free-stream velocities studied are listed in Table 10, along with the Reynolds number based on biggest cross-section diameter ($D = 0.4862 \text{ m}$) Re_D and free-stream velocity.

In addition to the force measurements on the model, particle image velocimetry (PIV) and single axis hot-wire probes were applied at selected wake locations to measure the wake velocity and frequency spectra of velocity fluctuations, respectively. The hot-wire anemometry data were provided in the form of velocities of the form V_h as well as perturbations of this velocity.¹ The quantity V_h is the effective hot-wire measured velocity given by

$$V_h = \sqrt{u^2 + k_v v^2 + w^2}, \quad (29)$$

where $k_v = 0.2$ as reported by experimental calibration [101].

¹Retrieved from <http://www.adl.gatech.edu/expaero/hubdrag/> on July 2, 2013, and verified through November 10, 2013.

Table 11: Matrix of flow conditions and analysis methods.

U_∞	8.941 m/s	13.41 m/s	22.35 m/s
120 rpm		Forces	
240 rpm	Forces, Velocities	Forces, Velocities	Forces




Figure 53: Hub model within John J. Harper wind tunnel test section.

7.1.2 Scissors Configuration

A modified model, which added a scissors assembly shown in Fig. 54, was also analyzed based on recommendations from industry. There was no PIV data obtained in this experiment, instead hot-wire anemometry was used to obtain fluctuating wake velocities to augment the load cell force data. Table 11 describes the different flow conditions analyzed and the analyses methods performed for the scissors configuration. Forces are computed for each simulation performed, while velocity fields (mean and fluctuating) were computed and correlated (wherever applicable) for $U_\infty = 8.941$ m/s and $U_\infty = 13.41$ m/s at 240 rpm.




Figure 54: Hub configuration with and without scissors.

7.2 *Simulation Approach*

Time-accurate predictions of the static hub and unsteady wake were simulated using the hybrid RANS-LES turbulence approach (GT-HRLES) [102] described in Section 2.1. Computations were performed *a priori* of correlations with experimental data. The background grid included the wind tunnel test section (see Fig. 53), where the tunnel walls were modeled as inviscid surfaces. The wind tunnel mount, as well as details of the hub, including pitch links, bolts and nuts were included in the configuration, as previously illustrated in Figs. 52 and 53. Each component was identified in the simulations separately, which resulted in sets of integrated loads for each component, as well as the entire assembly. This ability to group and monitor the unsteady history of loads on various surfaces is beneficial in order to assess the drag predictions against theory and experiment.

A baseline overset grid of 11.1 million nodes was adapted using the vorticity adaptation formulation discussed in Section 2.1.4 with $F_{tol} = 0.005$, such that at least 7-10 nodes are prescribed to resolve dominant vortices at one diameter downstream of the hub ($x = 1D$). Periodic grid adaptation was performed with the metric sampled

over one hub revolution. The boundary layer, consisting of 35 normal cell layers, had a maximum $y^+ = 0.35$ for all these simulations. The importance of the adaptation is emphasized in the shed vorticity of Figs. 55 and 56. The adapted grids provide significantly higher fidelity of the unsteady shed wake. This is particularly important in regions where the shed wake impacts on other hub components, as observed for the shed wake of the hub upstream of the main strut in Fig. 56.

To mimic the turbulence intensity of the wind tunnel, a free-stream turbulence intensity $Tu = 0.2\%$ was applied to the simulation. This value was chosen based on prior-reported turbulence in wind tunnels of the same generation as the Harper wind tunnel [16]. During 2012, there were major renovations to the wind tunnel, and it was reported that the wind tunnel turbulence intensity reduced to 0.03% [101]. However, to permit consistent correlations with the pre-renovation simulations, the original turbulence intensity was maintained.

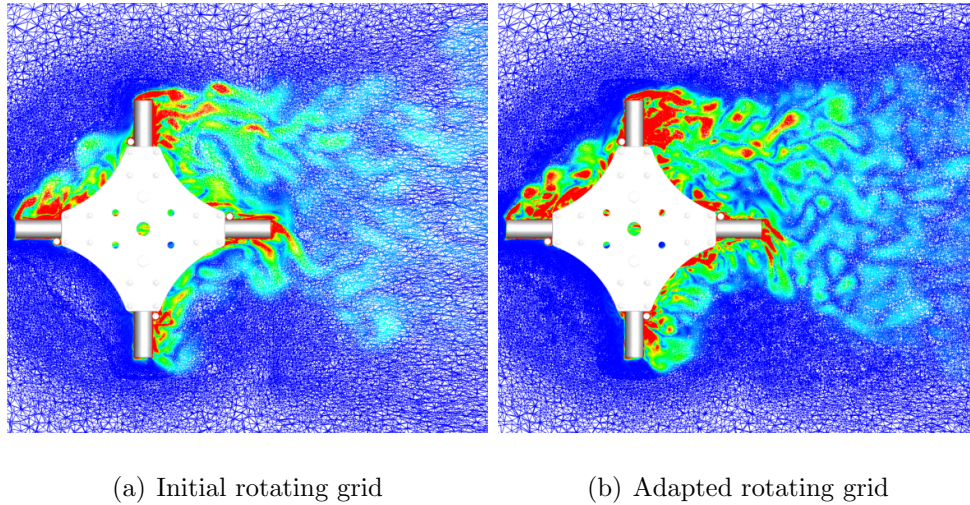


Figure 55: Planform (looking down) view of the vorticity magnitude for the initial and adapted grid.

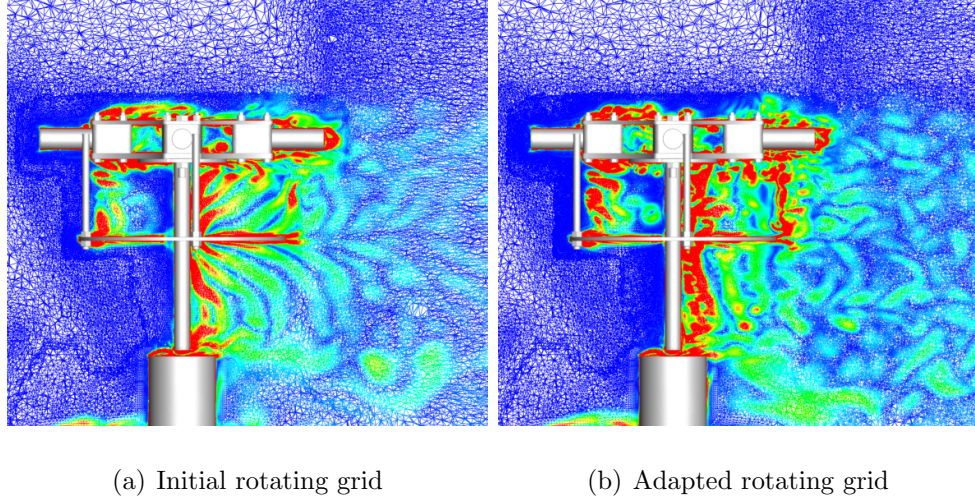


Figure 56: Left (looking forward) side view of the vorticity magnitude for the initial and adapted grid.

7.2.1 Temporal Resolution Study

Prior to conducting computations on the hub configurations, a temporal resolution study was performed to verify a sufficient time step for all analyses. Implicit time integration using the BDF2-opt scheme [79], described in Section 2.1, was applied for these simulations. A dimensional time step of 0.6944 ms was selected to be equivalent to a one degree azimuthal sweep of the 240 rpm rotating hub. Based on this rotation rate, each revolution corresponds to 4 Hz. Using this time step, a fluid particle traverses the characteristic hub length in about 80 time steps. Each time step was augmented with an average of 25 Newton subiterations (up to a maximum of 40 subiterations) to increase the temporal accuracy of the simulation. A temporal error controller maintained a specific residual reduction (1% of the estimated temporal error) so that the number of subiterations at each time step varies.

Convergence of the time step is demonstrated in Fig. 58, where velocity perturbation spectra are compared at four locations on the advancing side of the wake at $x = D/2$ (Fig. 57). In these plots, the frequency ranges up to 180/rev (720

Hz). The blending function is nearly zero here, which indicates that LES behavior is expected in this region of the wake. The time steps studied corresponded to 0.5° ($\Delta t = 0.3472 \times 10^{-3}$ s) and 1° ($\Delta t = 0.6944 \times 10^{-3}$ s) azimuthal sweeps. Prediction of the 4/rev (16 Hz) and 8/rev (32 Hz) features that are prominent at these wake locations is insensitive to time step refinement. Further, for frequencies higher than 100 Hz, a broad spectrum of scales is present showing evidence that the turbulence model is operating in LES mode. Based on these assessments, the dimensional time step was selected to be equivalent to a 1° azimuthal sweep of the rotating hub for all subsequent simulations.




