AERODYNAMICS AND DESIGN

FOR ULTRA-LOW REYNOLDS NUMBER FLIGHT

A DISSERTATION

SUBMITTED TO THE DEPARTMENT OF AERONAUTICS AND ASTRONAUTICS

AND THE COMMITTE ON GRADUATE STUDIES

OF STANFORD UNIVERSITY

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS

FOR THE DEGREE OF

DOCTOR OF PHILOSOPHY

Peter J. Kunz

June 2003

I certify that I have read this dissertation and in my opinion it is fully adequate, in scope and in quality, as a dissertation for the degree of Doctor of Philosophy.

Ilan M. Kroo (Principal Advisor)

I certify that I have read this dissertation and in my opinion it is fully adequate, in scope and in quality, as a dissertation for the degree of Doctor of Philosophy.

Juan J. Alonso

I certify that I have read this dissertation and in my opinion it is fully adequate, in scope and in quality, as a dissertation for the degree of Doctor of Philosophy.

Fritz Prinz

Approved for the University Committee on Graduate Studies:

Acknowledgments

This work was partially sponsored by the NASA Institute for Advanced Concepts, a unique program providing support for research that for one reason or another may sit outside the sphere of interest of mainstream funding resources. My tenure at Stanford was also supported by the Hugh H. Skilling Stanford Graduate Fellowship in Science and Engineering. This fellowship was made possible by the generosity of Mr. Frank Lynch. To both Stanford University and Mr. Lynch, I am eternally grateful.

I would like to thank my advisor and mentor at Stanford, Professor Ilan Kroo. His technical insight has been invaluable and the flexibility and enthusiasm with which he approaches graduate research and your students has wrought a wonderful environment in which to grow, professionally and personally. I would also like to thank another mentor, Professor Mark Maughmer. For almost 15 years he has been a teacher, a colleague, and most importantly a friend. The vigor with which I approach my work and life is in no small part inspired by him.

Finally and most importantly I'd like to thank my family, my father for giving me a goal, my mother for showing me how to achieve it, my brother for teaching me how to do it with honour and class, and my sister for helping me keep it all in perspective.

Abstract

Growing interest in micro-air-vehicles has created the need for improved understanding of the relevant aerodynamics. A reasonable starting point is the study of airfoil aerodynamics at Reynolds numbers below 10,000, here termed ultra-low Reynolds numbers. The effects of airfoil geometry on performance are explored using an incompressible Navier-Stokes solver. Variations in thickness, camber, and the shape of leading and trailing edges are studied. Results indicate an increase in maximum lift coefficient with decreasing Reynolds number, but the lift to drag ratio continues to decrease, making the power required for flight a more restrictive consideration than lift. This performance penalty can be mitigated by careful airfoil design. Contrary to the notion that viscous fairing reduces airfoil geometry effectiveness, the computational results indicate that geometry still has a profound effect on performance at ultra-low Reynolds numbers. To further explore this design space, the flow solver has been coupled with an optimizer, resulting in the first airfoils quantitatively designed for this flow regime and demonstrating that unconventional camberlines can offer significant performance gains.

Building on these results, tools are developed for ultra-low Reynolds number rotors, combining enhanced classical rotor theory with airfoil data from Navier-Stokes calculations. This performance prediction method is coupled with optimization for both design and analysis. Performance predictions from these tools are compared with three-dimensional Navier-Stokes analyses and experimental data for several micro-rotor designs. Comparisons among the analyses and experimental data show reasonable agreement both in the global thrust and power, but the spanwise distributions of these quantities exhibit deviations, partially attributable to three-dimensional and rotational effects that effectively modify airfoil section performance. While these issues may limit the applicability of blade-element type methods for detailed rotor design at ultra-low Reynolds numbers, such methods are still useful for evaluating concept feasibility and rapidly generating initial designs for prototyping and for further analysis and optimization using more advanced tools. Moving toward controlled powered flight at centimeter scales, several prototype rotorcraft have been fabricated and tested, exploring both the aerodynamics and system integration issues.

Nomenclature

A	Rotor disk area
B	Number of blades
C_d	Sectional drag coefficient ($D/ 0.5\rho V_\infty^2$)
C_f	Skin friction coefficient ($\tau/ 0.5\rho(\omega r)^2$)
C_l	Sectional lift coefficient ($L/ 0.5\rho U_{ref}^2$)
C_p	Pressure coefficient ($p-p_\infty / 0.5\rho(\omega r)^2$)
C_p’	Canonical pressure coefficient
C_Q	Rotor torque coefficient ($Q/ \pi\rho R^3(\omega R)$)
C_t	Rotor blade sectional thrust coefficient ($(dT/dR) / 0.5c\rho(\omega R)^2$)
C_T	Rotor thrust coefficient ($T/ \pi\rho R^2(\omega R)$)
D	Drag force
E	Modulus of elasticity (Young’s modulus)
G	Shear Modulus
J	Polar moment of inertia
L	Lift force
L/D	Lift-to-drag ratio
M	Rotor figure of merit (P_{ideal} / P_{actual})
P	Power

