

## ABSTRACT

Title of Dissertation: **INCORPORATION OF AIRFOIL-INTERACTIONAL  
DATA TO IMPROVE THE ACCURACY OF  
STACKED ROTOR PERFORMANCE PREDICTIONS  
IN THE DESIGN STAGE**

**Miranda Costenoble**  
**Doctor of Philosophy, 2023**

Dissertation Directed by: **Professor James Baeder**  
**Department of Aerospace Engineering**

In this dissertation, a methodology is presented for lower-fidelity modeling of stacked rotors, with improved accuracy compared to existing lower-fidelity rotor models. This methodology is built on a conventional prescribed vortex wake model of the rotor, with lifting-line airfoil blades. Such a lifting-line model cannot fully capture the aerodynamic interaction between the rotor blades, which are driven heavily by thickness and shape effects. To account for these effects, a high-fidelity 2D CFD code is used to model the airfoil-to-airfoil interaction along the span of the rotor. These 2D CFD loads are then injected into the lifting-line/prescribed wake rotor model, using an iterative technique to account for the changing deflections of the rotor blades.

To accurately determine the airfoil-interactional loads along the span of the rotor, it is necessary to have some way to relate the conditions along the span of the rotor in 3D to the airfoil conditions in 2D, and vice-versa. Methods for parameterization of the airfoil system are presented, which account for both the geometry of the rotor/airfoil system and their aerodynamic

conditions. Two different methods of relating the airfoil loads back to the rotor are presented, which offer different strategies depending upon the constraints of the underlying rotor model.

Any rotor design must include selection of the airfoils on the blades, and stacked rotors are no different. To that end, 2D airfoil simulations are presented, which demonstrate both the necessity of the current methodology, and offer suggestions for future stacked rotor design. 2D airfoil loads are pre-computed (prior to the lower-fidelity rotor simulations) using an established 2D CFD code. Automation of this code allows for rapid generation of large data sets with minimal user input.

The combined 2D CFD/3D prescribed wake methodology is presented and validated against recent experimental results. The baseline prescribed wake model is shown to significantly less variation of thrust and power with phase angle than the experiment. Inclusion of lifting-vortex airloads leads to improvements in thrust prediction, but with incorrect magnitude at small phase angles and incorrect power predictions. Only by including the 2D CFD airfoil-interactional airloads are the most accurate results achieved. Multiple inflow and coupling methods are also examined, with discussion of the strengths and limitations of each. Overall, it is believed that the current work should offer the potential for significantly faster and more accurate design of stacked rotors than was previously possible.

Incorporation of Airfoil-Interactional Data to Improve the Accuracy of Stacked  
Rotor Performance Predictions in the Design Stage

by

Miranda Costenoble

Dissertation submitted to the Faculty of the Graduate School of the  
University of Maryland, College Park in partial fulfillment  
of the requirements for the degree of  
Doctor of Philosophy  
2023

Advisory Committee:

Professor James Baeder, Chair/Advisor  
Associate Professor Anubhav Datta  
Professor Christopher Cadou  
Research Professor Roberto Celi  
Professor Howard Elman

© Copyright by  
Miranda Costenoble  
2023

## Acknowledgments

Many people and organizations have been involved, directly or indirectly, with supporting my research, and for that I must express my gratitude. First and foremost, I would like to thank my advisor Dr. Baeder, without whom none of this would have been possible. When I returned to the University of Maryland for my graduate studies I did so without a specific plan, but Dr. Baeder took me in as a research assistant, and set me on the research for my M.S. which would lead to the automated meshing methodology. Though I was initially unsure if I would stay on for my Ph.D. Dr. Baeder encouraged me to, and without his research insights, contacts with others, and keeping me on-task, I would not be here today.

I would also like to express my gratitude to the US Combat Capabilities Development Command Army Research Lab (ARL) and Oak Ridge Associated Universities (ORAU) for their support of my research efforts. Originally, I was granted a summer internship at ARL with support from ORAU back in 2018, but with continuous extensions that internship lasted until the end of last year. The support which that internship offered, both technical and financial, significantly eased the burden of conducting this research. Within ARL, I would like to thank Dr. Rajneesh Singh and Dr. Matthew Floros for serving as my research advisors, who pushed to keep me there for so long in support of my research. I would also like to thank Dr. George Jacobellis for his ongoing technical support and assistance in my research, as well as his prior work on the subject, which I have studied and learned from. Finally at ARL, I would like to offer my thanks

to Dr. Hao Kang, who further assisted me when I got stuck working with RCAS.

I would be remiss if I did not also thank Dr. Hossein Saberi of Advanced Rotorcraft Technologies, Inc. Dr. Saberi was the original developer of the CFD/CSD coupling methodology for RCAS, and despite his busy schedule he was able to take time to instruct me in its use. Without his help, it would not have been possible for me to fully develop the coupling portion of this work.

Finally, I would like to thank my friends and family who have supported me socially and emotionally during this process. Particularly, my parents Nancy and Steven Costenoble were the ones who pushed me to go back to school for my graduate studies, and my life is better for it. Since coming back to Maryland for school I began to develop a relationship with and eventually became engaged to my fiancée Quinn Campbell, who I moved in with at the start of the Covid pandemic. I don't know how I could have done any of this without her support.

# Table of Contents

|                                                               |      |
|---------------------------------------------------------------|------|
| Acknowledgements                                              | ii   |
| Table of Contents                                             | iv   |
| List of Tables                                                | vii  |
| List of Figures                                               | viii |
| List of Abbreviations                                         | x    |
| Chapter 1: Introduction                                       | 1    |
| 1.1 Terminology . . . . .                                     | 2    |
| 1.1.1 Unique Stacked Rotor Parameters . . . . .               | 3    |
| 1.2 Review of Prior Stacked Rotor Research . . . . .          | 5    |
| 1.2.1 Historical Efforts . . . . .                            | 5    |
| 1.2.2 Airfoil Interactions . . . . .                          | 7    |
| 1.2.3 Scissors Rotors . . . . .                               | 12   |
| 1.2.4 Towards a Modern Understanding . . . . .                | 14   |
| 1.2.5 Rediscovery and Recent Efforts . . . . .                | 16   |
| 1.2.6 Summary . . . . .                                       | 21   |
| 1.3 Rotor Design Methodology . . . . .                        | 22   |
| 1.3.1 Lower-Fidelity Rotor Models . . . . .                   | 24   |
| 1.3.2 Limitations and Application to Stacked Rotors . . . . . | 26   |
| 1.4 Objectives of the Present Work . . . . .                  | 27   |
| 1.5 Organization of Thesis . . . . .                          | 28   |
| Chapter 2: System Parameterization                            | 29   |
| 2.1 Aerodynamic Conditions . . . . .                          | 29   |
| 2.2 System Geometry . . . . .                                 | 30   |
| 2.2.1 Geometric Uniqueness . . . . .                          | 32   |
| 2.3 Consideration of Flow Conditions . . . . .                | 32   |
| 2.3.1 Modification of Geometry . . . . .                      | 34   |
| 2.3.2 Force Transformations . . . . .                         | 39   |
| 2.4 Summary . . . . .                                         | 40   |
| Chapter 3: 2D Methodology                                     | 42   |
| 3.1 2D Lifting-Vortex Model . . . . .                         | 43   |

|            |                                                       |     |
|------------|-------------------------------------------------------|-----|
| 3.1.1      | Periodicity . . . . .                                 | 45  |
| 3.2        | 2D CFD for Data Generation . . . . .                  | 50  |
| 3.2.1      | AutoMesh . . . . .                                    | 51  |
| 3.2.2      | HAM2D . . . . .                                       | 62  |
| 3.2.3      | Parallelization and Computing Resources . . . . .     | 68  |
| 3.3        | Data Interpolation . . . . .                          | 69  |
| 3.3.1      | Code Implementation and Usage . . . . .               | 72  |
| 3.3.2      | Parameter Selection . . . . .                         | 75  |
| 3.4        | Summary . . . . .                                     | 76  |
| Chapter 4: | 2D Results . . . . .                                  | 78  |
| 4.1        | System of Symmetric Airfoils . . . . .                | 78  |
| 4.1.1      | Symmetric Configuration . . . . .                     | 79  |
| 4.1.2      | Lifting Airfoil Comparison . . . . .                  | 84  |
| 4.1.3      | Angle of Attack and Pitch Offset . . . . .            | 89  |
| 4.2        | System of VR12 Airfoils . . . . .                     | 94  |
| 4.3        | Summary . . . . .                                     | 99  |
| Chapter 5: | 3D Methodology . . . . .                              | 101 |
| 5.1        | Considerations . . . . .                              | 102 |
| 5.2        | RCAS . . . . .                                        | 103 |
| 5.2.1      | Inflow Models . . . . .                               | 106 |
| 5.2.2      | Structural Model . . . . .                            | 121 |
| 5.2.3      | CFD/CSD Coupling Interface . . . . .                  | 123 |
| 5.3        | Multi-Airfoil Coupling Code . . . . .                 | 127 |
| 5.3.1      | Airfoil Positions and Flow Conditions . . . . .       | 128 |
| 5.3.2      | Coupling With CFD/CSD Interface . . . . .             | 130 |
| 5.3.3      | Coupling Without CFD/CSD Interface . . . . .          | 133 |
| 5.4        | Summary . . . . .                                     | 135 |
| Chapter 6: | 3D Results . . . . .                                  | 137 |
| 6.1        | Without Collective Offset . . . . .                   | 139 |
| 6.1.1      | Coupling Convergence . . . . .                        | 139 |
| 6.1.2      | Experimental Data Comparison . . . . .                | 144 |
| 6.1.3      | 2D CFD vs. 2D Lifting-Vortex . . . . .                | 146 |
| 6.1.4      | Selected Individual Cases . . . . .                   | 149 |
| 6.1.5      | Coupling With and Without CFD/CSD Interface . . . . . | 153 |
| 6.1.6      | Alternate CAMRAD Prescribed Wake Models . . . . .     | 156 |
| 6.2        | Collective Offset Cases . . . . .                     | 159 |
| 6.2.1      | Coupling With and Without CFD/CSD Interface . . . . . | 163 |
| 6.3        | Summary . . . . .                                     | 164 |
| Chapter 7: | Summary and Conclusions . . . . .                     | 166 |
| 7.1        | Key Contributions . . . . .                           | 171 |



## List of Tables

|     |                                                                                                       |     |
|-----|-------------------------------------------------------------------------------------------------------|-----|
| 4.1 | Summary of the operating range of $\alpha$ from Figure 4.9a . . . . .                                 | 93  |
| 4.2 | Maximum airfoil system efficiency and corresponding conditions from Figure 4.9 . . . . .              | 94  |
| 6.1 | Parameters of the stacked rotor modeled in this study . . . . .                                       | 138 |
| 6.2 | Parameters used for 2D CFD two-airfoil data generation . . . . .                                      | 138 |
| 6.3 | Change in upper and lower rotor thrust with CDS/CSD coupling for $\Delta\psi = 5.625^\circ$ . . . . . | 140 |

## List of Figures

|      |                                                                                                                             |    |
|------|-----------------------------------------------------------------------------------------------------------------------------|----|
| 1.1  | Examples of advanced air mobility concept vehicles utilizing stacked rotors . . . .                                         | 2  |
| 1.2  | Parameters relating to the relative positions of the rotors in a stacked rotor system                                       | 3  |
| 1.3  | Hamilton Standard “variable camber propeller” for tilt-wing/tilt-rotor applications                                         | 6  |
| 1.4  | Types of biplane stagger . . . . .                                                                                          | 7  |
| 1.5  | Illustration of some airfoil interactions described by Smith . . . . .                                                      | 9  |
| 1.6  | Example applications of “scissors” rotors . . . . .                                                                         | 13 |
| 1.7  | Stacked rotor test stand introduced by Johnson et.al. . . . .                                                               | 20 |
| 1.8  | Hierarchy of aerodynamic models commonly used in rotor design . . . . .                                                     | 23 |
| 2.1  | Geometric parameters of the stacked rotor blade and airfoil systems . . . . .                                               | 31 |
| 2.2  | Flow parameters of the stacked rotor blade and airfoil systems . . . . .                                                    | 34 |
| 2.3  | Notation for a generic vortical disturbance around the airfoil system . . . . .                                             | 38 |
| 3.1  | Phase angles between the closer blade pair and the farther blade pair . . . . .                                             | 48 |
| 3.2  | Strength and influence of the image vortices used to account for periodicity . . . .                                        | 49 |
| 3.3  | Effective angles of attack on nominal airfoils with and without accounting for periodicity . . . . .                        | 50 |
| 3.4  | Overview of the steps of the automated meshing procedure . . . . .                                                          | 52 |
| 3.5  | Adjustments for a wake surface or C-mesh using an O-mesh methodology . . . .                                                | 56 |
| 3.6  | Strand gathering and mesh unfolding, similarly to CLOVIS . . . . .                                                          | 57 |
| 3.7  | Overlap removal using point-in-polygon tests on the nodes of the outer boundary                                             | 58 |
| 3.8  | Bounds on the Laplacian smoothing from cells adjacent to a node . . . . .                                                   | 60 |
| 3.9  | Identification of the Hamiltonian paths through a representative unstructured mesh                                          | 63 |
| 3.10 | Strand directions and numbering scheme at a cell in a mesh . . . . .                                                        | 65 |
| 4.1  | Variation of forces for a system of two NACA 0008 airfoils at zero incidence angle                                          | 80 |
| 4.2  | CFD performed for a system of two NACA 0008 airfoils at zero incidence angle with $\frac{\Delta x}{c} = 0.75$ . . . . .     | 81 |
| 4.3  | Effect of circulation to change the apparent forces on each airfoil . . . . .                                               | 82 |
| 4.4  | Variation of total forces for a system of two NACA 0008 airfoils with $\Delta\theta = 4^\circ$ .                            | 85 |
| 4.5  | CFD performed for a system of two NACA 0008 airfoils with $\frac{\Delta x}{c} = 1.0$ and $\Delta\theta = 4^\circ$ . . . . . | 86 |
| 4.6  | Variation of individual forces for a system of two NACA 0008 airfoils with $\Delta\theta = 4^\circ$                         | 87 |
| 4.7  | Variation of forces with $\Delta\theta$ for a system of two NACA 0008 airfoils . . . . .                                    | 90 |
| 4.8  | Variation of forces with $\alpha_\infty$ for a system of two NACA 0008 airfoils . . . . .                                   | 92 |
| 4.9  | Effect of varying $\alpha_\infty$ and $\Delta\theta$ on the efficiency of a system of NACA 0008 airfoils                    | 93 |

|      |                                                                                                                                              |     |
|------|----------------------------------------------------------------------------------------------------------------------------------------------|-----|
| 4.10 | Efficiency and corresponding conditions for a system of two VR12 airfoils across a range of $\Delta x/c, \Delta z/c$ . . . . .               | 96  |
| 4.11 | Maximum lift coefficient and corresponding conditions for a system of two VR12 airfoils across a range of $\Delta x/c, \Delta z/c$ . . . . . | 98  |
| 5.1  | Relationship between the 2D and 3D models . . . . .                                                                                          | 101 |
| 5.2  | Superposition of horseshoe vortices along a lifting line . . . . .                                                                           | 108 |
| 5.3  | Vortex filaments forming a wake lattice grid . . . . .                                                                                       | 109 |
| 5.4  | Vortex particle (“blob”) representation of the wake . . . . .                                                                                | 114 |
| 6.1  | Convergence of thrust on each rotor in the stacked rotor system . . . . .                                                                    | 140 |
| 6.2  | Convergence of total performance for the stacked rotor system . . . . .                                                                      | 142 |
| 6.3  | Thrust on each rotor for $\Delta\theta_0 = 0^\circ$ compared to experimental data . . . . .                                                  | 144 |
| 6.4  | Total performance of the stacked rotor system for $\Delta\theta_0 = 0^\circ$ compared to experimental data . . . . .                         | 145 |
| 6.5  | Coupling based on 2D CFD and lifting-vortex airloads . . . . .                                                                               | 147 |
| 6.6  | Airfoil system conditions along the span of the blades for $\Delta\psi = 5.625^\circ$ . . . . .                                              | 149 |
| 6.7  | Airfoil-frame force coefficients along the span of the blades for $\Delta\psi = 5.625^\circ$ . . . . .                                       | 150 |
| 6.8  | Airfoil system conditions along the span of the blades for $\Delta\psi = 45^\circ$ . . . . .                                                 | 152 |
| 6.9  | Airfoil-frame Z-force coefficients along the span of the blades for $\Delta\psi = 45^\circ$ . . . . .                                        | 153 |
| 6.10 | Stacked rotor performance computed using different coupling methodologies . . . . .                                                          | 154 |
| 6.11 | Stacked rotor performance computed using the CAMRAD inflow model with coefficients given by Landgrebe . . . . .                              | 157 |
| 6.12 | Stacked rotor performance computed using the CAMRAD inflow model with coefficients given by Kocurek and Tangler . . . . .                    | 158 |
| 6.13 | Performance of the stacked rotor system for $\Delta\theta_0 = -4^\circ$ compared to experimental data . . . . .                              | 160 |
| 6.14 | Stacked rotor performance computed using different coupling methodologies for the $\Delta\theta_0 = -4^\circ$ configuration . . . . .        | 162 |

## List of Abbreviations

|        |                                                          |
|--------|----------------------------------------------------------|
| 2D     | Two-dimensional/dimensions                               |
| 3D     | Three-dimensional/dimensions                             |
| $a$    | Speed of sound in air ( <i>ft/s</i> , <i>m/s</i> , etc.) |
| AAM    | Advanced air mobility (formerly UAM)                     |
| BEMT   | Blade-element momentum theory                            |
| $c$    | Chord length ( <i>ft</i> , <i>in</i> , <i>m</i> , etc.)  |
| CFD    | Computational fluid dynamics                             |
| $C_d$  | Drag coefficient                                         |
| $C_l$  | Lift coefficient                                         |
| $C_m$  | Moment coefficient                                       |
| $C_P$  | Power coefficient                                        |
| $C_Q$  | Torque coefficient                                       |
| $C_T$  | Thrust coefficient                                       |
| $C_x$  | X-force coefficient                                      |
| $C_y$  | Y-force coefficient                                      |
| $C_z$  | Z-force coefficient                                      |
| $D$    | Drag force ( $lb_f$ or $N$ )                             |
| DoD    | United States Department of Defense                      |
| $F_x$  | X-force ( $lb_f$ or $N$ )                                |
| $F_y$  | Y-force ( $lb_f$ or $N$ )                                |
| $F_z$  | Z-force ( $lb_f$ or $N$ )                                |
| HAM2D  | 2D version of HAMSTR                                     |
| HAMSTR | Hamiltonian strand flow-solver                           |
| HPC    | High-performance computer                                |
| HPCMP  | DoD High-Performance Computing Modernization Program     |
| $L$    | Lift force ( $lb_f$ or $N$ )                             |

|                  |                                                                           |
|------------------|---------------------------------------------------------------------------|
| $M$              | Mach ( $M = u/a$ )                                                        |
| $r$              | Radial coordinate ( <i>ft, in, m, etc.</i> )                              |
| $R$              | Rotor radius ( <i>ft, in, m, etc.</i> )                                   |
| RCAS             | Rotorcraft Comprehensive Analysis System                                  |
| $u$              | Flow speed ( <i>ft/s, m/s, etc.</i> )                                     |
| UAM              | Urban air mobility                                                        |
| UAS              | Unmanned aerial system(s)                                                 |
| $\hat{x}$        | X-axis                                                                    |
| $\hat{y}$        | Y-axis                                                                    |
| $\hat{z}$        | Z-axis                                                                    |
| $\alpha$         | Angle of attack ( <i>deg or rad</i> )                                     |
| $\delta$         | An arbitrary small number                                                 |
| $\Omega$         | Rotor rotational speed ( <i>RPM, rad/s, etc.</i> )                        |
| $\phi$           | Inflow angle ( <i>deg or rad</i> )                                        |
| $\psi$           | Rotor azimuth angle ( <i>deg or rad</i> )                                 |
| $\theta_n$       | Pitch angle for blade $n$ ( <i>deg or rad</i> )                           |
| $\theta_0$       | Mean collective pitch angle of stacked rotor system ( <i>deg or rad</i> ) |
| $\Delta x$       | Airfoil horizontal offset ( <i>ft, in, m, etc.</i> )                      |
| $\Delta z$       | Airfoil or rotor vertical offset ( <i>ft, in, m, etc.</i> )               |
| $\Delta\psi$     | Index angle between rotors ( <i>deg or rad</i> )                          |
| $\Delta\theta$   | Airfoil differential pitch angle ( <i>deg or rad</i> )                    |
| $\Delta\theta_0$ | Stacked rotor differential collective angle ( <i>deg or rad</i> )         |

## Chapter 1: Introduction

A stacked rotor is a system of two rotors spinning in the same direction about a common rotational axis. This is in contrast to a more “conventional” counter-rotating coaxial rotor, in which two rotors share a common rotational axis but spin in opposite directions. Stacked rotors are a concept which has been investigated at various times throughout the history of rotorcraft, with a variety of different hypotheses and goals. This research has shown the potential for stacked rotors to offer benefits in performance, flexibility, and acoustics, as compared to more conventional rotors. Historically these benefits have not been fully realized, and stacked rotors have found only limited application in production vehicles. Recently however, emerging technologies such as distributed electric propulsion and novel concepts such as advanced air mobility—presenting a significant need for quieter, highly-efficient rotors—have once again stimulated interest in the use of stacked rotors. Some examples of AAM vehicle concepts utilizing stacked rotors are shown in Figure 1.1.

Despite their seemingly simple concept, stacked rotors present novel parameters to the rotor designer. Moreover, complex aerodynamic interactions occur between the blades of a stacked rotor. These interactions which are present in neither conventional single nor counter-rotating rotors, hence may not be accounted for by the traditional simplified rotor models used early in the rotor design process. For vehicle and rotor designers to truly unlock the potential of the



(a) Supernal S-A1, from [1]



(b) Uber eCRM and LaunchPoint EPS, from [2]

Figure 1.1: Examples of advanced air mobility concept vehicles utilizing stacked rotors

stacked rotor concept, novel design methodologies may be required to accurately represent these interactions. To develop such a methodology it is necessary to consider in detail the mechanisms of these interactions, which are largely covered by prior research efforts.

## 1.1 Terminology

Before proceeding to review the prior research on stacked rotors however, it is helpful to establish some terminology—not least because different terms have been used at different times. While this work follows modern practice (e.g. [3, 4]) in the term “stacked rotor” to describe the type of rotor system being studied, prior studies have used a variety of different terms. In the mid-to-late-2010s, the more-descriptive but less-catchy term “co-rotating coaxial rotor” was preferred (e.g. [5]). Older studies from the 1960s–70s use terms such as “variable camber” [6], “variable geometry” [7], and “scissors” [8], with the particular terms depending upon the authors’ design philosophies and research objectives. It is important to note however that from a modern perspective all of these are stacked rotors (except where noted), and may be referred to as such in this work, except where it is necessary to highlight differences in methodology.




Figure 1.2: Parameters relating to the relative positions of the rotors in a stacked rotor system

Another difference between prior studies is whether the stacked rotor system is considered to comprise one rotor (with offset between the *blades*) or two rotors (with offset between the *rotors*). Again there is a chronological distinction, with older studies (e.g. [6, 7]) typically referencing one rotor and more recent studies (e.g. [3, 4]) referencing two. From an *aerodynamic* perspective, this difference is only one of terminology, so long as both rotors are spinning at the same speed. From a *mechanical* perspective however, there may be a distinction depending upon whether the blades of the stacked rotor are attached to one rotor head or two, and if there are two rotor heads whether they share or each have a distinct drivetrain. This work will follow the nomenclature of two rotors, largely because this is the modern practice, but also because this nomenclature provides a framework to quantify the loads and performance of individual components of the stacked rotor system in addition to the system as a whole.

### 1.1.1 Unique Stacked Rotor Parameters

It is helpful at this stage to note the parameters unique to the stacked rotor system—i.e., those parameters which relate the positions of the two rotors to the system or to each other. A basic familiarity is assumed with the parameters and terminology for an isolated rotor<sup>1</sup>.

<sup>1</sup>If not familiar, it is suggested that the reader review any modern rotorcraft textbook, e.g. [9].

The **upper rotor** and **lower rotor** are defined such that “up” is the direction along the rotor’s rotational axis of positive thrust (or equivalently, opposite the direction of the mean inflow velocity). While this seems like a trivial definition, it is not strictly necessary that the rotor’s rotational axis be “vertical” or opposite to gravity (and indeed some studies [4,10] use a horizontal rotational axis), hence it may be necessary to disambiguate. Two parameters are then used to describe the positions of the two rotors relative to each other, as shown in Figure 1.2:

- **Vertical offset**  $\Delta z$ , positive when the upper rotor is above the lower rotor;  $\Delta z = 0$  implies a co-planar rotor (Fig. 1.2a)
- **Phase offset**  $\Delta\psi$ , positive when the lower rotor lags the upper rotor and negative when the lower rotor leads the upper rotor (Fig. 1.2b)

It is not necessary to consider values of  $\Delta z < 0$ , because a system where  $\{\Delta z, \Delta\psi\} = \{x, y\}$  is equivalent to a system where  $\{\Delta z, \Delta\psi\} = \{-x, -y\}$ . For a stacked rotor with  $N_b$  blades per rotor,  $\Delta\psi$  has a range of unique conditions of  $\frac{360^\circ}{N_b}$ . Typically this range is expressed as  $\Delta\psi \in \left[-\frac{180^\circ}{N_b}, \frac{180^\circ}{N_b}\right]$  to better examine effects near  $\Delta\psi = 0$ , but this range could alternately be expressed as  $\Delta\psi \in \left[0^\circ, \frac{360^\circ}{N_b}\right]$ .

In this study (as with other recent studies), the only rotor control parameter being examined is the collective (with zero cyclic control input). While it would be possible to express the collectives of the upper and lower rotors independently, it is useful to have a means to express the controls for the stacked rotor system. Those controls are defined thus:

- **Mean collective**  $\theta_0$ , where  $\theta_0 = \frac{1}{2} (\theta_{upper} + \theta_{lower})$
- **Collective offset**  $\Delta\theta_0$ , where  $\Delta\theta_0 = \frac{1}{2} (\theta_{upper} - \theta_{lower})$

With this definition, the individual rotor collectives may be easily determined as  $\theta_{upper} = \theta_0 + \Delta\theta_0$  and  $\theta_{lower} = \theta_0 - \Delta\theta_0$ . Note however that while this definition is consistent with recent work [4], some older works [7, 11] do not include the factor of  $\frac{1}{2}$  in  $\Delta\theta_0$ , and so care must be taken in comparing data sets.

## 1.2 Review of Prior Stacked Rotor Research

Research efforts towards what would now be recognized as stacked rotors extend at least as far back as the 1950's. These research efforts have had various hypotheses, and those hypotheses have not-infrequently influenced the researchers' interpretation of their results. Positively, because of the number of studies performed, a fuller understanding of the aerodynamic effects occurring in stacked rotors may be gained by examining multiple sources from throughout history, taking note of what each does and does not address.

### 1.2.1 Historical Efforts

One of the earliest efforts towards what would now be recognized as a stacked rotor was the Hamilton Standard "variable camber propeller" (Fig. 1.3), developed between 1958 and c. 1969 [6, 12]. Development of the variable camber propeller was driven by the unique needs of high-speed (0.6–0.8 Mach) tilt-wing and tilt-rotor vehicles, which must be able to produce high thrust with low/zero free-stream inflow (at takeoff), while maintaining high efficiency at lower thrust with high free-stream inflow (during cruise). Among the tradeoffs which must be considered when designing for both conditions is that the former condition is more easily met with a more-cambered airfoil, whereas the latter condition requires a less-cambered airfoil.




Figure 1.3: Hamilton Standard “variable camber propeller” for tilt-wing/tilt-rotor applications

Hamilton Standard addressed this compromise by designing a propeller<sup>2</sup> with pairs of closely-spaced blades<sup>3</sup> (Fig. 1.3a). By changing the pitch of the blades, their airfoils could comprise either an approximation of a single higher-camber<sup>4</sup> airfoil, or two “independent” lower-camber airfoils (Fig. 1.3b). The blades of the propeller were inversely tapered, which allowed for a consistent gap (and hence appropriate configuration) to be maintained between the airfoils along the span. Although the propeller apparently performed satisfactorily in both wind tunnel and flight tests the concept was ultimately abandoned, as it offered little advantage for flight speeds below 0.4 Mach [13]. In fact, with even the V-22 Osprey having a maximum flight speed in the range of “only” 0.4–0.5 Mach, modern VTOL rotorcraft have yet to achieve the flight speeds envisioned by Hamilton Standard in the 1960’s.

Another effort with a similar design philosophy was the Kaman K-16B, a modification of a Grumman Goose flying boat with a tilt-wing to enable vertical takeoff and landing [14].


<sup>2</sup>Lacking flapping and lead-lag hinges or flexures

<sup>3</sup>At various times, both 3-pair/6-blade and 4-pair/8-blade versions were tested.

<sup>4</sup>Or put another way, act like a deployed slotted or Junkers-type flap



(a) Positive stagger



(b) Negative stagger

Figure 1.4: Types of biplane stagger (public domain images)

Kaman recognized the same problem—the need for higher or lower airfoil camber in different flight conditions—as Hamilton Standard but addressed it in a slightly different way, using flaps attached directly to the blades. Of note, Kaman used true *rotors* in its design, with flapping and lead-lag hinges, and the flaps were intended to be used for both collective and cyclic control. Although the K-16B never flew outside of captive wind tunnel tests [15], it can be considered another stepping-stone towards the current work.

### 1.2.2 Airfoil Interactions

The Hamilton Standard and Kaman design teams clearly recognized the importance of airfoil-level effects to their designs, with Hamilton Standard even using a wind tunnel rig designed specifically to test pseudo-2D multi-airfoil configurations [6]. Before proceeding to discuss stacked rotor research further, it is useful to consider what these “airfoil interactions” actually are. Note that while most of the literature review is concerned with rotor (3D) simulations and experiments, this section will primarily be concerned with airfoil (2D) effects.


A system of two airfoils is at its most fundamental level like a biplane, and hence has been of interest since the earliest days of aviation. While the biplane wing configuration was typically chosen for structural (rather than aerodynamic) concerns, it was recognized as early as the 1920’s

that by positioning the wings carefully, aerodynamic benefits could be achieved. In particular, Norton in 1921 found that positive stagger (Fig. 1.4a) increased the lift of a biplane with little effect on the drag, while negative stagger (Fig. 1.4b) decreased the lift [16]. (This early study was limited however in the range of configurations tested, only looking at the effect of horizontal offset of  $\pm 1$  chord length.)


However, the seminal work on airfoil interactions must be A. M. O. Smith's 1975 paper *High-Lift Aerodynamics* [17]. While many works discuss the benefits and drawbacks of certain specific airfoil configurations, Smith lays out *in general* the mechanisms by which any system of airfoils (biplanes, flaps, slats, arbitrary numbers of elements etc.) will interact with each other, in addition to discussing other concerns for the design of high-lift airfoils. In particular, Smith identifies five effects which occur between adjacent airfoils:

- Slat effect
- Circulation effect
- Dumping effect
- Off-the-surface pressure recovery
- Fresh-boundary-layer effect


Understanding of these effects is important to understanding the results of other studies of stacked rotors (whether or not those studies' authors realized it at the time), and hence are discussed briefly in the following sections. Smith takes care to note however that none of these effects are related to jet-blowing, because without a means to artificially increase the total pressure of the flow (which would not be present for a "passive" system of airfoils) there can be no such effect.




(a) Slat effect



(b) Circulation effect due to a vortex



(c) Circulation effect due to an obstruction



(d) Dumping effect

Figure 1.5: Illustration of some airfoil interactions described by Smith, excerpted from [17]

### 1.2.2.1 Slat Effect

The slat effect is the result of a lifting element<sup>5</sup>—a “slat”—ahead of the “main” element being considered (Fig. 1.5a). The circulation created by this lifting element decreases the suction peak<sup>6</sup> on the main element, thus decreasing its lift. (The flow nearer the trailing edge of the main element remains relatively undisturbed.) However, the combined lift of the main element and the slat is greater than that of the main element in isolation. In addition, an interesting result of this effect is that the main element with a slat can sustain a higher angle of attack before stall than it could in isolation, allowing for further increases in lift. This is easy to understand by—as Smith suggests—thinking of the slat as a “turning aid” to the flow.

### 1.2.2.2 Circulation Effect

By moving the vortex to behind the “main” element, the circulation effect may be observed (Fig. 1.5b). Clearly, a vortex in this position would induce an upwards velocity—and hence positive angle of attack—at the trailing edge of the main element. For the Kutta condition<sup>7</sup> to be maintained, the circulation—and hence the lift—on the main element must increase. Smith observes that this increase in angle of attack at the trailing edge is analogous to the effect of a plain flap, except that rather than deflecting the airfoil into the flow, the flow is deflected into the airfoil.

What is important to cause the slat effect is the deflection of the flow, not the specific cause of that deflection. Indeed, Smith notes, such deflection could equivalently be caused by a

---

<sup>5</sup>Such as point vortex, or a finite-thick lifting airfoil

<sup>6</sup>Or alternately, increases the pressure on the main element’s upper surface

<sup>7</sup>I.E. that the circulation around the airfoil will be such that the rear stagnation point is held at the (sharp) trailing edge

non-lifting but finite-thickness obstruction to the flow placed near the trailing edge of the main element (as in Fig. 1.5c). As the flow navigates around the obstruction flow is induced at the trailing edge of the main element, hence the circulation effect. It should be apparent then that if a *finite-thickness lifting airfoil* is placed behind the main element, both thickness effects and lifting effects will be important to consider<sup>8</sup>.