Figure 57: HRLES blending function for $U_\infty = 13.41$ m/s (static configuration). Hot-wire sampling locations are plotted.

7.3 Adaptation Strategy Validation

The sensitivity of the drag prediction to the number adaptation iterations resulting from the adaptation strategy is detailed in Table 12. Three adaptation iterations were initially applied to the static hubs, where the hub is oriented at 0° and 45° . Substantial drag change (as much as 16%) was observed after the first adaptation




Figure 58: Comparison of normalized power spectral density (PSD) predicted using two different time steps for the $U_\infty = 13.41$ m/s rotating hub at 240 rpm. The wake locations (1-4) correspond to Fig. 57. (Note: Data are presented using the final adapted grid.)

iteration, but subsequent iterations yielded less than 1% change in drag. As the number of grid nodes increases with each adaptation cycle (as high as 50%), further iterations are deemed impractical for industry applications since prior literature [36, 37] from industry users have indicated that a 5% drag error is sufficient. Therefore, all hub simulations applied only one adaptive iteration. As detailed in Table 12, both orientations were adapted three times, but the 0° hub was adapted for two more iterations in order to confirm the trend of convergence. Figure 59 plots the drag coefficient over the five adaptation iterations. The drag for each adapted grid is within the 1% error bar of the final drag coefficient obtained at the fifth adaptation iteration. While small subsequent drag change is reported at every iteration, an asymptotic convergence to drag will not necessarily be achieved because adaptation

Table 12: Effect of adaptation on drag prediction for baseline static simulations (Note: Initial grid size is 11.7 million nodes.)

Adaptation Iteration	0° orientation		45° orientation	
	% Drag Change	Grid size (million nodes)	% Drag Change	Grid size (million nodes)
1	+3.3%	18.6	+15.8%	21.6
2	−0.47%	27.0	+0.77%	32.6
3	+0.99%	43.3	+0.13%	53.6
4	+0.72%	61.4	-	-
5	−0.53%	67.8	-	-

to local errors in vorticity bears no formal correlation to drag convergence.






Figure 59: Drag coefficient convergence with adaptation for the 0° static hub orientation. An error bar of 1% above and below the final drag coefficient is outlined (dashed line).

The adaptation strategy is validated by comparing time-averaged velocity profiles obtained via PIV at $x = 1D$ for the baseline configuration. The velocity profile is measured at the hub level, corresponding to $z = 0$, located at the hub's symmetry plane. By applying grid adaptation, good correlation between experimental and computational time-averaged wake velocity deficits is obtained, as plotted in Fig. 60.

The asymmetry of the wake for both the static and rotating hub simulations are



(a) Static hub at 0° orientation



(b) Rotating hub at 240 rpm

Figure 60: Comparison of time-averaged PIV and CFD hub centerline wake velocity deficits at 1 hub diameter downstream ($x = 1D$) along the tunnel axis. Experimental data obtained from Ref. [101].

apparent. As illustrated in the static hub at 0° orientation, shown in Fig. 60 (a), the slopes defining the wake deficits are well captured. The three primary velocity local minima ($y/D = -0.5, 0, 0.5$) are also predicted using the adapted grid. Figure 60 (b) illustrates that the rotating hub wake is predicted more accurately than the static hub wake. The application of grid adaptation in both these simulations is beneficial in accurately predicting the magnitudes and locations of these velocity local minima. The slope and features of the velocity deficit, exclusive of a small region of y/D between -0.3 and -0.1, are captured by computations. Overall the errors result in less than 5% difference in the drag from integrating the adapted grid's velocity profiles.

7.4 *Static & Dynamic Loads*

The first simulations undertaken were to evaluate the grid requirements and the ability of the computational simulations to predict the hub forces. With the requisite number of normal cell layers within the boundary layer, as discussed in Section 7.2 and Liggett et al. [103], the static hub drag is predicted very accurately by the computational approach using the initial unadapted grid, as depicted in Fig. 61. The drag is linear with the squared free-stream velocity. The maximum error occurs at the lowest free-stream velocity ($U_\infty = 5$ m/s) with an error of 5%. Most errors are within 1%-2% over the speed range.

Table 13: Drag coefficient predictions illustrating effect of fuselage on computed hub drag.

	C_D
Static hub at 0° orientation	1.2702
Static hub at 45° orientation	1.2271
Rotating hub at 240 rpm	1.2477




Figure 61: Correlation between experimental load cell and computational data for the 0° static hub model. Experimental data obtained from Ref. [101].

Comparison of the drag coefficient for the 0° and 45° static cases and the rotating case at 240 rpm results in drag coefficient values that are within 5% of one another. The maximum drag occurs at 0° static orientation, while the minimum occurs at the 45° static orientation. The drag changes are directly related to the minimization of the frontal flat plate area as the orientation changes. This change results in fewer high pressure (near stagnation) areas, that are instead now suction areas, as illustrated in Fig. 62. The rotating hub drag and pressure coefficient behavior falls between these two static results. Observing the pressures at the 0° azimuth between the static and rotating case (Figs. 62 (a) and (c)), the pressure distribution is fairly symmetric for the static case. For the rotating hub, the pressure progressively increases with radial distance on the advancing side (left side on the images), while the retreating side shank-block assembly shows a large low pressure region. The pressure variation with radial distance is expected, similar to the case of a full rotor.




Figure 62: Surface pressure coefficient distribution on the rotor hub ($U_\infty = 8.941$ m/s). Free-stream flow enters normal to the figure.

7.4.1 Component Drag Analysis

The different hub anatomy are defined in the Fig. 63 in order to differentiate the various hub components for drag contribution. Figure 64 breaks down the contribution of drag arising from these components from the static and rotating cases studied. The importance of modeling the hardware is apparent since it contributes nearly 10%-15% of the total drag for all cases. The component drag breakdown illustrates that

drag difference between the 0° and 45° static orientations occurs across nearly all components. One exception is the shanks, where a small increase is noticed because in the 45° azimuthal orientation all the shanks have more some exposure to the flow and oncoming wakes. Rotation at 240 rpm causes a significant decrease in drag from the blocks. There is more drag contribution from the pitch links and shanks due to rotation.




Figure 63: Definition of the different hub components.

Analysis of the change in the individual component contributions to drag (Fig. 64) provides some insight into the cause of the differences in the total drag, and are directly related to the surface pressure changes in Fig. 62. The drag from the cylindrical components for the 0° static orientation and the rotating case are presented in Tables 14 and 15, respectively. Using Fig. 52 to aid in the analysis of the four different assemblies, the differences in the physics is apparent. The drag contribution from the cylindrical components is compared with experimental results from Hoerner [104], where geometrically two-dimensional, interference-free drag data are reported. The ability of the HRLES turbulence model to accurately predict cylinder drag has been

evaluated extensively by Lynch and Smith [102]. Since the same methodology and recommended practices of that effort have been applied here, drag disagreement of these cylindrical components from their infinite cylinder experimental values can be attributed to interference effects and three-dimensional finite cylinder effects.




Figure 64: Component drag breakdown for the $U_\infty = 8.941$ m/s simulations.

The critical Reynolds number (Re_{cr}), marking the onset of transition to turbulence, for infinite cylinders in a cross-flow is approximately 300,000-400,000 [104]. At this Reynolds number ($Re = 0.29 \times 10^6$), the driveshaft Reynolds number is subcritical ($Re_d = 15,200$) and the drag coefficient is approximately 1.20. The driveshaft exhibits a drag value about 14% lower than the experimental drag prediction, attributable to interference effects from the swashplate and the pitch links. Similar lower drag trends are observed with the pitch links, which are due to interference drag arising from pitch link assemblies (see Fig. 52 (a)). Pitch link 4 obtains the closest cylinder drag in comparison to the experimental value at both Reynolds number scales (16% lower). This is expected given its orientation where minimal interference effects and a nearly unperturbed free-stream flow are encountered. Shank interference causes a reduction in drag for pitch links 1 and 2; pitch link 1 is affected by the presence of its shank and the assembly of pitch link 2 is on the leeward side of its shank.

Table 14: Drag coefficient tabulation for cylindrical components at 0° static azimuth orientations.

Component	Re_d	C_D
Driveshaft	15,200	1.0299
Pitchlink 1	7,600	0.7829
Pitchlink 2	7,600	0.8653
Pitchlink 3	7,600	0.4501
Pitchlink 4	7,600	1.003
Shank 2	21,000	0.5935
Shank 4	21,000	0.5129

Table 15: Average drag for rotating cylindrical components at $U_\infty = 8.941$ m/s.