Q	Rotor torque
RPM	Revolutions per minute
T	Rotor thrust
U_∞	Freestream velocity
c	chord length
p	Pressure
q	Dynamic pressure ($0.5\rho U_{ref}^2$)
r	Local section radius
t	maximum section thickness
u	Rotor vertical induced velocity or the u component of velocity (see context)
v	Rotor tangential induced velocity or the v component of velocity (see context)
v_i	Rotor inviscid tangential (swirl) induced velocity
v_v	Rotor viscous tangential (swirl) induced velocity
Re	Reynolds number based on chord length.
α	Geometric angle of attack in degrees
Γ	Potential vortex strength
ζ	Torsional deflection angle
θ	Jig twist angle or momentum thickness (see context)
κ	Rotor tip loss correction factor
ν	Kinematic viscosity
ϕ	Inflow angle relative to rotor plane
ρ	Freestream density
σ	Rotor solidity or standard deviation (see context)

τ	Shear stress
ω, Ω	Rotor angular velocity

Contents

List of Tables

List of Figures

Chapter 1

Introduction

1.1 Looking to Nature

Mankind has dreamed of flight for as long as he has dreamed. Flight in nature, in all its shapes and sizes, has been the seed of these dreams. Mankind has flown for barely over one hundred years, since the days of Lilienthal, then Chanute and the Wrights, but there has been flight in nature since before history. We have progressed, our technology has progressed, and our aircraft have progressed. Faster, and bigger, and bigger still, but rarely smaller and slower. Using the Reynolds number as a measuring stick, we have conquered flight first in the tens of thousands with models, progressed to the hundreds of thousands with gliders, and now operate in the millions and beyond, but nature still flies alone in realm of the hundreds and thousands: the ultra-low Reynolds number regime. We have done well, but man-made flight at the tiny scales that nature takes for granted has eluded us, but that is now changing.

The flight of small birds and insects has always held a fascination for mankind and entire lifetimes and careers have been devoted to the study of single species. We have studied their shapes and motions, theorized and modeled their mechanisms of flight, but we have been unable to emulate their success. A concise overview of both the challenges and possibilities for micro-air-vehicles has been provided by Spedding and Lissaman [1].

Major problems have included a technological inability to accurately manufacture very small-scale machines, an inability to power such devices, and the challenge of controlling them. Without an expectation of realization, there has been only marginal interest in the study of aerodynamics at tiny physical scales, but advances in technology have recently begun to make ultra-low Reynolds number flight a real possibility. The growing interest in micro-air-vehicle development has finally created a *need* for improved understanding of the relevant aerodynamics.

The flight regime of micro-aircraft poses numerous challenges for aerodynamic analysis and design, but little experimental or computational work exists for aerodynamic surfaces operating at ultra-low Reynolds numbers (Re) below 10,000. Interest in very small aircraft, operating in the Re=100,000 to Re=150,000 range, has rapidly grown and many such vehicles are currently under development, but research and development at much smaller scales is still in its infancy. Technological advances in micro-fabrication techniques and in the miniaturization of electronics are beginning to make mechanical micro-flight vehicles feasible from a systems and manufacturing standpoint. There are numerous potential applications for these vehicles, but first they must be capable of flight.

Much of the work presented here has been motivated by the Mesicopter Micro-Rotorcraft Development Program at Stanford University. The goal of this program, and a major focus of this thesis, is the design, development, and testing of micro-rotors and micro-rotorcraft. The broader goal of this work is to contribute substantially to the foundations of this essentially unexplored segment of applied aerodynamics and flight vehicle design.

Aerodynamic research at relevant Reynolds numbers has been severely limited. Experimental airfoil data has been published by Schmitz[2] and Althaus [3], among others, at chord Reynolds numbers as low as 20,000 to 30,000, but in terms of the nature of the viscous effects, this is considerably different from the range of interest below Re=10,000. This data also represents the extreme lower range of operation for many

wind-tunnel facilities and inconsistencies are often apparent in the resulting data. More recent work in experimental airfoil testing at ultra-low Reynolds numbers, some of it concurrent and subsequent to the initial publication of portions of this research [4, 5], has been published by Sunada and Kawachi [6] and as a comprehensive survey of relevant experiments and analyses by Azuma *et al.*[7]. These results support many of the conclusions put forward here regarding the effects of airfoil geometry on performance.