### 1.2.2.3 Dumping Effect

A lifting element downstream of the “main” element may also induce an increase in *tangential* velocity at the main airfoil’s trailing edge (this being in addition to the upwards velocity which results in the circulation effect). Hence, the flow leaving the trailing edge of (i.e. being dumped by) the main element will do so into a region of greater velocity than it would otherwise. Because of this, the required pressure recovery is reduced, and the main (forward) element may sustain greater suction over its length than it could in isolation, hence its lift is increased. Smith refers to this as the dumping effect (illustrated in Fig. 1.5d for a 3-element airfoil).

### 1.2.2.4 Off-the-Surface Pressure Recovery

In general, flow along an airfoil (from leading edge to trailing edge) accelerates to some peak velocity (negative pressure), then decelerates to some degree before reaching its trailing edge; after leaving the trailing edge the flow will accelerate back to the free-stream conditions. At the trailing edge, the flow has reduced energy due to the viscosity of the air. The flow along the airfoil after the peak velocity must decelerate such that its pressure is consistent with the zone into which it dumps, but if the amount of deceleration required to do so is too great, flow reversal

---

<sup>8</sup>This is also true of other airfoil interactional effects, but is easiest to understand for the circulation effect.

(separation) may occur. Hence, there is a limit imposed on the maximum suction which can be achieved.

In multi-element airfoils however, a leading element may be dumping into an area of lower pressure created by a trailing element. In this case, the flow may actually continue to decelerate *off the surface* of the leading element (and above the trailing element). Moreover, flow in the wake of an airfoil can sustain greater changes in pressure than can the flow in an (attached) boundary layer. Hence, with pressure recovered in the wake rather than on the surface, greater lift may be achieved for the multi-element airfoil's elements.

#### 1.2.2.5 Fresh-Boundary-Layer Effect


Smith notes that in a properly-designed multi-element airfoil, each element will develop its own boundary layer, meaning that the boundary layer at any given point will be thinner than would be the case for a single airfoil of equivalent chord. Since a thinner boundary layer can sustain a greater pressure gradient than a thicker one, separation on the multi-element airfoil will be delayed, and a greater lift may be achieved.

#### 1.2.3 Scissors Rotors

The “scissors rotor” is most famously used on the tail rotor of the Boeing (formerly Hughes) AH-64 Apache (Fig. [1.6a](#)), but in fact was developed earlier in an effort to reduce the noise produced by the tail rotor of a Hughes OH-6 [18]. This design is also used on the tail rotor of several helicopters designed by the Mil Design Bureau. There is no distinction between a scissors rotor and a stacked rotor, but it should be noted that all rotors *referred to* as scissors rotors are designed



(a) Boeing AH-64 tail rotor, from [20]



(b) Bell scissors main rotor, excerpted from [8]

Figure 1.6: Example applications of “scissors” rotors

with a fixed negative<sup>9</sup> phase angle. Hughes seems to have arrived at the configuration serendipitously (unequal spacing having been required for packaging reasons), whereas Mil suggests a more deliberate approach. Both Hughes and Mil claim acoustic benefits, while Mil also claims performance improvement, at least in hover [19].

The only apparent effort to use a scissors *main* rotor was by Bell Helicopter in the early 1970’s (Fig. 1.6b), and was largely unsuccessful. A modest performance *loss* in hover was noted compared to a conventional rotor. This loss of performance was compounded in forward flight by the increased drag of the scissors rotor hub. Acoustic performance was mixed—although the overall sound pressure was lower for the scissors rotor than for a two-bladed rotor, and comparable to that of a conventional four-bladed rotor, human observers rated the scissors rotor as noisier than either [8].

Ultimately, while at first blush the scissors rotors would seem to be a step towards modern stacked rotors, the dearth of published data—or indeed, evidence of a systematic approach to their design—limits the applicability of this earlier development to the current work. Moreover, the

---

<sup>9</sup>Being careful to note that this is negative according to the definition of phase angle used in this work. Older studies of scissors rotors typically define the phase angle in the opposite direction, i.e. positive when the lower blade *leads* the upper, or simply as the acute angle between the blades.

scissors rotor design concept is contrary to the idea of taking advantage of beneficial airfoil interactions, which would generally require that the upper blade be leading the lower. However, until recently these scissors rotors have been the only ones actually developed for use on production rotorcraft.

#### 1.2.4 Towards a Modern Understanding

A more systematic study on “variable geometry rotors” was performed by Landgrebe and Bellinger in 1974 [7]. Unlike the previous work of Hamilton Standard and Kaman which focused on airfoil effects, Landgrebe and Bellinger’s hypothesis was that rotor performance could be improved by altering *blade-vortex interactions* through the use of then-novel variations in rotor geometry. The specific geometry which they varied included vertical and azimuthal offsets between the rotor blades—hence, creating a stacked rotor—but also different rotor blade lengths and tip geometries.

While the latter geometric variation proved to have little effect, the stacked rotor parameters were noted with interest. In particular, an increase of nearly 13% in hover figure of merit was achieved by setting a vertical offset of 2 chord lengths and a phase angle of  $30^\circ$  (compared to an equal-spacing angle of  $60^\circ$  for the 6-bladed rotor being tested). No effect was noted due to collective offset, but only a fairly small range of  $\Delta\theta_0 = \pm 1^\circ$  was examined. The effect on forward flight performance was found to be negligible, in contrast to the contemporaneous effort by Bell/Sonneborn.

Examining the tip vortex positions, Landgrebe and Bellinger noted that in the cases with improved performance, the tip vortex from the upper/leading rotor blade would pass *above* the

lower/trailing rotor blade, before convecting down into the wake. (This is in contrast to the theory of the scissors rotor, which posits that the tip vortex from the leading blade passing *below* the trailing blade is beneficial [19].) Likewise, the tip vortices would pass further below the upper blades than in a coplanar rotor. Although they were able to make these observations, Landgrebe and Bellinger expressed doubt as to what mechanism was leading to the increased rotor performance.

Two years later in 1976, Rorke followed up on the model-sized (0.7 m radius) experiments of Landgrebe and Bellinger with tests of a full-scale (8.9 m radius) rotor [11]. Rorke's hypothesis was the same—that improvements in performance were due to beneficial changes in blade-vortex interactions—but the range of parameters tested was more limited than the earlier study. Still, Rorke was able to confirm that the benefits noted by Landgrebe and Bellinger did indeed scale to a full-size rotor. Rorke also noted acoustic gains of up to 4 dB, which corresponded to configurations of improved performance. (This is not inconsistent with Bell/Sonneborn, but Rorke did not investigate the subjective impact of rotor noise on human subjects.) Negatively, Rorke reported that the lower rotor blades of the stacked rotor would—apparently randomly—go out-of-track, an effect not observed on the small-scale rotor. This was attributed both to the small-scale rotor being much stiffer in flap, and to the full-scale rotor being tested in less-controlled outdoor conditions; the softer blades were thought to sensitize the full-scale rotor to slight perturbations in the wind.

### 1.2.5 Rediscovery and Recent Efforts

Over the course of the past decade there has been increasing interest in the stacked rotor concept, but this interest began almost by accident. In particular, several studies primarily focused on *counter-rotating* coaxial rotors sought quantify the amount of swirl velocity occurring therein by comparing them to co-rotating coaxial rotors. The validity of this comparison is dubious given the presence of coupled-airfoil and tip-vortex effects in stacked rotors which are different than those observed in counter-rotating coaxial rotors, but it would take several studies before this came to be understood.

In 2015, Ramasamy [21] was the first to attempt such a comparison (within a larger study covering coaxial, tandem, and tilt rotors), noting that at the time there had been no measurements made to validate or quantify the swirl recovery hypothesis. Ramasamy laid out the procedure for this comparison, which would be used in similar form by several later works—the collectives of the upper and lower rotors of the stacked rotor would be determined by trim such that the upper and lower rotors would have balanced torque<sup>10</sup>, and a total thrust equivalent to that of the counter-rotating coaxial rotor. Hence, it was assumed, the resulting difference in total power between the stacked and counter-rotating coaxial rotors could be attributed to swirl recovery in the latter.

Using this procedure, Ramasamy observed a 5% increase in power for the stacked rotor versus the counter-rotating coaxial rotor—which he attributed to swirl velocity recovery—but also that the stacked rotor was still 5% more efficient than an isolated rotor, which he noted as being consistent with the earlier work of Rorke [11]. (Although it should be pointed out that this is less than the 13% improvement in efficiency noted by Landgrebe and Bellinger [7], a differ-

---

<sup>10</sup>In the same direction, whereas for a counter-rotating coaxial rotor this would be equal and opposite.

ence which was not discussed.) Uniquely compared to both earlier and later studies, Ramasamy did not use synchronized rotors; rather, the two rotors had entirely independent drivetrains and motor controllers. Despite this, Ramasamy observed that the rotors self-synchronized to a small positive<sup>11</sup> phase angle, a synchronization that was maintained even after attempts to perturb the speed of one rotor (in which case the other rotor would adjust its speed to match).

Subsequently in 2017, Uehara and Sirohi [22] sought to examine swirl velocity recovery in a more systematic fashion. Modifying Ramasamy’s procedure slightly, their stacked rotor would have *torque* matched to the counter-rotating coaxial rotor<sup>12</sup>, such that differences in performance would appear in the *thrust* for each rotor system. Significantly, Uehara and Sirohi used a stacked rotor which was mechanically synchronized (like the earlier studies of the 1970’s); hence, a phase angle would need to be selected for comparison between the stacked and counter-rotating coaxial rotors.

In an attempt to determine an appropriate value, Uehara and Sirohi performed a sweep of the phase angle, trimming as described at each test condition. Two of those cases were eventually chosen for comparison to the counter-rotating coaxial rotor:  $\Delta\psi = 90^\circ$  where the upper rotor of both the stacked and counter-rotating rotors was closest (where swirl velocity recovery would account for a 4% increase in thrust), and  $\Delta\psi = -50^\circ$  where the upper and lower rotors of the stacked rotor generated approximately equal thrust (where swirl velocity recovery would account for a 2.7% increase in thrust). That these conditions are somewhat arbitrary should suggest the futility of this comparison.

Of particular interest however was that at  $\Delta\psi = 10^\circ$ , Uehara and Sirohi noted that the

---

<sup>11</sup>Ramasamy noted that “[...]the upper rotor led the lower by about 1/4 chord”, but did not specify a numerical value of the angle.

<sup>12</sup>With the upper and lower stacked rotors still having identical torque to match the equal-and-opposite torque of the counter-rotating coaxial rotor.

stacked rotor generated 4% *more* thrust than the counter-rotating coaxial rotor. Although Uehara and Sirohi did not examine the mechanisms which caused this increase in performance, they did acknowledge both the earlier work of Rorke and the idea that biplane theory would apply to a stacked rotor system.

Starting to move away from the swirl recovery hypothesis, Bhagwat [5] in 2018 sought to use RCAS [23] with a free-wake inflow model to gain greater insight into the aerodynamics of both stacked and counter-rotating coaxial rotors. Bhagwat concurred with Uehara and Sirohi that, given the parameters of their rotor system, the stacked rotor could be expected to have superior performance to a counter-rotating coaxial rotor in the range of  $\Delta\psi = 10^\circ$  to  $\Delta\psi = 15^\circ$ . Bhagwat also concurred with Ramasamy that their stacked rotor would have inferior performance to a counter-rotating coaxial rotor, although with a lesser difference than the 5% which Ramasamy had found. The reason for the closer performance of the stacked and counter-rotating coaxial rotors was not discussed; however, the differences in methodologies (experimental vs. computational) were not examined. Bhagwat did speculate that the difference between Uehara and Ramasamy's results could be due to differences in blade loading or in the relationship between the axial spacing and the number of rotor blades. Importantly, Bhagwat observed that the differences in self-induced and interference-induced inflow (i.e. the influence of the wake) are a more significant difference between stacked and counter-rotating coaxial rotors than any swirl recovery, agreeing with (but not explicitly acknowledging) the earlier works of Landgrebe and Bellinger and Rorke.

Other studies around this time also began examining the application of lower-fidelity rotor models to the stacked rotor. Misiorowski et.al. [24] sought to examine both the performance and acoustics of a stacked rotor (using the design of Uehara and Sirohi), using both Helios (a high-

fidelity 3D CFD code) and RCAS with viscous vortex particle method (VVPM, a lower-fidelity code). With regard to performance, good agreement was observed between Helios and the prior experimental data. However, it was observed that RCAS/VVPM predicted both a greater collective (for trim) and power loading than was seen in the experiment. Differences in spanwise load distribution were also observed between the Helios and RCAS/VVPM results. These differences were attributed to differences in the rotor wakes between the models, but the reasons for the differences in the wakes were not thoroughly examined.

The study with the most similar thought process to the current work however might be the one performed by Pollard et.al. [25], who sought to model a stacked rotor through the use of an unsteady 2D panel method code applied to the *system of two airfoils* along the span of the rotor blades. Indeed, Pollard et.al. appear to have been the first “modern” study which recognized the need for a comprehensive approach to model airfoil interactions in stacked rotors. However, their study had some major limitations which limit its usefulness to the current work. At the time of writing, their panel method did not include viscous effects, hence power could not be accurately determined. More fundamentally, Pollard et.al. did not make any attempt to model the rotor wakes—the airfoil system was treated as a “single” blade, and a standard Prandtl tip-loss factor was applied to account for wake effects. As many previous studies have examined the importance of the wake geometry of the stacked rotor, it is dubious how accurately isolated rotor assumptions (such as Prandtl tip-loss) could be applied to a stacked rotor.

The most recent work in the field has been focused around a new stacked rotor test stand, introduced in 2020 by Johnson et.al. [10] (Fig. 1.7) with computational work by Jacobellis et.al. [3]. This new test stand allows for the controlled variation of both phase offset *and* vertical offset (like Landgrebe, but unlike some other modern studies which were more limited). While broadly




Figure 1.7: Stacked rotor test stand introduced by Johnson et.al., excerpted from [10]

confirming the results of Landgrebe (regarding the positive effects of combining vertical and axial offset), Johnson tested a larger parameter space in finer resolution, and presented results in a way which demonstrated both the individual and coupled effects of the parameters. In particular, Johnson et.al. tested a wider range of collective offset ( $\Delta\theta_0 = \pm 4^\circ$ ) and found that, although best efficiency was achieved for  $\Delta\theta_0 = 0^\circ$ , setting negative  $\Delta\theta_0$  would increase the *sensitivity* of thrust to phase offset, which would be beneficial if  $\Delta\psi$  were to be used as a novel rotor control [4].

On the computational side, Jacobellis et.al. sought to further investigate the RCAS/VVPM methodology previously used by Misiorowski et.al. One key observation of Jacobellis et.al. was that the VVPM code had not been including the effects of *bound circulation* on the interaction between the blades of the stacked rotor—rather, all interaction occurred only due to the presence of the tip vortices. Such bound circulation effects would be significant components of<sup>13</sup> the slat and circulation effects noted by Smith [17]. After being modified to include these bound

---

<sup>13</sup>However, bound circulation interaction would *not* account for any thickness effects.

circulation effects, the VVPM code was shown to produce more accurate results [3]. However some disagreement between RCAS/VVPM and the experiment was noted, particularly at higher values of differential collective and at negative phase angles [4].

### 1.2.6 Summary

Many previous studies have confirmed that stacked rotors can offer greater efficiency than conventional/isolated rotors. With proper selection of vertical and phase offsets of the rotors, it is even possible for a stacked rotor to offer superior performance to a counter-rotating coaxial rotor. (On the other hand, with poor parameter selection, a conventional isolated rotor may offer superior performance to a stacked rotor.) Stacked rotors may also offer acoustic benefits as compared to other types of rotors<sup>14</sup>.

The performance benefits possible with stacked rotors are due to the complex aerodynamic interactions—both airfoil-to-airfoil and wake-to-blade—which they exhibit, interactions which are different from those found in either conventional isolated or counter-rotating rotors. Historically these interactions in stacked rotors have been insufficiently characterized, but recent research has made them increasingly well-understood.

Many previous studies have not recognized the importance of *both* types of aerodynamic interaction in stacked rotors, instead choosing to focus on one or the other (or neither, as apparently the case of some scissors rotors). With respect to airfoil-to-airfoil interactions, Hamilton Standard clearly recognized the potential and designed a propeller system accordingly, while (although not focused on rotary-wing applications specifically) Smith thoroughly described the

---

<sup>14</sup>Although rotor acoustics are outside the scope of the present work, rotor acoustic models typically require some fidelity of aerodynamic modeling to produce accurate results.

mechanisms behind these interactions. On the other hand, many studies—beginning with Landgrebe, through more recent work by Bhagwat—examined the wake structure of stacked rotors and its effect on performance, while at most only mentioning in passing the concept of biplane theory. Only more recently (e.g. [3]) have some studies recognized the need for both types of interaction.

One key limitation of almost all previous studies has been that these studies have focused mainly on characterizing some *particular* rotor. While it is easy to observe that a stacked rotor can offer superior performance with “some parameters”, it is more difficult to state for an *arbitrary* rotor which specific parameters will lead to increased performance without extensive investigation.

### 1.3 Rotor Design Methodology

The process of such investigation to determine the “best”<sup>15</sup> parameters for a rotor in a given application is simply the definition of the *rotor design process*. While the basics of the process should be immediately familiar to any professional in the field of rotorcraft it is nonetheless illustrative to describe them, as a means to observe where prior studies fit in, and the present limitations of the understanding of stacked rotor theory.

Rotor design relies on a hierarchy of aerodynamic models, an overview of which is shown in Figure 1.8. These models run on a scale from simple, inexpensive<sup>16</sup>, lower-fidelity models to more expensive, higher-fidelity models. Lower-fidelity models are used early in the design process, as they allow for rapid iteration of rotor design parameters<sup>17</sup>, examining potentially

---

<sup>15</sup>For some definition thereof

<sup>16</sup>In terms of computational cost, time spent by the designer, etc.

<sup>17</sup>Not just stacked rotor-specific parameters, but conventional rotor parameters such as chord, twist, number of




Figure 1.8: Hierarchy of aerodynamic models commonly used in rotor design

thousands of possible combinations in a matter of minutes to hours. The results of these lower-fidelity models allow for down-selection of the parameter space, yielding a much smaller number of configurations which can then be further studied and refined using higher-fidelity methods. These higher-fidelity methods, such as 3D CFD and physical models, may take days or weeks to analyze a single configuration, but their results can be expected to translate directly to the real world. In general, higher-fidelity methods may be applied directly to stacked rotors, as those methods simulate<sup>18</sup> fully the physics of the rotor system.

### 1.3.1 Lower-Fidelity Rotor Models

All lower-fidelity rotor models are able to provide faster results than higher-fidelity models by making simplifying assumptions about the physics involved. Particular lower-fidelity models are distinguished by which specific assumptions they make. For instance, the rotor's inflow could be computed by using momentum theory (as in BEMT) or by computing the influence of the wake vortices at the location of the rotor blades, with the geometry of the wake vortices determined by semi-empirical parameters (prescribed wake) or modeled as part of the simulation (free-wake). Lower-fidelity rotor models may be expected to give accurate results within the limitations of their assumptions, such as:

- Blade-element momentum theory is well-developed for a conventional, isolated rotor operating in hover.
- Canonical methods exist for prescribing the vortex wake geometry of an isolated rotor in hover and forward flight.

---

blades, etc.

<sup>18</sup>Computationally or sub-scale physically

- Tandem and multi<sup>19</sup>-rotors in forward flight will have interacting wakes, which can be captured by a free-vortex wake model.

Outside the limitations of their assumptions lower-fidelity rotor models may or may not be accurate; BEMT in forward flight for instance provides at best a rough estimate of performance even for an isolated rotor.

A major simplifying assumption of lower-fidelity rotor models is that the aerodynamic loads on their *blades* can be determined by integration of the aerodynamic loads on their *airfoils*, as is the case with lifting-line models<sup>20</sup>. The relationship between “free-stream” flow conditions and airfoil loads is given by an *airfoil table*. Such tables are freely available for many canonical airfoil geometries, or may be determined through comparatively-inexpensive 2D CFD methods. The goal of the lower-fidelity model is thus to determine the flow conditions—particularly the inflow—at a series of points along the span of the rotor blade. Once the flow conditions are known, the airfoil aerodynamic loads may be determined from an airfoil table, then integrated (over the length of the span) to determine the blade loads. These blade loads may be coupled back to the inflow, in which case the method is iterated until a converged aerodynamic solution is obtained.

Note that strictly speaking, “lower-fidelity” and “higher-fidelity” are a continuum, not a categorization of models; it would, for example, be correct to say that free-wake is a higher fidelity model than BEMT when comparing the two. For the purposes of this study however, it is useful to separate those models and methods which fully simulate the airflow around the rotor blades to determine the aerodynamic forces (3D CFD and physical models) from those which

---

<sup>19</sup>Quad-, hexa-, etc. as are often seen on commercially-available “drones”

<sup>20</sup>Prescribed and free-wake inflow rely on a lifting-line blade model to determine the trailed vortex strength; BEMT simply takes airfoil lift into account in its inflow calculation without using lifting-line specifically.

rely on airfoil tables to determine the loads (e.g. BEMT, prescribed and free-wake). In this study, the latter category will be referred to generally as “lower-fidelity models”, whereas the former category will be referred to as “higher-fidelity models.”

### 1.3.2 Limitations and Application to Stacked Rotors

Since lower-fidelity rotor models rely on airfoil tables, those rotor models will inherit whatever assumptions the airfoil tables include. Namely, airfoil tables are nearly always generated for an isolated airfoil in a steady, uniform free-stream flow. Plainly, these assumptions will not be strictly true even for an isolated, conventional rotor—the flow field near the rotor blade may be complex, and any cyclic control inputs or free-stream off-axial to the rotor imply unsteady airfoil behavior. However, these can be mitigated for a conventional rotor—the complexity in the flow around the rotor is generally of sufficiently large scale that the aerodynamic conditions along the blade can be taken as “free-stream” for its airfoils, and models exist (such as Theodorsen) to account for the effects of unsteady airfoil behavior.

The key difference between a stacked rotor and a conventional (single or counter-rotating) rotor is the presence of closely-spaced blades. While the interactions between one blade’s wake and the following blade are not-dissimilar from those found in conventional rotors, there is also significant interaction present at the airfoil level. Some of this interaction can be accounted for by lifting-line models so long as the influence of the bound circulation is included [3], which is not an inherently *difficult* calculation but may or may not actually be performed by a given rotor code. If such lifting-line interaction is included, it is important to note that the airfoils are being modeled as line (3D) or point (2D) vortices, and the influence is entirely due to the change in

the equivalent “free-stream” conditions around the airfoil. While this is clearly an improvement over no interaction whatsoever, the work of Smith [17] should make it clear that this “free-stream” interaction is insufficient to fully capture the interaction between the airfoils—effects due to finite thickness and *highly localized* changes in flow speed and direction are ignored.

## 1.4 Objectives of the Present Work

The objective then of the current work is to develop a methodology for lower-fidelity modeling of stacked rotors with improved accuracy by addressing the shortcomings of currently-available models. To do this, high-fidelity 2D CFD will be used to accurately model the interactions between the airfoils. Then, this 2D CFD data will be coupled to an existing lower-fidelity 3D rotor modeling tool. Through careful choice of lower-fidelity methodology and data coupling, it will be possible for this methodology to take into account both blade-to-blade (airfoil-level) and wake-to-blade (rotor-level) aerodynamic interactions.

A key difference between the methodology presented in this work and prior studies is that the methodology will be geared towards future use in the *design* of stacked rotors, not just *analysis*. Hence, in addition to validating the methodology, emphasis is placed on making the underlying tools modular (to fit together with existing design tools) and deployable (making the tools easily obtainable), thus making the tools and methodology practical for future rotor designers.

## 1.5 Organization of Thesis

With an understanding of the concept of stacked rotors, it is next necessary to understand how the 2D airfoil and 3D rotor systems relate to each other, before proceeding to discuss the 2D and 3D models themselves. This relationship is described and formulated as necessary in Chapter 2. All aspects of the 2D airfoil models are described in Chapter 3, including the lifting-vortex model and 2D CFD, along with mesh generation for the latter. In Chapter 4, the 2D methodology is used to explore the behavior of systems of two airfoils, and it is demonstrated under what conditions each model may or may not be valid. The 3D methodology, consisting of the rotor aerodynamic and structural models, and the specific means by which the 2D airloads are included, is presented in Chapter 5. Validation of the complete methodology is presented in Chapter 6, along with some experiments on alternate methodologies. Finally, Chapter 7 summarizes the accomplishments of the current work, and presents the intended direction for future research.

## Chapter 2: System Parameterization

Before proceeding to describe the solution methodologies, it is important to consider the relationship between the blades of the stacked rotor in 3D and their corresponding systems of airfoils in 2D. For conventional rotors, this relationship is simple—if the flow speed ( $M$ ) and direction ( $\alpha$ )<sup>1</sup> along the blades of the rotor<sup>2</sup> in 3D are known, then the aerodynamic forces ( $C_l, C_d, C_m$ ) on the airfoil in 2D may be determined directly from the airfoil tables<sup>3</sup>; the blade forces in 3D may then be computed by integrating the 2D aerodynamic forces. For a stacked rotor, the relative positions of the blades are of great importance, while  $\alpha$  and  $M$  must now be considered on *both* adjacent rotor blades. Careful consideration must be given to both the geometric and aerodynamic properties of the 2D and 3D systems to ensure that the modeled system accurately represents the physical system.

### 2.1 Aerodynamic Conditions

In this study, it is assumed that the rotor is operating in a flight condition in which the free-stream flow is radially symmetric, i.e. hover, axial climb, or axial descent. This assumption is coincidentally identical to the regime in which BEMT may strictly be considered valid [9]. For

---

<sup>1</sup>Where for a rotor  $\alpha = \theta - \phi$

<sup>2</sup>Hence,  $M$  and  $\alpha$  are functions of both  $r$  and  $\psi$

<sup>3</sup>It is assumed that as a condition of the rotor design, the airfoil profile along the span of the blades is specified.

stacked rotors in particular, this assumption may be justified by considering their proposed use cases, particularly as lifting-only rotors which are to be stowed during high-speed forward flight or as a technology to enable more highly efficient loitering.

Further, it is assumed that there is no crossflow on the blades of the stacked rotor. Even in an axially flight condition this will not be strictly true, due to 3D effects contained in the rotor wake. This assumption is not entirely unreasonable since the crossflow component does not contribute to lift [26], and for an axial flight condition the crossflow should only be significant near the edges of the blades. As long as the regions where crossflow is significant are sufficiently small—and in the case of the root, contribute little to the total rotor thrust and power—then the loss of accuracy in those regions can be considered acceptable.

## 2.2 System Geometry

Consider two blades of a stacked rotor, where the lower rotor is offset from the upper rotor by a phase angle  $\Delta\psi$  (positive when the lower rotor trails the upper) and vertically by  $\Delta z_{hub}$  (positive for the lower rotor below the upper). Assuming that there is no crossflow on the blades, the interaction between the rotor blades' airfoils can be represented by a section cut in the hub frame along a constant radius<sup>4</sup>, as shown in Figure 2.1a. Taking the section cut in the hub plane (denoted by the  $hp$  subscript), the two-airfoil system has four geometric parameters: The angle of the top ( $\theta_1$ ) and bottom ( $\theta_2$ ) blades with respect to the horizon, and the vertical ( $\Delta z_{hp}$ )<sup>5</sup> and

---

<sup>4</sup>If crossflow is significant, then this method could be performed with the section cut following the streamlines of the flow, which would require the directionality of the flow to be known *a priori*. This could be determined through reasonable assumptions or output from the rotor aerodynamic code. Consideration of significant crossflow is outside the scope of this work, but is an avenue for future research.

<sup>5</sup>Including the vertical separation at the hub, as well as flapping and bending.




Figure 2.1: Geometric parameters of the stacked rotor blade and airfoil systems

horizontal ( $\Delta x_{hp}$ )<sup>6</sup> separations between the airfoils' quarter-chords, as shown in Figure 2.1b.

These four parameters may be reduced to three by applying the following transformation from the hub frame to the upper airfoil's frame of reference:

$$\begin{aligned}
 \Delta x &= \Delta x_{hp} \cos(\theta_1) + \Delta z_{hp} \sin(\theta_1) \\
 \Delta z &= \Delta z_{hp} \cos(\theta_1) - \Delta x_{hp} \sin(\theta_1) \\
 \Delta \theta &= \theta_2 - \theta_1
 \end{aligned} \tag{2.1}$$

The result of this transformation is shown in Figure 2.1c. Note that in both the hub and airfoil reference frames, positive  $\Delta x$  denotes the lower airfoil behind the upper, positive  $\Delta z$  denotes the lower airfoil below the upper, and positive  $\theta$  or  $\Delta \theta$  denotes positive (clockwise) pitch.

<sup>6</sup>For straight blades, the horizontal separation is related to the phase angle by  $\Delta x_{hp} = r\Delta\psi$ . Lagging and bending effects must also be included.

### 2.2.1 Geometric Uniqueness

It should be clear that the choice of reference airfoil frame in Equation 2.1 is arbitrary. For any system of two airfoils described by  $\Delta x, \Delta z, \Delta \theta$  referenced to one airfoil, the system could instead be described by  $\Delta x', \Delta z', \Delta \theta'$  referenced to the other airfoil, where:

$$\begin{aligned}\Delta x' &= -\Delta x \cos(\Delta \theta) - \Delta z \sin(\Delta \theta) \\ \Delta z' &= -\Delta z \cos(\Delta \theta) + \Delta x \sin(\Delta \theta) \\ \Delta \theta' &= -\Delta \theta\end{aligned}\tag{2.2}$$

By this definition,  $\Delta x', \Delta z',$  and  $\Delta \theta'$  retain the same sign convention as  $\Delta x, \Delta z,$  and  $\Delta \theta$ . Assuming that both airfoils have the same profile<sup>7</sup>, this means that the choice of  $\Delta x, \Delta z,$  and  $\Delta \theta$  is not unique—the system defined by  $\Delta x', \Delta z',$  and  $\Delta \theta'$  is strictly identical. This observation allows the domain of  $\Delta x', \Delta z',$  and  $\Delta \theta'$  requiring 2D simulations to be limited (see Chapter 3.3.2), which helps reduce computational overhead<sup>8</sup>.

## 2.3 Consideration of Flow Conditions

Thus far only the geometry of the stacked rotor system has been considered, but the flow conditions around the rotor blades are clearly of vital importance. For a conventional rotor,  $\alpha$  and  $M$  along the blades would be determined by the inflow model used in the reduced-order 3D simulation (see Chapter 5.2.1), and those  $\alpha$  and  $M$  would be used as the free-stream (boundary)

---

<sup>7</sup>At present this is true of all known stacked rotors, but it is conceivable that a stacked rotor could be designed with different upper and lower rotor blade airfoils.

<sup>8</sup>If a stacked rotor were to be designed with different upper and lower rotor airfoils, then similar limitations on the domain could be imposed if the domain of the rotor parameters are limited, e.g. if it is specified that  $\Delta z_{hub} \geq 0$ .

conditions of the 2D airfoil simulation<sup>9</sup>. For the stacked rotor, the 3D inflow model will still compute the *local*  $\alpha$  and  $M$  along each of the blades, but the 2D airfoil simulation requires a single *global* free-stream boundary condition ( $\alpha_\infty$  and  $M_\infty$ ) for the airfoil system. For consistency with the system geometry,  $\alpha_\infty$  is referenced in this work with respect to the upper airfoil's chord line<sup>10</sup>.

The question now becomes how to determine a set of global flow conditions which will produce the correct local conditions on the rotor blades' airfoils. This is primarily a matter of flow directions; flow speed is dominated by the rotational component  $M_\Omega(r) = r\Omega/a$ , and varies quite little between the blades at any radial station. The flow *direction* however is much more sensitive to small perturbations. These perturbations may arise from such effects as:

- The mutual influence of each airfoil on the other [17], particularly due to bound circulation
- The effect of the trailed wake vortices, particularly the effect of the upper blade's tip vortex on the lower blade
- The contraction of the wake from the upper rotor to the lower rotor

From this list, it should be apparent that different sources of flow perturbation will affect each airfoil differently. It is important to match the angle of attack on both airfoils—hence two dependent parameters  $\alpha_1$  and  $\alpha_2$ —between the 2D and 3D models, since airfoil lift and drag depend significantly on  $\alpha$ . However, it is mathematically impossible to guarantee that the two dependent parameters can be matched if only one free parameter—in this case, the free-stream angle of attack  $\alpha_\infty$ —is available. Therefore, it is necessary to find a second free parameter.

---

<sup>9</sup>Or, more likely, as the free-stream lookup parameters for a pre-computed airfoil table.

<sup>10</sup>But could be equivalently referenced to the lower airfoil's chord line as  $\alpha'_\infty = \alpha_\infty + \Delta\theta$  (following the notation in Equation 2.2).