Component	Driveshaft	Pitchlinks	Shanks
Rotating at 240 rpm	1.0188	0.8681	0.5225

Pitch link 3 exhibits interference effects due to its location farthest downstream with respect to other components. Shanks 2 and 4, which should nominally compare with Hoerner [104] at drag values of 1.20 encounter interference and/or finite aspect ratio effects that result in significant reduced drag from the two-dimensional experimental values.

For the rotating hub (Table 15), the driveshaft drag again has close to the static value and is about 15% lower than the predicted value. As expected, the average drag of the pitch links and the shanks is much lower than that expected for interference-free behavior.

7.4.2 Time-Averaged and Vibratory Loads Analysis

The computational predictions of the effect the scissors on the loads is summarized in Table 16. The additional 10% drag increase at 120 rpm predicted by computational was also predicted by the experiment [101]. The net drag increase at 240 rpm varies with the free-stream velocity and does not correlate with the drag increase at 120

rpm. Also, the drag increase due to the scissors appears to decrease with increasing advance ratio or free-stream velocity. It is inconclusive whether the effect of the drag increase due to the scissors is a function of the free-stream velocity primarily since simulations were not performed at the lower rotor speed at different free-stream velocities.

The unsteady drag and side force coefficient harmonics for the scissors hub are plotted in Figs. 65 and 66, respectively. The effect of the scissors can be assessed via these plots on the contribution to these forces. It is clear that the scissors have a major contribution in the two-per-rev harmonic for both forces. On the other hand, the baseline configuration exhibits dominance in the four-per-rev harmonic. This is because the baseline configuration comprises four assemblies (shanks, blocks, pitch-links), which contribute in four-per-rev fashion. For both forces, the four-per-rev and eight-per-rev harmonics are prevalent and the four-per-rev energy is generally greater than that of the eight-per-rev. The exceptions are the drag harmonics at $U_\infty = 8.941$ m/s, where the eight-per-rev energy is stronger for both configurations. Frequencies greater than four-per-rev are caused by aerodynamic wake structures that impinge on the oncoming geometry; at $U_\infty = 8.941$ m/s, the advance ratio is smallest ($\mu = 1.463$) indicating that the wake influence and impingement on the oncoming blade on the advancing side is the greatest. This is illustrated by a qualitative comparison of the wakes of two simulations in Fig. 67. At the lower speed, there are significant wake shedding that impinges on the advancing strut (top) from the forward strut (left) compared to the wake shedding for the higher velocity. These structures have a greater contribution on the advancing blade higher harmonics, where a significant drag contribution is obtained. This explains the discrepancy observed in Fig. 65(a).

Table 16: Computationally predicted drag coefficients showing the effect of scissors at 240 rpm.

Flow Condition	Configuration	C_D	% increase
$U_\infty = 13.41$ m/s at 120 rpm	without scissors	1.265	10.6%
	with scissors	1.397	
$U_\infty = 8.941$ m/s at 240 rpm	without scissors	1.248	17.2%
	with scissors	1.463	
$U_\infty = 13.41$ m/s at 240 rpm	without scissors	1.291	11.3%
	with scissors	1.437	
$U_\infty = 22.35$ m/s at 240 rpm	without scissors	1.237	6.1%
	with scissors	1.312	






Figure 65: Comparison of unsteady drag harmonics of the rotating hub with scissors.




(a) $U_\infty = 8.941$ m/s




(b) $U_\infty = 13.41$ m/s

Figure 66: Comparison of unsteady side force of the rotating hub with scissors.



(a) $U_{\infty} = 8.941 \text{ m/s}$ ($\mu = 1.463$)



(b) $U_{\infty} = 13.41 \text{ m/s}$ ($\mu = 2.195$)

Figure 67: Q-criterion contours of the wake at $Z = 0.0 \text{ m}$ at 240 rpm.

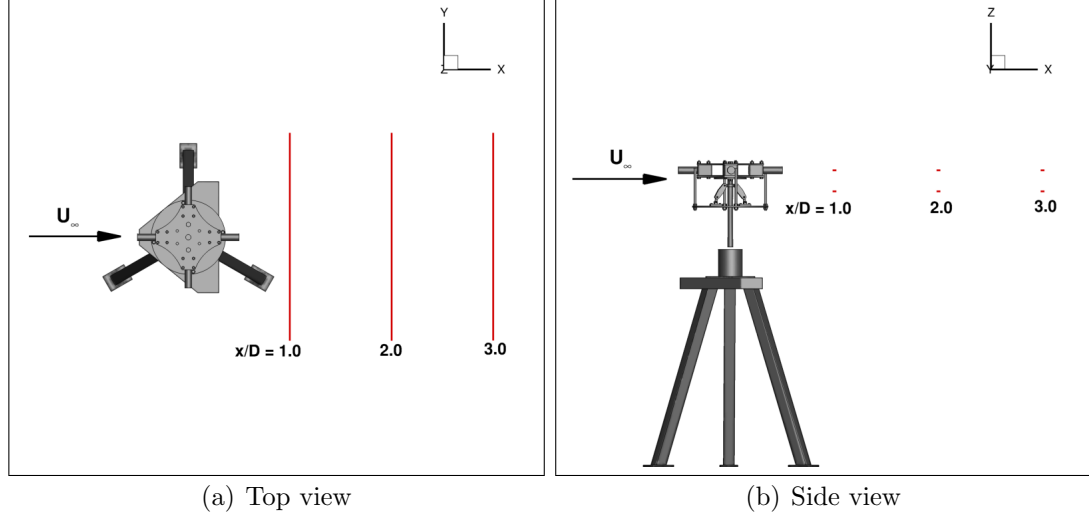


Figure 68: Wake velocity measurement locations in space. (Red line indicates traverse of each profile).

7.5 Wake Structure

The velocity profiles and frequency content is important to understand in terms how the flow field separates in the near-hub region. The effect of the scissors on the shape of the time-averaged velocity profile is assessed along with as well as the effect of the dominant frequencies. The dissipation of certain wake features with wake distance also needs to be understood from the stand point of understanding the adverse effect of the wake on the downstream geometry.

Time-averaged and fluctuating wake velocities were obtained from the computational simulations at $x = 1D$, $x = 2D$, and $x = 3D$ downstream of the hub for comparing trends and analyzing velocity spectra. The spatial locations of these probes are delineated in Fig. 68. In addition to the hub level at $z = 0.0$, data are obtained at wake locations directly behind the scissors at $z = -0.204 D$ so that the effect of the scissors on the downstream wake may be determined.

7.5.1 Wake Velocity Profiles

The time-averaged profiles occur over two revolutions since the velocity profile obtains convergence within that period. Figure 69 plots comparisons between averaging the velocity over one and two revolutions. The integrated velocity is within 1% of one another.






Figure 70: Comparison of computational wake velocity profiles for the different configurations at one hub diameter in the wake.

Since the scissor components are located near the hub center in the spanwise direction (Fig. 54), the predicted wake extent from the scissor components is hypothesized using engineering first principles to remain in the wake core ($-0.3 > y/D > 0.4$). The free-stream velocity near the tunnel walls should be recovered for the same free-stream conditions.


The computational data for both the baseline and scissors configurations are compared in Fig. 70. Time-averaged local streamwise velocity u data are presented. The differences in these time-averaged velocities lie primarily in the core of the wake. The velocity minimum is increased and translates to the left by $y/D = 0.1$ by the addition



(a) $x = 1D$



(b) $x = 2D$



(c) $x = 3D$

Figure 69: Time-averaged wake velocity comparison of averaging sampling size.

of the scissors. The superposition of the shear layers between the wake core and free stream indicate insignificant contributions from the addition of the scissors. The free stream outside of the wake is also identically recovered from both simulations, which is required if the free-stream velocities are identical. Comparisons with experimental data for time-averaged wake profiles were not made since inconsistencies were found between the experimental data sets provided by PIV, discussed in Section 7.3, and the hot-wire anemometry data provided for the scissors configuration. A thorough discussion of these inconsistencies is found in Ref. [105].

7.5.2 Time-Averaged Velocity Trends

Time-averaged velocity profiles resulting from the computational predictions at wake locations $x = 1D$, $x = 2D$, and $x = 3D$ for the different free-stream velocities at a rotation speed of 240 rpm is shown in Figs. 71 – 73. Several observations can be made. The velocity deficit generally decreases (the wake velocity approaches the free stream recovery) at a given spanwise location as the downstream distance from the hub increases. This decrease is not linear with respect to the distance, rather most recovery occurs during one to two hub diameters downstream. Additionally, the increase in downstream distance tends to diffuse or smooth the local extrema in the core. The wake bias toward the right (advancing side) decreases as the downstream distance increases, with the maximum changes before two hub diameters in the wake is reached. This wake skew is apparent by the larger velocity deficits on the wake’s advancing side.