Research into hovering flight at small physical scales has been dominated by studies of the hovering or near-hovering flight of insects and birds. The comprehensive work by Ellington [8] is an excellent example and provides significant insight into the mechanisms of micro-scale hovering flight in nature. Emphasis has typically been placed on the biomechanics of flight and the search for, and modelling of, unsteady aerodynamic mechanisms of high lift such as the 'clap and fling' mechanism proposed by Weis-Fogh [9, 10] or the unsteady wake mechanisms proposed by Rayner [11]. Common to all of these studies is a focus on inviscid aerodynamics, with, at most, simple accounting and estimates of viscous effects. As mentioned earlier, nature clearly has found a workable solution, but a key issue that currently prevents translating this large body of work to a man-made vehicle is the complexity of the required flapping motions and the difficulty associated with developing the necessary mechanisms at small scales without suffering significant mass and electromechanical efficiency penalties. This is also without any consideration of the relevant aerodynamic efficiencies. This dissertation seeks to explore the simplest possible solution: a non-articulating rotor operating under steady aerodynamic conditions.

1.2 Thesis Chapter Summary

A methodical and progressive approach has been taken to the exploration of the aerodynamic design space at ultra-low Reynolds numbers. Beginning with two-

dimensional steady aerodynamics, the work is then extended to rotor analysis and design, culminating in the development of several micro-rotorcraft prototypes.

- Chapter 1: Introduction, summary of work, and summary of contributions.
- Chapter 2: Overview of the utilized two-dimensional analysis methods, including limited validation and a grid sizing study.
- Chapter 3: Analysis and design of airfoils at ultra-low Reynolds numbers. This chapter explores the airfoil design space by parametrically varying the section geometry and Reynolds number. This culminates in the development of an automated optimal design tool for two-dimensional airfoils at ultra-low Reynolds numbers.
- Chapter 4: The theoretical development of a hybrid method for the automated design and analysis of micro-rotors. This approach combines classical rotor methods with several enhancement models and the judicious application of two-dimensional Navier- Stokes solutions.
- Chapter 5: Overview of the experimental and computational methodology used in the validation of the hybrid method, including a discussion of experimental error.
- Chapter 6: Discussion of results, including a description of several candidate micro-rotor designs and their theoretical and experimental performance. Several anomalous test results are explained via a hypothesis based on aero-structural deflections. The chapter discusses the effects on analysis and design of the different enhancement models for viscous swirl and vortex wake contraction.
- Chapter 7: Discussion and overview of several prototype micro-rotorcraft.
- Chapter 8: Conclusions and recommendations for future work.

1.3 Summary of Contributions

The sum total of this work provides new insight into both the aerodynamic challenges and the associated vehicle systems issues involved in realizing powered flight at ultra-low Reynolds numbers. Consequently, contributions have been made across a broad range of topic areas.

- First comprehensive and quantitative exploration of the airfoil design space at ultra-low Reynolds numbers.
- First quantitatively based optimal airfoil design at ultra-low Reynolds numbers.
- Demonstrated the applicability of classical rotor analysis methods, when properly augmented, for preliminary rotor design at ultra-low Reynolds numbers.
- Developed an automated, optimizer-based design and analysis tool for preliminary rotor design.
- The design of first rotors specifically tailored to this regime.
- Computationally analyzed and experimentally tested micro-rotors at unprecedented low Reynolds numbers
- Demonstrated the acute susceptibility of micro-rotors to aero-structural torsional deflections due to rotational effects.
- Developed, manufactured, and tested micro-rotorcraft prototypes providing new insight into system requirements and current technological limitations.

Chapter 2

Two-Dimensional Computational Analysis Methods

2.1 Motivation for Two-Dimensional Analysis

In many aerospace applications two-dimensional analysis is considered to be of limited value due to the existence of a large established knowledge-base and ever-increasing computational abilities that allow three-dimensional analysis of complex configurations and phenomena. These views *may* be justified when the end goal is the design of a specific device for a particular application, or to expand the level of understanding beyond the limits of two-dimensional study, but the present work encompasses something entirely different. The existing knowledge-base at ultra-low Reynolds numbers pertaining to applied aerodynamics, airfoils, and flight vehicle design is minimal. The need for a broader understanding of these areas, combined with limited time and resources, makes two-dimensional analysis a practical starting point. The analyses are relatively fast, permitting a large and varied test matrix. This is essential when the investigation is broadly exploratory both in geometry and flow properties. Two-dimensional analysis can also provide a more informative picture of fundamental behavior, free from three-dimensional effects such as cross-flow and induced drag which can be difficult to discern and isolate in both computational and experimental results.

What is being sought here, and in the next chapter, is the foundation of a new knowledge base for airfoils used in ultra-low Reynolds number applications, or, in the very least, confirmation that the reach of current knowledge extends to this regime. The exploration of this new regime requires the use of computational tools that are largely untested for this application. The goal of this chapter is to describe these tools, their implementation, and the necessary validation completed before the course of research could begin.