Figure 2.2: Flow parameters of the stacked rotor blade and airfoil systems

### 2.3.1 Modification of Geometry

Consider again the system of rotor blades shown previously in Figure 2.1a, but now including the flow around the rotor computed by the 3D inflow model (Figure 2.2a). Taking a section cut at a constant radius, the inflow model specifies the flow speed ( $u_{1,2}$ ) and direction ( $\phi_{1,2}$  or  $\alpha_{1,2}$ ) at each airfoil (Figure 2.2b); the geometry is also specified as described in the prior section<sup>11</sup>. Notwithstanding the *source* of any perturbations in the flow, the *effect* of perturbations in the flow is to change the direction of the free-stream flow.

Let a new conceptual reference frame be defined which follows the streamlines of the flow. Such a reference frame would be suitable for modeling using 2D CFD, since a uniform free-stream (boundary) condition could be defined. The geometry of the airfoil system in the

<sup>11</sup>Geometric specification is omitted in Figure 2.2 for clarity.

flow-referenced frame would look similar to the geometry in the physical system, but with some rotation of the airfoils relative to each other. The differential pitch between the two airfoils in the flow-referenced frame represents the *effective* angle between the two airfoils; let that angle be denoted  $\Delta\theta_{eff}$ . In practice, the direction of the streamlines of the flow is neither necessarily known nor easily computed. However, by using  $\Delta\theta_{eff}$  as the second free parameter with  $\alpha_\infty$ , it is possible to obtain a model system (Figure 2.2c) which allows  $\alpha_1$  and  $\alpha_2$  to be matched to the conditions predicted by the 3D inflow model.

Exactly how to determine  $\alpha_\infty$  and  $\Delta\theta_{eff}$  depends on whether or not the mutual effects of bound circulation on the 2D airfoils are included in the 3D rotor inflow model. In either case, the inflow model must include all relevant 3D effects—particularly the wake—since these cannot be determined strictly in 2D.

### 2.3.1.1 Rotor Models Including the Influence of Bound Circulation

Some rotor inflow models treat the bound vorticity of the rotor blades as no different than the shed (wake) vorticity. In this case, the mutual effect of the airfoils on each other must be isolated, otherwise the effect would be double-counted. To accomplish this, a 2D lifting-vortex model of the airfoils (see Chapter 3.1)—which serves as a 2D approximation of a 3D lifting-line model—must be used. Assuming that the inputs to the 2D lifting-vortex model include  $\alpha_\infty$  and  $\Delta\theta$ , a simple iterative scheme may be used to obtain the correct  $\alpha_{1,2}$ . Rather than analyzing specifically how  $\alpha_\infty$  and  $\Delta\theta$  effect  $\alpha_{1,2}$  in the 2D lifting-vortex model, the following simplifying assumptions are made about the relationship between the system inputs and outputs:

- $\alpha_\infty \propto \text{mean}(\alpha_1, \alpha_2)$

- $\Delta\theta_{eff} \propto (\alpha_2 - \alpha_1)$
- Changing  $\Delta\theta_{eff}$  by an amount  $\delta$  results in a proportional change in  $\alpha_2$ , hence a change of  $0.5\delta$  in  $\text{mean}(\alpha_1, \alpha_2)$

It will be shown in Chapter 4.1 that these assumptions are highly similar to the actual behavior of the airfoil system prior to stall<sup>12</sup>, particularly when only the bound circulation effects are considered. Although these assumptions are not exactly true, they are sufficient to allow convergence of the iterative scheme.

With the approximate relationship between the inputs and the outputs defined, a simple under-relaxed fixed-point iteration is used to converge both  $\alpha_\infty$  and  $\Delta\theta_{eff}$  simultaneously. Based on the assumed relationships given above, the mean and difference of  $\alpha_{1,2}$  are used as the convergence criteria rather than  $\alpha_{1,2}$  directly, converging either pair also converges the other pair since:

$$\begin{bmatrix} \text{mean} \\ \text{diff} \end{bmatrix} = \begin{bmatrix} 0.5 & 0.5 \\ -1 & 1 \end{bmatrix} \begin{bmatrix} \alpha_1 \\ \alpha_2 \end{bmatrix} \longleftrightarrow \begin{bmatrix} \alpha_1 \\ \alpha_2 \end{bmatrix} = \begin{bmatrix} 1 & -0.5 \\ 1 & 0.5 \end{bmatrix} \begin{bmatrix} \text{mean} \\ \text{diff} \end{bmatrix} \quad (2.3)$$

The initialization and iteration of this scheme is shown in Algorithm 2.1.

### 2.3.1.2 Rotor Models Not Including the Influence of Bound Circulation

Other rotor inflow models<sup>13</sup> may ignore the effect of the bound vorticity entirely. This is not an unreasonable assumption for conventional rotors, where the blades are always<sup>14</sup> far enough

---

<sup>12</sup>Allowing that  $y \propto x$  indicates a relationship such as  $y = ax + b$  for some values  $a > 0$  and  $b$

<sup>13</sup>Particularly RCAS's prescribed and free wake models, which are of interest in this work.

<sup>14</sup>Or almost always, in the case of a counter-rotating coaxial rotor

---

**Algorithm 2.1:** Defining the parameters for the 2D airfoil model, where the corresponding 3D inflow model includes the influence of bound circulation

---

```

Input :  $\alpha_1^{(\text{inflow})}, \alpha_2^{(\text{inflow})}$  from 3D inflow model
mean =  $0.5 \left( \alpha_1^{(\text{inflow})} + \alpha_2^{(\text{inflow})} \right)$ ;          /* Set target criteria */
diff =  $\alpha_2^{(\text{inflow})} - \alpha_1^{(\text{inflow})}$ ;
tol =  $\delta$ ;
 $\omega = 0.95$ ;                      /* Relaxation parameter */
 $\alpha_\infty = \text{mean}$ ;          /* Set initial conditions */
 $\Delta\theta_{eff} = \text{diff}$ ;
for a number of iterations do
     $\left[ \alpha_1^{(\Gamma)}, \alpha_2^{(\Gamma)} \right] = \text{vortexAirfoils2D}(\alpha_\infty, \Delta\theta_{eff}, \dots)$ ; /* See Chapter 3.1 */
     $\varepsilon_{\Delta\theta} = \text{diff} - \left( \alpha_2^{(\Gamma)} - \alpha_1^{(\Gamma)} \right)$ ;          /* Define error terms */
     $\varepsilon_\alpha = \text{mean} - 0.5 \left( \alpha_1^{(\Gamma)} + \alpha_2^{(\Gamma)} \right) - 0.5\varepsilon_{\Delta\theta}$ ;
    if  $|\varepsilon_\alpha| < \text{tol}$  then
        break;                      /* Converged */
    else
         $\alpha_\infty = \alpha_\infty + \omega\varepsilon_\alpha$ ;          /* Update parameters */
         $\Delta\theta_{eff} = \Delta\theta_{eff} + \omega\varepsilon_{\Delta\theta}$ ;
    end
end
return  $\alpha_\infty, \Delta\theta_{eff}$ ;

```

---




Figure 2.3: Notation for a generic vortical disturbance around the airfoil system

apart for their bound circulation to have little influence on each other. For stacked rotors however this interaction between the airfoils is a significant concern.

The difficulty with trying to select  $\alpha_\infty$  and  $\Delta\theta_{eff}$  for this case arises from the fact that there is now no completely reliable reference condition—the airfoils do not interact with each other so the 2D lifting-vortex model provides no useful information, and it is not known *a priori* how perturbations in the flow affect each airfoil individually. However, considering the sources of perturbations in the flow discussed previously, it can be seen that most perturbations *near* the airfoils—excluding mutual bound circulation effects—are due to the effect of the upper rotor on the lower rotor not vice-versa, and hence will mostly affect  $\Delta\theta_{eff}$ . As perturbations such as the tip vortices convect further into the wake they move farther away from *both* airfoils; hence, using the notation given in Figure 2.3:

$$\begin{aligned}\vec{d}_2 &= \vec{d}_1 + \vec{\delta d} \\ \text{if } |\delta d| &\ll |d_1| \text{ then } \vec{d}_2 \approx \vec{d}_1\end{aligned}$$

And since, for a generic vortical disturbance with strength  $\Gamma$  at a location  $\vec{d}$  relative to the point

where the velocity is resolved:

$$\begin{bmatrix} u \\ v \end{bmatrix} = \frac{\Gamma}{2\pi|d|^2} \begin{bmatrix} d_y \\ -d_x \end{bmatrix}$$

It follows that the velocities experienced by each airfoil due to the far disturbance will be approximately the same, hence the effect of the far disturbance is to change the free-stream which the airfoils experience.

Combining the observations in the prior paragraph, the free-stream direction is set as  $\alpha_\infty = \alpha_1$  and the offset angle is set as  $\Delta\theta_{eff} = \alpha_2 - \alpha_1$ . While it is important to stress that there is limited rigor to this approximation, it will be shown in Chapter 6 that this approximation allows for accurate results to be obtained.

### 2.3.2 Force Transformations

The primary focus of this chapter has so far been on obtaining the parameters of the 2D airfoil system from the 3D rotor system. Relating the 2D force coefficients back to the 3D rotor system is a simpler exercise, but some consideration of the reference frames involved is required.

In particular, note that:

- 2D CFD results will give  $C_l$  and  $C_d$  in the free-stream ( $\alpha_\infty$ ) wind axes
- The 3D rotor model will give  $C_l$  and  $C_d$  in the local ( $\alpha_{1,2}$ ) wind axes
- For rotor thrust and power, and potentially force coupling, it may be necessary to specify the horizontal ( $C_x$ ) and vertical ( $C_z$ ) components of the force in the rotating hub reference frame

Moreover, although  $\Delta\theta_{eff}$  changes the geometric angle of the lower airfoil in the *modeled* airfoil system, it is important to remember that the angle of the airfoil in the *physical* system has not changed.

Taking these facts into account, it is convenient to first define an airfoil-fixed coordinate system, with  $\hat{x}$  along the airfoil's chord and  $\hat{z}$  upwards. For each airfoil, the transformation from the free-stream wind axes to the airfoil axes is:

$$\begin{aligned} \begin{bmatrix} C_x \\ C_z \end{bmatrix}_1^{(\text{airfoil})} &= \begin{bmatrix} \cos(\alpha_\infty) & \sin(\alpha_\infty) \\ -\sin(\alpha_\infty) & \cos(\alpha_\infty) \end{bmatrix} \begin{bmatrix} C_d \\ C_l \end{bmatrix}_1^{(\alpha_\infty)} \\ \begin{bmatrix} C_x \\ C_z \end{bmatrix}_2^{(\text{airfoil})} &= \begin{bmatrix} \cos(\alpha_\infty + \Delta\theta_{eff}) & \sin(\alpha_\infty + \Delta\theta_{eff}) \\ -\sin(\alpha_\infty + \Delta\theta_{eff}) & \cos(\alpha_\infty + \Delta\theta_{eff}) \end{bmatrix} \begin{bmatrix} C_d \\ C_l \end{bmatrix}_2^{(\alpha_\infty)} \end{aligned} \quad (2.4)$$

Then for both airfoils, the transformation from the airfoil axes to the hub axes is:

$$\begin{bmatrix} C_x \\ C_z \end{bmatrix}_{1,2}^{(\text{hub})} = \begin{bmatrix} \cos(\theta_{1,2}) & -\sin(\theta_{1,2}) \\ \sin(\theta_{1,2}) & \cos(\theta_{1,2}) \end{bmatrix} \begin{bmatrix} C_x \\ C_z \end{bmatrix}_{1,2}^{(\text{airfoil})} \quad (2.5)$$

Where  $\theta$  is particular for each airfoil, and may vary along the span due to blade twist. Any other necessary rotations may be determined trivially from these reference axes.

## 2.4 Summary

In this chapter, a methodology has been shown to relate the conditions (geometric and aerodynamic) of the rotor-blade system to those of the corresponding airfoil system, and likewise

to relate the airfoil forces in the free-stream frame to the forces in the blade and hub frames. Since it is necessary for the blade-to-airfoil methodology to account for what aerodynamics are present in the rotor model, two variations have been presented, depending upon whether the rotor model does or does not include bound circulation interaction. The methodology in either case is straightforward, but relies on simplifying assumptions to account for the differences between what parameters are available in each model. The airfoil-to-blade methodology is a simple coordinate transformation of the forces, but it is necessary to understand which reference frames are preserved in order to correctly perform the transformation.

While this chapter has discussed the relationships between the parameters and forces in different systems, it has not been concerned with how the parameters and forces are actually obtained. Determination of the forces on the 2D airfoil system will be discussed in the next Chapter. The geometric and aerodynamic parameters of the rotor-blade system are determined from the lower-fidelity rotor model, which will be discussed in [Chapter 5](#).

## Chapter 3: 2D Methodology

As has been discussed in Chapter 1.3, lower-fidelity rotor models typically rely on 2D *airfoil tables* for isolated airfoils. These airfoil tables give the pre-computed force and moment coefficients on the airfoil at a variety of angles of attack and Mach numbers. Canonical two-parameter interpolation schemes are then used to determine the force and moment coefficients based on the tabulated data for angles of attack and Mach numbers between the given values. Various canonical formats (such as C81) exist to express these airfoil tables.

The goal of this chapter is to present a methodology analogous to 2D airfoil table interpolation to determine the force and moment coefficients for isolated airfoils, but relevant to the problem of the two-airfoil systems found in stacked rotors. This methodology will necessarily have two significant parts: data *generation* and data *interpolation*. Data generation is the process of pre-computing the loads on the two-airfoil system for some configurations and aerodynamic conditions, which will be performed using 2D CFD. A 2D lifting-vortex model of the airfoils is also required, but will only be executed at runtime (i.e. the lifting-vortex model will not contribute to the pre-computed airfoil data). Data interpolation for the two-airfoil system requires interpolating across five parameters, and a method for doing so using previously-developed radial basis functions will be presented.

### 3.1 2D Lifting-Vortex Model

A lifting-vortex model of the airfoil system is a 2D analogue to the 3D lifting-line model. The lifting-vortex model is able to capture a significant portion of the mutual interaction between the airfoils, particularly where the gap between the airfoils exceeds the chord length<sup>1</sup>. Where the airfoils are spaced more closely, only high-fidelity 2D CFD is able to accurately capture their interactions. However, the 2D lifting-vortex model does have utility for:

- Determining the equivalent free-stream conditions for the airfoil system if the 3D inflow model includes bound circulation interaction (see Algorithm 2.1)
- Adding in the effects of bound circulation interaction if not included in the 3D inflow model, without the need for high-fidelity 2D CFD

The model itself is based on potential flow theory. Each airfoil is represented as a simple<sup>2</sup> potential flow vortex located at the airfoil's quarter-chord. These vortices are superimposed on a uniform flow, which represents the free-stream (see Chapter 2.3). The angle of attack on each airfoil is resolved at the airfoil's quarter-chord<sup>3</sup>. Pre-computed airfoil tables for the isolated airfoil(s) are used to relate the angle of attack to the lift (coefficient) generated by the airfoil.

The Kutta-Joukowski theorem states that the lift per unit span is related to the bound circulation by:

$$L' = \rho u \Gamma \quad (3.1)$$

---

<sup>1</sup>This measurement is useful for scale, but there is no particular significance to a 1:1 ratio of chord length to gap.

<sup>2</sup>While it would be possible to use a vortex core model, physical constraints on the system (IE finite airfoil thickness and chord) prevent the airfoils from being sufficiently close for a vortex core model to be necessary. It is however assumed that the velocity induced at the exact center of the vortex is zero.

<sup>3</sup>IE the center of its representative vortex

And by definition, the lift per unit span is related to the lift coefficient by:

$$L' = \frac{1}{2} \rho u^2 c C_l \quad (3.2)$$

Combining Equations 3.1 and 3.2, the lift coefficient is related to the bound circulation by:

$$\Gamma = \frac{1}{2} u c C_l \quad (3.3)$$

Since the lift coefficient is typically given in airfoil tables as a function of Mach ( $M$ )<sup>4</sup> rather than the dimensional flow speed ( $u$ ) and normalized per unit chord, it is useful to define a similarly normalized bound circulation from Equation 3.3, where:

$$\Gamma' \equiv \frac{\Gamma}{ac} = \frac{1}{2} M C_l \quad (3.4)$$

Next, consider a potential vortex with strength  $\Gamma$  located at a point<sup>5</sup>  $[x_v, y_v]$ . At a point in the flow with coordinates  $[x, y]$ , the velocity due to the vortex is given by:

$$\vec{u}^{(v)} = \begin{bmatrix} u_x \\ u_y \end{bmatrix}^{(v)} = \frac{\Gamma}{2\pi r^2} \begin{bmatrix} y - y_v \\ x_v - x \end{bmatrix} \quad (3.5)$$

Where  $r = \sqrt{(y - y_v)^2 + (x_v - x)^2}$  is the straight-line distance between the vortex and the point.

Equation 3.5 may be normalized by dividing both sides by  $a$ , and factoring  $c$  from  $\Gamma$ ,  $x$ ,  $y$ , and  $r$

---

<sup>4</sup>And angle of attack ( $\alpha$ )

<sup>5</sup>In Cartesian coordinates used here,  $x$  is positive right and  $y$  is positive up. This is in contrast to the notation defined in Section 2.2 where  $\Delta x$  is positive right but  $\Delta z$  is positive down. It is therefore important to note that  $y = -\Delta z$  when considering the system of two airfoils.

to obtain:

$$\begin{bmatrix} M_x \\ M_y \end{bmatrix}^{(v)} = \frac{\Gamma'}{2\pi \left(\frac{r}{c}\right)^2} \begin{bmatrix} \frac{y}{c} - \frac{y_v}{c} \\ \frac{x_v}{c} - \frac{x}{c} \end{bmatrix} \quad (3.6)$$

And note that the notation  $x/c$ ,  $y/c$ , and  $r/c$  is consistent with airfoil data referenced to a unit chord length.

In a system of two airfoils represented as lifting vortices in a freestream, each airfoil will influence the direction of the local flow around the other. Since for real airfoils there may be an arbitrary relationship between flow speed/direction and lift<sup>6</sup>, it is necessary to solve Equations 3.4 and 3.6 iteratively. A simple under-relaxed fixed-point iteration to achieve this is shown in Algorithm 3.1, which invokes Algorithm 3.2 as a convenient implementation Equation 3.6. Note that Algorithm 3.1 assumes that the geometric inputs  $\Delta x$  and  $\Delta z$  are given normalized by the airfoil chord length.

### 3.1.1 Periodicity

Thus far in this chapter, the 2D two-airfoil system has been considered in isolation. However, it is necessary to briefly consider the geometry of the stacked rotor system as a whole, in particular the meaning of the phase angle. Ignoring bending, the horizontal offset in the hub frame between the pair of closest airfoils is  $\Delta x_{hp} = r\Delta\psi$ ; hence as the magnitude of the index angle  $\Delta\psi$  increases, the airfoils move further apart. It follows from Equation 3.6 that as the airfoils move further apart, their influence on each other will decrease asymptotically towards zero. However, as the closer pair of blades moves further apart ( $\Delta\psi$  increases), the blades of one rotor

---

<sup>6</sup>The typical approximation is  $C_l(\alpha) = C_{l,\alpha} \cdot \alpha + C_{l,0}$ , but this does not take into account the effect of flow speed, is not true near and post-stall, and may not be true pre-stall particularly at low Reynolds numbers.

---

**Algorithm 3.1:** Computation of local flow conditions for two mutually-interacting airfoils represented as 2D lifting vortices

---

```

Function vortexAirfoilSystem( $AF_1, AF_2, \Delta x, \Delta z, \Delta \theta, \alpha_\infty, M_\infty$ )
  Input : Airfoil data tables  $AF$  for each airfoil, containing lookup for  $C_l(\alpha, M)$ 
  Input : Geometry of two-airfoil system  $\Delta x, \Delta z, \Delta \theta$ 
  Input : Free-stream flow conditions  $\alpha_\infty, M_\infty$ 
  tol =  $\delta$ ;
   $\omega = 0.5$ ;                                /* Relaxation parameter */
   $\alpha_1 = \alpha_\infty$ ;                      /* From free-stream conditions */
   $\alpha_2 = \alpha_\infty + \Delta \theta$ ;
   $M_{\infty,1}^{(u)} = M \cos \alpha_1$ ;          /* Free-stream flow components */
   $M_{\infty,1}^{(v)} = M \sin \alpha_1$ ;
   $M_{\infty,2}^{(u)} = M \cos \alpha_2$ ;
   $M_{\infty,2}^{(v)} = M \sin \alpha_2$ ;
   $M_1 = M$ ;
   $M_2 = M$ ;
  for a number of iterations do
     $\alpha_1^{(0)} = \alpha_1$ ;                    /* Store for convergence check */
     $\alpha_2^{(0)} = \alpha_2$ ;
     $C_{l,1} = AF_1.C_l(\alpha_1, M_1)$ ;          /* Update lift coefficients */
     $C_{l,2} = AF_2.C_l(\alpha_2, M_2)$ ;
     $\Gamma'_1 = 0.5M_1C_{l,1}$ ;                  /* Update vortex strengths */
     $\Gamma'_2 = 0.5M_2C_{l,2}$ ;
     $\Delta M_1^{(u)}, \Delta M_1^{(v)} = \text{vort2d}(\Gamma'_2, \Delta x, \Delta z, 0, 0)$ ;          /* AF 2 on AF 1 */
     $M_1^{(u)} = M_{\infty,1}^{(u)} + \Delta M_1^{(u)}$ ;
     $M_1^{(v)} = M_{\infty,1}^{(v)} + \Delta M_1^{(v)}$ ;
     $M_1 = \sqrt{(M_1^{(u)})^2 + (M_1^{(v)})^2}$ ;
     $\alpha_1 = \omega \tan^{-1}(M_1^{(v)}/M_1^{(u)}) + (1 - \omega)\alpha_1$ ;
     $\Delta M_2^{(u)}, \Delta M_2^{(v)} = \text{vort2d}(\Gamma'_1, 0, 0, \Delta x, \Delta z)$ ;          /* AF 1 on AF 2 */
     $M_2^{(u)} = M_{\infty,2}^{(u)} + \Delta M_2^{(u)}$ ;
     $M_2^{(v)} = M_{\infty,2}^{(v)} + \Delta M_2^{(v)}$ ;
     $M_2 = \sqrt{(M_2^{(u)})^2 + (M_2^{(v)})^2}$ ;
     $\alpha_2 = \omega \tan^{-1}(M_2^{(v)}/M_2^{(u)}) + (1 - \omega)\alpha_2$ ;
    if  $|\alpha_1 - \alpha_1^{(0)}| < \text{tol}$  and  $|\alpha_2 - \alpha_2^{(0)}| < \text{tol}$  then
      | break;                                /* Solution converged */
    end
  end
  return  $\alpha_1, \alpha_2, M_1, M_2$ ;
end

```

---

---

**Algorithm 3.2:** Implementation of Equation 3.6 used by Algorithm 3.1

---

```
Function vort2d( $\Gamma, x_v, y_v, x_p, y_p$ )  
  Input : Vortex strength  $\Gamma$   
  Input : Vortex location  $[x_v, y_v]$   
  Input : Measurement point  $[x_p, y_p]$   
   $d_x = x_p - x_v$ ; /* Offset between vortex and measurement */  
   $d_y = y_p - y_v$ ;  
   $r^2 = d_x^2 + d_y^2$ ;  
   $u = \Gamma d_y / (2\pi r^2)$ ; /* Velocity components */  
   $v = -\Gamma d_x / (2\pi r^2)$ ;  
  return  $u, v$ ;  
end
```

---

move closer to the next adjacent blades on the other rotor ( $\Delta\psi' = \Delta\psi - \frac{2\pi}{N_b}$  decreases), as shown in Figure 3.1. For individual rotors with  $N_b$  blades each, the distance to the next blade<sup>7</sup> is given by:

$$\Delta x'_{hp} = r\Delta\psi' = r\Delta\psi \pm \frac{2\pi r}{N_b} \quad (3.7)$$

In practice, even where the blades are evenly-spaced<sup>8</sup> the interaction between the airfoils of both blades cannot be assumed to be near-zero, nor can it be assumed that this interaction is limited to the inboard sections of the blade. Thus, some modification to Algorithm 3.1 is necessary to account for the presence of the “farther” blades’ airfoils.

Modification of the algorithm entails the placement of two “image vortices” in the airfoil system, as shown in Figure 3.2. Each image vortex represents the airfoil on the next adjacent blade, but because the flow is radially symmetric it is assumed that all blades have the same lift, hence the strength of the image vortex can be set as the strength of the corresponding “physical” vortex in the two-airfoil system. Using hub-frame coordinates (see Fig. 2.1b), each image vortex is placed at a location corresponding to an angle  $\frac{2\pi}{N_b}$  from its physical vortex and opposite the

---

<sup>7</sup>Again ignoring bending.

<sup>8</sup> $\Delta\psi = \Delta\psi' = \pm\pi/N_b$




Figure 3.1: Phase angles between the closer blade pair ( $\Delta\psi$ ) and the farther blade pair ( $\Delta\psi'$ )

physical vortex from the other airfoil, thus<sup>9</sup>:

$$\begin{aligned} p^{(\Gamma_1)} &= \left( \frac{2\pi r}{N_b} \cdot \text{sgn}(\Delta\psi), 0 \right) \\ p^{(\Gamma_2)} &= \left( \left[ \frac{2\pi r}{N_b} - r\Delta\psi \right] \cdot \text{sgn}(\Delta\psi), \Delta z_{hp} \right) \end{aligned} \quad (3.8)$$

The hub-frame coordinates given by Equation 3.8 may then be transformed to airfoil coordinates by Equation 2.1—hence several rotor-related parameters are required to determine the image vortex positions, but these are known at runtime. The velocity induced by each image vortex (Eq. 3.6) is applied *only* to the adjacent airfoil on the “other” rotor, *not* to the following airfoil on the “same” rotor<sup>10</sup>. Once the vortex locations are determined, Algorithm 3.1 is modified by introducing two additional calls to Algorithm 3.2 with appropriate vortex strengths ( $\Gamma'_1, \Gamma'_2$ ) and

<sup>9</sup>Note that if  $\Delta\psi = 0$ , then  $\text{sgn}(\Delta\psi) = 1$  for purposes of vortex placement.

<sup>10</sup>For a two-bladed rotor this is physically reasonable, since the blades on the same rotor are 180° apart, hence the influence of one blade on the other should be zero. Stacked rotors with more than two blades per rotor have not been investigated in this work, but it is conceivable that more pairs of image vortices could be added as necessary.




Figure 3.2: Strength and influence of the image vortices used to account for periodicity

locations (Eq. 3.8, 2.1), and the resulting induced velocities are added to  $\Delta M$  as appropriate.

An example of the effect of correcting for periodicity using image vortices is shown in Figure 3.3. At this notional point along the span of a rotor, the blade airfoils on the *same* rotor are separated by 16 chords ( $\frac{2\pi r}{N_b c} = 16$ ), hence a separation of  $\Delta x/c = 8$  implies equal spacing of the blades. The free-stream angle of attack is  $4^\circ$ , and since  $c_l$  is given as a function of  $\alpha$  from the airfoil tables, changes in local  $\alpha$  correspond directly to changes in lift. Without the image vortices (red lines), as the airfoils move further apart their influence on each other diminishes, but there is still a significant change of nearly  $\pm 0.5^\circ$  from freestream when the airfoils are evenly-spaced. Moreover, the airfoils show different influence on each other at  $\Delta x/c = 8$  and  $\Delta x/c = -8$ , despite these representing physically identical conditions on the rotor. With the image vortices however (blue lines), the airfoils experience identical conditions at  $\Delta x/c = \pm 8$ , which also match the free-stream due to the influence of the physical and image vortices being equal and opposite.

Thus far the discussion of periodicity has focused only on the 2D lifting vortex model, which is intended as a supplement to—not alternative for—the 2D CFD. Is it necessary to account for periodicity in the 2D CFD simulations, requiring four (or more) airfoils and rotor-dependent parameters? An examination of Figure 3.3 shows that where the blades are close (small  $\Delta x/c$ ), the dominant influence is that of the adjacent blades, and the image vortices have relatively little




Figure 3.3: Effective angles of attack on nominal airfoils with and without accounting for periodicity

effect. Moreover, it will be shown in Chapter 4.1 that 2D CFD is of little benefit compared to the 2D lifting vortex model where the airfoils are spaced more than a few chords apart. Therefore, the effects of periodicity are safely ignored when performing the 2D CFD simulations.

### 3.2 2D CFD for Data Generation

With the focus on modularity in this work, it must be emphasized that any reasonable high-fidelity<sup>11</sup> 2D CFD code could be used to simulate the flow around the two-airfoil system. Different codes—and the various solution methods which they employ—will have their own benefits and drawbacks, and a significant benefit which should not be overlooked is familiarity to

<sup>11</sup>Navier-Stokes, including turbulence and transition models

the researcher. An important consideration however is that thousands of individual simulations must be performed, hence it is important that the mesh generation and flow solution processes be automated, as without automation case setup and execution would require an inordinate amount of time and effort.


In particular then, this study uses a 2D version of the Hamiltonian-strand flow solver (HAMSTR, HAM2D) developed at the University of Maryland [27,28], with an automated meshing methodology (AutoMesh) developed previously by the current author [29]. With HAM2D and AutoMesh, a point-continuous mesh (which avoids the interpolation errors inherent to overset meshes) is generated around the airfoil(s), then the flow is solved using a “structured-like” algorithm. This results in a process which is highly accurate—avoiding the interpolation errors inherent with overset grids or those with hanging nodes—while also nearly as fast as more conventional structured flow solvers. Both AutoMesh and HAM2D have been previously validated, but the basic methodology of each is described in this sections.

### 3.2.1 AutoMesh


The automated meshing methodology is a significant enabler of the automated airfoil data generation as a whole. All geometric configurations (unique combinations of  $\Delta x$ ,  $\Delta z$ , and  $\Delta \theta$ ) require their own mesh to be generated, and building a data set sufficient for rotor parameter variation typically takes several hundred airfoil configurations. Hence, it is necessary to have a methodology which can be trusted to consistently produce sufficiently good-quality meshes without needing much or any input on the part of the user. AutoMesh was developed with the specific intent of offering (as much as possible) a “hands-off” experience for the user, taking




(a) Geometrically-determined distribution of the points along the airfoil surface




(b) Near-body (boundary layer) structured mesh



(c) Removal of overlapping regions for multiple airfoil elements



(d) Hybrid mesh, with unstructured triangles for background



(e) All-quadrilateral unstructured mesh after quadrangulation

Figure 3.4: Overview of the steps of the automated meshing procedure

airfoil coordinates in and returning a mesh out, and hence is well-suited for the task.

AutoMesh is comprised of several different individual routines which are scripted together (using Python) to produce a complete mesh. An overview of the output of each routine is shown in Figure 3.4. This section will discuss the routines, highlighting some of their important concepts. While AutoMesh can handle airfoils with an arbitrary number of elements, the surface point allocation and near-body mesh generation are performed separately for each element, so discussion of them will refer to a single airfoil. All other routines deal with the full system of airfoils at once.

### 3.2.1.1 Surface Point Allocation

The airfoil coordinates provided by the user to AutoMesh must have sufficient resolution to capture the *geometric* features of the airfoil, but they need not be distributed along the airfoil's surface in any specific way<sup>12</sup>. Rather, the first step of AutoMesh is to redistribute points along the surface of the airfoil element in a way which achieves the following objectives (Fig. 3.4a):

- Maintain a small maximum angle between adjacent points
- Maintain a certain maximum distance between adjacent points
- Cluster near points with discontinuous slope (i.e. sharp points)

The first objective ensures that points will cluster around areas of high curvature (e.g. the leading edge), where the flow will be turning rapidly and there may be relatively rapid changes in flow speed or direction. The second objective ensures that sufficient resolution is maintained in low-

---

<sup>12</sup>The user-supplied airfoil geometry may even be open or closed; if open then AutoMesh will draw a straight line between the endpoints to close it.

curvature regions. The third objective ensures that points are clustered in areas (e.g. near the trailing edge) where flow separation is likely to occur.

Within the code, the user-supplied coordinates are parameterized in terms of the linear arc length  $s$  along the airfoil surface from the starting point. The coordinates  $x, y$  and all derived properties are assumed to be strictly functions of  $s$ . Also computed at each point are the turning angle  $\theta_i$  between the incoming and outgoing surfaces, and the geometric curvature  $\kappa_i$ :

$$\begin{aligned}\theta_i &= \tan^{-1} \left( \frac{y_{i+1} - y_i}{x_{i+1} - x_i} \right) - \tan^{-1} \left( \frac{y_i - y_{i-1}}{x_i - x_{i-1}} \right) \\ \kappa_i &= \frac{y''x' - x''y'}{(x'^2 + y'^2)^{3/2}}\end{aligned}\tag{3.9}$$

Where the derivatives of  $x$  and  $y$  are taken with respect to  $s$ . Likewise, both  $\theta_i$  and  $\kappa_i$  are considered to be functions of  $s$ . Individual points with significantly larger turning angles than their adjacent points<sup>13</sup> are defined as sharp points.

The first point, last point, and any sharp points along the airfoil surface are considered to be fixed points. Between fixed points, regions of high- and low-curvature are determined according to a program-default or used-supplied cutoff value of  $\kappa$ . With the regions defined, the code determines the coarsest number of points which could be assigned to each region, while maintaining the requirements for a maximum angle (high-curvature) or distance (low-curvature) between points. The user may define any number of points to be distributed along the surface, so each region is then allocated a total number of points proportional to that of its coarsest distribution.