As the free-stream velocity increases (higher advance ratios), the asymmetry of the wake diminishes in the wake locations examined. This implies that at higher advance ratios the wake skew tends to disappear, which is expected since the effect of rotation is limited due to the dominance of the free-stream energy.




Figure 71: Time-averaged wake velocity comparison at $Z = 0.0$ m at different downstream locations. The rotor hub conditions are at 240 rpm and $U_{\infty} = 8.941$ m/s.




Figure 72: Time-averaged wake velocity comparison at $Z = 0.0$ m at different downstream locations. The rotor hub conditions are at 240 rpm and $U_\infty = 13.41$ m/s.




Figure 73: Time-averaged wake velocity comparison at $Z = 0.0$ m at different downstream locations. The rotor hub conditions are at 240 rpm and $U_\infty = 22.35$ m/s.


7.5.3 Velocity Spectral Analysis

Wake velocity spectra from experiment and computation are plotted and compared in Figs. 74 - 79 at the hub level ($Z = 0.0$ m). The contours indicate the power spectral density (in logarithmic scale). The frequency axis is in the vertical direction, marked by the per-rev frequency so that wake structures may be traced to their originating geometry. Noting that the contour ranges are the same for both computation and experiment, a general trend is observed that the computational predictions for the power spectra content are greater than their experimental counterparts. Also, since the computational data are averaged for many more data points over the span and over two revolutions, the data appear to be noisy. However, clear correlations are observed particularly in the data at the four-per-rev and two-per-rev levels.


The four-per-rev levels structures are generally present over the range the of the wake, but they are much stronger in the retreating portion of the wake (left side

of the plots). These structures emanate primarily from the blade shanks and pitch links. These structures also persist to the $X = 3D$ downstream location, the farthest point measured by experiment. The eight-per-rev and sixteen-per-rev content are much stronger for $U_\infty = 13.41$ m/s, but they diminish more rapidly at the downstream locations compared to the four-per-rev structures. Most of these structures are strongest in the retreating portion of the wake, where the spanwise mixing of the flow is much less compared to the advancing rotor wake due to the skewed wake shape. A considerable two-per-rev shedding is also observed in both sets of data and is intermittently prominent across the wake.

Additionally, traces of six-per-rev shedding are observed. This corresponds to a frequency of 24 Hz, which can be attributed to the vortex shedding off the blade shanks ($D = 0.0349$ m). Sakamoto and Arie [106] have observed that for cylinders of aspect ratio ≈ 2.5 , which corresponds to the aspect ratio of the blade shanks, the Strouhal number ($St = \frac{fD}{U}$) can vary between 0.11 and 0.14, depending on the portion of the cylinder submerged in the boundary layer. Using the tip speed V_{tip} , as the characteristic velocity of shedding, a Strouhal number of 0.135 is obtained for this configuration. Therefore, the six-per-rev structures here can be traced to the rotating blade root. This is similar to the findings reported by Reich et al. [39], where a six-per-rev wake structure was attributed to a Strouhal shedding due to the hub arms of a four-bladed rotor hub.




(a) Experiment




(b) Computational Simulation

Figure 74: Wake PSD comparison at $z = 0.0$ m and $x = 1D$. The rotor hub conditions are at 240 rpm and $U_\infty = 8.941$ m/s. Experimental data obtained from Ref. [101].




(a) Experiment




(b) Computational Simulation

Figure 75: Wake PSD comparison at $z = 0.0$ m and $x = 2D$. The rotor hub conditions are at 240 rpm and $U_\infty = 8.941$ m/s. Experimental data obtained from Ref. [101].




(a) Experiment




(b) Computational Simulation

Figure 76: Wake PSD comparison at $z = 0.0$ m and $x = 3D$. The rotor hub conditions are at 240 rpm and $U_\infty = 8.941$ m/s. Experimental data obtained from Ref. [101].




(a) Experiment




(b) Computational Simulation

Figure 77: Wake PSD comparison at $z = 0.0$ m and $x = 1D$. The rotor hub conditions are at 240 rpm and $U_\infty = 13.41$ m/s. Experimental data obtained from Ref. [101].




(a) Experiment




(b) Computational Simulation

Figure 78: Wake PSD comparison at $z = 0.0$ m and $x = 2D$. The rotor hub conditions are at 240 rpm and $U_\infty = 13.41$ m/s. Experimental data obtained from Ref. [101].



(a) Experiment




(b) Computational Simulation


Figure 79: Wake PSD comparison at $z = 0.0$ m and $x = 3D$. The rotor hub conditions are at 240 rpm and $U_\infty = 13.41$ m/s. Experimental data obtained from Ref. [101].

Figures 80 - 82 plot the computational velocity spectra in the wake at $z = -0.204D$, illustrating the wake character at behind the scissors geometry. Unlike the wake at the hub level ($z = 0.0$), the wake extent here is narrower, but the rightward bias is still observed, which is expected. It is clear that in addition to the four-per-rev structures, there are also two-per-rev structures that are not as prominently at the hub level. The two-per-rev structure persists strongly up until the $x = 3D$ location

whereas the four-per-rev features tend to diminish.




(a) $U_\infty = 8.941 \text{ m/s}$




(b) $U_\infty = 13.41 \text{ m/s}$

Figure 80: Wake PSD comparison at $z = -0.204D$ and $x = 1D$ at 240 rpm.




(a) $U_\infty = 8.941$ m/s




(b) $U_\infty = 13.41$ m/s

Figure 81: Wake PSD comparison at $z = -0.204D$ and $x = 2D$ at 240 rpm.



(a) $U_\infty = 8.941$ m/s



(b) $U_\infty = 13.41$ m/s

Figure 82: Wake PSD comparison at $z = -0.204D$ and $x = 3D$ at 240 rpm.

CHAPTER VIII

CONCLUSIONS

An enhanced overset grid adaptation methodology targeted for complex rotating systems associated with external aerodynamic flows has been developed in the form of two contributions to the community:

- Feature-based grid adaptation on overset, dynamic grids for complex, time-dependent wake flows has been developed and successfully demonstrated with an unstructured solver using URANS and hybrid RANS-LES turbulence approaches. A unique feature of this contribution is that all component grids of the overset system are adapted, enabling features to be captured as they arise from their originating surfaces. Additionally, this metric-based adaptation method is compatible with BDF-based implicit time integration schemes.
- An efficient and robust parallel localization algorithm for interpolation of unstructured overset grid systems has been developed for massively distributed systems. This permits second-order accurate solution transfers for adaptive grids resulting in low interpolation errors.

Temporal and spatial design order accuracy has been verified for this improved methodology. These improvements have been applied and demonstrated on flows for a rotor-fuselage interaction configuration and a generic four-bladed rotating hub.

The methodology has been successfully demonstrated on a canonical rotor-fuselage interaction case that has been widely used for validating computational methods. In particular, the following observations can be made:

- A single grid adaptation framework has been extended to overset grid systems

to allow for large relative motion. This required a strategy to identify, track, and reconstruct the component grid where each node is located.

- Correlations with experiment of time-averaged and instantaneous pressures illustrates solution dependency on the feature selection and formulation.
- Flow field features and fuselage pressures converge within one or two grid adaptation cycles. The most accurate adaptation scheme for the inviscid simulations is the *vorticity-mixed scheme*. For the HRLES simulations, both the *vorticity-mixed scheme* and *Q-criterion-mixed scheme* show similarly accurate results.
- The HRLES turbulence model is significantly superior than the $k\omega$ -SST RANS model in the preservation of rotor wake features. The high eddy viscosity prediction resulting from the $k\omega$ -SST model significantly dissipates the features contributing to vortex-fuselage interaction.
- The ability of the *Q-criterion-mixed scheme* to capture some features in viscous near-body flows is limited by its inability to adapt in the vortex sheet. The Q-criterion targets tip vortices and disregards adaptation in the vortex sheet, which in many instances require refinement for accurate wake predictions.
- Modeling the flow as either inviscid or fully turbulent affects predicted tip vortex location with respect to experiment. The *vorticity-mixed scheme* from the inviscid simulation shows a spatial lag of the tip vortex location (2% rotor radius), while the same for the HRLES simulation is within ($\pm 1\%$ rotor radius).
- The sensitivity of the method to the adaptation interval size is observed in the prediction of the secondary interaction, where the smaller interval computations provide marginal improvements against experiment. The prediction of the secondary interaction location also improves with the smaller intervals. The

smaller intervals predict consistent physics as that of the periodic interval but with substantial cost increases.

- The aerodynamics analysis of rotor-fuselage interaction indicate that the presence of the rotor blades cause the strongest high frequency vibratory response in the fuselage's side force, rolling moment, and yawing moment. These vibratory loads in a realistic helicopter would adversely affect the vehicle's handling qualities.