2.2 Incompressible Navier-Stokes Solver

The two-dimensional computational analyses make extensive use of the INS2d incompressible Navier-Stokes solver developed by Rogers [12, 13]. This code utilizes the artificial compressibility method, first introduced by Chorin [14], to deal with incompressible flows. It is capable of both steady and time-accurate computations using single grids, Chimera over-set grids, and moving grids for relative motion of multiple solid bodies. Only the single grid, steady and time-accurate capabilities are used in the current work.

2.2.1 Artificial Compressibility

Navier-Stokes solvers for the compressible flow equations generally require some form of preconditioning at very low Mach numbers. As the Mach number is decreased, the flow field approaches the incompressible limit. The character of the governing equations goes from hyperbolic to elliptic and the speed of sound becomes infinite. Typical artificial dissipation schemes introduce an amount of spurious dissipation that is proportional to the speed of sound. Because of this, hyperbolic flow solvers are not well suited to accurately solve elliptic problems.

The artificial compressibility method offers a straightforward and efficient means of preconditioning to allow for the solution of an incompressible homogeneous flow field. The incompressible conservation of mass equation,

$$\frac{\partial u}{\partial x} + \frac{\partial v}{\partial y} = 0 \tag{2.1}$$

is modified by the addition of a pseudo-time derivative of density:

$$\frac{\partial \rho}{\partial \tau} + \frac{\partial u}{\partial x} + \frac{\partial v}{\partial y} = 0 \tag{2.2}$$

Density is related to pressure via an artificial equation of state, where δ is the artificial compressibility:

$$p = \frac{\rho}{\delta} \tag{2.3}$$

This introduces an artificial and finite acoustic speed governed by the selection of the δ parameter. This addition to the conservation of mass equation, combined with the conservation of momentum equation, results in a hyperbolic system of equations which may be marched in pseudo-time. As the solution converges to a steady state, the artificial compressibility term drops out and a divergence-free solution is attained. The artificial compressibility acts in a manner similar to a relaxation parameter.

2.2.2 Computation of Forces and Moments

Forces and moments are calculated within INS2d by numerical integration of pressure and shear forces on the airfoil surface. This method is commonly used and considered acceptable in most applications, but given the unique flow conditions and the need to analyze extremely thin sections, an additional method of force calculation has been developed and implemented as a stand-alone post-processor.

Conventional airfoils have round leading edges and trailing edges that approximate a cusp. The sections considered here are typically only 2% thick with the added complexity of manufacturing constraints requiring the consideration of blunt leading and trailing edges due to minimum gage constraints. As a result, a significant portion of the drag is being calculated in regions of very high surface curvature. Discretizing these airfoils creates the possibility of incurring discretization and subsequent force integration errors in the leading and trailing edge regions.

A second issue is the possibility of excessive numerical and spatial dissipation in the solution affecting the results. INS2d has not previously been used for airfoil analysis at ultra-low Reynolds numbers and excessive dissipation could introduce errors, and in extreme cases result in solutions that erroneously appear to be steady-state.

Both of these issues have been assessed, and the forces have been verified by utilizing an off-body control-volume approach. The method applies the two-dimensional conservation of momentum equation which in integral form may be expressed as:

$$\vec{F}' = \oint \rho \vec{V}(\vec{V} \cdot \hat{n}) dS + \oint p \hat{n} dS \qquad (2.4)$$

This far-field calculation is analogous to the surface force integration with the viscous body force term replaced by the momentum flux across the control volume. As an additional check, the post-processor calculates the total mass flux for the control volume. Instead of the forces exerted on the airfoil by the fluid, the effect on the fluid from the airfoil is calculated. The large rounded control volume does not have the regions of high local curvature found at the surface, rather the integration area has been expanded, reducing local sensitivities due to any particular region. Several issues arise with this approach. These include numerical dissipation within the flow field, which the method attempts to assess, and application too far from the airfoil, where the cell sizes generally are larger and spatial discretization error and increased dissipation can become issues.

The control volume method has been applied to INS2d analyses of a NACA 4402 aifoil operating at two degrees and six degrees angle of attack. The control volume has been applied at various distances from the airfoil surface for several Reynolds numbers. The resulting drag coefficients, are compared to the results from surface force integration in Figure 2.1. The maximum variation at any single Reynolds number is less than one percent for control volumes out to a radius of approximately six chordlengths. Beyond this point the variations increase due to the rapid growth in cell size, reaching 25% at 12 chordlengths. Dissipative effects may be visible in the systematic increase in the drag variation seen with decreasing Reynolds number at two degrees, but the maximum variation across the Reynolds number range of interest is only two percent and the variation is absent from the six degree cases. Variations in the predicted lift are smaller, averaging 0.5% out to six chordlengths.