Each region of the airfoil now has a *number* of points assigned to it, but the *locations* of those points along the surface ( $s$ ) must still be determined. How the locations are determined

---

<sup>13</sup>Comparing the turning angle between adjacent points helps to distinguish sharp-but-shallow angles from coarsely-defined but smooth curves.

depends upon the type of region:

- In low-curvature regions, assign locations to maintain constant intervals of  $s$
- In high-curvature regions, assign locations to maintain a constant turning angle

The first computation is trivial. The second computation requires calculating the rolling integral of curvature over arclength, i.e.:

$$f(s) = \frac{1}{2} \sum_{s_0}^s (\kappa_i + \kappa_{i-1}) (s_i - s_{i-1}) \quad (3.10)$$

Then,  $s_i$  are chosen such that  $f(s_i) - f(s_{i-1})$  is constant for all  $i$ . While this distribution of points would satisfy the need to maintain the maximum angle and distance criteria, the distribution near sharp points must be addressed. This is done by using hyperbolic tangent stretching, such that the spacing at the sharp point is reduced to the intended target value, while ensuring a smooth distribution of points across the airfoil surface.

With the locations of the points along the surface determined, the last step in the surface point allocation is to reconstruct the actual airfoil surface. The reconstruction is preformed using monotone cubic spline interpolation on  $x(s)$  and  $y(s)$ . To ensure correct reconstruction near corners, these splines only extend between (i.e. do not cross) fixed points.

### 3.2.1.2 Near-Body Mesh Generation

The near-body mesh is a structured mesh which extends a sufficient distance out from the surface of the airfoil to capture the airfoil's boundary layer (Fig. 3.4b). A strand/advancing front approach is used, and the code has common routines (with slight modifications where appropriate,




Figure 3.5: Adjustments for a wake surface or C-mesh using an O-mesh methodology

as shown in Figure 3.5) to construct either an O-mesh, C-mesh, or O-mesh with C-mesh-type wake strand.

The thickness of the first layer of the mesh may be defined manually, or computed based on flat-plate boundary layer theory as:

$$\frac{\Delta y}{c} = \frac{L}{c} y^+ \left( \frac{L}{c} Re_L \right)^{13/14} \quad (3.11)$$

Where a value of  $y^+ = 1$  is desired in the final mesh to ensure accuracy without an excessive number of points. Similarly, the target near-body mesh cutoff distance may be specified by the user, or computed based on flat-plate theory as:

$$\frac{\delta}{c} = 0.37 \frac{L}{c} \left( \frac{L}{c} Re_L \right)^{-(1/5)} \quad (3.12)$$

Since a real airfoil will have a thicker boundary layer than a flat plate, a cutoff distance of  $2\delta$  is generally used, but larger values may be necessary to fully capture the boundary layer at higher angles of attack.




Figure 3.6: Strand gathering and mesh unfolding, similarly to CLOVIS

After the first layer, the next few layers have equal spacing, then the spacing is allowed to grow according to a default or user-defined stretching ratio. Within the equally-spaced region, the strands are held exactly orthogonal to the surface, to ensure that the cells have as little skew as possible. Outside of this region, the direction of the strands may be adjusted such that the wraparound-direction mesh spacing will become more regular as the mesh grows.

Some consideration of how to handle concave regions is also necessary, since in these regions the strands of the mesh may intersect. The technique used to account for this is based on CLOVIS [30] with the addition of “unfolding”, as shown in Figure 3.6. As with CLOVIS, if any strands would intersect (Fig. 3.6a) then the surface is extended by a distance greater than the intended layer thickness and the intersecting strands are gathered to a point. The unfolding algorithm adds some extra distance such that the resulting surface is less concave than the previous layer (Fig. 3.6b). The surface is then trimmed by a fixed amount over all strands, resulting in a valid (non-intersecting) layer with reduced concavity (Fig. 3.6c). Finally, some smoothing is performed, and this process is repeated until the structured mesh is at least as large as the specified cutoff distance.




Figure 3.7: Overlap removal using point-in-polygon tests on the nodes of the outer boundary

### 3.2.1.3 Overlap Removal

Thus far, the mesh generation procedure has only been concerned with the individual airfoil elements. However, if the airfoil elements are closely spaced (as in Fig. 3.4c) then it is possible that the boundary layer meshes may overlap. Since the goal of the meshing methodology is to achieve a point-continuous mesh, it is necessary for this overlap to be removed.

Initially a scheme was tried which used an overset mesh tool to detect the overlapping cells [31], but this required additional software, and was found to have difficulty in some marginal cases. Instead, the current work uses simple point-in-polygon tests to identify the nodes on the outer boundary of one mesh which exist inside the boundary of another (Fig. 3.7a). All of these nodes are first tagged, then removed from the mesh (along with their corresponding cells). The boundary is updated (Fig. 3.7b), and if necessary the process is repeated until no more nodes exist inside another mesh's boundary. One limitation of this method is that it cannot handle a case in which one mesh exists *entirely* within the boundary of another mesh, but this is a non-issue for any practical airfoil case.

#### 3.2.1.4 Off-Body Mesh Generation

Next, it is necessary to construct a background mesh which both fills the gaps between the airfoil element meshes, and extends to the far-field boundary. The process for generating the background mesh begins with Delaundo, an open-source unstructured-triangle grid generator using a frontal Delaunay method [32, 33]. The outer boundaries of the airfoil element near-body meshes and the intended far-field boundary are passed to Delaundo, which fills the intermediate space. The resulting mesh is a hybrid, consisting of both quadrilaterals and triangles, unstructured but contains significant *regions* of structure, and point-continuous by virtue of Delaundo connecting to the nodes of the near-body mesh boundaries.

This hybrid mesh must then be transformed into an all-quadrilateral mesh. One simple way to achieve this would be to use Catmull-Clark subdivision, which transforms every initial triangle into three quadrilaterals and every initial quadrilateral into four quadrilaterals, with one additional node per each edge and each cell in the initial mesh. However, a major drawback of using Catmull-Clark subdivision directly is that subdividing triangles necessarily produces quadrilaterals with uneven angles (i.e. poor quality quadrilaterals), even with a high-quality triangular mesh.

To produce an all-quadrilateral mesh with better quality than can be achieved through subdivision alone, the HAMSTRAN method [34] is used. HAMSTRAN operates on the triangles in the mesh, but does not affect the quadrilaterals in the initial mesh. As part of the HAMSTRAN method all of the boundaries of the triangular region are subdivided, including the structured/unstructured (i.e. triangle/quadrilateral) boundary. To avoid the need for templates to subdivide only the boundary of the quadrilateral region, all of the previous operations are conducted




Figure 3.8: Bounds on the Laplacian smoothing from cells adjacent to a node

on a coarsened mesh. Once these operations have been performed, the intended mesh is injected and its boundary is stitched to the boundary subdivided by HAMSTRAN. To achieve a greater fineness in the final mesh the HAMSTRAN mesh was further subdivided, using Catmull-Clark subdivision. Since the mesh produced by HAMSTRAN is composed entirely of quadrilaterals of reasonable quality, the quality of the subdivided mesh is preserved. Additionally, the final mesh has somewhat fewer points and quadrilaterals compared to those quadrangulated by subdivision alone, reducing the computational requirements of the flow solver.

### 3.2.1.5 Mesh Smoothing

When constructing the background mesh, Delaundo attempts to generate triangles as close to isomorphic as possible. However, the near-body mesh generator has no strict requirements to achieve all isomorphic quadrilaterals and, in fact, generates high-aspect ratio quadrilaterals both near the surface of the airfoil element and in concave regions. As a result, the combined triangle/quadrilateral mesh may exhibit large changes in cell size and/or aspect ratio near the

transition region. (This problem is particularly acute when the meshes have been trimmed.) As a result, some degree of mesh smoothing is required. This smoothing is performed after all the previous steps have been completed.

It was found that an ordinary Laplacian smoothing algorithm could result in cells becoming self-intersecting in highly skewed regions, and also could not prevent cells from becoming concave. Additionally, applying smoothing naively could result in unintended changes in the cell sizes and growth ratio in the (structured) boundary layer mesh region. To resolve the former issue, limits were applied to define a region in which a given point must exist to produce valid cells (Fig. 3.8). These limits are calculated from the other points belonging to each cell to which the point belongs. If the point were to cross the diagonal between adjacent nodes on a cell, or the line formed by either adjacent node and the opposite node, then the cell would become concave. Crossing any more lines would result in a negative cell volume (Fig. 3.8a). Thus, any given node must remain outside the triangle defined by the remaining three nodes as well as the extension of that region away from the node, for all of the cells containing that node (Fig. 3.8b). Enforcing this limit is a computationally efficient calculation, which relies on only the four basic mathematical operations, applied to each set of points to be tested. To resolve the latter issue, smoothing strength was varied based on both the distance from the airfoil surface, and the area jump at the nodes. The resulting mesh shows improved smoothness of the cell size transition.

### 3.2.1.6 Mesh Quality Metrics

Several quality metrics were implemented to judge the quality of the resulting meshes. The two which were found to be most useful (which have already been mentioned in passing) were

the nodal area ratio and the cell skew:

$$\begin{aligned}
(\text{area ratio})_{\text{node}} &= \frac{\max(\text{area of cells containing node})}{\min(\text{area of cells containing node})} \in [1, \infty) \\
(\text{skew})_{\text{cell}} &= \frac{1}{90^\circ} \max(\theta_{\max} - 90^\circ, 90^\circ - \theta_{\min}) \in [0, 1]
\end{aligned} \tag{3.13}$$

Where in this case  $\theta$  are the interior angles of the cell. For both metrics, a smaller value implies better quality and a larger value implies worse quality. Some amount of skew however is inherent where the structured mesh turns rapidly, for instance near sharp corners, and some amount of skew is also inherent to an unstructured mesh where the cells grow more rapidly than those of a structured mesh. Likewise, the nodal area ratio may be thought of as an unstructured equivalent of the stretching ratio in the structured mesh, hence it is expected that the nodal area ratio will be greater than 1.0 in much of the mesh.

### 3.2.2 HAM2D

HAMSTR is a line-based Reynolds-averaged Navier-Stokes (RANS) unstructured flow solver, which was developed at the University of Maryland. A 2D version (HAM2D) is used in this work, but a 3D version has also been developed. The key concept behind HAMSTR is the ability to find lines (Hamiltonian paths) in an unstructured grid by connecting opposite edges (or faces in 3D), as shown in Figure 3.9—this loop-finding imposes the requirement for an all-quadrilateral mesh. Along these Hamiltonian paths, the flow solution is computed similarly to a structured mesh. The underpinnings of HAMSTR are described in greater detail in other research [28, 35], but are summarized in this section.




Figure 3.9: Identification of the Hamiltonian paths (colored and dashed lines) through a representative unstructured mesh, excerpted from [28]

### 3.2.2.1 Governing Equations and Solution Scheme

The governing equation of HAM2D is the unsteady, compressible Navier-Stokes equation, which may be expressed in integral form as:

$$\frac{\partial}{\partial t} \int_V \mathbf{Q} dV + \int_s (\mathbf{F}_c - \mathbf{F}_v) \cdot ds = 0 \quad (3.14)$$

Where  $\mathbf{Q}$  is the vector of conserved variables:

$$\mathbf{Q} = [\rho, \rho u, \rho v, E]^T \quad (3.15)$$

And  $\mathbf{F}_c$   $\mathbf{F}_v$  are the convective and viscous fluxes, respectively:

$$\begin{aligned}\mathbf{F}_c &= [\rho \mathbf{V}, \rho u \mathbf{V} + n_x p, \rho v \mathbf{V} + n_y p, \rho H \mathbf{V}]^T \\ \mathbf{F}_{v_x}(\mathbf{Q}) &= [0, \tau_{xx}, \tau_{yx}, u\tau_{xx} + v\tau_{xy} - \bar{q}_x]^T \\ \mathbf{F}_{v_y}(\mathbf{Q}) &= [0, \tau_{xy}, \tau_{yy}, u\tau_{yz} + v\tau_{yy} - \bar{q}_y]^T\end{aligned}\tag{3.16}$$

Where  $\rho$  is the density,  $(u, v)$  are the velocity components,  $E$  is the total energy,  $H$  is the enthalpy,  $p$  is the pressure,  $\tau$  is the viscous stress tensor, and  $\bar{q}$  is the rate of thermal conduction of the flow.

Reynolds-averaging of the governing equations is used to avoid the need to directly capture turbulence in the flow. In the Reynolds-averaging approach, the dependent variables ( $f$ ) are broken down into time-averaged ( $\bar{f}$ ) over some time ( $\Delta t$ ) and fluctuating ( $f'$ ) components, hence:

$$u = \bar{u} + u', \quad v = \bar{v} + v', \quad \rho = \bar{\rho} + \rho', \quad p = \bar{p} + p' \tag{3.17}$$

Where, by definition:

$$\begin{aligned}\bar{f} &\equiv \frac{1}{\Delta t} \int_{t_0}^{t_0 + \Delta t} f dt \\ \overline{f'} &\equiv \frac{1}{\Delta t} \int_{t_0}^{t_0 + \Delta t} f' dt = 0, \text{ but } \overline{f'f'} \neq 0\end{aligned}\tag{3.18}$$

The decomposed variables (Eq. 3.17) are substituted into the governing equation (Eq. 3.14) which is then time-averaged, resulting in a new governing equation known as the Reynolds-Averaged Navier-Stokes (RANS) equation. This also introduces an additional term, which is




Figure 3.10: Strand directions and numbering scheme at a cell in a mesh, excerpted from [35]

known as the Reynolds stress tensor:

$$(\bar{\tau}_{ij})_{turb} = -\overline{\rho u'_i u'_j} \quad (3.19)$$

This additional term is approximated using a turbulence model, as will be discussed below.

A cell-centered finite volume method is used to discretize the governing RANS equation. As part of the discretization, the fluxes (Eq. 3.16) must be reconstructed at the cell boundaries (faces). This reconstruction is performed for the inviscid fluxes using a stencil-based scheme, with direction defined along the Hamiltonian paths, as shown in Figure 3.10. The particular stencil used in this work is the fifth-order weighted essentially-non-oscillatory (WENO) finite difference scheme:

$$q_{i+\frac{1}{2}}^L = \frac{\omega_1}{3} \bar{q}_{i-2} - \frac{1}{6} (7\omega_1 + \omega_2) \bar{q}_{i-1} + \frac{1}{6} (11\omega_1 + 5\omega_2 + 2\omega_3) \bar{q}_i$$

$$\begin{aligned}
& + \frac{1}{6} (2\omega_2 + 5\omega_3) \bar{q}_{i+1} - \frac{\omega_3}{6} \bar{q}_{i+2} \\
q_{i-\frac{1}{2}}^R &= \frac{\omega_1}{3} \bar{q}_{i+2} + \frac{1}{6} (7\omega_1 + \omega_2) \bar{q}_{i+1} + \frac{1}{6} (11\omega_1 + 5\omega_2 + 2\omega_3) \bar{q}_i \\
& + \frac{1}{6} (2\omega_2 + 5\omega_3) \bar{q}_{i-1} - \frac{\omega_3}{6} \bar{q}_{i-2}
\end{aligned} \tag{3.20}$$

Where  $\omega_{1,2,3}$  are weighting parameters. The inviscid flux is then reconstructed from the left and right states using Roe's flux difference splitting scheme. Velocity and temperature gradients are computed using second-order accurate methods, and those gradients are used to compute the viscous fluxes.

The discretized and linearized version of Equation 3.14 forming the system of equations which must be solved at each time step is:

$$\left[ \frac{I}{\Delta t} + \frac{\partial \mathbf{R}}{\partial \mathbf{Q}} \right] \Delta \mathbf{Q} = -\mathbf{R}(\mathbf{Q}^n) \tag{3.21}$$

Where  $\mathbf{R}(\mathbf{Q}^n)$  is the residual evaluated at time step  $n$ , i.e.:

$$\mathbf{R}(\mathbf{Q}^n) = - \sum_i \hat{\mathbf{F}}_i(\mathbf{Q}^n) \cdot ds_i \tag{3.22}$$

HAM2D includes several iterative schemes to solve Equation 3.21, including diagonally-dominant ADI (DDADI) and line-Gauss-Seidel (LGS) [36]. In this study however, the equation is solved for all cases using the generalized minimum residual (GMRES) method [37]. In particular, a right-preconditioned GMRES method is used, of the general form:

$$AM^{-1}(Mx) = B \tag{3.23}$$

Where from Equation 3.21  $x$  is the vector of variables  $\Delta Q$ ,  $A$  is the the left-hand side matrix,  $B$  is the right-hand side of the equation (i.e. the residual, as defined by Eq. 3.22), and the preconditioner  $M$  is an approximate factorization of the  $A$  matrix.

### 3.2.2.2 Turbulence and Transition Models

The goal of the turbulence model is to approximate the Reynolds stress tensor (Eq. 3.19) in the RANS equation. The basic turbulence model used in this study is the one-equation Spalart-Allmaras (SA) model:

$$\frac{D\tilde{\nu}}{Dt} = c_{b1}\tilde{S}\tilde{\nu} + \frac{1}{c_\sigma} [\nabla \cdot ((\nu + \tilde{\nu}) \nabla \tilde{\nu}) + c_{b2} (\nabla \tilde{\nu})^2] - c_{w1}f_w \left[ \frac{\tilde{\nu}}{\tilde{d}} \right]^2 \quad (3.24)$$

Where:

$$\nu_t = \tilde{\nu}f_{v1}, \text{ with } f_{v1} = \frac{\chi^3}{\chi^3 + c_{v1}^3} \text{ and } \chi = \frac{\tilde{\nu}}{\nu} \quad (3.25)$$

While the SA turbulence model is recognized for producing accurate results for high-Reynolds number external flows, it cannot predict laminar-turbulent boundary layer transition on its own. A separate transition model is required, which in this study is the  $\gamma - \overline{Re_{\theta t}} - SA$  model [28]. This transition model begins with the transport equation for the intermittency ( $\gamma$ ):

$$\frac{D(\rho\gamma)}{Dt} = P_\gamma - D_\gamma + \frac{\partial}{\partial x_j} \left[ (\mu + \mu_t) \frac{\partial \gamma}{\partial x_j} \right] \quad (3.26)$$

Computation of the terms  $P_\gamma$  and  $D_\gamma$  requires the momentum thickness Reynolds number ( $\overline{Re_{\theta t}}$ ),

the transport equation for which is:

$$\frac{D(\overline{\rho Re_{\theta t}})}{Dt} = P_{\theta t} + P_{CF} + \frac{\partial}{\partial x_j} \left[ 2(\mu + \mu_t) \frac{\partial \overline{Re_{\theta t}}}{\partial x_j} \right] \quad (3.27)$$

Equations 3.27 and 3.26 are solved separately using DDADI. Finally, the intermittency is used to modify the production term of the SA equation (Eq. 3.24) as:

$$\frac{D\tilde{\nu}}{Dt} = \gamma P_v - D_v + \frac{1}{\sigma} [\nabla \cdot ((\nu + \tilde{\nu}) \nabla \tilde{\nu}) + c_{b2} (\nabla \tilde{\nu})^2] \quad (3.28)$$

### 3.2.3 Parallelization and Computing Resources

A unique computational mesh must be generated for each airfoil geometry (or, in the case of the stacked rotor, for each set of geometric conditions), and HAM2D computes the flow for one set of aerodynamic (free-stream) conditions at a time. On modern hardware, mesh generation takes less than a minute, while computing the flow may take a few hours<sup>14</sup> per condition. While it would be feasible to run a few cases on a local workstation, airfoil table generation may require hundreds or even thousands of cases to be simulated. (The number of cases is discussed in more detail below in Section 3.3.2.) Hence, the need for parallel computing is significant.

In this study, almost all<sup>15</sup> 2D airfoil cases were run on the DoD HPCMP’s Centennial high-performance computer (HPC) [38]. Running on an HPC allowed for all of the desired cases to be run in parallel<sup>16</sup>. Hence, “production” groups of cases for the two-airfoil system typically took around 7 hours, independent of how many cases were included in the group.

---

<sup>14</sup>Depending upon factors including the complexity of the geometry and resulting flow field, solver options, desired convergence tolerance, etc.

<sup>15</sup>A small number of individual cases were run on a local workstation, for testing purposes.

<sup>16</sup>Within the limitations imposed on maximum cores- and jobs-per-user.

### 3.3 Data Interpolation

One of the key concepts of this methodology for lower-fidelity modeling of stacked rotors is the ability to perform all of the “expensive” high-fidelity (2D CFD) computations prior to and independently of the rotor simulations, similarly to how existing lower-fidelity models of isolated rotors use conventional airfoil tables. Since the particular geometric and aerodynamic conditions of the rotor will thus not be known at the time that the 2D CFD computations are performed, it is necessary to test a range of conditions and then to be able to interpolate that data onto the rotor when the conditions become known.

For an isolated rotor, the aforementioned is simple—lift, drag, and moment data is generated for a canonical range<sup>17</sup> of aerodynamic conditions ( $\alpha$  and  $M$ ), forming an airfoil table. Then simple two-dimensional interpolation schemes (bilinear, bicubic, etc.) may be used to interpolate the data, allowing a reasonably accurate estimation of lift, drag, and moment to be obtained for any arbitrary aerodynamic condition. A limitation of these simple methods however is that they often require the input data to have been generated on a rectilinear grid. Generating data over a wide range of cases—for instance, including high- $M$  low- $\alpha$  conditions near the blade tip, low- $M$  high- $\alpha$  conditions near the root, and negative- $\alpha$  low- $M$  conditions near the root on the retreating blade—may require generating data for some cases that are not just unnecessary, but unlikely to be physically possible—such as high- $M$  negative- $\alpha$ . Also, the data must exist and be valid at all sampled points—a single missing point would break the assumption of a rectilinear grid.

For the stacked rotor, a methodology which does not have the limitations of the canonical methods for an isolated rotor is not just desirable, but essential. As discussed earlier, it will be

---

<sup>17</sup>Or within a reasonable range expected to occur on the rotor

necessary to interpolate with five input variables—the aerodynamic conditions  $\alpha$  and  $M$  as with an isolated airfoil, but also the geometric conditions  $\Delta x$ ,  $\Delta z$ , and  $\Delta \theta$ . Certain ranges of those variables will be necessary to sample at higher resolution, but limiting the resolution where the airfoil forces do not change rapidly would significantly help to control computational costs. Also, with many thousands of individual cases it is likely that some simulations may fail, and it would be undesirable to need to repeat each case if they could instead be ignored. Both of these goals indicate the need for a mesh-free interpolation method.

Radial basis function (RBF) interpolation [39] is well-suited to this application, being mesh-free and compatible with  $N$ -dimensional data. RBF interpolation relies on *radial basis functions*, which are functions ( $\phi$ ) whose value depends only upon the magnitude of their argument, i.e.:

$$\phi(\|\mathbf{x}\|) = \phi(r), \text{ for } \mathbf{x} \in \mathbb{R}^N \quad (3.29)$$

Where  $r$  is the length (radius) of the  $N$ -dimensional vector  $\mathbf{x}$ .

Suppose that exact values of the data ( $d_i$ ) are known at some specific points ( $\mathbf{c}_i$  where  $i \in [1, n]$ ), and an interpolant ( $f(\mathbf{x})$ ) is sought which would provide the data at an arbitrary point ( $\mathbf{x}$ ). An RBF interpolant is comprised of a linear combination of RBFs and polynomials, i.e.:

$$f(\mathbf{x}) = \sum_{i=1}^n a_i \phi(\|\mathbf{x} - \mathbf{c}_i\|) + \sum_{j=1}^m b_j p_j(\mathbf{x}) \quad (3.30)$$

Where  $a_i$  and  $b_j$  are weighting coefficients, and  $p_j(\mathbf{x})$  is a polynomial of degree no greater than  $m$ . Observe as well that  $\mathbf{c}_i$  are the *centers* of the radial basis functions. The weighting coefficients are determined by applying the constraint that the interpolant must return exact values at the

sample points, i.e.:

$$f(\mathbf{c}_i) = d_i, \forall i \in [1, n] \quad (3.31)$$

And the constraint that:

$$\sum_{i=1}^n a_i p_j(\mathbf{c}_i) = 0, \forall j \in [1, m] \quad (3.32)$$

Applying the former constraint (Eq. 3.31) to Equation 3.30 leads to a system of equations:

$$\Phi \mathbf{a} + P \mathbf{b} = \mathbf{d} \quad (3.33)$$

Where  $\mathbf{a}$ ,  $\mathbf{b}$ , and  $\mathbf{d}$  are the column vectors of  $a_i$ ,  $b_i$ , and  $d_i$  respectively, and

$$\Phi = \begin{bmatrix} \phi(\|\mathbf{c}_1 - \mathbf{c}_1\|) & \phi(\|\mathbf{c}_2 - \mathbf{c}_1\|) & \cdots & \phi(\|\mathbf{c}_n - \mathbf{c}_1\|) \\ \phi(\|\mathbf{c}_1 - \mathbf{c}_2\|) & \phi(\|\mathbf{c}_2 - \mathbf{c}_2\|) & \cdots & \phi(\|\mathbf{c}_n - \mathbf{c}_2\|) \\ \vdots & \vdots & \ddots & \vdots \\ \phi(\|\mathbf{c}_1 - \mathbf{c}_n\|) & \phi(\|\mathbf{c}_2 - \mathbf{c}_n\|) & \cdots & \phi(\|\mathbf{c}_n - \mathbf{c}_n\|) \end{bmatrix}$$

$$P = \begin{bmatrix} p_1(\mathbf{c}_1) & \cdots & p_m(\mathbf{c}_1) \\ \vdots & \ddots & \vdots \\ p_1(\mathbf{c}_n) & \cdots & p_m(\mathbf{c}_n) \end{bmatrix} \quad (3.34)$$

Similarly, the second constraint (Eq. 3.32) results in a system of equations:

$$P^T \mathbf{a} = 0 \quad (3.35)$$

Hence, the weighting coefficients may be determined as the solution to the system of linear

equations (combining Eq. 3.33 and 3.35):

$$\begin{bmatrix} \Phi & P \\ P^T & 0 \end{bmatrix} \begin{bmatrix} \mathbf{a} \\ \mathbf{b} \end{bmatrix} = \begin{bmatrix} \mathbf{f} \\ 0 \end{bmatrix} \quad (3.36)$$

In this study, the particular RBF used in the interpolation scheme is the thin plate spline:

$$\phi_i(r) = r^2 \log(r), \text{ where } r = \|\mathbf{x} - \mathbf{c}_i\| \quad (3.37)$$

The minimum degree of the polynomial term is one less than the degree of the RBF. Thus, since the thin plate spline has degree 2, the polynomial degree used is  $m = 1$  (linear).

### 3.3.1 Code Implementation and Usage

The implementation of RBF interpolation used in this study is the `RBFInterpolator` class included as part of Python’s SciPy<sup>18</sup> library [40]. Using a pre-written interpolation class avoids the need to “reinvent the wheel” for what is an otherwise well-established method. Importantly, Python is not just free and open-source, but is also most likely already installed on any computer system used for scientific computing. Hence, any code written in Python should be easily able to be shared with other researchers, to be deployed for their own needs.

With tens of thousands of individual data points, it typically takes several minutes to generate the `RBFInterpolator` class, while the code computes the matrices in Eq. 3.34 and solves Eq. 3.36 for the weighting coefficients. Moreover, separate interpolator classes are required for each desired data set (output), of which there are six—lift ( $c_l$ ), drag ( $c_d$ ), and moment ( $c_m$ ) coef-

---

<sup>18</sup>Since SciPy version 1.7.0

ficients, for each of the two airfoils in the system. While it would not necessarily be *prohibitive* to generate the interpolators at the time that the lower-fidelity rotor simulations are to be performed, doing so could represent a significant fraction of the entire solution time. Conveniently however, observe that the entire RBF interpolation methodology as described up to Equations 3.36 and 3.37 does not rely on any test points, only the sampling points in the given data set `d`. This suggests that it should be possible to build the interpolator once the airfoil data has been generated, prior to any rotor simulations being performed. Indeed, `RBFInterpolator` creates a class for the *interpolator*, which may subsequently be treated like any other Python object. Thus, the `RBFInterpolator` classes need only be generated once for each data set<sup>19</sup>; then once they have been generated, Python’s `pickle` module [41] is used to save the classes to files. When the rotor simulations are to be performed, `pickle` is used to load the already-generated `RBFInterpolator` classes from the files, which may be used as normal. The file load operations and the calculation of the interpolated values at the test points are relatively inexpensive operations, adding little computational overhead to the overall methodology.

Because RBF interpolation is a mesh-free methodology, there are no specific limitations imposed on the ordering or format of the data points. The Python code developed in this study expects the input data to be contained in a file with the columns corresponding to the particular independent ( $\alpha$ ,  $M$ ,  $\Delta x$ ,  $\Delta z$ ,  $\Delta \theta$ ) and dependent ( $c_l$ ,  $c_d$ ,  $c_m$  for each airfoil) variables, separated by spaces, and the rows corresponding to the particular case. These files are ASCII-formatted and nominally human-readable, making them easy to inject additional data into if refinement is needed. It would however be trivial to modify the code to accept any reasonably comparable format (e.g. CSV), so long as the relationship between the independent and dependent variables

---

<sup>19</sup>If the data sets are changed or modified, the `RBFInterpolator` classes must be regenerated.

is noted.

Some care must be taken in the computation of the aerodynamic moment about the second (lower) airfoil. HAM2D can easily separate the lift and drag (both referenced to the free-stream flow) for the two airfoils. It can do likewise for the aerodynamic moment, but that moment is computed about the fixed point  $(x, y) = (0, 0.25)$  regardless of the location of the airfoil in the computational mesh. This is reasonable for the upper airfoil, but for the lower airfoil this results in a moment computed by HAM2D which contains a significant contribution of the offset between the quarter-chord and the calculation point, i.e.:

$$c_m^{(\text{HAM2D})} = c_m^{(\text{airfoil})} + \bar{r} \times (c_l, c_d) \quad (3.38)$$

Where  $\bar{r}$  is the offset between the airfoils<sup>20</sup>. It is necessary to compute and subtract the second term, to obtain the pitching moment about the second airfoil's quarter-chord. Initially it was attempted to do this using the *interpolated* data, but it was found that the interpolation errors in the lift, drag, and moment did not correlate, resulting in incorrect values for the moment even after correction<sup>21</sup>. This problem was resolved by instead correcting the moment *before* building the interpolators, i.e. using only data at the known points.

Further information regarding how the table lookup/interpolation methodology for the airloads integrates with the lower-fidelity rotor modeling code will be discussed in Chapter 5.3.

---

<sup>20</sup>I.E.  $(\Delta x, \Delta z)$  rotated to be referenced relative to the free-stream axes.

<sup>21</sup>In fact, since  $\bar{r} \times (c_l, c_d) \gg c_m^{(\text{airfoil})}$ , any lack of correlation in the interpolation errors would tend to be exaggerated.

### 3.3.2 Parameter Selection

Conventional airfoil tables for isolated airfoils generally have two parameters ( $\alpha$  and  $M$ ), and  $O(10^2)$  individual data points. Two-airfoil data tables necessarily have five parameters ( $\alpha$ ,  $M$ ,  $\Delta x$ ,  $\Delta z$ , and  $\Delta\theta$ ), and so it would be tempting—and concerning—to think that such tables would need  $O(10^5)$  data points to be collected. However, some useful observations can be made (this will be discussed further in Chapter 4.1) which help to significantly limit the number of data points required:

- For larger values of  $\Delta x$  (greater than about 5 chords) and likewise  $\Delta z$ , 2D CFD is of little to no benefit compared to the 2D lifting-vortex model
- Within the aforementioned range, lift and drag only change rapidly when  $\Delta x$  is within  $\pm 2c$
- Post-stall behavior of the airfoils is a 3D phenomenon (hence high-fidelity 2D CFD is of little practical use), limiting the necessary range of  $\alpha$  and  $\Delta\theta$
- The range of  $M$  may be selected based on the conditions expected to actually occur on the rotor

In practice, “production” cases in this study used  $O(10^4)$  2D CFD simulations—still a large number, but a tractable task on HPC systems. If 2D airfoil loads were required for conditions outside of those tested, the code reverted to the 2D lifting-vortex model, which is fast enough to be executed during the lower-fidelity rotor modeling process. It will be shown in Chapter 6 that these constrained parameters were still sufficient to produce accurate results in the lower-fidelity rotor model.

### 3.4 Summary

This chapter has presented all of the methods necessary for generation of the 2D loads of the two-airfoil system. One method which may be used to determine the two-airfoil loads is a simple lifting-vortex model, based on pre-computed loads for the isolated airfoils. The lifting-vortex model is analogous to the bound circulation of 3D lifting-lines, and it will be shown in the next Chapter that the lifting-vortex model is sufficiently accurate when the airfoils are separated by more than a few chord lengths. More accurate—but more expensive—is the use of high-fidelity 2D CFD, which can capture highly localized changes in the flow around the two-airfoil system.

In particular, this study uses HAM2D as its flow solver and AutoMesh as its corresponding mesh generator, the background for both of which has been described in this chapter. HAM2D is a Navier-Stokes flow solver, which implements both turbulence and transition models. The key concept of HAM2D is that the solution is performed along lines in an all-quadrilateral unstructured mesh, using methods similar to those of a structured mesh. AutoMesh was designed to create such a mesh, with the required user input being decreased through automation and default parameters. AutoMesh covers all stages of mesh generation, from surface definition through background mesh generation and quadrangulation. The combination of AutoMesh and HAM2D has been validated for accuracy in prior studies, and is currently used in academia and government.