Using a combination of experimental and computational methods, the characterization of the integrated loads and complex wake field of a scaled helicopter main rotor hub has been obtained. From this investigation, it can be concluded that

- Large regions of separated flow contribute to considerable drag from the complex hub. The azimuthal variation of the static hub model shows variation in drag corresponding to the projected frontal area of the hub. For the four-bladed hub configuration, the 45° azimuthal orientation yields similar drag measurements to the mean values obtained for the rotating hub.
- Strong correlations between *a priori* computational results and experimental data permit computational predictions to explain the fundamental physics of complex rotating hubs with high confidence. The overset adaptation capability allows grid modifications across both background and near-body grids, which are needed to accurately resolve the wake interactions that occur in both the near-body and far-field grids.
- The simulation of the hub static configuration at 0° quantified the interference drag on components such as the driveshaft, pitch links and shanks. Divergence from canonical experimental drag values for cylinders in cross-flow implied interference effects from nearby geometry. These interference effects were confirmed by analysis of computational surface and flow field features.

- The addition of scissors results in a dominant contribution to the loads from the two-per-rev frequency. Four-per-rev structures were observed in both configurations and their persistence is strong on the retreating hub side up to three diameters downstream of the hub. Higher harmonic structures are much more pronounced at the higher free-stream flow condition but dissipate sooner than the four-per-rev features.
- An additional six-per-rev shedding was attributed to the vortex shedding from the blade shanks, confirming physics that were independently observed for a similar hub in a separate experimental effort [39].

:

8.1 Recommendations for Future Work

Several areas of improvement for both accurate and efficient computations of rotating turbulent flows are recommended. For grid adaptation methods, future work should include:

- Unsteady-overset adjoint-based grid adaptation

The uncertainty associated with the selection of the adaptation feature and interval size for complex multiscale flow phenomena imply that a more rigorous adaptation process, such as adjoint-based adaptation, may be more appropriate than feature-based adaptation. Future efforts should focus on the development of an adaptation process that will circumvent the issues identified with feature-based adaptation, as well as extension of the method to other rotorcraft-related applications. Recent extension of the adjoint-solver toward unsteady overset dynamic grid turbulent problems by Nielsen et al. [107] provides further impetus for a complimentary adaptation methodology. The extension of the metric-based adaptation methodology for overset dynamic grids in this thesis may be

leveraged toward this next goal.

- FUN3D boundary layer grid adaptation

Currently, the metric-based grid adaptation is confined to tetrahedral elements external to the boundary layer. The boundary layer height is either prescribed to the solver so that adaptation may be frozen within that region. Alternately, mixed-element boundary layers can be used since non-tetrahedral elements are frozen by default. Full mixed-element adaptation, either by generic element splitting or deformation needs to be developed so that boundary layer elements may be refined, particularly along the wall. Surface grid refinement, already available via GridEx [108] and CAPRI [109], should be used for full boundary layer adaptation.

- Adaptation to eliminate orphan nodes

In some scenarios, grid adaptation may lead to coarsened grids that may result in overset fringe nodes that do not have a suitable donor. Therefore, grid coarsening based on the formulation should be limited such that orphans are eliminated. This improvement would be fully realized if the solver has an internal overset capability, without requiring external libraries such as DiRTlib and SUGGAR++.

Recommended improvements for overset methods and general grid-to-grid interpolations include:

- Increase robustness of localizations

Surface and boundary layer grid adaptation may introduce nodes outside of the original domain in the new grid because the underlying CAD surface and discretized (planar) boundary faces are not coincident. This would require the ability to identify that those nodes need to be localized using data from the

neighboring surface nodes rather than full trilinear barycentric interpolation. Furthermore, these surface nodes may need to be localized using linear or bilinear (triangles) barycentric coordinates. This enhanced robustness feature would allow for the same localization method to be utilized toward grid sequencing for realistic viscous grids.

- FUN3D internal solver overset capability

Internal overset capability, without requiring linking to DiRTlib and SUGGAR++, would be advantageous for the FUN3D methodology since a generalized interpolation (and localization) method may be used toward both unsteady grid adaptation as well as overset moving-body assembly and interpolation and grid-sequencing.

- Conservative interpolation methods

Conservative interpolation is recommended for both the solution interpolation between overset component grids and solution transfers between adaptive grids. Additionally, incorporating the conservation property to the overset technique would ensure a consistent finite volume method. A possible method to overcome this is the mesh intersection method [63]. Possible challenges include the cost and memory requirement to perform the intersection in three dimensions.

Future work recommendations for rotating systems applications include:

- Advances in Hybrid-RANS LES methods

Enhancements to the existing HRLES model for more accurate modeling of turbulence kinetic energy dissipation are underway. The application of the newly implemented local dynamic k^{SGS} model (LDKM) is required since this would generalize the model, making it less dissipative due to the grid. Also, a

formal study of the consistency of the model to moving geometry and adaptive, overset grids is necessary.

- Analysis with different configurations

In addition to the rotor-fuselage interaction and the rotor hub examined in this thesis, the overset adaptive grid methodology should be applied toward more realistic geometries including aeroelastic applications where an extensive dataset of loads and wake velocities is available. This would enhance the fundamental knowledge of rotor and hub physics for the rotorcraft and wind turbine communities.

REFERENCES

- [1] SHERIDAN, P. and SMITH, R., “Interactional Aerodynamics-A New Challenge to Helicopter Technology,” *Journal of the American Helicopter Society*, vol. 25, no. 1, pp. 3–21, 1980.
- [2] STREMEL, P. M., “The Calculation of Rotor/Fuselage Interaction for Two-Dimensional Bodies,” Tech. Rep. TM-102855, NASA, August 1990.
- [3] 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, no. 1, pp. 5–21, 2006.
- [4] LYNCH, C. E., *Advanced CFD Methods for Wind Turbine Analysis*. PhD thesis, Georgia Institute of Technology, 2011.
- [5] ABRAS, J. N., *Enhancement of Aeroelastic Rotor Airload Prediction Methods*. PhD thesis, Georgia Institute of Technology, 2009.
- [6] O’BRIEN, D. M., *Analysis of Computational Modeling Techniques for Complete Rotorcraft Configurations*. PhD thesis, Georgia Institute of Technology, 2006.
- [7] FEJTEK, I. and ROBERTS, L., “Navier-Stokes Computation of a Wing/Rotor Interaction for a Tilt Rotor in Hover,” *AIAA Journal*, vol. 30, no. 11, pp. 2595–2603, 1992.
- [8] CHAFFIN, M. S. and BERRY, J. D., “Helicopter Fuselage Aerodynamics Under a Rotor by Navier-Stokes Simulation,” *Journal of American Helicopter Society*, vol. 42, no. 3, pp. 235–242, 1997.
- [9] BETTSCHART, N., “Rotor Fuselage Interaction: Euler and Navier-Stokes Computations with an Actuator Disk,” in *Proceedings of the 55th American Helicopter Society Annual Forum*, (Montreal, Canada), May 1999.
- [10] TADGHIGHI, H., “Simulation of Rotor-Body Interactional Aerodynamics: An Unsteady Rotor Source Distributed Disk Model,” in *Proceedings of the 57th Annual Forum*, (Washington, DC), American Helicopter Society, May 2001.
- [11] POTSDAM, M. and STRAWN, R., “CFD Simulations of Tiltrotor Configurations in Hover,” in *Proceedings of the 58th American Helicopter Society Annual Forum*, (Montreal, Canada), June 11–13, 2002.