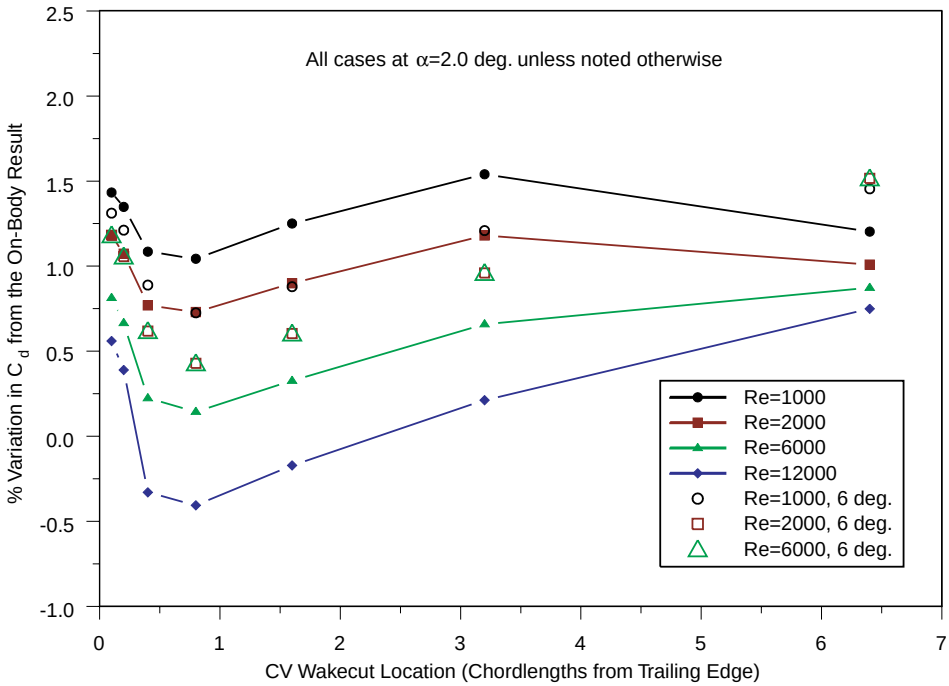

FIGURE 2.1 Comparison of INS2d on-body and control volume drag values for the NACA 4402.

2.2.3 Grid Sizing Study

All calculations use a C-grid topology with either 256 by 64 cells or 512 by 128 in the leading edge studies and the optimized designs. The airfoil is paneled using 70% of the streamwise cells, with 10% of these clustered at the leading edge. Constant initial normal spacing provides for approximately 25 cells in the boundary layer at 10% chord for the 256 by 64 grids. The outer grid radius is placed at 15 chord-lengths. The results of a grid-sizing study using the NACA 0002 are displayed in Figure 2.2. Relative error values are based on analysis with a 1024 by 256 grid. The relative error convergence with grid size is close to quadratic and values for lift and drag are essentially grid independent with a 0.2% variation in C_l and a 0.7% variation in C_d over three levels of grid refinement.

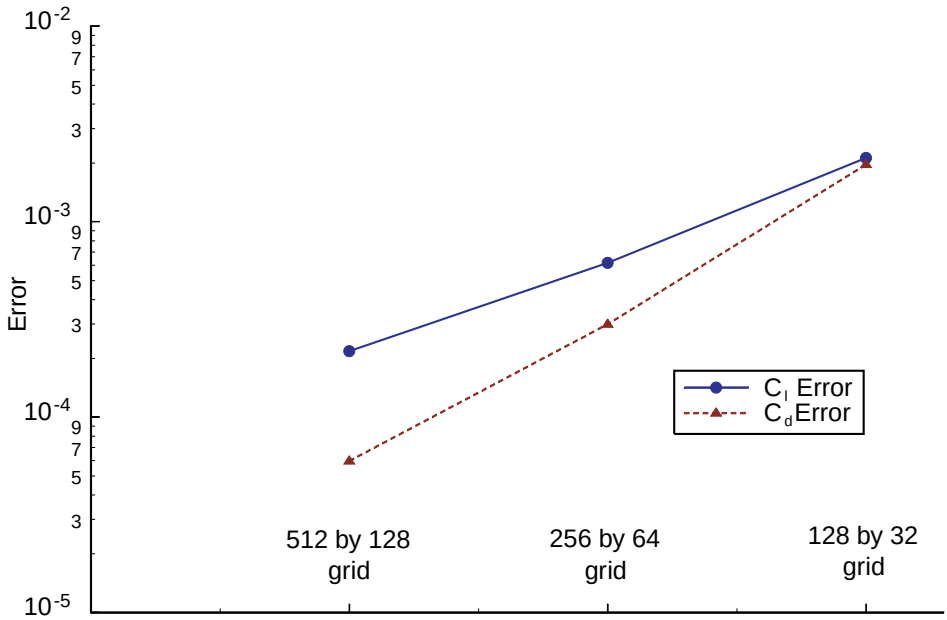

FIGURE 2.2 Results of a grid-sizing study for the NACA 0002 at Re=6000.