Lower-fidelity stacked rotor simulations require 2D airfoil data in a variety of configurations. To avoid the need to perform 2D CFD during the rotor simulations themselves, the CFD airfoil loads are pre-computed over a range of configurations, similarly to how isolated airfoil

loads are pre-computed for ranges of  $\alpha$  and  $M$ . When the lower-fidelity stacked rotor simulations are performed, these pre-computed loads are interpolated to correspond with the conditions found along the span of the blades. For the stacked rotor, this requires an interpolation scheme which is capable of interpolating over the five independent parameters of the two-airfoil system (in contrast to the two independent parameters of an isolated airfoil table). This is achieved in this study by using radial basis function interpolation, building on classes offered by `Python` for ease of deployment. Another significant benefit of RBF interpolation is that it is a mesh-free method, hence the airfoil data may be collected only where it is necessary, with off-design conditions ignored and decreased resolution where the forces do not change rapidly.

## Chapter 4: 2D Results

Since the mesh generation and flow solver methodology has been validated previously [29], the results presented in this chapter will focus mainly on the behavior of two-airfoil systems. While some of this could be derived from prior work such as Smith [17], the airfoil data presented here is formatted specifically to be relevant to the airfoil systems found in stacked rotors. Hence, both a qualitative and quantitative understanding of these results will demonstrate both the need for this methodology and ways in which stacked rotors could be designed to take the most advantage of beneficial airfoil interactions.

### 4.1 System of Symmetric Airfoils

As a symmetric NACA airfoils is often used as a simple test case for an isolated airfoil methodology, so too is a system of two symmetric NACA 0008 airfoils used here as a simple test case. Nonetheless, some key features of the airfoil interaction can still be observed. For all cases,  $M = 0.2$  and  $\frac{\Delta z}{c} = 0.1667$ . Note that since Mach (Reynolds number) effects on the airfoil forces are fairly insignificant in the range relevant to the stacked rotor rig ( $M \approx 0.08$  to  $M \approx 0.4$ ), the effect of varying Mach number will be ignored.

### 4.1.1 Symmetric Configuration

Arguably the simplest configuration of two symmetric airfoils is to place both airfoils at zero incidence angle to the free-stream flow ( $\alpha_\infty = 0^\circ$  and  $\Delta\theta = 0^\circ$ ). Under these conditions, if each airfoil was modeled in isolation, it would produce zero lift. If the airfoils were modeled using a 2D lifting-vortex model then they would still produce zero lift—and in fact experience zero interaction of any sort—regardless of their placement relative to each other, because (following Algorithm 3.1):

1. Both airfoils are initialized to the free-stream condition, i.e.  $\alpha_{(1,2)} = \alpha_\infty = 0^\circ$
2. From (isolated) airfoil table lookup,  $c_l(\alpha = 0^\circ) = 0$
3. Zero lift implies zero circulation ( $\Gamma$ ) and hence zero velocity induced on either airfoil by the other ( $\Delta M$ ) per Eq. 3.6, regardless of their relative positions
4. With zero velocity induced on either airfoil, the angle of attack on neither airfoil changes, hence  $\alpha_{(1,2)} = 0^\circ$  represents a converged solution

Likewise, since  $c_d$  is strictly a function of  $\alpha$  (which does not change), the drag on the airfoils would be expected to be the same regardless of the placement of the airfoils.

However, when the airfoil system is modeled for a range of  $\Delta x$  using high-fidelity 2D CFD, significant changes in both the total and individual forces can be seen depending on the configuration (Figure 4.1). In particular, a net lifting force on the two-airfoil system is generated in the range  $|\Delta x/c| < 2$ , with a peak magnitude<sup>1</sup> of  $c_l = 0.0665$  at  $\Delta x/c = 0.75$  (Fig. 4.1a). The

---

<sup>1</sup>As should be expected due to symmetry, mean system  $c_l(\Delta x) = -c_l(-\Delta x)$  with the individual elements having mirrored forces, hence it is only necessary to discuss  $\Delta x > 0$ .






Figure 4.1: Variation of forces for a system of two NACA 0008 airfoils at zero incidence angle

drag on each element also changes significantly (and in fact is “negative” in some cases) in the same range, but the total drag of the system is increased when  $|\Delta x / c| < 0.75$  (Fig. 4.1b).

Why should net lift be generated by a seemingly symmetric system of airfoils? Consider the case  $\Delta x / c = 0.75$  corresponding to the peak lift, the flow field about which is shown in Figure 4.2b (the corresponding computational mesh for which is shown in Fig. 4.2a). The lower airfoil element represents a blockage to the flow near the trailing edge of the upper airfoil, which results in the circulation effect (see Chapter 1.2.2.2) as the flow is deflected around the obstruction, inducing a trailing-edge angle of attack on the upper element. Hence, the upper element will have a positive lift. This positive lift is somewhat offset by the lower element having negative lift, as a consequence of the lifting upper element reducing the suction on the lower element’s upper surface through the slat effect (see Chapter 1.2.2.1). However, the reduction in lift on the lower element due to the slat effect is less than the amount of lift on the “slat” (upper element) itself, hence the net result is a positive lifting force.




(a) Computational mesh




(b) Dynamic pressure

Figure 4.2: CFD performed for a system of two NACA 0008 airfoils at zero incidence angle with  $\frac{\Delta x}{c} = 0.75$



(a) Effect of circulation on the angle of attack seen by each element



(b) Force components resolved in different reference frames

Figure 4.3: Effect of circulation to change the apparent forces on each airfoil

The changes in drag on each airfoil element (particularly the negative drag values) may seem non-physical at first glance, but actually are a consequence of the change in direction of the resultant force<sup>2</sup> on each airfoil element. As a simplified model, consider each lifting airfoil to be a point vortex (following the Kutta-Joukowski theorem, Eq. 3.1). A positive circulation (lift) about the leading airfoil will induce a downwards<sup>3</sup> velocity on the trailing airfoil, resulting in the trailing airfoil seeing an angle of attack less than that of the free-stream (as shown in Fig. 4.3a). Similarly, positive circulation about the trailing airfoil will induce an upwards velocity on the leading airfoil, resulting in the leading airfoil seeing an increased angle of attack. It should be expected that the lift and drag on each element resolved with respect to that airfoil's *local* angle of attack ( $L'$ ,  $D'$ ) should be positive. However, HAM2D reports the drag (as shown e.g. in Fig.

<sup>2</sup>The vector force on the airfoil

<sup>3</sup>Because  $u_\infty \gg v_\infty$ , the contribution of the *horizontal* velocity induced by the lifting vortex to the change in flow direction will be small compared to that of the *vertical* velocity.

4.1b) with respect to the *global* (free-stream) angle of attack. Hence (as noted in Fig. 4.3b), the apparent lift and drag with respect to the free-stream ( $L, D$ ) will be:

$$\begin{aligned} L &= L' \cos(\Delta\alpha) + D' \sin(\Delta\alpha) \\ D &= D' \cos(\Delta\alpha) - L' \sin(\Delta\alpha) \end{aligned} \quad (4.1)$$

Where  $\Delta\alpha$  is the offset between the local and global angles of attack. Hence, since generally  $L' \gg D'$ , it is possible that  $D < 0$  on an airfoil element in a multi-element system. In no case however is the *total* drag of the airfoil system ever less than zero.

While for the most part the total drag of the airfoil system remains constant<sup>4</sup>, the total drag is increased in the range  $|\Delta x/c| < 0.75$ , i.e. when the airfoils are close to evenly spaced. In these configurations, the airfoils behave like a converging–diverging duct, which presents a greater obstruction to the flow than the sum of the two isolated airfoils. Additionally, this duct-like effect creates a suction between the two airfoils, as can be noted by examining the individual airfoil lifts in Figure 4.1a.

Although the configuration of two symmetric airfoils at zero incidence angle is not likely to be used in a stacked rotor, some important observations relating to stacked rotor design can still be made. Clearly, there exist configurations in which the airfoil system exhibits beneficial (increased lift, constant drag) or detrimental (decreased lift, increased drag) interactions, which would be relevant to the performance (thrust, efficiency) of a stacked rotor system. Greater than the change in the total forces however is the change in the individual forces on each airfoil. These individual airfoil forces directly contribute to the rotor blade deflections. Moreover, neither the change in

---

<sup>4</sup>Note that this is consistent with the earlier observations of Norton [16]

total force nor the change in the individual forces would be observed (for these configurations) without using high-fidelity CFD tools (2D or 3D).

#### 4.1.2 Lifting Airfoil Comparison

It would be reasonable to question at this point whether the effects seen in the symmetric airfoil configuration correlate to those seen in a configuration of greater relevance to a stacked rotor. After all, the total lift generated was fairly small, and the airfoils experienced relatively little interaction which could be attributed *strictly* to bound circulation. A more relevant case would conceivably involve airfoils which *would* generate lift even in isolation.

To test the interaction between lifting airfoils, the previous case was modified by introducing an offset to the incidence angle of  $\Delta\theta = 4^\circ$ . All other parameters were identical to the previous case (i.e.  $\alpha = 0^\circ$ ,  $\Delta x$  varied). In this configuration, only the lower airfoil would be expected to experience lift if modeled in isolation. However, a 2D lifting-vortex (or 3D lifting-line) model *would* demonstrate mutual interaction between the airfoils<sup>5</sup>.

The comparison between the 2D lifting vortex model and HAM2D is shown explicitly for this case in Figures 4.4 and 4.6. In these figures,  $\Delta c_n$  is the difference in airfoil force coefficient between HAM2D and the 2D vortex model. In terms of the mean forces on the airfoil system, it can be seen that while the 2D vortex model is able to broadly capture the *trend* in variation of lift with position, it significantly diverges in *magnitude* when compared to HAM2D (Fig. 4.4a). At  $\Delta x/c = 1$  where the lift of the system is at a maximum value of  $c_l = 0.418$  per HAM2D, the 2D vortex model under-predicts the lift by just under 27%. Examining the flow field in

---

<sup>5</sup>On the first step of Algorithm 3.1 only the lower airfoil would influence the upper, but because the upper airfoil would then be generating lift it would in turn influence the lower on subsequent steps.





Figure 4.4: Variation of total forces for a system of two NACA 0008 airfoils with  $\Delta\theta = 4^\circ$

this configuration (Fig. 4.5b, with corresponding mesh shown in Fig. 4.5a), the presence of the slat<sup>6</sup>, circulation<sup>7</sup>, and dumping<sup>8</sup> effects should be apparent, and as such the changes to the flow around each airfoil are highly localized. As the airfoils move farther apart ( $|\Delta x|$  increases) the difference between the models ( $\Delta C_l$ ) decreases to nearly zero, even as the total lift has not yet converged to that of the airfoils in isolation. Physically (e.g. recalling Eq. 3.5), as the airfoils move farther apart, the influence of one airfoil on the other varies relatively less along the length of each airfoil’s chord (since  $r \gg c$ )—that is, the influence between the airfoils may be accurately approximated by a change only in the free-stream conditions around each airfoil, which is precisely how the 2D lifting-vortex model represents the interactions. With respect to the mean drag of the airfoil system (Fig. 4.4b), the 2D vortex model is entirely unable to capture the “channel” drag, but for  $|\Delta x/c| > 1$  the drag varies fairly little, and both models are well-matched.


<sup>6</sup>Lifting upper element ahead of the lower element

<sup>7</sup>Both lifting and real-thickness lower element below and behind the trailing edge of the upper element

<sup>8</sup>As flow leaving the trailing edge of the upper element does so into the suction region of the lower element




(a) Computational mesh




(b) Dynamic pressure


Figure 4.5: CFD performed for a system of two NACA 0008 airfoils with  $\frac{\Delta x}{c} = 1.0$  and  $\Delta\theta = 4^\circ$




(a) Upper element lift



(b) Upper element drag



(c) Lower element lift



(d) Lower element drag

Figure 4.6: Variation of individual forces for a system of two NACA 0008 airfoils with  $\Delta\theta = 4^\circ$

Examining the lift on each airfoil, the channel effect in the HAM2D simulations is more apparent. The upper element experiences a strong downwards force (Fig. 4.6a) in the range  $|\Delta x/c| < 1$ , while the lower element experiences a strong upwards force in the same range (Fig. 4.6c). Outside of this range, the 2D vortex model tends to under-predict the upper airfoil’s lift and (to a lesser degree) over-predict the lower airfoil’s lift. That the difference in the individual element lifts is greater than the difference in total lift is similar to what was observed in the previous example (Fig. 4.1a), but the 2D vortex model is unable to accurately predict the magnitude of this difference.

Similarly, the lifting-vortex model is not able to capture the large differences in drag on each element (Figs. 4.6b and 4.6d) when the airfoils are closely spaced. Even outside of the “channel” region, for  $1 < \Delta x/c < 2$  it can be seen that HAM2D predicts slightly lower drag on the upper element, and slightly higher drag on the lower element. It should of course be noted that these differences in “drag” are not unrelated to the differences in lift, as described previously.

Some useful conclusions can be drawn from this case. In particular, it is reasonable to limit the range of geometric configurations which require high-fidelity 2D CFD simulation, since for “large” spacing between the airfoils the 2D vortex model produces accurate results. The range of configurations where 2D CFD is required is driven mainly by the difference in *lift* between the models, which for these airfoils is most significant where  $|\Delta x/c| < 4$ . In contrast, differences in drag between the models are only observed for  $|\Delta x/c| < 2$ , and are most significant when the airfoils are very closely spaced. While closely-spaced airfoils acting as a “channel” would not likely be a design condition over the span of the blades of a stacked rotor, it is possible for this condition to occur near the blade root (where the blades are necessarily closer together, since  $\Delta x \approx r \Delta \psi$ ).

### 4.1.3 Angle of Attack and Pitch Offset

Before examining the effect of the pitch offset ( $\Delta\theta$ ) in more detail, it is instructive to consider how changing  $\Delta\theta$  relates to the characteristics of a “conventional” airfoil. Imagine that the two airfoils in the system are placed such that the leading edge of the lower/trailing element is near the trailing edge of the upper/leading element. Changing the angle  $\Delta\theta$  between the airfoils would be physically similar to a conventional airfoil deflecting a plain or split flap downwards. At an even higher level, this is similar to the effect of changing the *camber*<sup>9</sup> of a single airfoil of equivalent chord. It is trivial to observe that prior to stall, a cambered airfoil has a lift curve of the form:

$$\begin{aligned} c_l(\alpha) &= c_{l_\alpha} \cdot (\alpha - \alpha_0), \text{ or} \\ c_l(\alpha) &= c_{l_\alpha} \cdot \alpha + c_{l_0} \end{aligned} \tag{4.2}$$

Where  $\alpha_0$  is the zero-lift angle of attack or  $c_{l_0}$  is the lift at  $\alpha = 0$ . In either case, the effect of the camber shows up primarily as the lift increment  $\alpha_0$  or  $c_{l_0}$ . Moreover, it can be shown using thin-airfoil theory that the relationship between the flap deflection angle and the the lift increment will be roughly linear for small-angle deflections, i.e.:


$$c_l(\alpha, \delta) = c_{l_\alpha} \cdot \alpha + c_{l_\delta} \cdot \delta \tag{4.3}$$

Where  $\delta$  is the flap deflection angle and  $c_{l_\delta}$  is the flap deflection lift curve slope.


Does this theory hold true for changing the pitch offset in the two-airfoil system? To

---


<sup>9</sup>Exactly as Hamilton Standard proposed, see Chapter [1.2.1](#)



(a) Lift coefficient



(b) Drag coefficient



(c) Lift-to-drag ratio

Figure 4.7: Variation of forces with  $\Delta\theta$  for a system of two NACA 0008 airfoils

test this supposition, a system of two airfoils was simulated in HAM2D with a fixed horizontal position of  $\Delta x/c = 1.25$  and free-stream angle of attack  $\alpha_\infty = 0^\circ$ , the results of which are shown in Figure 4.7. HAM2D shows that there is in fact a nearly-linear relationship between the pitch offset and the mean lift of the airfoil system (Fig. 4.7a). Interestingly, despite the upper and lower airfoils being at increasingly different angles to the flow as  $\Delta\theta$  increases, each airfoil contributes a nearly equal amount of lift to the total system, with the upper element producing slightly more lift than the lower at low  $\Delta\theta$  and the lower producing more at higher  $\Delta\theta$ .

The relationship between drag and  $\Delta\theta$  (Fig. 4.7b) is clearly not linear however, with the mean drag on the airfoil system exhibiting a “drag bucket” type behavior. Once again the individual element drags are heavily influenced by coupling with the lift, but it is apparent that the increase in mean drag on the system is driven more by the increase in the drag on the lower element. Since the lift increases linearly with  $\Delta\theta$  but the drag does not it would be reasonable to suggest that some  $\Delta\theta$  would correspond to the best efficiency of the airfoil system ( $c_l/c_d$ ), a suggestion which is borne out by Figure 4.7c. In this particular case, the maximum efficiency is  $c_l/c_d \approx 72$  which occurs at a pitch offset of  $\Delta\theta = 6^\circ$ . This efficiency is for one particular  $\alpha_\infty$ , but per the general form of Equation 4.3 it should be expected that  $\alpha_\infty$  and  $\Delta\theta$  will have similar effects.

Before examining the variation of  $\alpha_\infty$  and  $\Delta\theta$  together, consider the effect of  $\alpha_\infty$  on its own. In this case, the pitch offset was fixed at  $\Delta\theta = 6^\circ$  (which was the condition for maximum efficiency with  $\alpha_\infty = 0^\circ$ ), and the horizontal offset remained fixed at  $\Delta x/c = 1.25$ . Results for this case are shown in Figure 4.8. As expected, the relationship between  $\alpha_\infty$  and the mean lift on the two-airfoil system is linear (Fig. 4.8a). Compared to an isolated NACA 0008 airfoil, the two-airfoil system exhibits a nearly identical lift curve slope, but is offset by a zero-lift  $\alpha$  of slightly




Figure 4.8: Variation of forces with  $\alpha_\infty$  for a system of two NACA 0008 airfoils

over  $-4^\circ$ . Likewise, lift stall occurs about  $4^\circ$  sooner for the two-airfoil system ( $\alpha_{stall} = 4^\circ$ ) than for the isolated airfoil ( $\alpha_{stall} = 8^\circ$ ). However, the maximum lift just prior to stall is slightly greater for the two-airfoil system ( $C_{l,max} = 0.895$ ) than for the isolated airfoil ( $C_{l,max} = 0.834$ ). Despite the two-airfoil system having a similar lift curve slope to the isolated airfoil, the upper element has a relatively steeper lift curve slope while the lower element has a flatter slope. This is reasonable considering the slat (Chapter 1.2.2.1) and circulation (Chapter 1.2.2.2) effects—the upper airfoil will tend to see increased lift, while the lower airfoil will see decreased lift, with the net result being lift enhancement.

Drag variation with  $\alpha_\infty$  (Fig. 4.8b) follows similar trends to those observed with respect to  $\Delta\theta$ . However, it can be seen that when stall occurs, it does so primarily on the upper element. By contrast, the drag on the lower element increases smoothly, even beyond the angle of attack at which the isolated airfoil has stalled. Again, this broadly follows what would be expected due to the slat effect.




Figure 4.9: Effect of varying  $\alpha_\infty$  and  $\Delta\theta$  on the efficiency of a system of NACA 0008 airfoils

| Case                     | $\alpha$ before stall | $c_l/c_d$ before stall | Equivalent low- $\alpha$ | Range of $\alpha$ |
|--------------------------|-----------------------|------------------------|--------------------------|-------------------|
| Isolated airfoil         | $8^\circ$             | 44                     | $2^\circ$                | $6^\circ$         |
| $\Delta\theta = 0^\circ$ | $6^\circ$             | 51                     | $2.25^\circ$             | $3.75^\circ$      |
| $\Delta\theta = 3^\circ$ | $5^\circ$             | 55                     | $0^\circ$                | $5^\circ$         |
| $\Delta\theta = 6^\circ$ | $4^\circ$             | 54                     | $-2^\circ$               | $6^\circ$         |

Table 4.1: Summary of the operating range of  $\alpha$  from Figure 4.9a

The effect of combining variation in  $\alpha_\infty$  and in  $\Delta\theta$  on airfoil system efficiency is shown in Figure 4.9. Overall, it can be seen that increasing  $\Delta\theta$  results in similar trends in efficiency, but trends which occur at progressively lower  $\alpha_\infty$  (Fig. 4.9a). Additionally, increasing  $\Delta\theta$  increases the range of  $\alpha$  at which  $c_l/c_d$  is high, an effect summarized in Table 4.1. For illustrative purposes, the “operating range” was defined as the range of  $\alpha_\infty$  over which  $c_l/c_d$  is at least as great as its value immediately prior to stall. Observe that by this definition however, the two-airfoil system tends to have a narrower (or at least no wider) operating range than the isolated airfoil.

Realistically, a lifting rotor on a vehicle would be trimmed to achieve some thrust—hence

| Case                     | $(c_l/c_d)_{max}$ | Corresponding $\alpha$ | Corresponding $c_l$ |
|--------------------------|-------------------|------------------------|---------------------|
| Isolated airfoil         | 64                | $4^\circ$              | 0.48                |
| $\Delta\theta = 0^\circ$ | 61                | $4^\circ$              | 0.49                |
| $\Delta\theta = 3^\circ$ | 71                | $2^\circ$              | 0.50                |
| $\Delta\theta = 6^\circ$ | 73                | $0.5^\circ$            | 0.55                |

Table 4.2: Maximum airfoil system efficiency and corresponding conditions from Figure 4.9

some lift along its blades—with however much power (drag) is required to do so. Hence, Figure 4.9b shows the airfoil efficiency plotted versus lift coefficient. From this plot, it can be seen that the two-airfoil system is of little benefit (or detriment) for  $c_l$  less than about 0.3. Above this value, the different airfoil systems differentiate themselves. While the two-airfoil system with  $\Delta\theta = 0^\circ$  has worse maximum efficiency than an isolated airfoil, the systems with  $\Delta\theta = 3^\circ$  and  $\Delta\theta = 6^\circ$  show more than a 10% *increase* in maximum efficiency (summarized in Table 4.2). Moreover, the peak efficiency for the two-airfoil system occurs at a greater value of  $c_l$  than for the isolated airfoil. Hence, if a rotor was designed to generate a certain amount of thrust, gains in rotor efficiency due to improved airfoil efficiency could be compounded by a decrease in the required blade area.

## 4.2 System of VR12 Airfoils

Now that some of the advantages and limitations of two-airfoil systems have been shown, consider what would be required for the design of a rotor for a production vehicle. Realistically, some of the advantage demonstrated for two-airfoil systems could be achieved using isolated cambered airfoils (although it will be shown in this section that the benefits of the two-airfoil system do translate to systems of cambered airfoils). Additionally, it would likely not be possible


to precisely control the separation between the airfoils along the entire span of the blade—even ignoring blade motions, straight untapered blades will have a separation between their airfoils of  $\Delta x = r\Delta\psi$ . While it could be possible to alter this relationship using inverse taper (as with Hamilton Standard) or curved blades, such blades would introduce unique structural and aerodynamic concerns, which may not make them suited for rotorcraft applications<sup>10</sup>. Hence, it is more likely that  $\Delta x$  (and  $\Delta z$ ) will vary along the span of the blade. Can a system of two airfoils still offer a benefit to the rotor system in this case?

To examine some parameters relevant to a design-like scenario, a system of VR12 airfoils was modeled for a range of  $\Delta x$ ,  $\Delta z$ ,  $\Delta\theta$ , and  $\alpha$ , with a fixed  $M = 0.2$ . The VR12 airfoil was chosen specifically because, aside from being a typical rotorcraft airfoil, it is the airfoil used on the test stand and rotor blades developed by Johnson et.al. [10], against whose data the results from the current study will be examined in Chapter 6. Rather than attempt to show the variation with each parameter, the goal of this section is to show what could be done to maximize the previously-demonstrated positive interactions between the airfoils.


For the expected nominal operating condition of a rotor, it would be desired to minimize the required power while achieving the required thrust. At an airfoil-level, this means maximizing the  $c_l/c_d$  (or  $L/D$ ) of the airfoil—or system of airfoils—along the span of the rotor. Figure 4.10 shows the improvement possible with a coupled airfoil system, along with what conditions would be required to achieve this. Clearly, there is a narrow zone of configurations centered around  $(\Delta x/c, \Delta z/c) = (1.0, 0.2)$  which offer significant benefit over an isolated airfoil (Fig. 4.10a, where the iso-line represents  $(c_l/c_d)_{max}$  for an isolated VR12 airfoil). Whether such a configuration of airfoils is possible at all depends upon the rotor hub geometry and expected

---


<sup>10</sup>At least not without further research which would be outside the scope of the current work.



(a) Maximum  $c_l/c_d$ , with iso-line representing an isolated airfoil



(b) Corresponding  $\alpha_\infty$



(c) Corresponding  $\Delta\theta$

Figure 4.10: Efficiency and corresponding conditions for a system of two VR12 airfoils across a range of  $\Delta x/c, \Delta z/c$

blade motions, but likely could not be achieved over the entire rotor span. However, it can also be seen that there is a broad region of at-least-modest improvement to the airfoil system efficiency even as  $\Delta x$  and  $\Delta z$  are increasingly positive. (This should be expected, based on the results seen previously in Figure 4.4.) It can also be seen that for  $\Delta x/c < 0.75$  (varying weakly with  $\Delta z$ ), the two-airfoil system is actually less efficient than an isolated airfoil.


Importantly, it can also be shown what conditions are required to achieve the maximum possible airfoil efficiency. Consider a stacked rotor with straight untapered blades, having a configuration of airfoils around  $(\Delta x/c, \Delta z/c) = (1.0, 0.2)$  somewhere near the blade root. Moving out along the span towards the tip of the blades,  $\Delta x$  will increase. Figure 4.10b suggests that while zero or slightly negative upper-airfoil angle of attack would be desired near the root, a greater angle of attack of around  $5^\circ$  would be desired outboard on the span. Additionally, Figure 4.10c shows that while a high  $\Delta\theta$  of up to  $13^\circ$  would lead to better efficiency near the root, a lower  $\Delta\theta$  of  $6 - 7^\circ$  would lead to superior efficiency outboard on the rotors. Combining the conditions of  $\alpha$  and  $\Delta\theta$  suggests that not only would some twist along the blades of a stacked rotor lead to superior efficiency, but in fact some *differential twist* between the rotors may also be desirable. To date, no study<sup>11</sup> has yet been published examining the effect of differential blade twist on a stacked rotor.

Another concern in rotor design might be the maximum load that the rotor can support without causing stall<sup>12</sup>, which at an airfoil level would relate to  $c_{l,max}$ . Results for maximum lift of the two-airfoil VR12 system are shown in Figure 4.11, and are broadly similar to the results shown for maximum  $c_l/c_d$ . Again, there is a significant increase in maximum possible


---

<sup>11</sup>Of which the author is aware


<sup>12</sup>Or exceeding the available power, but engine dynamics are outside the scope of the current work



(a) Maximum  $c_l$ , with iso-line representing an isolated airfoil



(b) Corresponding  $\alpha_{\infty}$



(c) Corresponding  $\Delta\theta$

Figure 4.11: Maximum lift coefficient and corresponding conditions for a system of two VR12 airfoils across a range of  $\Delta x/c, \Delta z/c$

$c_l$  in a narrow region centered around  $(\Delta x/c, \Delta z/c) = (1.0, 0.2)$ , but a broad area of moderate improvement for  $\Delta x/c > 0.4$  and positive  $\Delta z$  (Fig. 4.11a, where the iso-line represents  $c_{l,max}$  for an isolated VR12). The corresponding free-stream angle of attack for  $c_{l,max}$  (Fig. 4.11b) is greater than that for maximum  $c_l/c_d$ , with  $\alpha = 7^\circ$  in the region of greatest  $c_{l,max}$  improvement and  $\alpha = 10^\circ$  as  $\Delta x$  is increased. Likewise,  $\Delta\theta = 13^\circ$  would be required in the region of greatest improvement, decreasing to  $\Delta\theta = 7^\circ$  as  $\Delta x$  increases. These results suggest that twist of the blades (both together and differential) could be of benefit not just to efficiency, but also to the maximum thrust. Notice also that as  $\Delta z$  is increased, the required  $\alpha$  increases and the required  $\Delta\theta$  decreases for  $c_{l,max}$ , an effect which was not observed for maximum  $c_l/c_d$ .

### 4.3 Summary

It has long been established that a system of multiple airfoils may generate more lift and be more aerodynamically efficient than an isolated airfoil. In this chapter in particular then, the behavior of two-airfoil systems has been examined to demonstrate both the need for the current methodology, and how certain parameters of the airfoil system will affect the system's overall performance.

Regarding the need for the current methodology, it has been shown in this Chapter that when the airfoils are separated by less than about 4–5 chord lengths, high-fidelity 2D CFD will capture aerodynamic interactions which a simple lifting-vortex model is unable to represent. (Further, it will be shown in Chapter 6 that offsets of 5 chord lengths or less are highly relevant to the airfoil systems seen in a stacked rotor.) Among the effects which the lifting-vortex model cannot represent are those due to the real thickness of the airfoils, and those which are due to

localized changes in the flow around the airfoils. As the airfoils move further apart however the CFD and lifting-vortex models agree, suggesting that limits may be placed on the range of 2D CFD necessary for lower-fidelity rotor simulations.

When designing a rotor, it is necessary to understand how changing the parameters of the airfoil system will change the performance of the rotor. In isolation, changing the free-stream angle of attack results in changes in forces analogous to those of an isolated airfoil, while changing  $\Delta\theta$  is analogous to changing the camber or flap deflection angle of an airfoil. Moreover, there is a specific combination of  $\alpha$  and  $\Delta\theta$  for a system of airfoils which will lead to the greatest improvement in efficiency. While it is clear that the greatest improvement in efficiency also requires a specific close spacing between the two airfoils, it may not be possible<sup>13</sup> to maintain this particular spacing along the span of the blades. As the separation between the airfoils changes, the corresponding  $\alpha$  and  $\Delta\theta$  for best efficiency improvement also change, suggesting that both blade twist and differential twist in stacked rotors may lead to further gains in efficiency.

---

<sup>13</sup>Without resorting to such novel designs as curved or swept blades

## Chapter 5: 3D Methodology

While the previous chapters of this work have been concerned primarily with extraction of the flow conditions (Chapter 2) and data generation for a two-airfoil system (Chapters 3 and 4), it is now time to consider the lower-fidelity rotor (3D) model itself in greater detail. The goal of this chapter is not to develop an entire lower-fidelity rotor model, but to consider which models may or may not be suitable for use with stacked rotors, and to show how suitable models can be improved using interactional airfoil data. Since the particular implementation of the rotor model will generally dictate how the 2D airloads can be adjusted, the coupling between the selected rotor model and the 2D airfoil code will also be discussed in this section.

A general overview of the relationship between the 2D and 3D models is shown in Figure 5.1. Regardless of which specific lower-fidelity rotor model is used, it is the job of the rotor model to compute the aerodynamic conditions (inflow and wake) around the rotor; these conditions are




Figure 5.1: Relationship between the 2D and 3D models

necessary inputs to the 2D airfoil model. The rotor model must also account for blade motions, including rigid body/hinge motion and bending (beam) deflections. These motions are especially important for stacked rotors, since the airfoil loads may be sensitive to changes in relative blade positions.

## 5.1 Considerations

With respect to any lower-fidelity rotor model, there are a few key considerations:

1. Accuracy/physicality
2. Ease of use for the rotor designer
3. Prerequisites

The first consideration is a seemingly trivial statement that the results of using the lower-fidelity rotor model to simulate a rotor should generally agree with observations and experiments on actual rotors. However, many lower-fidelity rotor models make simplifying assumptions about the physics of the rotor wake, assumptions which may be “good enough” for conventional rotors but do not hold true for stacked rotors. In particular, previous studies (e.g. [5, 7]) have shown that the interaction on the lower rotor of the tip vortex from the upper rotor is important to the performance of the stacked rotor as a whole. Therefore, any rotor model which does not include a representation of the tip vortices would not be suitable for modeling a stacked rotor. Since there are some peculiarities regarding how the selected rotor code implements the inflow models, the inflow models are described in greater detail alongside the rotor code in the following section.

The second consideration is important since the goal of this work is to develop a method-

ology which can be *used* in future rotor design, not just a theoretical background. Ideally the methodology should interface with one or more codes already in use by rotor designers, and novel components of the methodology should be easy for the rotor designer to learn how to use. Additionally, the particular lower-fidelity rotor model chosen should require a minimum of “tuning” or adjustment by the rotor designer to produce accurate results—both for ease-of-use, and out of concern that if the model is “tuned” correctly for one case, that tuning may be incorrect if the case is changed.

The third consideration is somewhat similar to the second, but relates to how difficult it is to *obtain* and execute the code for the lower-fidelity rotor model. While it would be nice if the rotor model could be free and open-source, the time and labor required to develop and validate such specialty<sup>1</sup> software does not necessarily make this a reasonable proposition. Additionally, although the point of lower-fidelity rotor models is that they obtain results at a lower computational cost than high-fidelity models, different lower-fidelity models may require different degrees of computing power; when it is necessary to run a large number of rotor cases, the computing power required will be scaled accordingly.