- [12] RENAUD, T., POTSDAM, M., O'BRIEN, JR, D. M., and SMITH, M. J., "Evaluation of Isolated Fuselage and Rotor-Fuselage Interaction Using CFD," in *Proceedings of the 60th American Helicopter Society Annual Forum*, (Baltimore, MD), June 7–10, 2004.
- [13] RUFFIN, S. M., O'BRIEN, JR, D. M., SMITH, M. J., HARIHARAN, N. S., LEE, J. D., and SANKAR, L., "Comparison of Rotor-Airframe Interaction Utilizing Overset and Unstructured Grid Techniques," in *AIAA 42nd Aerospace Sciences Meeting*, no. AIAA-2004-0046, (Reno, NV), January 5–8, 2004.
- [14] NAM, H., PARK, Y., and KWON, O., "Simulation of Unsteady Rotor-Fuselage Interactions Using Unstructured Adaptive Sliding Meshes," *Journal of the American Helicopter Society*, vol. 51, no. 2, pp. 141–149, 2006.
- [15] STEIJL, R. and BARAKOS, G., "Sliding Mesh Algorithm for CFD Analysis of Helicopter Rotor-Fuselage Aerodynamics," *International Journal for Numerical Methods in Fluids*, vol. 58, no. 5, pp. 527–549, 2008.
- [16] BRAND, A. G., *An Experimental Investigation of the Interaction Between a Model Rotor and Airframe in Forward Flight*. PhD thesis, Georgia Institute of Technology, 1989.
- [17] LIOU, S. G., KOMERATH, N. M., and MCMAHON, H. M., "Velocity Measurements of Airframe Effects on a Rotor in Low Speed Forward Flight," *Journal of Aircraft*, vol. 26, no. 4, pp. 340–348, 1989.
- [18] O'BRIEN, D. M. and SMITH, M. J., "Analysis of Rotor-Fuselage Interactions Using Various Rotor Models," in *AIAA 43rd Aerospace Sciences Meeting*, no. AIAA-2005-0468, (Reno, NV), January 10–13, 2005.
- [19] O'BRIEN, JR, D. M. and SMITH, M. J., "Understanding the Physical Implications of Approximate Rotor Methods Using an Unstructured CFD Method," in *Proceedings of the 31st European Rotorcraft Forum*, (Florence, Italy), September 13–15, 2005.
- [20] LEE, J. K. and KWON, O. J., "Predicting Aerodynamic Rotor-Fuselage Interactions by Using Unstructured Meshes," *Transactions of the Japan Society for Aeronautical and Space Sciences*, vol. 44, no. 146, pp. 208–216, 2005.
- [21] STEIJL, R. and BARAKOS, G., "Computational Study of Helicopter Rotor-Fuselage Interactions," *AIAA Journal*, vol. 47, no. 9, pp. 2143–2157, 2009.
- [22] JUNG, M. S. and KWON, O. J., "Numerical Simulation of Unsteady Rotor Flow Using an Unstructured Overset Mesh Flow Solver," *International Journal of Aeronautical and Space Sciences*, vol. 10, no. 1, pp. 104–111, 2009.
- [23] XU, H. and YE, Z., "Euler Calculation of Rotor-Airframe Interaction Based on Unstructured Overset Grids," *Journal of Aircraft*, vol. 48, no. 2, pp. 707–713, 2011.

- [24] WILLIAMS, R. and MONTANA, P., “A Comprehensive Plan for Helicopter Drag Reduction,” in *Proceedings of the American Helicopter Society Symposium on Helicopter Aerodynamic Efficiency*, (Hartford, CT), March 6–7, 1975.
- [25] KEYS, C. and WIESNER, R., “Guidelines for Reducing Helicopter Parasite Drag,” *Journal of the American Helicopter Society*, vol. 20, no. 1, pp. 31–41, 1975.
- [26] HOFFMAN, J., “The Relationship Between Rotorcraft Drag and Stability and Control,” in *Proceedings of the 31st Annual Forum of the American Helicopter Society*, (Washington, DC), May 13–15, 1975.
- [27] KERR, A., “Effect of Helicopter Drag Reduction on Rotor Dynamic Loads and Blade Life,” in *Proceedings of the American Helicopter Society Symposium on Helicopter Aerodynamic Efficiency*, (Hartford, CT), March 6–7, 1975.
- [28] WAKE, B., HAGEN, E., OCHS, S., and MATALINIS, C., “Assessment of Helicopter Hub Drag Prediction with an Unstructured Flow Solver,” in *Proceedings of the 65th Annual Forum of the American Helicopter Society*, (Grapevine, TX), May 27–29, 2009.
- [29] SALTZMAN, E. and AYERS, T., “A Review of Flight-to-Wind Tunnel Drag Correlation,” in *AIAA 1st AIAA Flight Testing Conference*, no. AIAA-1981-2475, (Las Vegas, NV), November 11–13, 1981.
- [30] BUSHNELL, D., “Scaling: Wind Tunnel to Flight,” *Ann. Rev. of Fluid Mechanics*, vol. 38, no. 1, pp. 111–128, 2006.
- [31] GREGORY, J., PORTER, C., and McLAUGHLIN, T., “Circular Cylinder Wake Control using Spatially Distributed Plasma Forcing,” in *4th AIAA Flow Control Conference*, no. AIAA-2008-4198, June 23–26, 2008.
- [32] ROSHKO, A., “Perspectives on Bluff Body Aerodynamics,” *Journal of Wind Engineering and Industrial Aerodynamics*, vol. 49, no. 1-3, pp. 79–100, 1993.
- [33] SHEEHY, T. and CLARK, D., “A Method for Predicting Helicopter Hub Drag,” Tech. Rep. USAARMDL-TR-75-48, United Technologies Corp., Stratford, CT, Sikorsky Aircraft Div., 1976.
- [34] SHEEHY, T., “A General Review of Helicopter Rotor Hub Drag Data,” *Journal of the American Helicopter Society*, vol. 22, p. 2, 1977.
- [35] OCHS, S., MATALINIS, C., WAKE, B., and EGOLF, T., “Evaluation of Helios CFD Toolset for Faired Rotor-Hub Drag Prediction,” in *Proceedings of the 67th Annual Forum of the American Helicopter Society*, (Virginia Beach, VA), May 3–5, 2011.

- [36] BRIDGEMAN, J. and LANCASTER, G., “Predicting Hub Drag on Realistic Geometries,” in *Proceedings of the American Helicopter Society Aeromechanics Specialists’ Conference*, (San Francisco, CA), January 20–22, 2010.
- [37] BRIDGEMAN, J. and LANCASTER, G., “Physics-Based Analysis Methodology for Hub Drag Prediction,” in *Proceedings of the 66th Annual Forum of the American Helicopter Society*, (Phoenix, AZ), May 11–13, 2010.
- [38] HILL, M. and BRIDGEMAN, J., “Rotating Hub Drag Prediction Methodology,” in *Proceedings of the American Helicopter Society Future Vertical Lift Aircraft Design Specialists’ Conference*, (San Francisco, CA), January 18–20, 2012.
- [39] REICH, D., ELBING, B., BEREZIN, C., and SCHMITZ, S., “Water Tunnel Flow Diagnostics of Wake Structures Downstream of a Model Helicopter Rotor Hub,” in *American Helicopter Society 69th Annual Forum Proceedings, Phoenix, AZ*, May 21–23, 2013.
- [40] KANG, H. and KWON, O., “Effect of Wake Adaptation on Rotor Hover Simulations Using Unstructured Meshes,” *Journal of Aircraft*, vol. 38, no. 5, pp. 868–877, 2001.
- [41] KANG, H. and KWON, O., “Unstructured Mesh Navier-Stokes Calculations of the Flow Field of a Helicopter Rotor in Hover,” *Journal of the American Helicopter Society*, vol. 47, no. 2, pp. 90–99, 2002.
- [42] DINDAR, M., SHEPHERD, M., FLAHERTY, J., and JANSEN, K., “Adaptive CFD Analysis for Rotorcraft Aerodynamics,” *Computer Methods in Applied Mechanics and Engineering*, vol. 189, no. 4, pp. 1055–1076, 2000.
- [43] POTSDAM, M. and MAVRIPLIS, D., “Unstructured Mesh CFD Aerodynamic Analysis of the NREL Phase VI Rotor,” in *47th AIAA Aerospace Sciences Meeting*, no. AIAA-2009-1221, (Orlando, FL), January 5–8, 2009.
- [44] CAVALLO, P. A., SINHA, N., and FELDMAN, G., “Parallel Unstructured Mesh Adaptation Method for Moving Body Applications,” *AIAA Journal*, vol. 43, no. 9, pp. 1937–1945, 2005.
- [45] BERGLIND, T. and TYSELL, L., “Time-Accurate CFD Approach to Numerical Simulation of Store Separation Trajectory Prediction,” in *AIAA 29th Applied Aerodynamics Conference*, no. AIAA-2011-3958, (Honolulu, HI), June 27–30, 2011.
- [46] MEAKIN, R., “On Adaptive Refinement and Overset Structured Grids,” in *AIAA 13th Computational Fluid Dynamics Conference*, no. AIAA-1997-1858, (Snowmass Village, CO), June 29–July 2, 1997.
- [47] MEAKIN, R., “Adaptive Spatial Partitioning and Refinement for Overset Structured Grids,” *Computer Methods in Applied Mechanics and Engineering*, vol. 189, no. 4, pp. 1077–1117, 2000.