2.2.4 Comparison with Experiment

There are a small number of relevant experiments in the literature that provide a reasonable basis for comparison with INS2d at the Reynolds numbers of interest. One interesting result of the current work, to be discussed in greater detail in Chapter 3, that is supported by experiment is an increase in attainable lift coefficient as the Reynolds number is reduced. Thom and Swart [15] tested a small R.A.F. 6a airfoil model in an oil channel and water channel at Reynolds numbers below 2000. They observed large increases in lift coefficient at fixed angles of attack as the Reynolds number was reduced from 2000 to almost one.

Validation of the computational analyses is difficult due to the almost complete absence of experimental data at relevant Reynolds numbers. The Thom and Swart experiment is based on a 1.24cm chord airfoil with manufacturing deviations from the R.A.F. 6a. This small test piece was hand filed to shape, causing the measured geometry to vary across the span. An exact validation is not possible due to the unknowns in the section geometry, but comparison with computations for the R.A.F. 6 airfoil with a 256 by 64 grid show reasonable agreement with experiment. No coordinates for the R.A.F. 6a could be located, but the R.A.F. 6 appears to be nearly identical. The results are shown in Figure 2.3. The Reynolds number varies from point to point and ranges from Re=650 to Re=810. The computed drag is on average 7.5% lower than experiment, but the trends in C_d with angle of attack agree. Corresponding C_l data is only given for α=10.0. The computational result matches the experimental value of C_l=0.52 within 3.0%.

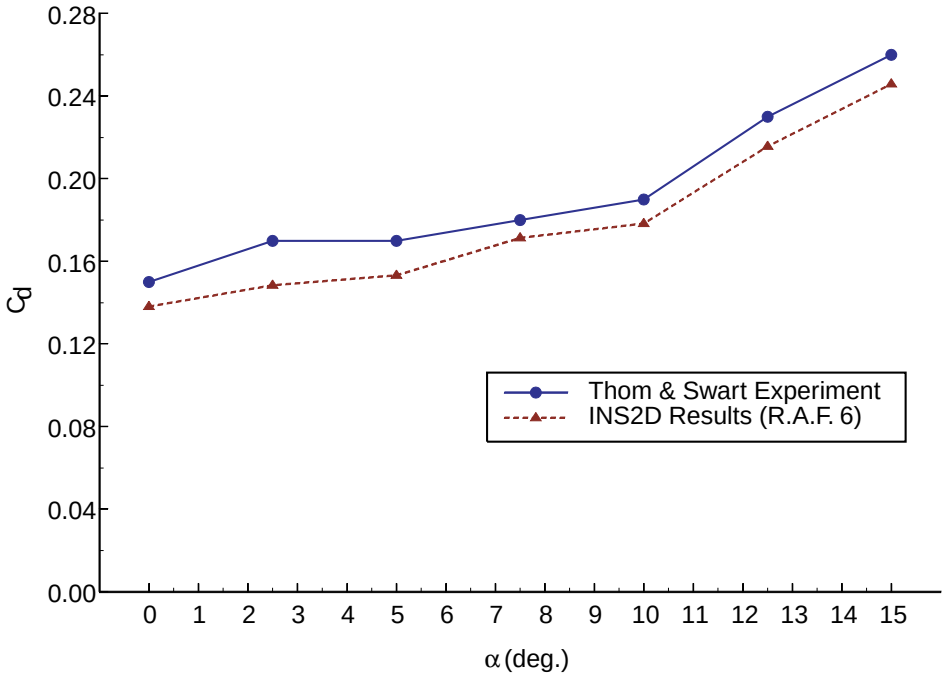

FIGURE 2.3 Comparison of computed and experimental C_d for the R.A.F. 6 and R.A.F. 6a airfoils.

2.3 Viscous-Inviscid Interaction Methods

Using integral boundary layer formulations in conjunction with inviscid flow field solutions offers the potential for significant computational savings over full viscous flow solvers. Many potential flow panel methods prove problematic when attempting to model very thin airfoils due to the close proximity of singularity elements, but Euler solvers have no such issues and still represent a significant computational savings. The common problem for these methods is the unknown extent to which the limitations and assumptions of integral boundary layer theory may effect the quality of the solution.