## 5.2 RCAS

The Rotorcraft Comprehensive Analysis System (RCAS) is an advanced software system for complete, reduced-order modeling of arbitrary rotorcraft systems [23]. RCAS began development with the US Army Aviation and Missile Research, Development, and Engineering’s (AMRDEC) Aeroflight Dynamics Directorate (AFDD), and is currently developed and maintained

---

<sup>1</sup>Compared to something more broadly applicable like Python, rotorcraft design and analysis is a much smaller field.

by Advanced Rotorcraft Technology, Inc. (ART). While RCAS requires a proprietary license, it is widely used in academic, government, and industry applications, making it a reasonable candidate to serve as the basis of the 3D methodology in this work.

RCAS uses a hierarchical model for both structure and aerodynamics, in which each component consists of a number of primitive elements, and those components are connected together to form the complete rotorcraft model [42]. For instance, a rotor’s aerodynamic model would contain the blades, which would include aerodynamic segments having properties of chord, airfoil section, etc.; the rotor’s structural model would contain the blades, which would include structural segments having properties of stiffness, etc. RCAS also offers models for wings, fuselages, and engines, which are not of interest to the current work.

In the current work, a periodic solution analysis is used in RCAS, which allows the rotor loads and blade deflections to be determined for a given set of control inputs<sup>2</sup>. (The data from these simulations is time-averaged in post-process, for comparison with similarly time-averaged experimental data.) However, RCAS is also capable of performing several other types of rotorcraft analysis:

- *Trim analysis*<sup>3</sup>, in which an iterative procedure is used to determine the controls (trim variables) required to achieve a certain flight condition (trim constraints).
- *Maneuver analysis*, in which the nonlinear transient response is computed for a set of specified controls or motions (position, velocity, or acceleration).
- *Stability analyses* at one or more equilibrium conditions, where linearization of the con-

---

<sup>2</sup>In this case, upper and lower collective

<sup>3</sup>“Periodic solution” is grouped with “trim” in the RCAS analysis hierarchy, but it should be noted that the analysis performed in this study is *not* a “trim” analysis as it would be typically understood for a rotorcraft.

trol/response equations are performed, allowing for analysis of the eigenvalues (stability) or linearized response.

These analyses are not unique to rotors, but can be performed on the full model including wings, engines, and fuselages with coupled effects.

Output data is selected through several menus, which may be attached to the RCAS model as TOPR scripts. Further data manipulation in RCAS (including output-to-file) may be performed with RSCOPE, and several options for RSCOPE-formatted user-additional scripts may be attached to an RCAS model.

While RCAS clearly offers a significant depth and breadth of modeling and analysis options, much of this is superfluous to the methodology of the present work. In this study, the rotor system is examined in isolation (i.e. without modeling the support structure). Because only the aerodynamics—and structural effects on the aerodynamics—are of interest, the rotors are simulated at a specified speed, with motor dynamics being ignored. Similarly, analysis is performed using fixed values of collective and examining the resulting thrust and power, i.e. not trimming to any particular target. While this may seem like a simplistic analysis compared to the capabilities of RCAS, the analysis follows the procedure of recent experimental and computational studies (e.g. [3,4]). Moreover, the point of the current work is not to analyze any specific rotor or rotorcraft, but to validate the methodology of including 2D airfoil-interactional data in lower-fidelity rotor simulations. The capabilities of RCAS which are unused in this study are considered a plus, as once the methodology is validated these capabilities should be able to be used alongside the methodology with little to no further modification required.

With respect to the lower-fidelity rotor analysis used in the current work, the most important

models in RCAS are the aerodynamic (inflow and wake) model, the structural (beam) model, and the CFD/CSD coupling interface. These models are described in greater detail in the following sections. Except as noted, descriptions of the models (Sections 5.2.1 and 5.2.2) are given from [43].

### 5.2.1 Inflow Models

By default, RCAS offers four models for computing the self-induced inflow of a rotor, ranging from simplest to most complex: Momentum theory (uniform or annular, the latter as in BEMT), dynamic inflow, prescribed vortex wake, or free vortex wake [43]. Some recent studies have also used a viscous vortex particle method (VVPM) coupled to RCAS , which serves the same function.

Within the RCAS model, the upper and lower rotors of the stacked rotor system are modeled as separate subsystems/supercomponents<sup>4</sup>. As stated previously, considering the upper and lower rotors to be “separate” provides a useful framework for analyzing both the loads of each part of the stacked rotor system, and the system as a whole. However, this also requires that an aerodynamic interference (interaction) model be specified between the rotors.

#### 5.2.1.1 Momentum Theory and Dynamic Inflow

The simplest possible inflow model is uniform momentum theory. Let  $\lambda = \lambda_m - \lambda_c$ , where  $\lambda$ ,  $\lambda_m$ , and  $\lambda_c$  are the total, thrust-induced, and climb-induced components of the inflow respectively, and  $\mu$  is the in-plane velocity, all of which are normalized by the tip speed. Uniform

---

<sup>4</sup>The RCAS structural model refers to each rotor as a “subsystem”, while the RCAS aerodynamic model refers to each rotor as a “supercomponent”. Regardless of notation, this means that each rotor is “independent” within the solver hierarchy.

momentum theory states that:

$$\lambda_m = \frac{C_T}{2\sqrt{\lambda^2 + \mu^2}} \quad (5.1)$$

This equation for  $\lambda_m$  must be solved iteratively, which is often performed by Newton's method. The equation may also be extended for annular inflow, where inflow may vary along the span of the rotor.

The second-simplest possible inflow model is dynamic inflow. RCAS specifically uses the model of Peters and He [44]. Dynamic inflow is based on unsteady potential flow theory, and from there derives a coupled set of equations for lift (blade loading) and inflow represented in terms of azimuthal harmonics and radial distributions. This coupled set of equations forms a system of ordinary differential equations (ODEs), which may be solved using typical methods for ODEs. This model is suitable for unsteady computations, and has been shown to produce accurate results in forward flight.

A major limitation of both momentum theory and dynamic inflow is that neither one models the localized distribution of flow (i.e. vortices) in the wake—at most, the the inflow may be used to determine an approximate mean wake velocity. While this is reasonable for isolated rotors, and may even be sufficient for computing the download on a blunt body (e.g. a fuselage or pylon), it is known that interaction between the wake vortices and the blades is important to stacked rotor performance.




Figure 5.2: Superposition of horseshoe vortices along a lifting line, excerpted from [45]

### 5.2.1.2 Prescribed Wake

Prescribed vortex wake and free vortex wake both derive generally from lifting-line theory. In lifting-line theory [45], a finite-length lifting wing<sup>5</sup> is represented as a series or superposition of lifting vortex segments bound to the wing's quarter chord (the *lifting line*). At each point along the wing, the bound vortex strength is given by the Kutta-Joukowski theorem ( $L = \rho V \Gamma$ ). According to Helmholtz's theorem a vortex filament cannot end in a fluid, hence the vortex filaments along the wing are considered to be horseshoe vortices, which continue unbound behind the wing. The vortices behind the wing form a trailed vortex sheet (as shown in Figure 5.2). (In a typical flow behind a wing, this vortex sheet will roll up to form discrete tip (and in the case of a blade rotor, root) vortices.) The induced downwash on the wing is computed by summing the influence of each trailed vortex along the wing's span<sup>6</sup>.

For fixed-wing aircraft, the wake sheet (or discrete tip vortex) would convect back, away from the wing, hence a series of straight, infinite-length vortex wake filaments would be a reasonable approximation of the wake geometry when computing the downwash. However, in a

<sup>5</sup>Including a rotor blade, which is simply a rotating wing.

<sup>6</sup>A closed-form solution may only be determined if the lift distribution along the wing is known *a priori*.




Figure 5.3: Vortex filaments forming a wake lattice grid, excerpted from [9]

rotor flow, the wake will convect downwards<sup>7</sup> while remaining comparatively “near” the rotor. Hence, further consideration of the wake geometry is required beyond what would be required for a fixed-wing lifting line.

Rather than infinite-length straight lines, consider each filament of the trailed wake to be made up of a series of finite-length segments. These vortex segments would follow the direction of the wake, potentially including rolling up to form the tip vortex, as shown in Figure 5.3. The question then is, how can “the direction of the wake” be determined?

Prescribed vortex wake<sup>8</sup> is one answer to that question. As its name might suggest, in a prescribed wake model the geometry of the wake is specified either *a priori* or based on a few semi-empirical parameters. The advantage of prescribing the wake geometry is that it is cheaper to do so than to determine the geometry dynamically. Plainly however, the prescribed geometry can only be an approximation of the actual wake flow, and if the assumptions of this approximation are violated then so too will its validity.

<sup>7</sup>And in the direction of the free-stream flow, but because this study is only focused on hover, the only component of the flow will be the rotor’s own inflow.

<sup>8</sup>Or simply “prescribed wake”, as “vortex” is understood to be implied.

RCAS implements a few different means of prescribing the wake geometry, the simplest of which is known as classical prescribed wake. In the classical prescribed wake model, the whole wake is modeled as a vortex lattice grid (Fig. 5.3a). Let  $\mathbf{r}^{(ll)}$  be a bound lifting-line vortex segment extending between points  $\{r_1^{(ll)}, z_1^{(ll)}, \psi_1^{(ll)}\}$  and  $\{r_2^{(ll)}, z_2^{(ll)}, \psi_2^{(ll)}\}$ , with coordinates defined in the hub-frame polar coordinate system. In hover, classical prescribed wake defines the coordinates of the resulting wake filament  $\mathbf{r}^{(w)}$  as:

$$\begin{aligned} r_i^{(w)} &= \sqrt{\left(r_1^{(ll)}\right)^2 + \left(r_2^{(ll)}\right)^2} \\ z_i^{(w)} &= v_i t_i + \sqrt{\left(z_1^{(ll)}\right)^2 + \left(z_2^{(ll)}\right)^2} \\ \psi_i^{(w)} &= \omega t_i \end{aligned} \tag{5.2}$$

Where  $v_i$  is the mean induced velocity at the rotor disk and  $t_i = \Delta t(i - 1)$  is the time at step  $i$  after the wake is shed based on the time step size  $\Delta t$ . One limitation of the classical prescribed wake model is that it specifies a *cylindrical* wake (i.e. one without any contraction) since  $r_i^{(w)}$  is constant with time. Additionally, no rollup is specified in the classical prescribed wake model. Still, classical prescribed wake can produce accurate results, and is suitable for modeling the flow field near the blade before the age at which the tip vortex would be expected to coalesce.

RCAS also offers the CAMRAD prescribed wake model, as an alternative to the classical model. The CAMRAD model was developed to account for cases where there may be regions of both positive and negative lift (hence  $\Gamma$  may have both positive and negative peaks) along the span of the rotor blade, such as in high-speed forward flight and some tiltrotor applications. It may however still be used in more conventional cases. Aside from being able to account for negative circulation, another key distinction between the CAMRAD and classical wake models

is that the CAMRAD model represents the roll-up of the trailed vortex sheet into the far wake (as in Fig. 5.3b).

In the CAMRAD prescribed wake model, the wake is divided into three modeling zones:

- The *near wake* immediately emanating from the blade, which is modeled using the full vortex lattice grid to ensure that radial and azimuthal variations of the wake near the blade can be fully captured.
- The *rolling-up wake*, which is modeled with a line segment for the tip vortex and a wake sheet for the inboard wake, and primarily serves as a transition region rather than a rigorous physical model.
- The *far wake*, which models the strong tip vortex as a line segment, but also includes an inboard wake sheet.

The CAMRAD model gives the geometry of the tip vortex as a vector  $\vec{r}_w$  in the tip-path plane reference frame as:

$$\vec{r}_w(\psi, \phi) = \vec{r}_b(\psi, \phi) + \phi \vec{\mu} + \vec{D}(\psi, \phi) \quad (5.3)$$

Where  $\psi$  is the current azimuth angle of the blade and  $\phi$  is the wake age. The distortion  $\vec{D}(\psi, \phi)$  is defined as:

$$\vec{D}(\psi, \phi) = D_r \cos(\psi - \phi) \hat{b}_1^T + D_r \sin(\psi - \phi) \hat{b}_2^T + D_z \hat{b}_3^T \quad (5.4)$$

Where  $\hat{b}_{(1,2,3)}^T$  are the base unit vectors in the tip-path plane. The axial distortion factor is:

$$D_z = \begin{cases} -K_1 \phi & \phi \leq \frac{2\pi}{N} \\ -K_1 \frac{2\pi}{N} - K_2 \left( \phi - \frac{2\pi}{N} \right) & \phi \geq \frac{2\pi}{N} \end{cases} \quad (5.5)$$

Where:

$$K_1 = f_1 \lambda_i \text{ and } K_2 = f_2 \lambda_i \quad (5.6)$$

And the radial (contraction) distortion factor is:

$$D_r = -(1 - K_4) (1 - e^{-K_3 \phi}) r_i \quad (5.7)$$

It is necessary to specify the constant factors  $f_1, f_2, K_3, K_4$  (or alternately,  $K_1, K_2, K_3, K_4$ ) for the CAMRAD prescribed wake model<sup>9</sup>. While RCAS allows user definition of these factors, it also offers semi-empirical calculation of these parameters based on prior studies. In particular, Landgrebe's method gives the parameters as:

$$\begin{aligned} K_1 &= 0.25 (C_T + 0.001\theta_1) \\ K_2 &= f_2 (1 + 0.01\theta_1) \sqrt{C_T} \\ K_3 &= 0.143 + 27C_T \\ K_4 &= 0.78 \end{aligned} \quad (5.8)$$

Where  $\theta_1$  is the blade twist rate in degrees. Kocurek and Tangler offered a modification based on experimental data on low-aspect-ratio rotors as:

$$\begin{aligned} K_1 &= B + C \left( \frac{C_T}{N^n} \right)^m \\ K_2 &= f_2 \sqrt{C_T - N^n \left( -\frac{B}{C} \right)^{\frac{1}{m}}} \end{aligned}$$

---

<sup>9</sup>Observe that  $f_1 = f_2 = 1$  and  $K_3 = K_4 = 0$  would recover the classical prescribed wake model.

$$K_3 = 4.0\sqrt{C_T} \quad (5.9)$$

Where

$$\begin{aligned} B &= 0.000729\theta_1 \\ C &= 2.3 - 0.206\theta_1 \\ m &= 1.0 - 0.25e^{0.040\theta_1} \\ n &= 0.5 - 0.0172\theta_1 \end{aligned} \quad (5.10)$$

RCAS also offers the option to compute the parameters of either the Landgrebe or Kocurek and Tangler models based on the maximum bound circulation, rather than  $C_T$ .

### 5.2.1.3 Free Wake

Like prescribed wake, free vortex wake<sup>10</sup> models begin with a lifting-line representation of the blades. However, free wake models seek to compute “the direction of the wake” at runtime. This computation is based on the vorticity transport theorem which (with some simplifications, including the assumption of potential flow) states that elements of vorticity will convect along with the flow, hence:

$$\frac{d\vec{r}}{dt} = \vec{V}_{total} = \vec{V}_{induced} + \vec{V}_{freestream} \quad (5.11)$$

---

<sup>10</sup>Again “vortex” may be understood to be implied and the model would be referred to simply as “free wake”.




Figure 5.4: Vortex particle (“blob”) representation of the wake, excerpted from [9]

Where  $\vec{V}_{induced}$  arises from the other vortices in the flow field. In the blade-fixed coordinate system, the above equation can be represented as:

$$\frac{\partial \vec{r}}{\partial \psi} + \frac{\partial \vec{r}}{\partial \zeta} = \frac{\vec{V}_{local}}{\Omega} \quad (5.12)$$

Where  $\psi$  and  $\zeta$  are the blade azimuth and wake age, respectively. For a numerical solution, the partial derivatives are discretized into finite differences, hence:

$$D_{\psi} + D_{\zeta} = \frac{\vec{V}_{local}}{\Omega} \quad (5.13)$$

The difference operators may be determined using a periodic or time-accurate method. Periodic difference operators are relevant only to steady flight conditions<sup>11</sup>, while time-accurate operators make the free wake method suitable for maneuver as well.

#### 5.2.1.4 VVPM

A key limitation of conventional/vortex-line-based free wake models is the assumption of potential flow, which does not allow for accurate determination of higher-order or viscous effects. As an alternative to vortex-line models, vortex-particle models have been developed (e.g. Fig. 5.4) which allow for a more physical, mesh-free representation of the wake. The particular formulation of the viscous-vortex particle method (VVPM) available for RCAS is an add-on to FLIGHTLAB, and is described in [46]. The governing equation for the VVPM is the vortex velocity form of the Navier-Stokes equation:

$$\frac{d\omega}{dt} = \omega \cdot \nabla \mathbf{u} + \nu \nabla^2 \omega \quad (5.14)$$

Where  $\omega$  is the vorticity field and  $\nu$  is the kinematic viscosity. In the vortex-particle method, the vorticity field is represented by a set of  $N$  vortex particles, i.e.:

$$\omega(\mathbf{x}, t) = \sum_{i=1}^N \xi_{\sigma}(\mathbf{x} - \mathbf{x}_i) \alpha_i = \sum_{i=1}^N \xi_{\sigma}(\mathbf{x} - \mathbf{x}_i) \omega_i V_i \quad (5.15)$$

Where  $i$  is the index of the particle having position  $\mathbf{x}_i$ , vector-valued total vorticity  $\alpha_i$ , and volume  $V_i$ , and  $\xi_{\sigma}(\mathbf{x})$  is a distribution function. With the vortex particle representation (Eq. 5.15), the governing equation 5.14) may be rewritten as:

$$\frac{d\alpha}{dt} = \alpha \cdot \nabla \mathbf{u} + \nu \nabla^2 \alpha, \text{ where } \mathbf{u}(\mathbf{x}, t) = - \sum_{j=1}^N \frac{1}{\sigma_j^3} K(\rho)(\mathbf{x} - \mathbf{x}_j) \times \alpha_j \quad (5.16)$$

---

<sup>11</sup>Such as hover, as is being examined in this study.

Where  $\rho$  is a nondimensional distance parameter and  $\sigma$  is a smoothing parameter. Similarly to the line-based free wake model, the movement of the vortex particles is governed by the convection of the flow, i.e.:

$$\frac{d\mathbf{x}_i}{dt} = \mathbf{u}(\mathbf{x}_i, t) \quad (5.17)$$

Unlike the line-based free wake model however, VVPM is able to account for vortex stretching (the  $\omega \cdot \nabla \mathbf{u}$  term in Eq. 5.14) and viscous diffusion (the  $\nu \nabla^2 \omega$  term in Eq. 5.14). The vortex core stretching is computed directly, i.e.:

$$\left. \frac{d\alpha_i}{dt} \right|_{ST} = [\nabla \mathbf{u}(x_i, t)] \cdot \alpha_i \quad (5.18)$$

Where the velocity gradient is determined from the particle discretization. The viscous diffusion is computed using the particle strength exchange method, where the Laplacian is represented approximately as an integral, i.e.:

$$\nabla^2 \omega(\mathbf{x}) \approx \frac{2}{\sigma^2} \int \eta_\sigma(\mathbf{x} - \mathbf{y}) [\omega(\mathbf{y}) - \omega(\mathbf{x})] d\mathbf{y} \quad (5.19)$$

The above expression is discretized using the vortex particle discretization, then integrated over the volume of the particle.

In the particular implementation available for RCAS, VVPM is an external/add-on code and behaves slightly differently to the other inflow models. Additional scripts are required to configure the model for the VVPM code. Additionally, the interface between RCAS and the VVPM code precludes the use of RCAS's CFD/CSD coupling interface—hence, it is necessary to use the method described in Section 5.3.3 if the 2D interactional loads are to be injected into

the VVPM simulation.

### 5.2.1.5 Interference Models

While technically a separate component from the inflow models, it is useful at this stage to consider the aerodynamic interference model, which accounts for interaction between the upper and lower rotors. (Without an interference model, RCAS would consider each rotor in isolation, which is plainly incorrect for the components of a stacked rotor.) RCAS offers four<sup>12</sup> interference models relevant to rotor-to-rotor interactions:

1. Simple<sup>13</sup> rotor-to-rotor
2. Simple horseshoe vortex
3. Cylindrical vortex sheet
4. Discrete vortex wake

In the *simple rotor-to-rotor* model, it is assumed that the “lower” rotor is far enough separated from the “upper” rotor that the upper rotor’s wake is fully developed, hence (following momentum theory) the mean induced velocity in the wake is twice the mean induced inflow velocity. The wake velocity of the upper rotor is imposed as a contribution to the mean inflow on the lower rotor; the contribution is scaled as a function of the overlap between the rotors. This offers about as good an approximation of the rotor-to-rotor interference as momentum theory does of the rotor inflow, which is to say a “first order” estimate at best.

---

<sup>12</sup>Two more interference models are implemented in RCAS, but are relevant only to solid bodies or ground effect, not to rotors.

<sup>13</sup>The use of “simple” in the names of these interference models is RCAS’s terminology, not commentary on the part of the author.

The *simple horseshoe vortex* model represents the entire lifting rotor as a single horseshoe vortex. The strength of the vortex is set by replacing lift in the Kutta-Joukowski theorem with the thrust, i.e.  $T = \rho VT$ . The bound vortex line is placed on the rotor disk perpendicular to the direction of the free-stream flow, and the trailed vortices emanate perpendicular to the bound vortex at an angle corresponding to the wake skew. This is a reasonable approximation of a rotor in forward flight, where the blade tip vortices will tend to roll up on the edges of the wake, but not so much for hover where this roll-up does not occur.

As its name suggests, the *cylindrical vortex sheet* model represents the wake of a rotor as a cylindrical sheet of vortices, emanating from the rotor. The rotor itself is represented as having an infinite number of blades, such that the vorticity along the surface of the cylinder is constant, which enables a closed-form solution for the velocity induced by the wake to be found at any point in the flow.

Finally, the *discrete vortex wake* model determines the influence of one rotor on another by using the wake computed by an appropriate inflow model—i.e. prescribed or free vortex wake. It should be clear that the discrete vortex wake interference model is the only one which can account for the azimuthal variation in the wake, since all other models represent the rotor as a uniform disk. Hence, discrete vortex wake is the only interference model suitable for representing rotor-to-rotor interactions in the stacked rotor. One limitation of this model is that the interference is computed based only on the *wake* vortices, not the *bound* vortices. Hence if using this interference model, some way of accounting for this bound circulation interaction would be required (see Chapters [2.3.1.2](#) and [3.1](#)).

Note that unlike the other inflow models included in RCAS, the FLIGHTLAB implementation of VVPM models vortex interactions on the entire system, hence it does not require an

interference model. At one point the VVPM implementation did not include bound-circulation interaction between the blades, but this was corrected in a later version [3].

#### 5.2.1.6 Down-Selection

Having described the inflow and interference models, it is now time to consider which models will be used in this study. Plainly, momentum theory and dynamic inflow would not be suitable since they do not model the wake of the individual blades. Alternately, it could be stated that the simple rotor-to-rotor, simple horseshoe vortex, and cylindrical vortex sheet interference models are unsuitable for the same reason, leaving only discrete vortex wake as a suitable interference model, which is not compatible with either momentum theory or dynamic inflow. Thus, the available models which could adequately represent the physics of the stacked rotor are:

- Prescribed vortex wake inflow with discrete vortex wake interference
- Free vortex wake inflow with discrete vortex wake interference
- Viscous vortex particle method

Ideally, the free wake model should represent more of the physics of the system than the prescribed wake. And indeed, RCAS with free wake inflow has been used in prior studies to model the stacked rotor system [5]. In practice however, it was found that it was difficult to adjust the free wake model in a way that would achieve a meaningful result. Hence, free wake would not be preferred for ease-of-use, although future researchers could consider experimenting with this model themselves.

VVPM comes the closest of the available models to representing the actual physics of the system, including modeling bound circulation effects (which are not included in the discrete vortex wake interference model). VVPM has also been used in prior studies on stacked rotors [3,4], including by the author [47]. A major limitation of VVPM however is that it is incompatible with RCAS’s CFD/CSD coupling interface (this interface will be described further in Section 5.2.3). Additionally, the VVPM module is separate from RCAS, and requires its own license. Finally, the VVPM module is implemented as a GPU code which helps to reduce the computational time required, but hence when running multiple rotor simulations in parallel requires multiple GPUs (which is a less common configuration than multiple CPUs) to avoid significant slowdown.

Instead, the prescribed vortex wake model was chosen for use in the current work. While prescribed wake is the simplest possible model which could be used to represent the physics of the stacked rotor system, this could be seen as a positive—in the next chapter, it will be shown that the current methodology using 2D interactional airloads is able to produce accurate results even with the comparatively simple underlying model. Further, during testing it was noted that the cylindrical prescribed wake actually produced *more accurate* results than the CAMRAD model with either Landgrebe or Kocurek and Tangler deformation coefficients<sup>14</sup>. While the cylindrical wake model does not model the roll up of the tip vortex, it should be noted that retaining the vortex lattice field near the blades (as the CAMRAD model does) is necessary to ensure that the wake distribution immediately behind the blade is correctly modeled. Since the most interaction between the blades in the stacked rotor will occur when they are close together, the vortex lattice wake is thus reasonable.

---

<sup>14</sup>In both cases, it was found that basing the models on maximum bound circulation produced “better” results than basing the models on  $C_T$ , but still worse results than the cylindrical case.

### 5.2.2 Structural Model

The current study is concerned only with the aerodynamic design of a stacked rotor's blades, not with their structural design. However, some representation of the structure of the blades is necessary—even for aerodynamic design—so that the relative positions of the blades may be properly determined. This is especially true for hingeless rotors<sup>15</sup>, where both the stiffness of the flexure and the blades themselves are important to the blade motions. Hence, while the structural model will not be investigated, it is still pertinent to consider its setup and solution.

RCAS fundamentally uses a finite-element model for structural analysis. Like the aerodynamic model the structural model is hierarchical, with each subsystem having a primitive structure, and the primitive structure being assembled from various nodes and elements. Within RCAS, each element's equations of motion are expressed in the general form:

$$q_e(u, w, t, x_e, \dot{x}_e, \ddot{x}_e) = 0 \quad (5.20)$$

Where  $q_e$  is the residual (zero at equilibrium) for the element  $e$ ,  $u$  is the external forcing,  $w$  is the wind,  $t$  is time, and  $x_e, \dot{x}_e, \ddot{x}_e$  are the degrees of freedom and their first and second derivatives with time. For linear elements, the above equation can be expressed as:

$$q_e = F_e - M_e \ddot{x}_e - C_e \dot{x}_e - K_e x_e \quad (5.21)$$

Where  $F_e$  is the element force vector, and  $M_e$ ,  $C_e$ , and  $K_e$  are the mass, damping, and stiffness

---

<sup>15</sup>Hingeless rotors are somewhat common in design modern concepts for stacked rotors where distributed propulsion demands rotors in a smaller size than “conventional” main lifting rotors, see e.g. Figure 1.1.

matrices respectively. RCAS considers each element to be rigidly bound to some “parent” nodes and to have degrees of freedom only at its “free” nodes, hence the “assembly” process for the residuals and degrees of freedom consists only of collecting those for each element:

$$\{Q\} = \begin{Bmatrix} q_e^{(1)} \\ q_e^{(2)} \\ \vdots \\ q_e^{(n)} \end{Bmatrix} \text{ and } \{X\} = \begin{Bmatrix} x_e^{(1)} \\ x_e^{(2)} \\ \vdots \\ x_e^{(n)} \end{Bmatrix} \quad (5.22)$$

If there are nonlinear elements, the matrices for the system must be re-computed using numerical perturbation, where:

$$M = \frac{\partial Q}{\partial \ddot{X}} \text{ and } C = \frac{\partial Q}{\partial \dot{X}} \text{ and } K = \frac{\partial Q}{\partial X} \quad (5.23)$$

Such that the coupled system of structural equations to be solved is:

$$Q = F - M\ddot{X} - C\dot{X} - KX \quad (5.24)$$

RCAS defines several types of structural elements in its library (including springs, dampers, slides, hinges, and rigid bars or masses) each with their own equations of motion as appropriate [42]. In particular, rotor blades are modeled (in RCAS generally and this study specifically) using nonlinear beam elements. The nonlinear beam element connects two nodes, one of which is the parent node and one of which is the free node—hence over the span of a rotor blade, there will be one node at the root with defined position, while all other nodes have equations of motion determined by one segment. Each free node in the nonlinear beam element has six degrees

of freedom: three positional (enlongation, flap position, lead/lag position) and three rotational (torsion, flap angle, lead/lag angle). The beam sectional properties ( $EA$ ,  $EI$ ,  $GJ$ , mass, etc.) must be specified for each element<sup>16</sup>. At runtime, RCAS uses these properties and a specified set of shape functions and number of Gaussian quadrature points to compute the mass, damping, and stiffness matrices from the strain-displacement relationship.

It should be noted that RCAS does not require the aerodynamic and structural nodes to be located at the same coordinates<sup>17</sup>. In the structural model, RCAS applies the airloads to the beam by extending massless rigid bar elements from the aerodynamic control points to the structural Gauss points. In the aerodynamic model, the deflections of the aerodynamic control points are determined by evaluating the beam deflections at the appropriate radial station.

### 5.2.3 CFD/CSD Coupling Interface

In addition to its own aerodynamic models, RCAS offers the ability to couple with external CFD codes. The intent of the coupling is to allow the airloads (particularly interactional airloads) to be computed with greater fidelity than is possible through RCAS's own aerodynamic models by using a high-fidelity 3D CFD code. Since most CFD codes do not include a structural model, RCAS acts as a computational structural dynamics (CSD) code, computing the deflections of the structure, which are passed to the CFD code to be used as mesh motions and deformations. An appropriate number of iterations of CFD and CSD are performed until a converged CFD/CSD solution is achieved. Several implementations of the RCAS CFD/CSD coupling interface have been developed at various times [42], but FILEIO2018 is the particular version used in this

---

<sup>16</sup>In the current work, the beam sectional properties are constant along the span of the blade, hence all segments have the same properties.

<sup>17</sup>Although the aerodynamic and structural node coordinates should, for obvious reasons, both refer to points *along the same line*.

work [48], and the description of the interface is taken from its documentation.

Although the CFD/CSD coupling interface was intended for use with high-fidelity 3D CFD codes, two key features of the interface make it a useful handle for injection of the 2D interactional airloads:

- The `FILEIO2018` interface is designed to work with a *generic* 3D CFD code, and as such treats the 3D code itself as a “black box”.
- Despite expecting CFD airloads, RCAS still computes its own aerodynamic solution first, then uses the CFD airloads as a *correction* on subsequent iterations.

Further, the `FILEIO2018` interface is included as part of the basic RCAS distribution, meaning that any current RCAS user should be able to utilize tools and models invoking it. Of course, an understanding of the interface data and routines is important to develop tools which utilize the interface.

The CFD/CSD coupling consists of three phases:

1. Setup/initialization
2. RCAS  $\rightarrow$  CFD position exchange
3. CFD  $\rightarrow$  RCAS airloads exchange

The first phase is performed only at the start of the rotor simulation, while the second and third phases are repeated for some number of coupling iterations.

When the coupling is initialized, RCAS and the CFD code cross-check to ensure that each model is correctly configured, with the same number of coupled aerodynamic systems (rotors/blades, wings, etc.) and the same expected number of coupling iterations. RCAS and CFD

also exchange the spatial and temporal points at which each will compute their solutions. The coupling interface is designed such that these points do not need to match, since it is expected (but not required) that the CFD code will use much finer spatial and temporal discretization than RCAS.

In the position exchange step, RCAS expresses the motions of the rigid bodies (reference frame motion) and deformations (within those reference frames), at the control points specified by the CFD code. Displacement is expressed as a 3-component vector  $(x, y, z)$ , while rotational deformation is expressed redundantly in terms of Euler angles and their corresponding rotation matrix. Accelerations of the rigid body are also exchanged, but are not used in this study. For rotor coupling, the rigid body motions include the motion of the non-rotating rotor reference frame with respect to the inertial frame, and the motion of the blade reference frame with respect to the rotor reference frame. Deformations of the blades are expressed within the blade reference frame. The user may also manually specify a reference frame for the data exchange.

Airloads exchange is performed at the RCAS aerodynamic control points (ACPs). The airloads exchanged are dimensional forces and moments (3 components each, 6 components total), using the unit system agreed upon by RCAS and the CFD code<sup>18</sup>. It is important to note that the reference frame for both the loads and moments is the undeformed blade reference frame, but that the reference frame is attached to the *deformed* position of the blades—i.e. there is no moment computed due to the force acting on the blade at a position offset from the undeformed coordinates.

Two coupling modes are offered as part of the interface: loose coupling (LCP) and tight coupling (TCP). Loose coupling is suitable for steady conditions and is used exclusively in this

---

<sup>18</sup>In this study the unit system is US Customary/ft-slug-sec, hence the airloads are expressed in slugs.

study, while tight coupling is suitable for cases with maneuver. In loose coupling, the aerodynamic and structural solutions are computed to a converged periodic steady-state, then the data exchange is performed, and the process is repeated for several iterations. Data is exchanged at the same time for all time steps and control points. In contrast, data is exchanged in tight coupling at either every time step or sub-iteration. Tight coupling (such as for a maneuver) may be preceded by some number of loose coupling iterations (e.g. to establish a steady-state trim solution).

RCAS computes an initial set of motions and deformations for the CFD/CSD coupling by following its normal aerodynamic and structural solution routines, as described previously. On subsequent LCP iterations, RCAS computes the difference between its airloads and the CFD airloads (called  $\Delta$ -airloads), then applies the  $\Delta$ -airloads as an “external” force to its ACPs. Hence, despite taking account of the CFD airloads, RCAS still must compute its own aerodynamic solution at each iteration. In TCP however, RCAS entirely replaces its own airloads with those computed by the CFD code. Hence, *only* LCP will be suitable for the injection of the two-airfoil 2D CFD data.