- [48] HENSHAW, W. and SCHWENDEMAN, D., “Parallel Computation of Three-Dimensional Flows Using Overlapping Grids with Adaptive Mesh Refinement,” *Journal of Computational Physics*, vol. 227, no. 16, pp. 7469–7502, 2008.
- [49] KANNAN, R. and WANG, Z., “Overset Adaptive Cartesian/Prism Grid Method for Stationary and Moving-Boundary Flow Problems,” *AIAA Journal*, vol. 45, no. 7, pp. 1774–1779, 2007.
- [50] CANONNE, E., BENOIT, C., and JEANFAIVRE, G., “Cylindrical Mesh Adaptation for Isolated Rotors in Hover,” *Aerospace Science and Technology*, vol. 8, no. 1, pp. 1–10, 2004.
- [51] DUQUE, E., BISWAS, R., and STRAWN, R., “A Solution Adaptive Structured/Unstructured Overset Grid Flow Solver with Applications to Helicopter Rotor Flows,” in *AIAA 13th Applied Aerodynamics Conference*, no. AIAA-95-1766, (San Diego, CA), June 19–22, 1995.
- [52] SANKARAN, V., SITARAMAN, J., WISSINK, A., DATTA, A., 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, FL), January 4–7, 2010.
- [53] 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, IL), June 28–July 1, 2010.
- [54] WISSINK, A., POTSDAM, M., SANKARAN, V., SITARAMAN, J., YANG, Z., and MAVRIPLIS, D., “A Coupled Unstructured-Adaptive Cartesian CFD approach for Hover Prediction,” in *Proceedings of the 66th Annual American Helicopter Society Forum*, (Phoenix, AZ), May 11–13, 2010.
- [55] WISSINK, A., JAYARAMAN, B., DATTA, A., SITARAMAN, J., POTSDAM, M., KAMKAR, S., MAVRIPLIS, D., YANG, Z., JAIN, R., LIM, J., and STRAWN, R., “Capability Enhancements in Version 3 of the Helios High-Fidelity Rotorcraft Simulation Code,” in *AIAA 50th Aerospace Sciences Meeting and Exhibit*, no. AIAA-2012-0713, (Nashville, TN), January 9–12, 2012.
- [56] CHADERJIAN, N. M. and AHMAD, J. U., “Detached Eddy Simulation of the UH-60 Rotor Wake Using Adaptive Mesh Refinement,” in *Proceedings of the 68th Annual American Helicopter Society Forum*, (Fort Worth, TX), May 1–3, 2012.
- [57] KAMKAR, S. J., WISSINK, A., SANKARAN, V., and JAMESON, A., “Feature-driven Cartesian Adaptive Mesh Refinement for Vortex-Dominated Flows,” *Journal of Computational Physics*, vol. 230, pp. 6271–6298, July 2011.

- [58] CAVALLO, P. A., ARUNAJATESAN, S., SINHA, N., and BAKER, T. J., “Transient Mesh Adaptation Using and Error Wake and Projection Method,” in *AIAA 44th Aerospace Sciences Meeting and Exhibit*, no. AIAA-2006-1149, (Reno, NV), Jan 9–12, 2006.
- [59] NAJAFIYAZDI, M., HABASHI, W. G., and FOSSATI, M., “Improved Transient-Fixed-Point Mesh Adaptation Using Orthogonality-Preserving Metric Intersection,” in *AIAA 20th Computational Fluid Dynamics Conference*, no. AIAA-2011-3691, (Honolulu, HI), June 27–30, 2011.
- [60] ALAUZET, F., FREY, P. J., GEORGE, P. L., and MOHAMMADI, B., “3D Transient Fixed Point Mesh Adaptation for Time-Dependent Problems: Application to CFD Simulations,” *Journal of Computational Physics*, vol. 222, no. 2, pp. 592–623, 2007.
- [61] ZHOU, C. H. and AI, J. Q., “Mesh Adaptation for Simulation of Unsteady Flow with Moving Immersed Boundaries,” *International Journal for Numerical Methods in Fluids*, vol. 72, no. 4, pp. 453–477, 2013.
- [62] ALAUZET, F. and OLIVIER, G., “Extension of Metric-Based Anisotropic Mesh Adaptation to Time-Dependent Problems Involving Moving Geometries,” in *AIAA 49th Aerospace Sciences Meeting and Exhibit*, no. AIAA-2011-896, (Orlando, FL), January 4–7, 2011.
- [63] ALAUZET, F. and MEHRENBARGER, M., “ P^1 -Conservative Solution Interpolation on Unstructured Triangular Meshes,” *International Journal of Numerical Methods in Engineering*, vol. 84, no. 13, pp. 1552–1588, 2010.
- [64] MAVRIPLIS, D., “Multigrid Techniques for Unstructured Meshes,” in *26th Computational Fluid Dynamics Lecture Series Program of the von Karman Institute (VKI) for Fluid Dynamics*, (Rhode-Saint-Genèse, Belgium), March 13–17, 1995.
- [65] NOACK, R., “Resolution Appropriate Overset Grid Assembly for Structured and Unstructured Grids,” in *AIAA 16th Computational Fluid Dynamics Conference*, no. AIAA-2003-4123, (Orlando, FL), June 23–26, 2003.
- [66] LÖHNER, R. and AMBROSIANO, J., “A Vectorized Particle Tracer for Unstructured Grids,” *Journal of Computational Physics*, vol. 91, no. 1, pp. 22–31, 1990.
- [67] LYNCH, C., PROSSER, D., and SMITH, M., “An Efficient Actuating Blade Model for Unsteady Rotating System Wake Simulations,” *Computers and Fluids*, in press.
- [68] ROGET, B. and SITARAMAN, J., “Robust and Scalable Overset Grid Assembly for Partitioned Unstructured Meshes,” in *AIAA 51st Aerospace Sciences Meeting*, no. AIAA-2013-0797, (Grapevine, TX), January 7–10, 2013.

- [69] PLIMPTON, S., HENDRICKSON, B., and STEWART, J., “A Parallel Rendezvous Algorithm for Interpolation Between Multiple Grids,” *Journal of Parallel and Distributed Computing*, vol. 64, no. 2, pp. 266–276, 2004.
- [70] LEE, P., YANG, C., and YANG, J., “Fast Algorithms for Computing Self-Avoiding Walks and Mesh Intersections over Unstructured Meshes,” *Advances in Engineering Software*, vol. 35, no. 2, pp. 61–73, 2004.
- [71] BUNING, P. and PULLIAM, T. H., “Initial Implementation of Near-Body Grid Adaption in OVERFLOW.” Presentation at 11th Symposium on Overset Composite Grid and Solution Technology, Dayton, OH, October 15–18, 2012.
- [72] BONHAUS, D. L., “An Upwind Multigrid Method for Solving Viscous Flows on Unstructured Triangular Meshes,” Master’s thesis, Georgia Washington University, 1993.
- [73] ANDERSON, W. K., RAUSCH, R. D., and BONHAUS, D. L., “Implicit/Multigrid Algorithms for Incompressible Turbulent Flows on Unstructured Grids,” *Journal of Computational Physics*, vol. 128, no. 2, pp. 391–408, 1996.
- [74] PARK, M. and DARMOFAL, D., “Parallel Anisotropic Tetrahedral Adaptation,” in *AIAA 46th Aerospace Sciences Meeting and Exhibit*, no. AIAA-2008-0917, (Reno, NV), January 7–10, 2008.
- [75] BIBB, K., GNOFFO, P., PARK, M., and JONES, W., “Parallel, Gradient-Based Anisotropic Mesh Adaptation for Re-entry Vehicle Configurations,” in *9th AIAA/ASME Joint Thermophysics and Heat Transfer Conference*, no. AIAA-2006-3579, (San Francisco, CA), June 5–8, 2006.
- [76] JONES, W., NIELSEN, E., and PARK, M., “Validation of 3D Adjoint-Based Error Estimation and Mesh Adaptation for Sonic Boom Predictions,” in *AIAA 44th Aerospace Sciences Meeting and Exhibit*, no. AIAA-2006-1150, (Reno, NV), January 9–12, 2006.
- [77] LEE-RAUSCH, E., PARK, M., JONES, W., HAMMOND, D., and NIELSEN, E., “Application of Parallel Adjoint-Based Error Estimation and Anisotropic Grid Adaptation for Three-Dimensional Aerospace Configurations,” in *AIAA 23th Applied Aerodynamics Conference*, no. AIAA-2005-4842, (Toronto, Canada), June 6–9, 2005.
- [78] CHORIN, A., “A Numerical Method for Solving Incompressible Viscous Flow Problems,” *Journal of Computational Physics*, vol. 135, no. 2, pp. 118–125, 1997.
- [79] VATSA, V. N., CARPENTER, M. H., and LOCKARD, D. P., “Re-evaluation of an Optimized Second Order Backward Difference (BDF2OPT) Scheme for Unsteady Flow Applications,” in *AIAA 48th Aerospace Sciences Meeting and Exhibit*, no. AIAA-2010-0122, (Orlando, FL), January 4–7, 2010.