The MSES program developed by Drela [16] has been applied in this study with limited success. This is a two-dimensional Euler solver, coupled with an integral boundary layer formulation. It appears to give reasonable drag predictions over a narrow range of angles of attack, but the limitations of the boundary layer formulation cause the solution

to diverge if significant regions of separated flow exist. This is a general limitation of these methods. Unfortunately, this is often the case with airfoils operating in this regime, even at moderate angles of attack.

A comparison of results from MSES and INS2d for NACA 4402 and NACA 4404 airfoils at Re=1000 are presented in Figures 2.4 and 2.5. The upper end of each curve represents the maximum angle of attack for which a steady-state solution was attainable. The most obvious feature is the failure of the MSES analysis at about 3 degrees geometric angle of attack, much earlier than the INS2d results which continue to converge past 10 degrees. Over the range that MSES does converge to a solution, the trends in the results agree with the INS2d calculations, and drag values are close to each other, although the curves appear offset to some degree. In both figures, the effects of increasing thickness agree. Both analyses indicate similar reductions in the lift curve slope and equivalent increases in drag. The MSES solutions predict a lower lift curve slope and a slightly higher zero lift angle of attack, resulting in a significant deviation in predicted lift, and approximately 5% lower drag than the equivalent INS2d results. These results indicate that under the limitations of low angles of attack, the much faster inviscid/integral boundary layer codes provide a functional alternative to full viscous flow field solutions, but would not be applicable for detailed design where analysis of the entire operating range is needed.

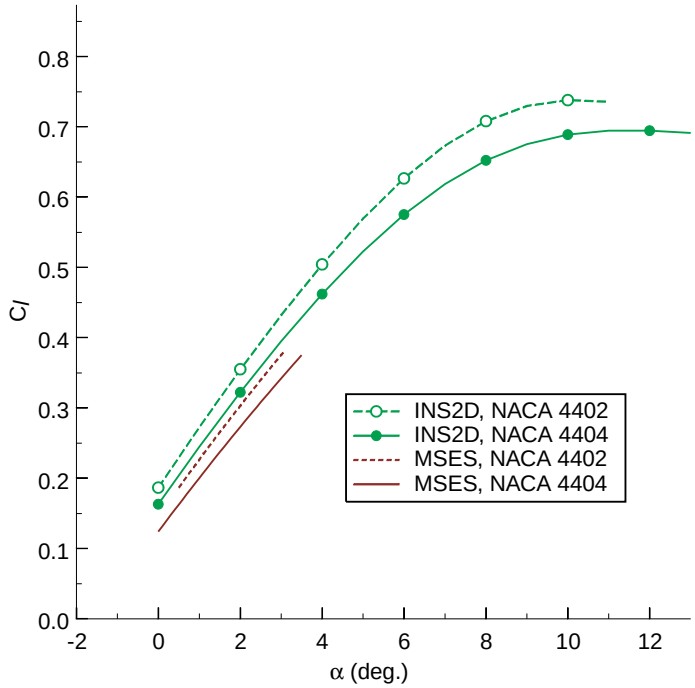

FIGURE 2.4 MSES and INS2d predicted lift curves for the NACA 4402 and NACA 4404. Flow is fully laminar and Re=1000.

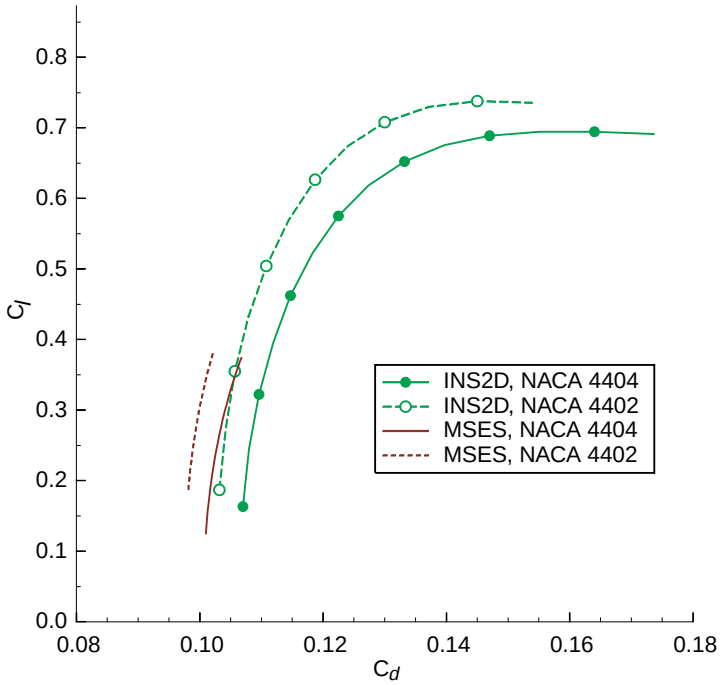

FIGURE 2.5 MSES and INS2d predicted drag polars for the NACA 4402 and NACA 4404. Flow is fully laminar and Re=1000.