As might be suggested by its name, data exchange using the `FILEIO2018` interface is performed through ASCII files, with the particular file formats specified by RCAS [48]. In addition a series of flag files are used, the presence of which signifies the readiness of the CFD and CSD codes to have their data read by the other.

Not that while RCAS and the CFD code exchange positional and loads data, they do not exchange any information about the flow around the rotor itself—under the assumption that the high-fidelity CFD code does not need any low-fidelity aerodynamic information from RCAS, and that RCAS only needs updated loads to update the structural deformations. If information about the flow is required to be exchanged, that exchange must be defined by the user through custom

RSCOPE scripts.

### 5.3 Multi-Airfoil Coupling Code

Previous chapters have discussed how the blade positions and aerodynamic conditions are obtained for a “generic” rotor model (Chapter 2), and how that data can be used to determine the 2D loads for the corresponding system of two airfoils (Chapter 3). This section will discuss how in particular how the data from RCAS is used to obtain the 2D parameters, and how the resulting 2D airloads are coupled back to RCAS. With appropriate scripts, the necessary data can be obtained *from* RCAS with or without the CFD/CSD coupling interface. The choice of coupling methodology has a more significant impact on the method by which the loads are provided back *to* RCAS. Aside from those functions particular to input parsing and output formatting<sup>19</sup>, the internal routines of the code are identical regardless of the coupling method being used.

For convenience, the 2D airfoil data in the multi-airfoil coupling code is computed at positions along the span of the rotor blade which match the aerodynamic control points (ACPs) in the RCAS model. This is not strictly necessary for either coupling method, and in fact coupling without the CFD/CSD interface requires interpolating the resulting data onto a different series of exchange points. However, when using RCAS’s CFD/CSD coupling interface, this ensures that neither is resolution lost if the 2D loads were to be computed on a coarser set of points, nor false resolution gained if the loads were computed on a finer set of points (for which RCAS output data would need to be interpolated). To ensure consistency between the RCAS model and the multi-airfoil code, the multi-airfoil code directly reads the positions of the aerodynamic nodes and segments of the blades (and the relative positions and orientations of the two rotors) from the

---

<sup>19</sup>Not just data formatting, but specifying and computing *which* data is necessary to output

RCAS input file.

All of the computations are also performed at every azimuthal (time) step of the rotor simulation. Time-averaging is only performed as a post-processing step. Note however that any relative motion of the blades is ignored, as the 2D CFD simulations are taken in steady conditions. This is reasonable within the scope of this study, as the unsteady blade motions should be close to zero in hover.

### 5.3.1 Airfoil Positions and Flow Conditions

To obtain the relative position of the airfoils it is necessary to obtain the positions of the rotor blades in a common reference frame. A logical reference frame for this is the hub-fixed polar coordinate system, with coordinates defined as  $(r, \psi, z)^{(hp)}$  where  $r$  is zero at the root and positive outboard,  $\psi$  is positive in the direction of rotation, and  $z$  is positive up. The horizontal separation between the blades is taken along a line of constant radius<sup>20</sup> (as in Fig. 2.1a) and the vertical separation must be defined positive with the lower rotor's airfoil below the upper, hence:

$$\begin{aligned}\Delta x^{(hp)} &= \bar{r} \Delta \psi^{(hp)} = \frac{r_{upper}^{(hp)} + r_{lower}^{(hp)}}{2} \left( \psi_{upper}^{(hp)} - \psi_{lower}^{(hp)} \right) \\ \Delta z^{(hp)} &= z_{upper}^{(hp)} - z_{lower}^{(hp)}\end{aligned}\tag{5.25}$$

While it is assumed that the upper and lower rotor blades' ACPs are specified at the same radial stations along the blades, the deformed radial positions may not be exactly identical. Hence to avoid biasing one rotor over the other, the mean radial position  $\bar{r}$  is used. Once the positions of the airfoils are obtained in the hub frame (Eq. 5.25), they may be rotated to the upper airfoil

---

<sup>20</sup>I.E. *not* the shortest linear distance.

frame (Eq. 2.1) for use in the interpolator as described.

For both coupling methods, it is necessary to add a script to the RCAS model to write the aerodynamic conditions ( $\alpha$  and  $M$ ) along the blade. This is done simply by saving the necessary data using a user-defined RSCOPE script. In particular, this user-defined script was implemented as `useradditionalendstep.exe` because it was found that this script would execute prior to the CFD/CSD coupling iteration. Any user-defined script which would execute later in the RCAS solution process would do so after CFD/CSD coupling was completed. The user-defined script `useradditionalendstep.exe` actually executes several times as RCAS iterates before the CFD/CSD coupling is performed, but the script was written such that updated output files would overwrite files from the prior iteration, such that when the coupling is performed only the files from the current (final, converged) iteration would exist.

Based on the  $\alpha$  on each blade, the free-stream angle of attack ( $\alpha_\infty$ ) and effective pitch offset between the airfoils ( $\Delta\theta$ ) in the two-airfoil system is determined using the method described in either Chapter 2.3.1.1 or 2.3.1.2 as appropriate for the inflow and interference models. The flow speed varies fairly little between the airfoils, being dominated by the rotational speed at a given radius ( $r\Omega$ ), but to avoid bias due to slight variations in flow speed, the free-stream flow is set to the mean velocity, i.e.  $M_\infty = \bar{M} = (M_{upper} + M_{lower}) / 2$ .

If the parameters of the airfoil system ( $\alpha$ ,  $M$ ,  $\Delta x$ ,  $\Delta z$ ,  $\Delta\theta$ ) are within the range for which 2D CFD has been computed<sup>21</sup>, then the 2D CFD airfoil loads will be used. It was found that extrapolation of the 2D CFD data was unreliable, so outside of that range the airfoil loads revert to either the 2D lifting-vortex model or the loads computed by RCAS, again as appropriate for

---

<sup>21</sup>For simplicity, each parameter is checked separately against its own minimum and maximum values in the CFD data.

the inflow and interference models.

It is left as an exercise to the researcher to ensure that the two-airfoil 2D CFD data is generated for a set of airfoils corresponding to those actually used on the rotor blades. The multi-airfoil code also requires an airfoil table for the isolated airfoil, to correctly compute the lifting-vortex interaction. This isolated airfoil table should be consistent with both the isolated airfoil table provided to RCAS and the two-airfoil 2D CFD data.

### 5.3.2 Coupling With CFD/CSD Interface

When coupling using RCAS's CFD/CSD interface, the multi-airfoil lookup code is initialized at the same time as RCAS is started. On start-up, the multi-airfoil code waits for RCAS to signal that it is ready to perform CFD/CSD coupling, reads RCAS's input and data exchange files to set its rotor blade parameters consistently with RCAS, then reports to RCAS that it is ready to proceed. If at any point the multi-airfoil lookup code times out while waiting for RCAS to be ready, the code will send a termination signal to RCAS before exiting.

ACP positions in the hub frame are determined straightforwardly from the model definition and the RCAS coupling data:

$$(x, y, z)_{ACP}^{(hp)} = R_\psi \left( (x, y, z)_{ACP}^{(bf)} + (x, y, z)_{hub} + \text{POSDEF} \right) \quad (5.26)$$

Where  $R_\psi$  is the rotation matrix due to a z-axis rotation by the index angle  $\psi$ ,  $(x, y, z)_{ACP}^{(bf)}$  is the undeformed ACP position in the blade frame,  $(x, y, z)_{hub}$  is the hub (origin) position, and POSDEF is the *position deformation* supplied by RCAS during the coupling. The Cartesian

hub-frame coordinates are converted to polar coordinates using the standard transformation<sup>22</sup>

$$\begin{aligned} r_{ACP}^{(hp)} &= \sqrt{\left(x_{ACP}^{(hp)}\right)^2 + \left(y_{ACP}^{(hp)}\right)^2} \\ \psi^{(hp)} &= \tan^{-1} \left( \frac{y_{ACP}^{(hp)}}{x_{ACP}^{(hp)}} \right) \end{aligned} \quad (5.27)$$

These values may be used in Equation 5.25, and the airfoil sectional forces ( $c_l, c_d, c_m$ ) may then be determined using the interpolation scheme described previously.

To couple the loads back to RCAS, it is necessary to compute the dimensional airloads ( $F, M$ ) in the correct reference frame from the airfoil sectional loads. First, the force coefficients<sup>23</sup> are expressed in the airfoil local frame as:

$$\begin{aligned} \begin{bmatrix} c_x \\ c_z \end{bmatrix}_{upper} &= \begin{bmatrix} \cos(\alpha_\infty) & -\sin(\alpha_\infty) \\ \sin(\alpha_\infty) & \cos(\alpha_\infty) \end{bmatrix} \begin{bmatrix} c_d \\ c_l \end{bmatrix}_{upper} \\ \begin{bmatrix} c_x \\ c_z \end{bmatrix}_{lower} &= \begin{bmatrix} \cos(\alpha_\infty + \Delta\theta) & -\sin(\alpha_\infty + \Delta\theta) \\ \sin(\alpha_\infty + \Delta\theta) & \cos(\alpha_\infty + \Delta\theta) \end{bmatrix} \begin{bmatrix} c_d \\ c_l \end{bmatrix}_{lower} \end{aligned} \quad (5.28)$$

Next, the forces and moments in the airfoil/blade-fixed (i.e. deformed) reference frame are computed for each aerodynamic segment as:

$$F_{seg}^{(def)} = \begin{bmatrix} f_x \\ f_y \\ f_z \end{bmatrix}_{seg} = \frac{1}{2} \rho u_\infty^2 S_{seg} \begin{bmatrix} 0 \\ -c_d \\ c_l \end{bmatrix}_{seg}$$

---

<sup>22</sup>Specifically, the Python function `math.atan2()` is used instead of simply `math.atan()` to ensure that  $\psi$  is correctly defined in the range  $[-\pi, \pi]$ .

<sup>23</sup>The sectional moment coefficient is not effected by the change in the reference frame

$$M_{seg}^{(def)} = \begin{bmatrix} m_x \\ m_y \\ m_z \end{bmatrix}_{seg} = \frac{1}{2} \rho u_\infty^2 S_{seg} c_{seg} \begin{bmatrix} c_m \\ 0 \\ 0 \end{bmatrix}_{seg} \quad (5.29)$$

Where  $S_{seg}$  is the planform area of the aerodynamic segment, i.e. the chord ( $c_{seg}$ ) multiplied by the span of the particular segment. Notice that in the above calculation, the reference frame was changed from the airfoil (x/drag back, z/lift up) to the blade (x outboard, y forward, z up) frame. Finally, the forces and moments are transformed from the deformed frame to the undeformed blade frame as:

$$\begin{aligned} F_{seg} &= R_{ANGDEF} R_{\beta_p}^{-1} F_{seg}^{(def)} \\ M_{seg} &= R_{ANGDEF} R_{\beta_p}^{-1} M_{seg}^{(def)} \end{aligned} \quad (5.30)$$

Where  $R_{ANGDEF}$  is the transformation matrix defined by the *angular deformation* supplied by RCAS during the coupling (ANGDEF), and  $R_{\beta_p}^{-1}$  is the inverse of the transformation matrix defined by a rotation by the precone angle ( $\beta_p$ ) about the y-axis. Note that because ANGDEF defines the rotation of the *structure*, the force transformation involves the forwards (*not* inverse) transformation matrix.

The loads resulting from Equation 5.30 at each segment and azimuthal step are written to the required coupling file, then a flag file is generated to indicate to RCAS that the CFD airloads are ready to be read. This completes one LCP iteration. Unless the multi-airfoil lookup code expects the current iteration to be the last, the code then returns to the beginning of the coupling loop to wait for the updated deflections. After the final LCP iteration, the multi-airfoil lookup

code exits. Rotor performance metrics (thrust, power, etc.) are computed by RCAS and output at the completion of the simulation.

### 5.3.3 Coupling Without CFD/CSD Interface

It is also possible to inject the 2D CFD airloads without going through RCAS's CFD/CSD interface. This procedure requires additional computational time and may not be as physical, but it does enable coupling to be performed with any RCAS inflow model, including VVPM.

This coupling procedure begins by running an "standard" RCAS simulation to convergence, using airfoil tables for the isolated airfoil as normal. The positions and rotations of the blade structural nodes are selected as an output in RCAS, along with the aerodynamic conditions along the blades<sup>24</sup>. In this case, the blade position data includes the specified structure and the deflections, so it may be plugged directly into Equation 5.27.

After the first RCAS simulation is completed, the coupling code is called. The coupling code reads the positions and aerodynamic conditions, interpolates the structural positions onto the ACPs, and determines the airfoil sectional forces from the 2D CFD data as described previously. The question thus is how these loads may be used by RCAS on subsequent iterations.

Consider that in hover, there is no in-plane velocity across the rotor. Hence, the flow speed seen by the airfoils is strictly related to their radial (*not* azimuthal) position, and vice-versa:

$$M_{\infty} = \frac{r\Omega}{a} \Leftrightarrow r = \frac{M_{\infty}a}{\Omega} \quad (5.31)$$

Hence,  $M_{\infty}$  may be used as a proxy for the radial position along the blades. Based on this,

---

<sup>24</sup>The same script as for the CFD/CSD coupling case may be used to output these , but there are other means of specifying this data to be output from a converged RCAS simulation.

the distribution of the airfoil forces along the span of the blades may be expressed instead as a function of the free-stream Mach number:

$$c_{l,d,m}(r) \Leftrightarrow c_{l,d,m}(M_\infty) \quad (5.32)$$

This formulation should be familiar, since for an isolated airfoil the sectional forces are expressed in the airfoil table as a function of the angle of attack *and the Mach number*. Hence, the theory of this coupling methodology is that adjustment to the forces along the span of the blade can be achieved by adjusting the forces in the airfoil table at the corresponding Mach numbers. Specifically, the procedure followed is thus:

1. Compute the expected airfoil forces on each element as described, along the span of the blade ( $c_{l,d,m}(r)$ )
2. Express the variation of the airfoil forces along each element in terms of the Mach number ( $c_{l,d,m}(M)$ ), and interpolate<sup>25</sup> to obtain the airfoil forces at the particular Mach numbers contained in the isolated airfoil table ( $M_{AT}$ )
3. Compute the difference between the expected airfoil loads and the airfoil loads actually used by RCAS ( $\Delta c_{l,d,m}(M_{AT})$ )
4. Define a new airfoil table *for each element*, with the difference in airfoil loads added to the previous airfoil table at the corresponding Mach numbers for every angle of attack
5. Write a new RCAS input file containing the updated airfoil tables

---

<sup>25</sup>Or extrapolate near the root and tip.

RCAS is then run again using the updated input file. This process of running RCAS, post-processing, and generating a new RCAS input, is run for a specified number of iterations. A shell script (created automatically by a pre-processor) is used to automate the iteration process.

Obviously, this procedure has some significant limitations. Because several RCAS simulations must be performed, with each simulation being fully converged, this process takes more time than coupling with the CFD/CSD interface. At a physical level, this process also assumes that all of the lift in the system is circulatory about each airfoil uniquely—while this assumption is not necessarily unreasonable, it is also not rigorous. Whether these limitations are a significant drawback is debatable, and it will be shown in the next Chapter that this procedure can produce accurate results in some cases but inaccurate results in others.

## 5.4 Summary

This chapter has laid out the methods and procedures associated with lower-fidelity rotor modeling for the stacked rotor, including coupling with the 2D airfoil interactional loads which were described in Chapters 3–4. RCAS is used as the baseline lower-fidelity rotor modeling code, both due to its ability to include structural and several aerodynamic models, as well as its common use in academia, government, and industry. However, the current methodology could be adapted to other lower-fidelity rotor codes, so long as those codes model the rotor at an appropriate level and include suitable data interfaces.

Previous studies (see Chapter 1) have shown that the aerodynamics of the blades of a stacked rotor are significantly influenced by vorticity trailed from other blades on the rotor. The lower-fidelity aerodynamic models which model (as opposed to approximate) this shed vorticity

include prescribed vortex wake, free vortex wake, and vortex particle wake methods. All three methods use a lifting-line representation of the rotor blades, and are generally suitable for both hover and forward flight, but differ in how specifically they model the geometry of the rotor wake. This study focuses on prescribed wake methods, due both to their easier implementation in RCAS compared to free wake, and being able to use RCAS's CFD/CSD coupling interface unlike the vortex particle wake method. While RCAS offers both classical (cylindrical) and deformed prescribed wake methods, it will be shown in the next Chapter that deformed prescribed wake methods which are suitable for isolated rotors do not produce accurate results for stacked rotors.

RCAS's CFD/CSD coupling interface provides a way for the aerodynamic loads on the blades to be adjusted while the lower-fidelity rotor simulation is performed. Because RCAS considers CFD to be a "black box", and RCAS continues to perform its own aerodynamic calculations, this interface is suitable for inclusion of the 2D airfoil-interactional airloads. In this study, the `Python` code described in Chapter 3.3 was extended to work with this interface. At each coupling iteration, the `Python` code reads the aerodynamic and geometric conditions output from RCAS, then interpolates to determine the 2D aerodynamic loads, and integrates those loads to obtain dimensional 3D airloads, which are finally passed back to RCAS. A similar code was also written which performs coupling through modification of the airfoil tables. However, it will be shown in the next Chapter that airloads coupling through the CFD/CSD interface is generally more reliable than this alternate method.

## Chapter 6: 3D Results

The main focus of this chapter is validation of the lower-fidelity stacked rotor model with the inclusion of the high-fidelity 2D CFD two-airfoil loads. These airfoil loads will mainly be included using RCAS's CFD/CSD coupling interface, and the basic aerodynamic model used in RCAS is the cylindrical prescribed wake. However, comparison will be shown with other models in RCAS. For demonstration, some behaviors of the coupling model will also be shown.

All of the validation is performed using experimental data from Johnson, et.al. [4], hence the RCAS rotor model (modified from Jacobellis, et.al. [3] with the scripts necessary for coupling) is configured to match the experimental setup. A summary of the geometric and aerodynamic parameters of the rotor is given in Table 6.1, and the structural properties of the rotor are described in the experimental studies [49]. Observe that since each rotor has two blades, a phase offset of  $\Delta\psi = \pm 90^\circ$  represents an equal-spaced rotor<sup>1</sup>.

2D CFD data was gathered for the two-airfoil system using the methodology presented in Chapter 3, over the range of parameters shown in Table 6.2. Within the given ranges the distribution of points was not necessarily uniform<sup>2</sup>. A tip Reynolds number of 750,000 was estimated<sup>3</sup>

---

<sup>1</sup>While  $\Delta\psi = +90^\circ$  and  $\Delta\psi = -90^\circ$  are physically the same case and hence should be redundant, data points for both are shown to demonstrate the correct behavior of the methodology.

<sup>2</sup>The distribution of data points was based on observation of the trends in Figures 4.10 and 4.11, with additional simulations performed as necessary. While this was sufficient for the current work, greater efficiency in performing the 2D CFD simulations could perhaps be obtained through the application of an appropriate optimization algorithm.

<sup>3</sup>Lacking specific atmospheric conditions for the lab in [4], sea level conditions were assumed.

| Parameter                      | Value    |
|--------------------------------|----------|
| Radius                         | 3.637 ft |
| Blades per rotor               | 2        |
| Precone                        | 3°       |
| Rotational speed               | 1200 RPM |
| Root cut                       | 18.76% R |
| Chord                          | 3.15 in  |
| Airfoil                        | VR12     |
| Tip Mach number (approx.)      | 0.405    |
| Root cut Mach number (approx.) | 0.076    |
| Vertical offset at hub         | 2.3 in   |

Table 6.1: Parameters of the stacked rotor modeled in this study, from [4, 10]

| Parameter      | Minimum Value | Maximum Value |
|----------------|---------------|---------------|
| $\Delta x/c$   | -5            | 5             |
| $\Delta z/c$   | -1.25         | 1.75          |
| $\Delta\theta$ | -2°           | 13°           |
| $\alpha$       | -2°           | 9°            |
| $M$            | 0.05          | 0.4           |

Table 6.2: Parameters used for 2D CFD two-airfoil data generation

for the rotor, corresponding to the tip Mach number of 0.405. For all other 2D CFD cases, the Reynolds number was scaled linearly with the Mach number, to ensure that the Reynolds number would vary correctly along the span of the rotor blade. The range of  $\Delta x$  was chosen based upon where the 2D CFD and 2D lifting-vortex model were expected to give different results (see e.g. Chapter 4.1.2), while the ranges of the other parameters were chosen to encompass the range of conditions expected to be observed along the span of the rotor.

## 6.1 Without Collective Offset

For the first set of cases, the collective is set to  $\theta_0 = 8^\circ$  with  $\Delta\theta_0 = 0^\circ$  (i.e. for both the upper and lower rotors, collective is  $8^\circ$ ). Prior experiments have shown that, with proper selection of the phase angle, having equal collective on both rotors leads to the greatest efficiency of the rotor system. Since this is a relatively simple case, it was chosen both for validation of the methodology and for a more detailed examination of the methodology's behaviors.

### 6.1.1 Coupling Convergence

Before discussing the accuracy of the methodology, it should be demonstrated that the solution is converged. Through trial and error, it was found that a slight relaxation factor of  $\omega = 0.8$  in the force update would assist with convergence for some cases (at particular  $\Delta\psi$ ), and this relaxation factor is used in all cases.

Convergence of the thrust with coupling iteration on each rotor in the stacked rotor system for various  $\Delta\psi$  is shown in Figure 6.1. Clearly, there is a significant change in the forces between




Figure 6.1: Convergence of thrust on each rotor in the stacked rotor system

| Iteration | Upper Rotor |        | Lower Rotor |        |
|-----------|-------------|--------|-------------|--------|
|           | $C_T$       | Change | $C_T$       | Change |
| Baseline  | 0.003865    | –      | 0.002851    | –      |
| 1         | 0.004533    | 17.3%  | 0.001742    | -38.9% |
| 2         | 0.004608    | 1.67%  | 0.001654    | -5.04% |
| 3         | 0.004626    | 0.38%  | 0.001641    | -0.79% |
| 4         | 0.004630    | 0.10%  | 0.001637    | -0.24% |
| 5         | 0.004631    | 0.01%  | 0.001635    | -0.08% |

Table 6.3: Change in upper and lower rotor thrust with CDS/CSD coupling for  $\Delta\psi = 5.625^\circ$

the baseline model<sup>4</sup> and even the first coupling iteration<sup>5</sup>. While the baseline model shows some variation of the thrust with respect to  $\Delta\psi$ , the variation is much less than that shown after coupling with the two-airfoil loads. In addition, the phase angles at which the thrust minima and maxima change slightly between the models, with for instance the baseline model showing a “wide” peak centered around  $\Delta\psi = 45^\circ$  while the model including the two-airfoil loads shows a much sharper peak at  $\Delta\psi = 5.625^\circ$ .


The amount of this change (both absolute and in relative terms) decreases after the first coupling iteration, but takes some number of iterations before it can be considered converged. As an example, the thrust coefficient on each rotor at  $\Delta\psi = 5.625^\circ$ —corresponding to the upper rotor thrust peak (Fig. 6.1a) and lower rotor thrust minimum (Fig. 6.1b), and a condition where the models sharply diverge—is shown in Table 6.3, along with the relative change between iterations. The lower rotor thrust changes more in relative terms than the upper rotor on each iteration, but both rotors see thrust changing by less than 1% after 3 iterations. Observe however that as the phase angle increases, the amount of change decreases both between the baseline and adjusted models, and between coupling iterations.

While the individual upper and rotor loads are important to understand the blade motions and load paths when designing a rotor, the net behavior of the stacked rotor system is more pertinent to the design of the vehicle as a whole. To this end, the convergence of the total performance of the stacked rotor system is shown in Figure 6.2. Once again, it’s clear that the largest change is between the baseline model and the first coupling iteration, but the total thrust (Fig. 6.2a) changes somewhat more on subsequent iterations than the individual element forces (that is, the changes


---

<sup>4</sup>I.E. cylindrical prescribed wake, with only isolated airfoil loads


<sup>5</sup>In prior work published by the author [47], it was also seen that even a single “coupling” iteration could offer improvement over the baseline case.



(a) Total thrust



(b) Total power



(c) Efficiency

Figure 6.2: Convergence of total performance for the stacked rotor system

in element forces on each iteration compound). A similar change between iterations is noted in the total power (Fig. 6.2b). In both thrust and power, the model including 2D CFD two-airfoil forces shows two distinct minima, one around  $\Delta\psi = -5^\circ$  and another around  $\Delta\psi = 15^\circ$ , with the negative-index angle minimum being a larger deviation than the positive-index minimum. Combining the changes in thrust and power, an increase in efficiency (Fig. 6.2c) is shown at small phase angles, with maxima at both positive and negative  $\Delta\psi$ . (That is, the power is predicted to decrease by “more than” the thrust.) Compare this result from the current methodology to the baseline model, which shows relatively little change of rotor efficiency with phase angle.

An interesting consequence of the methodology demonstrated here is the theoretical ability to separate the wake (tip vortex) and bound circulation (airfoil) components of the interaction between the rotors, since the baseline model includes only the former while the adjusted loads include both. Clearly there is some variation in the rotor loads with  $\Delta\psi$  in the baseline model, hence the wake interactions cannot be ignored. However, the bound circulation interaction is clearly significant to both the individual loads and overall performance when the phase angle is small. This contrasts with the hypothesis of some earlier studies that the vortex interaction is the primary effect in the stacked rotor system (e.g. [7, 22]), but comports with more recent work [3].

Based on the results presented in this section, 4 CFD/CSD coupling iterations will be used for all other results presented, unless noted otherwise. Using 4 iterations ensures that the rotor thrust and power will vary by significantly less than 1% (see e.g. Table 6.3), with some margin for error if changing certain rotor parameters causes slower system convergence. While it would be possible to run more iterations to achieve more complete convergence, it is important to remember that this is a *lower-fidelity method*, hence higher-level convergence is not likely to significantly improve the accuracy of the methodology itself.





Figure 6.3: Thrust on each rotor for  $\Delta\theta_0 = 0^\circ$  compared to experimental data from [4]


### 6.1.2 Experimental Data Comparison

Validation of the current methodology was performed by comparing the predicted thrust and power to the experimental results of Johnson et.al. [4].<sup>6</sup> To demonstrate the improvement which the current methodology offers compared to the prior state of the art, the baseline (cylindrical prescribed wake) predictions are also shown. Results of this comparison for the thrust on the upper and lower rotor blades are shown in Figure 6.3. Overall, significantly better agreement with the experiment is seen in terms of the thrust *trends* for the model with multi-airfoil loads than for the baseline model. While (at positive phase angles) the thrust on the upper rotor is over-predicted and the thrust on the lower rotor is under-predicted with the multi-airfoil loads compared to the experiment, the amount of this over- or under-prediction is nearly constant. Additionally, the multi-airfoil lookup code is able to capture exactly the phase angles at which the


<sup>6</sup>While the experimental results do include some uncertainty, that uncertainty was reported to be on the order of 3% of the mean values [49]. Because this represents a relatively small interval, only the mean experimental values are shown in this study.



(a) Total thrust



(b) Total power



(c) Efficiency

Figure 6.4: Total performance of the stacked rotor system for  $\Delta\theta_0 = 0^\circ$  compared to experimental data from [4]

thrust maxima and minima occur. While the baseline model in some configurations comes closer in *magnitude* to the experimental data this agreement is only coincidental<sup>7</sup>, and indeed it is farther in magnitude in other configurations.


Predicted and experimental results for the total performance of the stacked rotor system are shown in Figure 6.4. Again, it is quite evident that the trends in thrust (Fig. 6.4a) are more accurately predicted with the multi-airfoil loads than with the baseline model. With the multi-airfoil loads a thrust minimum is actually predicted around  $\Delta\psi = -5.625^\circ$  to  $0^\circ$ , whereas the baseline model shows a thrust *maximum* in this configuration. Total power (Fig. 6.4b) is under-predicted with the multi-airfoil loads, but observe that the present methodology agrees with the experiment that the *configurations* corresponding to the power minima are  $\Delta\psi = -5.625^\circ$  and  $\Delta\psi = 11.25^\circ$ . Again, the baseline model is unable to accurately predict the trends in power variation with azimuth angle. Combining the thrust and power results, the rotor efficiency (Fig. 6.4c) matches well between the experiment and multi-airfoil model at positive phase angles. At negative phase angles however the multi-airfoil model diverges from the experimental results, as a consequence of the under-prediction of power. In contrast, the baseline model predicts hardly any variation of rotor efficiency with azimuth angle.

### 6.1.3 2D CFD vs. 2D Lifting-Vortex


The multi-airfoil coupling code uses both 2D lifting-vortex and 2D CFD airfoil loads, with the former used where the airfoils are far apart and the latter used where the airfoils are close. One question not answered by examining the accuracy of the complete methodology however is

---


<sup>7</sup>As in the common phrase “a broken clock is still right twice a day.”




(a) Upper rotor thrust



(b) Lower rotor thrust



(c) Total thrust



(d) Total power

Figure 6.5: Coupling based on 2D CFD and lifting-vortex airloads

how different actually are the rotor results with or without the 2D CFD loads<sup>8</sup>, considering that previous studies have indicated a significant improvement just by including the bound circulation interaction between the blades [3].

For sake of comparison, the multi-airfoil coupling code was run with both the combined CFD and lifting-vortex loads, and with the lifting-vortex loads only, the results of which are shown<sup>9</sup> in Figure 6.5. The rotor thrust (Figs. 6.5a and 6.5b) sees similar trends in adjustment, but somewhat different magnitudes when the blades are closely-spaced. In particular, using only the lifting-vortex loads results in about an 11% under-prediction of the thrust peak on the upper rotor (at  $\Delta\psi = 5.625^\circ$ ) compared to the 2D CFD loads. While the *magnitude* of the peak predicted by using only the lifting-vortex loads is nearly identical to the experiment, the *variation* of the upper rotor thrust with  $\Delta\psi$  near the peak is much less than that shown by the experiment or prediction obtained using the 2D CFD loads, which suggests that the agreement in peak magnitude when using only the lifting-vortex loads is coincidental. On the lower rotor, the different 2D loads result in different predictions for both the magnitude and phase angle of the thrust peak.

More significant differences between the models are seen regarding the total thrust (Fig. 6.5c) and power (Fig. 6.5d). Using only the 2D lifting-vortex loads, the total thrust decreases smoothly as the magnitude of the phase angle is decreased, although there is a thrust *maxima* at  $\Delta\psi = 0^\circ$ . In contrast, both the experiment and the prediction obtained by using 2D CFD airloads indicate that the most significant minimum occurs only at  $\Delta\psi = -5.625^\circ$ , and that the total thrust at  $\Delta\psi = 0^\circ$  will be close to that *minimum* value. Similar trends as in total thrust are shown in total power when using only the lifting-vortex loads. With the 2D CFD loads however, as the

---

<sup>8</sup>Or alternatively, how much of the span of the rotor blade is operating in a condition where the 2D CFD loads are significantly different from the 2D lifting-vortex loads.

<sup>9</sup>The baseline cylindrical prescribed wake results are not shown for clarity, since those results have been shown in Figures 6.3–6.4.





Figure 6.6: Airfoil system conditions along the span of the blades for  $\Delta\psi = 5.625^\circ$

phase angle decreases from a positive value, one power minima is predicted at  $\Delta\psi = 11.25^\circ$  in a region where the power comes close to a plateau, and another, stronger minima is predicted at  $\Delta\psi = -5.625^\circ$ . Since accurate prediction of the power is vital to the overall modeling of the rotor, this demonstrates the necessity of the 2D CFD airfoil loads.


#### 6.1.4 Selected Individual Cases

Thus far, all results in this chapter have been shown in terms of the rotor performance (i.e. thrust and power). To understand the physics of the stacked rotor however, it is helpful to consider what is occurring *along* the blades themselves. Two configurations are presented in this section: One with  $\Delta\psi = 5.625^\circ$  corresponding to the upper rotor thrust maximum, and one with  $\Delta\psi = 45^\circ$  where the blades are further apart but still experiencing interaction.


The relative positions and angles of the airfoil system along the blades for the configuration with  $\Delta\psi = 5.625^\circ$  are shown in Figure 6.6. Regardless of the aerodynamic model, the horizontal




(a) Upper blade Z-force



(b) Upper blade X-force



(c) Lower blade Z-force



(d) Lower blade X-force

Figure 6.7: Airfoil-frame force coefficients along the span of the blades for  $\Delta\psi = 5.625^\circ$

position between the airfoils ( $\Delta x$ ) is dictated mostly by the angle between the blades ( $\Delta x^{(hp)} = r\Delta\psi$ ). The vertical position between the airfoils ( $\Delta z$ ) however varies between the models, with the baseline aerodynamic model showing a nearly-constant separation and the adjusted models (2D lifting-vortex and CFD) predicting that the separation will increase along the span (Fig. 6.6a). This would be expected however, considering that the vertical position is significantly influenced by the flap-bending of the blades which itself is dictated by the blade thrust, and the adjusted models predict a greater difference in thrust between the upper and lower blades than the baseline model (recall Fig. 6.5). In contrast to the positions of the airfoils, the angles in the airfoil system (Fig. 6.6b) vary only slightly between the different aerodynamic models. Observe that the effective difference in pitch between the airfoils ( $\Delta\theta$ ) is slightly negative, since both blades are operating at the same collective ( $\theta_0 = 8^\circ$ ,  $\Delta\theta_0 = 0^\circ$ ) and the lower airfoil is in the downwash of the upper airfoil. Comparing the positions and angles of the airfoil system to the range of CFD parameters (Table 6.2), it can be seen that the 2D CFD airloads should be able to be applied along the entire span of the blades in this configuration.