- [80] ROE, P. L., “Approximate Riemann Solvers, Parameter Vectors, and Difference Schemes,” *Journal of Computational Physics*, vol. 43, no. 2, pp. 357–372, 1981.
- [81] NIELSEN, E., LU, J., PARK, M., and DARMOFAL, D., “An Implicit, Exact Dual Adjoint Solution Method for Turbulent Flows on Unstructured Grids,” *Computers and Fluids*, vol. 33, no. 9, pp. 1131–1155, 2004.
- [82] HAIRER, E. and WANNER, G., *Solving Ordinary Differential Equations II: Stiff and Differential-Algebraic Problems*. Springer-Verlag, Berlin Heidelberg, 2nd ed., 1996.
- [83] MENTER, F. R., “Two-Equation Eddy-Viscosity Turbulence Models for Engineering Applications,” *AIAA Journal*, vol. 32, no. 8, pp. 1598–1605, 1994.
- [84] LYNCH, C. E. AND SMITH, M. J., “Hybrid RANS-LES Turbulence Models on Unstructured Grids,” in *38th AIAA Fluid Dynamics Conference and Exhibit*, no. AIAA-2008-3854, (Seattle, WA), June 23–26, 2008.
- [85] SANCHEZ-ROCHA AND M., KIRTAS, M., MENON, S., “Zonal Hybrid RANS-LES Method for Static and Oscillating Airfoils and Wings,” in *44th AIAA Aerospace Sciences Meeting and Exhibit*, no. AIAA-2006-1256, (Reno, Nevada), January 9-12, 2006.
- [86] PARK, M. A., *Anisotropic Output-Based Adaptation with Tetrahedral Cut Cells for Compressible Flows*. PhD thesis, Massachusetts Institute of Technology, 2008.
- [87] VENDITTI, D. A., *Grid Adaptation for Functional Outputs of Compressible Flow Simulations*. PhD thesis, Massachusetts Institute of Technology, 2002.
- [88] HECHT, F., “The mesh adapting software: Bamg..” url-
<http://www.ann.jussieu.fr/hecht/ftp/bamg/>, 1998. [Online; accessed 19-December-2013].
- [89] NOACK, R. W., BOGER, D. A., KUNZ, R. F., and CARRICA, P. M., “Sug-gar++: An Improved General Overset Grid Assembly Capability,” in *AIAA 19th Computational Fluid Dynamics Conference*, no. AIAA-2009-3992, (San Antonio, TX), June 22-25, 2009.
- [90] NOACK, R., “DiRTlib: A Library to Add an Overset Capability to Your Flow Solver,” in *17th AIAA Computational Fluid Dynamics Conference*, no. AIAA-2005-5116, (Toronto, Canada), June 6–9, 2006.
- [91] SMITH, M. J., SHENOY, R., KENYON, A. R., and BROWN, R. E., “Vorticity-Transport and Unstructured RANS Investigation of Rotor-Fuselage Interactions,” in *Proceedings of the 35th European Rotorcraft Forum*, (Hamburg, Germany), September 22–25, 2009.

- [92] LÖHNER, R., “Robust, Vectorized Search Algorithms for Interpolation on Unstructured Grids,” *Journal of Computational Physics*, vol. 118, no. 2, pp. 380–387, 1995.
- [93] DOMPIERRE, J., LABBÉ, P., VALLET, M.-G., and CAMARERO, R., “How to Subdivide Pyramids, Prisms, and Hexahedra into Tetrahedra,” in *Proceedings of the 8th International Meshing Roundtable*, (South Lake Tahoe, CA), pp. 195–204, October 10–13, 1999.
- [94] SHERER, S. and SCOTT, J., “High-Order Compact Finite-Difference Methods on General Overset Grids,” *Journal of Computational Physics*, vol. 210, no. 2, pp. 459–496, 2005.
- [95] HAY, A., ETIENNE, S., PELLETIER, D., and GARON, A., “Code Verification of an Unsteady RANS Finite-Element Solver,” in *AIAA 21st Computational Fluid Dynamics Conference*, no. AIAA-2013-3088, (San Diego, CA), June 24–27, 2013.
- [96] ETIENNE, S., GARON, A., and PELLETIER, D., “Code Verification for Unsteady Flow Simulations with High Order Time-Stepping Schemes,” in *AIAA 47th Aerospace Sciences Meeting and Exhibit*, no. AIAA-2009-0169, (Orlando, FL), January 5–8, 2009.
- [97] MAVRIS, D. N., *An Analytical Method for the Prediction of Unsteady Rotor / Airframe Interactions in Forward Flight*. PhD Thesis: Revised Edition, Georgia Institute of Technology, 1991.
- [98] HELLSTEN, A., “Some Improvements in Menter’s $k-\omega$ SST Turbulence Model,” in *29th AIAA Fluid Dynamics Conference*, no. AIAA-1998-2554, (Albuquerque, NM), June 15–18, 1998.
- [99] SMITH, M., COOK, J., SANCHEZ-ROCHA, M., SHENOY, R., and MENON, S., “Improved Prediction of Complex Rotorcraft Aerodynamics,” in *Proceedings of the 69th Annual Forum of the American Helicopter Society*, (Phoenix, AZ), May 21–23, 2013.
- [100] RAGHAV, V., SHENOY, R., SMITH, M. J., and KOMERATH, N. M., “Investigation of Drag and Wake Turbulence of a Rotor Hub,” *Aerospace Science and Technology*, vol. 28, no. 1, pp. 164–175, 2013.
- [101] KOMERATH, N., “Deconstructing Hub Drag, Part I: Experiments,” Tech. Rep. ONR Final Report: Grant No. N0001409-1-1019, September 2013.
- [102] LYNCH, C. E. and SMITH, M. J., “Extension and Exploration of a Hybrid Turbulence Model on Unstructured Grids,” *AIAA Journal*, vol. 49, no. 11, pp. 2585–2591, 2011.

- [103] LIGGETT, N. D. and SMITH, M. J., “Temporal Convergence Criteria for Time-Accurate Viscous Simulations of Separated Flows,” *Computers & Fluids*, vol. 66, pp. 140–156, 2012.
- [104] HOERNER, S., *Fluid-Dynamic Drag*. Hoerner Fluid Dynamics, Brick Town, NJ, 1965.
- [105] SHENOY, R. and SMITH, M. J., “Deconstructing Hub Drag, Part II: Computational Deconstruction of Hub Drag,” Tech. Rep. ONR Final Report: Grant No. N0001409-1-1019, September 2013.
- [106] SAKAMOTO, H. and ARIE, M., “Vortex Shedding from a Rectangular Prism and a Circular Cylinder Placed Vertically in a Turbulent Boundary Layer,” *Journal of Fluid Mechanics*, vol. 126, pp. 147–165, 1983.
- [107] NIELSEN, E. J. and DISKIN, B., “Discrete Adjoint-Based Design for Unsteady Turbulent Flows on Dynamic Overset Unstructured Grids,” vol. 51, no. 6, pp. 1355–1373, 2013.
- [108] JONES, W., “GridEx – An Integrated Grid Generation Package for CFD,” in *AIAA 16th Computational Fluid Dynamics Conference*, no. AIAA-2003-4129, (Orlando, FL), June 23–26, 2003.
- [109] AFTOSMIS, M., DELANAYE, M., and HAIMES, R., “Automatic Generation of CFD-Ready Surface Triangulations from CAD Geometry,” in *AIAA 37th Aerospace Sciences Meeting and Exhibit*, no. AIAA-99-0776, (Reno, NV), January 11–14, 1999.

VITA

Rajiv Shenoy is the son of Rajendra and Pramila Shenoy of Hampton, VA. From an early age, he had a fascination with airplanes and loved attending air shows. As he grew up, he enjoyed playing tennis and took up an interest in geography.

While attending Bethel High School in Hampton, VA, he developed an increased affinity toward the physical sciences and wanted to pursue a future degree and career in engineering. He graduated high school in 2003 as class salutatorian and chose to attend Georgia Tech for his undergraduate studies. After obtaining his Bachelor of Science degree in Aerospace Engineering in 2007, he went on to work with Analytical Services & Materials in Hampton, VA as a CFD Engineer. There, he became interested in learning more about computational methods and decided to pursue a graduate degree to enable him to become an active contributor in the field. In 2009, he began his graduate program at Georgia Tech under the guidance of Prof. Marilyn Smith. He participated in two Langley Aerospace Research Summer Scholars tours in 2011 and 2012, where he gained more exposure in computational methods and helped him further his thesis goals.

He remained an active member of both the Georgia Tech American Institute of Aeronautics and Astronautics and American Helicopter Society (AHS) chapters as a graduate student. As AHS membership chair for the Georgia Tech chapter, he was involved in one of the highest membership increases in the chapter's history, which earned the chapter an award at the 2013 AHS annual forum. Upon completion of his PhD, he joined Combustion Research & Flow Technology, Inc. (CRAFT Tech) in Pennsylvania.