Force distributions along the rotor blades are shown for the  $\Delta\psi = 5.625^\circ$  configuration in Figure 6.7, where  $C_x$  is the chordwise force and  $C_z$  is the chord-normal force in each airfoil's own frame of reference. Compared to the baseline model, the 2D lifting-vortex loads result in a relatively constant increment of  $C_z$  along the span of the blades (Figs. 6.7a and 6.7c). The 2D CFD loads however result in a different load distribution, with the outboard section of the upper blade being more heavily loaded and the outboard section of the lower blade being more lightly loaded. This different distribution of loads could be an important consideration to the structural design of the blades (particularly the spar). Due to the angle between the blades and the flow,  $C_x$  is largely dominated by the component due to  $C_l$  (see Eq. 5.28), and hence resemble a “mirror




Figure 6.8: Airfoil system conditions along the span of the blades for  $\Delta\psi = 45^\circ$

image” of  $C_z$  at a lower magnitude.

For the configuration with  $\Delta\psi = 45^\circ$ , the positions and angles of the airfoil system along the span of the blades are shown in Figure 6.8. With this phase angle the blades are further apart, but some degree of interaction between the blades is expected. Based on the positions of the blades’ airfoils (Fig. 6.8a), the 2D CFD airfoil loads will be relevant inboard of about 45% span. Outboard of this location, the multi-airfoil lookup code reverts to the 2D lifting-vortex airloads. The limiting parameter for the 2D CFD airloads is the range of horizontal position ( $\Delta x$ ), but it has been shown that CFD is not necessary beyond the specified range (see eg. Fig. 4.4). Observe that the vertical offset between the airfoils ( $\Delta z$ ) is negative over part of the span due to the change in reference frame with collective on the upper rotor (Eq. 2.1). Still, the range of the 2D CFD data collected would allow for the vertical offset at the hub to be increased if desired. Likewise, the angles along the blades (Fig. 6.8b) are also within the range of the 2D CFD data.

Z-force coefficients<sup>10</sup> for the  $\Delta\psi = 45^\circ$  configuration are shown in Figure 6.9. As ex-

<sup>10</sup>X-force coefficients are not shown, since they follow the same pattern shown in Figure 6.7 of being “mirrored





Figure 6.9: Airfoil-frame Z-force coefficients along the span of the blades for  $\Delta\psi = 45^\circ$


pected, outboard of 45% span, the multi-airfoil lookup code reverts to using lifting-vortex airloads, while inboard of this location 2D CFD loads are used. The transition between the CFD and lifting-vortex loads is not smooth (possibly due to the presence/absence of the periodicity model, see Chapter 3.1.1), suggesting that there may be some room for improvement in the model. However, no artificial smoothing is performed, to avoid introducing loads other than those expected by the particular physics models.

### 6.1.5 Coupling With and Without CFD/CSD Interface


While it is not *necessary* to use the airfoil table-based coupling with the prescribed wake model, it was still desired to examine the validity of this alternate coupling methodology, should future researchers need it for their own studies. To compare the airfoil table-based and CFD/CSD coupling, both methodologies use the same baseline model (as described previously). Four iterations of the CFD/CSD coupling are performed while six iterations of the airfoil table-based and reduced magnitude”.




(a) Upper rotor thrust



(b) Lower rotor thrust



(c) Total thrust



(d) Total power

Figure 6.10: Stacked rotor performance computed using different coupling methodologies

coupling are performed, but in both cases the solutions are considered iteration-converged. Both coupling methodologies were run using the same two-airfoil CFD and isolated airfoil (for lifting-vortex) data sets.

Results for the different coupling methodologies are shown in Figure 6.10. Overall trends match between both coupling methodologies, but there are some differences in magnitude. Thrust on the lower rotor (Fig. 6.10b) is very close between the two models, while thrust on the upper rotor (Fig. 6.10a) is nearly the same except in the region of the thrust peak, where CFD/CSD coupling predicts a slightly greater thrust. Combining the thrusts of each individual rotor (Fig. 6.10c), the airfoil table coupling more accurately captures the magnitude of the thrust minimum around  $\Delta\psi = -5.625^\circ$ , but then significantly under-predicts thrust at nearly all positive phase angles compared to both the experiment and CFD/CSD coupling. A similar trend is noted in the total power (Fig. 6.10d), where the airfoil table coupling comes slightly closer to the experimental power minimum at  $\Delta\theta = -5.625^\circ$ , but under-predicts power at the second minimum at  $\Delta\theta = 11.25^\circ$  much more significantly than the CFD/CSD coupling.

Based on these results, it may be suggested that while the CFD/CSD coupling is preferred, the airfoil table-based coupling is not irrelevant and may still produce meaningful results. However, care should be taken with both models, as although they are able to capture the trends in the change of thrust and power with configuration, neither is entirely able to match the experimental magnitudes.


### 6.1.6 Alternate CAMRAD Prescribed Wake Models

Thus far, all results presented have used the cylindrical (i.e. without any contraction) prescribed wake for the baseline inflow model, as defined by Equation 5.2. In a real rotor however, the wake will contract as it convects past the rotor. The CAMRAD model (Eq. 5.3–5.7) theoretically offers a more physical model by also allowing for the contraction and distortion of the wake to be prescribed. Both Landgrebe (Eq. 5.8) and Kocurek and Tangler (Eq. 5.9–5.10) have proposed semi-empirical parameters for the CAMRAD model, based on experimental observation of certain rotors.


Attempts were made to model the stacked rotor using the CAMRAD prescribed wake with both Landgrebe (Fig. 6.11) and Kocurek and Tangler (Fig. 6.12) parameters. Not that in each figure, the “baseline” model is the CAMRAD model with appropriate parameters, but without multi-airfoil loads. It was found in both cases that basing the parameters on  $C_T$  produced non-physical output, so the results shown are for the parameters based on the maximum bound circulation ( $\Gamma_{max}$ ). Even with the parameters based on  $\Gamma_{max}$ , it is clear that both models significantly under-predict the thrust and power for the stacked rotor by a factor of around one-half<sup>11</sup>. The CAMRAD models do predict more variation of thrust and power with azimuth angle than is predicted by the cylindrical prescribed wake, but this variation does not correlate to the variation observed in the experimental data. With the addition of the multi-airfoil loads, some similar trends as the experiment (particularly in the individual rotor thrusts) are observed, but the magnitudes of thrust and power are still significantly below those seen in the experiment. This should

---


<sup>11</sup>Other than the particular inflow model, the rest of the RCAS rotor model is identical for the CMARAD and cylindrical prescribed wakes, so it is considered unlikely that the difference with the experimental data was due to an error in the configuration of the rotor model.




(a) Upper rotor thrust



(b) Lower rotor thrust




(c) Total thrust




(d) Total power


Figure 6.11: Stacked rotor performance computed using the CAMRAD inflow model with coefficients given by Landgrebe




(a) Upper rotor thrust



(b) Lower rotor thrust



(c) Total thrust



(d) Total power

Figure 6.12: Stacked rotor performance computed using the CAMRAD inflow model with coefficients given by Kocurek and Tangler


not be interpreted as a flaw in the multi-airfoil coupling code, but rather proof that the code needs a somewhat-physical wake model as a starting point.

The aforementioned results do not preclude the possibility that the CAMRAD prescribed wake model could be made to produce more accurate results than the cylindrical prescribed wake model for a stacked rotor. The parameters developed by Landgrebe and Kocurek and Tangler are based on conventional (equal blade spacing), isolated rotors. Aside from the effect of blade spacing, there may be modeling differences if the stacked rotor is represented as “one” or “two” rotors. If these factors were taken into account and appropriate experimental results were examined, it may be possible to develop a novel set of semi-empirical parameters which would make the CAMRAD prescribed wake model relevant to stacked rotors. However, the development of such parameters is outside the scope of the current work.


## 6.2 Collective Offset Cases

While previous studies have suggested that best efficiency is obtained by using equal collective on the upper and lower rotors, recent studies [4] have suggested that some differential collective could enable novel control schemes to be implemented for the stacked rotor. To test the accuracy of the multi-airfoil coupling code, simulations were performed with a mean collective of  $\theta_0 = 8^\circ$  and a differential collective of  $\Delta\theta_0 = -4^\circ$ , such that the upper rotor collective is  $4^\circ$  and the lower rotor collective is  $12^\circ$  per the definitions given in Chapter 1.1.1. Since the CAMRAD prescribed wake model was shown to give inaccurate results, the baseline inflow model used in this section is again the cylindrical prescribed wake.


Results for the thrust on each rotor are shown in Figure 6.13. On the upper rotor (Fig.




(a) Upper rotor thrust



(b) Lower rotor thrust



(c) Total thrust




(d) Total power


Figure 6.13: Performance of the stacked rotor system for  $\Delta\theta_0 = -4^\circ$  compared to experimental data from [4]

6.13a, the baseline model predicts hardly any variation of thrust at positive phase angles, but the multi-airfoil coupling code shows almost exactly the same trend as the experiment, albeit with a slight and constant over-prediction. Neither the baseline nor the multi-airfoil load model is able to fully capture the thrust minimum at  $\Delta\psi = -5.625^\circ$ , but the multi-airfoil loads come closer in both magnitude and trend to the experiment than the baseline model. Thrust on the lower rotor (Fig. 6.13b) varies only slightly with phase angle and all models are lose in magnitude, but the multi-airfoil loads result in more accurate prediction of the thrust minima at  $\Delta\psi = \pm 5.625^\circ$  than the baseline model. The baseline model did not capture the thrust maximum on the lower rotor at  $\Delta\psi = 0^\circ$ , and the coupling code did not converge for that configuration.


Total thrust of the stacked rotor system (Fig. 6.13c) is predicted to vary little by the baseline model, but the inclusion of the multi-airfoil loads results in a generally correct trend (except at  $\Delta\psi = -25^\circ$  specifically), in particular the locations of the thrust minimum and maximum at  $\Delta\psi = -5.625^\circ$  and  $\Delta\psi = 5.625^\circ$  respectively, with some over-prediction of the magnitude. Both the variation of the thrust with azimuth angle and the over-prediction of thrust are driven mainly by the upper rotor. Results for total rotor power (Fig. 6.4b) are somewhat more ambiguous—while the multi-airfoil loads are able to predict a thrust maximum at positive  $\Delta\psi$  which the baseline model is not able to predict, neither model shows particularly positive results at negative phase angles. It should be noted however that the main region of interest in stacked rotor design would be small magnitude and/or positive phase angles, and the multi-airfoil coupling code is overall reasonably accurate in this regime.




(a) Upper rotor thrust



(b) Lower rotor thrust



(c) Total thrust



(d) Total power

Figure 6.14: Stacked rotor performance computed using different coupling methodologies for the  $\Delta\theta_0 = -4^\circ$  configuration

### 6.2.1 Coupling With and Without CFD/CSD Interface

It was sought to further validate the airfoil table-based coupling by employing it on the rotor configuration with collective offset. The comparison between the coupling methodologies and the experimental data is shown in Figure 6.14. The airfoil table-based coupling more accurately predicts the total thrust (Fig. 6.14c) than the CFD/CSD coupling at positive phase angles. However this apparent improvement in accuracy is somewhat coincidental, since both coupling methodologies over-predict the thrust on the upper rotor (Fig. 6.14a) but the airfoil table-based coupling more significantly under-predicts thrust on the lower rotor at small positive phase angles (Fig. 6.14b). Additionally, while the airfoil-table based coupling accurately predicts the magnitude of the total thrust peak, it shows no improvement over the CFD/CSD coupling in predicting the magnitude of the thrust minimum. This means that the airfoil table coupling significantly under-predicts the *variation* of thrust at small  $\Delta\psi$  compared to using the CFD/CSD coupling interface. The airfoil table-based coupling does appear to be able to damp some of the irregularity in the solution at negative phase angles, despite that irregularity having been present in the baseline model.

A more significant difference between the coupling methodologies is seen in the total power (Fig. 6.14d). While the CFD/CSD-based coupling is able to predict the magnitude of the total power with reasonable accuracy, the airfoil table-based coupling results in a power prediction twice as large as the actual value in the range  $|\Delta\psi| < 30^\circ$ . The specific cause of this discrepancy is not well understood, given that the same error was not seen in the  $\Delta\theta_0 = 0^\circ$  configuration (Fig. 6.10d). However, recall that the alternate coupling methodology works by changing the lift and drag on the airfoil tables used by RCAS. It is possible that, due to the twisting of the blades or

the changing angle of the flow, some component of the lift along the blades is being improperly translated into drag (power). Ultimately at this stage, despite being able to predict the thrust well, the airfoil table-based coupling cannot be recommended in general due to a lack of robustness when predicting the rotor power.

### 6.3 Summary

Overall, the inclusion of the multi-airfoil loads has been shown in this chapter to significantly improve the accuracy of the stacked rotor thrust and power predicted by the lower-fidelity rotor model compared to the baseline cylindrical prescribed wake model. When varying the phase angle of the stacked rotor, the baseline model shows some slight variation in thrust (both total and on each rotor) and power, which is expected due to the change in the positions of the trailed vorticity from each blade relative to the other blades. However, the variation of the thrust and power with phase angle in the baseline model is much less than that seen in the experiment. Including just the 2D lifting-vortex airloads (as an analogue for the bound-circulation interaction) significantly increases the variation of thrust and power with phase angle. However, only by including the 2D CFD airfoil-interactional airloads is it possible to fully capture the variation of thrust and power with phase angle, particularly in the important regime where the blades are closest together. In addition to the total loads, the distribution of the loads along the span of the blades is also different between the baseline, lifting-vortex, and 2D CFD airloads, which may have implications for blade structural design. While final agreement with the experiment is acceptable, room remains for improvement, particularly in prediction of the total power.

Accurate agreement with the experiment requires an acceptably-accurate baseline model.

While the cylindrical prescribed wake works well enough, a model with wake contraction would theoretically be more physical. However, the contracted wake models available in RCAS were found to produce significantly inaccurate results, which was theorized to be due to these models having been developed for isolated rotors. It is possible that the current methodology could be improved through the development of a contracted prescribed wake model particular to the case of stacked rotors.

While coupling using RCAS's CFD/CSD coupling interface is reasonably reliable, the results of coupling coupling the airloads by adjusting the airfoil tables are more mixed. For the case with equal upper and lower rotor collective, both coupling methodologies performed reasonably well. With greater lower rotor collective however the airfoil table-based coupling produced a significantly inaccurate prediction of total power. The reason for this discrepancy is currently unknown, but given the robustness of the CFD/CSD coupling interface, the airfoil table adjustment is considered to be a secondary concern.

## Chapter 7: Summary and Conclusions

Prior research has established that, with proper selection of configuration, stacked rotors may offer benefits in absolute performance, acoustics, and flexibility as compared to conventional isolated and counter-rotating rotors<sup>1</sup>. These benefits are of keen importance to modern AAM vehicle concepts, which must be quieter to operate in proximity to people on the ground, while requiring higher performance out of smaller rotors in distributed electric propulsion concepts. However, those studies which have examined stacked rotors have focused mainly on testing and validating the concept experimentally, or using high-fidelity methods. To fully realize the benefits of the stacked rotor concept in a production vehicle, a design methodology beginning with lower-fidelity rotor models is required.

Existing lower-fidelity rotor models typically rely on a lifting-line approximation of the rotor blades, and are differentiated by how they represent the vorticity in the wake (e.g. prescribed vortex lattice grid, deforming lattice, vortex particles, etc.). While these models may offer a first-order approximation of the interactions between adjacent rotor blades, prior research on multi-element airfoils has shown that airfoil interaction is a complex phenomenon, involving highly localized changes to the flow which are effected by factors such as the shape and thickness of real airfoils. A true stacked rotor design methodology must offer some means of taking these

---

<sup>1</sup>On the other hand, with *improper* selection of configuration, stacked rotors may perform significantly worse.

airfoil interactions into account.

In this work, a methodology has been presented for lower-fidelity modeling of stacked rotors which would be suitable for a design scenario. At the rotor-level, this methodology uses a prescribed vortex wake model, but is not limited to this and could be generalized to any wake model which adequately models the localized variations in the wake around the rotor blades (i.e. any vortex wake method). To address the shortcomings of lifting-line blade-to-blade interaction, 2D CFD is used to model the interaction between the blades' airfoils with high fidelity, but at lower cost than would be required for 3D CFD modeling of the entire rotor system.

Pre-computation of the 2D airfoil-interactional loads allows the lower-fidelity rotor simulations to be conducted with little computational overhead. However, pre-computation of the loads also means that it is necessary to have an interpolation scheme which can take multiple inputs, since the conditions of the airfoil system will not be known to the 2D CFD solver. Radial basis function interpolation is one such scheme, and in this study the implementation of RBF interpolation offered by Python's SciPy library was used. Building on a free, open-source tool such as Python enables simple deployment of the developed code to researchers working with any system already used for scientific computing. The 2D CFD simulations themselves are performed using the HAMSTR/HAM2D solver with automated mesh generation, which allows for accurate results to be obtained with minimal user input. The mesh generation and flow solver codes have been previously developed at the University of Maryland, and are currently in use in both academia and government.

A key step in the methodology is establishing the relationship between the conditions along the rotor blades in 3D, and the corresponding airfoil systems in 2D. While at first this would seem trivial since such relationship has been established for conventional rotors, the assumptions

which underlie that relationship do not necessarily hold true for stacked rotors. In particular, it is not possible for both the geometric and aerodynamic relationship between the airfoils in 3D to be maintained in 2D, nor should this be expected given the assumption in 2D simulations of uniform free-stream. Instead, a strategy has been presented of relaxing the geometric similarity in order to achieve aerodynamic similarity between the 3D and 2D systems. While this is not necessarily a rigorous strategy, it appears to provide adequate results.

The 2D mesh generation and flow solution methodology has been validated by previous studies. However, examination of the behavior of the 2D airfoil system is instructive to the behavior and modeling of the stacked rotor system. It is shown that the 2D CFD is able to capture interactional effects which a simple lifting-vortex model based on isolated airfoil tables could not, and that these interactional effects may be significant to the overall performance of the airfoil system. It is also shown that 2D CFD is only necessary when the airfoils are arbitrarily close to each other, and that the lifting-vortex model is sufficient when the airfoils are far apart. This observation allows for the necessary range of the 2D CFD simulations—and thus the overall cost of this methodology—to be significantly limited. With further examination of the 2D airfoil results, it is possible to make suggestions which may lead to superior performance in future stacked rotor designs, in particular blade twist and differential twist between the upper and lower rotors, which has not been examined by any prior study in stacked rotors.

Past studies have variously sought to explain the difference between stacked and conventional rotors in terms of either wake interactions or airfoil interactions. In this study however (and in some recent research) it has been shown that neither effect can be considered in isolation. If it were the case that the only significant effect was the interactions between the wake and

the blades, then the baseline prescribed wake<sup>2</sup> model would sufficiently capture the interaction between the rotors. On the contrary, the results presented in this study show that while the prescribed wake model on its own predicts some variation of thrust and power with phase angle, the amount of variation is much less than the experiment, and the trends of variation either disagree entirely or match only in the broadest interpretation. With the inclusion of the 2D lifting-vortex loads these trends are improved, such that the phase angles at which thrust minima and maxima occur are predicted with reasonable accuracy. However the predicted trends in power remain inaccurate, and the magnitude of the predicted thrust is under-predicted in the critical regime of small phase angles where the stacked rotor offers best efficiency and control potential. In this regime, the trends and magnitudes of the thrust are further improved with the inclusion of the 2D CFD airfoil-interactional loads. Thrust matches well with the experiment overall, showing correct locations of minima and maxima, and coming reasonably close to the experimental magnitudes. Trends in power are correctly shown, particularly the double-minima of power for the  $\Delta\theta_0 = 0^\circ$  configuration. The magnitude of the power is under-predicted however in the  $\Delta\theta_0 = 0^\circ$  configuration, although it is reasonably accurate in the  $\Delta\theta_0 = -4^\circ$  configuration.

Aside from the rotor loads, differences between the baseline, lifting-vortex, and 2D CFD models were noted in loads along each blade. In particular, while the lifting-vortex resulted in a roughly constant increment of force along the span of the blades, the inclusion of the 2D CFD airloads resulted in a different spanwise loading distribution, with the upper and lower blades being more heavily and lightly loaded towards their tips, respectively. These differences in load distribution may be relevant to both blade deflections, and to rotor blade structural design.

Two different coupling methodologies were presented, one based on RCAS's CFD/CSD

---

<sup>2</sup>Or in prior work, vortex particle wake

coupling interface, and one in which the airfoil tables in the rotor model were adjusted. The methodology based on the CFD/CSD coupling interface was found to be robust in both the  $\Delta\theta_0 = 0^\circ$  and  $\Delta\theta_0 = -4^\circ$  configurations. While the airfoil table-based coupling worked about as well in the former configuration, a large discrepancy in power was noted in the latter configuration. The reason for this discrepancy is currently unknown, but it is suggested that the CFD/CSD coupling interface should be preferred.

Of course, the 2D CFD model cannot on its own compute the aerodynamic conditions along the rotor blades in 3D. To compute these conditions, it is necessary to build on an existing inflow/wake model. In this study, a cylindrical prescribed vortex wake methodology was used, which is computationally efficient and produced reasonable results. However, the cylindrical prescribed wake does not model contraction or distortion of the wake. While alternate prescribed wake methodologies were tried, these methodologies were based on experimental results for isolated rotors, and proved not to be suitable for the stacked rotor. If however an adequate distortion model could be developed for the stacked rotor, it is possible that the accuracy of the 2D airload-coupling methodology could be further improved.

It is hoped and intended that the methodology presented in this study will lead to future improvements in stacked rotor design. By using the methodology presented in this study, rapid parameter variation and design iterations of the stacked rotor may be performed. For instance, it is trivial to change the twist of the blades in this model, whereas to do so on an experimental test stand would require constructing entirely new sets of rotor blades. Of course, further testing may be warranted to further validate this methodology for a wider variety of configurations than have recently been tested experimentally. It is hoped however that this methodology may be used to inform future experimental work, which in turn will produce data that may be used to further

validate and refine this model.

## 7.1 Key Contributions

The key contributions of the current work may be summarized as follows:

- Explicitly defined a reference system for systems of airfoils, which is not a unique choice but does preserve the meaning of 2D CFD inputs and outputs (Chapter 2.2)
- Described how to relate the aerodynamic conditions along the blades stacked rotor computed by the lower-fidelity inflow model to the equivalent 2D airfoil system, by changing the system geometry to account for flow non-uniformity (Chapter 2.3)
- Developed the pre- and post-processing routines necessary for building the 2D airfoil-interactive tables using 2D CFD (automating the methods described in Chapter 3.2)
- Explicitly defined a 2D lifting-vortex model for airfoil interactions (Chapter 3.1), and demonstrated that such a model may be used in place of 2D CFD without much loss of accuracy if and only if the airfoils are sufficiently far apart (Chapter 4.1.2)
- Demonstrated the need for higher-fidelity modeling of 2D airfoil interactions (Chapters 4.1.1 and 4.1.2)
- Shown the ability to make design suggestions for stacked rotors based on 2D airfoil-interactive data (Chapter 4.2)
- Identified key features and limitations of existing lower-fidelity rotor models, as applied to stacked rotors (Chapter 5.2.1)

- Identified methods to include 2D multi-airfoil airloads in lower-fidelity rotor models (Chapters [5.3.2](#) and [5.3.3](#))
- Developed software specifically intended to interface with RCAS, to implement inclusion of the 2D multi-airfoil airloads (implementing the methods described in Chapter [3.3](#) and either Chapter [5.3.2](#) or [5.3.3](#))
- Demonstrated that the described methodology produces performance predictions which more accurately match experimental results than does the current state-of-the-art (Chapters [6.1.2](#) and [6.2](#))
- Demonstrated the need for high-fidelity, as opposed to lifting-vortex, modeling of the airfoil interactions in stacked rotors (Chapters [6.1.3](#) and [6.1.4](#))

## Bibliography

- [1] Vertical Flight Society. Electric VTOL news: Supernal (Hyundai Motor Company) S-A1. Online: <https://evtol.news/hyundai-s-a1/>, 2022.
- [2] Graham Warwick. Electric propulsion start-up LaunchPoint EPS targets drones, EVTOLs. Online: <https://launchpointeps.com/placement-aviation-week-111020-electric-propulsion-start-up/>, 10 November 2020.
- [3] George Jacobellis, Rajneesh Singh, Chloe Johnson, Jayant Sirohi, and Rob McDonald. Investigation of stacked rotor performance in hover pt. ii: Computational validation. In *VFS Transformative Vertical Flight*, January 2020.
- [4] Chloe Johnson, Jayant Sirohi, George Jacobellis, and Rajneesh Singh. Performance and acoustics of a stacked rotor with differential collective pitch. In *Vertical Flight Society 78th Annual Forum and Technology Display*, May 2022.
- [5] Mahendra Bhagwat. Co-rotating and counter-rotating coaxial rotor performance. In *AHS Technical Conference on Aeromechanics Design for Transformative Vertical Flight*, January 2018.
- [6] Barry Tully. Variable camber prop studied for vtol. *Aviation Week*, pages 98–101, 15 August 1960.
- [7] Anton Landgrebe and E. Dean Bellinger. Experimental investigation of model variable-geometry and ogee tip rotors. *NASA Contractor Report*, CR-2275, February 1974.
- [8] Walter Sonneborn and Jan Drees. The scissors rotor. *Journal of the American Helicopter Society*, 20(3):18–27, July 1975.
- [9] J. Gordon Leishman. *Principles of Helicopter Aerodynamics*. Cambridge University Press, 32 Avenue of the Americas, New York, NY 10013-2473, USA, second edition, 2006.
- [10] Chloe Johnson, Jayant Sirohi, George Jacobellis, Rajneesh Singh, and Rob McDonald. Investigation of stacked rotor performance in hover pt. i: Experimental measurements. In *VFS Transformative Vertical Flight*, January 2020.
- [11] James Rorke. Hover performance tests of full scale variable geometry rotors. *NASA Contractor Report*, CR-2713, August 1976.

- [12] Hamilton Standard. Advanced propellor concepts...developed today. Online: [http://www.enginehistory.org/Propellers/HamStdVC/HS\\_VariableCamber.pdf](http://www.enginehistory.org/Propellers/HamStdVC/HS_VariableCamber.pdf), 1965.
- [13] Tom Fey. The hamilton standard variable camber propeller. Online: <http://www.enginehistory.org/Propellers/HamStdVC/HamStdVC.shtml>, 27 June 2021.
- [14] J. O. Emmerson. Letters: Variable camber prop. *Aviation Week*, page 164, 12 September 1960.
- [15] All The World's Rotorcraft. Kaman K-16. Online: [http://www.aviastar.org/helicopters\\_eng/kaman\\_k-16.php](http://www.aviastar.org/helicopters_eng/kaman_k-16.php).
- [16] F. H. Norton. The effect of staggering a biplane. *NACA Technical Notes*, TN-70, September 1921.
- [17] A. M. O. Smith. High-lift aerodynamics. *Journal of Aircraft*, 12(6):501–530, 1975.
- [18] James Chiles. Air america's black helicopter. *Smithsonian Air and Space Magazine*, Online: <https://www.smithsonianmag.com/air-space-magazine/air-americas-black-helicopter-24960500/>, March 2008.
- [19] Ray Prouty. *Helicopter Aerodynamics*, volume II, chapter 18, pages 69–72. Eagle Eye Solutions, 2009.
- [20] Isaag Gershman. Prime Portal: Israeli AH-64A Apache walk around. Online: [http://www.primeportal.net/hangar/isaac\\_gershman/ah-64a/index.php?Page=1](http://www.primeportal.net/hangar/isaac_gershman/ah-64a/index.php?Page=1), 2008.
- [21] Manikandan Ramasamy. Hover performance measurements toward understanding aerodynamic interference in coaxial, tandem, and tilt rotors. *Journal of the American Helicopter Society*, 60(3), July 2015.
- [22] Daiju Uehara and Jayant Sirohi. Quantification of swirl recovery in a coaxial rotor system. In *AHS International 73rd Annual Forum and Technology Display*, May 2017.
- [23] Hossein Saberi, Maryam Khoshlahjeh, Robert Ormiston, and Michael Rutkowski. Overview of RCAS and application to advanced rotorcraft problems. In *AHS 4th Decennial Specialist's Conference on Aeromechanics*, January 2004.
- [24] Matthew Misorowski, Farhan Gandhi, and Phuriwat Anusonti-Inthra. Comparison of acoustic predictions using distributed and compact airloads. In *Vertical Flight Society 75th Annual Forum and Technology Display*, May 2019.
- [25] Beau Pollard, Michael Patterson, and Siena Whiteside. A combined blade element and vortex panel method for stacked rotor analysis. In *AIAA Aviation Forum*, June 2019.

- [26] Robert T. Jones and Doris Cohen. *Aerodynamic Components of Aircraft at High Speeds*, volume VII, chapter A,4: Effect of Sweepback, pages 36–48. Princeton University Press, 1957.
- [27] Yong Su Jung, Bharath Govindarajan, and James Baeder. Turbulent and unsteady flows on unstructured line-based hamiltonian paths and strands grids. *AIAA Journal*, 55(6):1986–2001, 2017.
- [28] Yong Su Jung and James Baeder.  $\gamma - \overline{Re_{\theta_t}}$  spalart-almaras with crossflow transition model using hamiltonian-strand approach. *Journal of Aircraft*, 56(3):1040–1055, 2019.
- [29] Miranda Costenoble, James Baeder, and Yong Su Jung. Automated mesh generation and solution analysis of arbitrary airfoil geometries. *Journal of Aircraft*, 59(5):1120–1136, September–October 2022.
- [30] Jayanarayanan Sitaraman, Vinod K. Lakshminarayan, Beatrice Roget, and Andrew M. Wissink. Progress in strand mesh generation and domain connectivity for dual-mesh cfd simulations. In *55th AIAA Aerospace Sciences Meeting*, 9–13 January 2017.
- [31] Alexander Costenoble, Yong Su Jung, Bharath Govindarajan, and James Baeder. Automated mesh generation and solution analysis of arbitrary airfoil geometries. In *AHS Technical Conference on Aeromechanics Design for Transformative Vertical Flight*, January 2018.
- [32] Jens-Dominik Müller. *On Triangles and Flow*. PhD thesis, University of Michigan, 1996.
- [33] Jens-Dominik Müller. Grid generation tools. Online: <http://www.ae.metu.edu.tr/tuncer/ae546/prj/delaundo/>.
- [34] Zhu Zihao. HAMSTRAN, an indirect method to create all-quadrilateral grids for the HAMSTR flow solver. Master’s thesis, University of Maryland, 2017.
- [35] Yong Su Jung. *Hamiltonian Paths and Strands for Unified Grid Approach for Computing Aerodynamic Flows*. PhD thesis, University of Maryland, 2019.
- [36] P.E.O. Buelow, S. Venkateswaran, and C.L. Merkle. Stability and convergence analysis of implicit upwind schemes. *Computers and Fluids*, 30(7):961–988, 2001.
- [37] Youcef Saad and Martin Schultz. GMRES: A generalized minimum residual algorithm for solving nonsymmetric linear systems. *SIAM Journal on Scientific and Statistical Computing*, 7(3):856–869, July 1986.
- [38] Department of Defense High Performance Computing Modernization Program. HPC centers: Home. Online: <https://centers.hpc.mil/index.html>.
- [39] Wilna du Toit. Radial basis function interpolation. Master’s thesis, University of Stellenbosch, 2008.
- [40] The SciPy Community. SciPy API reference. Online: <https://docs.scipy.org/doc/scipy/reference/generated/scipy.interpolate.RBFInterpolator.html>, January 2022.

- [41] Python Software Foundation. `pickle`—Python object serilization. Online: <https://docs.python.org/3/library/pickle.html>, 2022.
- [42] US Army Research, Development, and Engineering Command. *RCAS User's Manual: Volume I*, February 2019.
- [43] US Army Research, Development, and Engineering Command. *RCAS Theory Manual*, February 2019.
- [44] David Peters and Cheng Jian He. Correlation of measured induced velocities with a finite-state wake model. *Journal of the Americal Helicopter Society*, (36):59–70, July 1991.
- [45] John D. Anderson Jr. *Fundamentals of Aerodynamics*. McGraw Hill, 3rd edition, 2001.
- [46] Chengjian He and Jinggen Zhao. Modeling rotor wake dynamics with viscous vortex particle method. *AIAA Journal*, 47(4):902–916, April 2009.
- [47] Miranda Costenoble, James Baeder, and George Jacobellis. Incorporation of airfoil-interactive data to improve the accuracy of stacked rotor performance predictions in the design stage. In *VFS Transformative Vertical Flight*, 25–27 January 2022.
- [48] Advanced Rotorcraft Technology, Inc. *File-based CFD Coupling for RCAS*, 12 July 2019.
- [49] George Jacobellis, Rajneesh Singh, Chloe Johnson, Jayant Sirohi, and ROb McDonald. Experimental and computational investigation of stacked rotor performance in hover. *Aerospace Science and Technology*, 116, September 2021.