2.4 Flow Field Assumptions

The INS2d analyses make use of three assumptions about the flow field: the flow is incompressible, fully laminar, and steady. Incompressibility is well justified for this application since the highest Mach number encountered by the rotors to be described in Chapter 6 was Mach 0.3. For a broad range of applications, the Mach number will be considerably below this value and the flow is essentially incompressible.

The fully laminar flow assumption is more uncertain. Schlichting [17] indicates that transition on a flat plate at zero incidence occurs at Reynolds numbers greater than 350,000, this is in the abscence of any adverse pressure gradient however. The introduction of an adverse pressure gradient as seen in the pressure recovery of an airfoil introduces two possibilities, laminar separation or turbulent transition. The degree of separation that might result in transition and the transition length are the unclear issues, but the alternatives to a fully laminar assumption are even less satisfactory. The flow field could be assumed fully turbulent, which is surely not the case, or transition could be artificially and rigidly imposed at a specified location. Of these three, the fully laminar assumption is the least restrictive and most physically accurate in the range of Reynolds number and angle of attack of interest here.

The steady-state assumption represents tremendous computational savings over time-accurate analyses and has been verified with time-accurate computations. Airfoil polars have been generated by increasing the angle of attack until the steady-state solution fails to converge. Analyses were completed at Reynolds numbers of 1000, 2000, 6000, and 12,000. For analyses at Re=6000 and above, failure to converge is taken as an indication of unsteady phenomena in the flow field. For the Re=1000 and Re=2000 cases, the polars have been computed assuming a steady-state flow field, but then a time-accurate analysis has been completed at or above the indicated maximum lift-to-drag ratio angle of attack. This assures the absence of significant unsteady effects in the presented data,

but does not represent a rigidly defined upper limit. It is possible that higher angles of attack would still exhibit steady behavior.

The steady-state results for the NACA 4402 airfoil at Re=6000 have been compared with data generated using the time-accurate mode of INS2d. The results of this study are illustrated in Figure 2.6. The time-accurate computations consist of an impulsive start with 20 chord lengths of total travel. Each time step represents 0.02 chord lengths of travel. The steady-state solutions are indicated by the horizontal reference lines. The symbols on the α=8.0 case represent a halving of the time step. These demonstrate that the 0.02 chord length time step is adequate for resolution of the temporal variations. Steady-state analysis failed to converge at α=5.0. Time accurate solutions agree with the steady-state result up past α=4.0. At α=6.0, one degree past the point where the steady-state analysis stops, small amplitude periodic behavior is visible in the time accurate solution. At this point, the lift and drag of the section are still reasonably well defined and quasi-steady. At α=8.0, the amplitude of the oscillations observed in the time accurate solutions has increased considerably and the effects of multi-frequency shedding are visible.

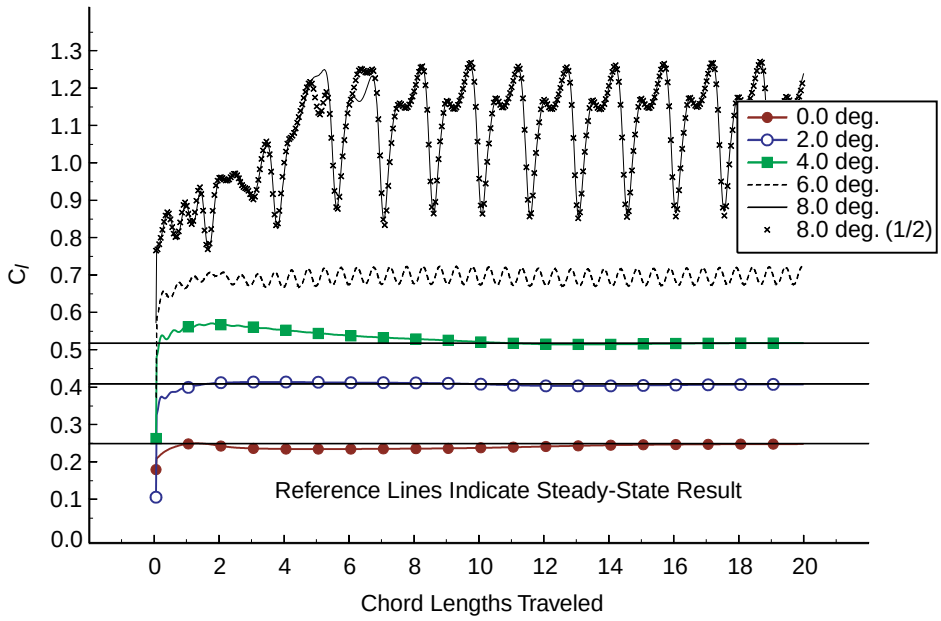

FIGURE 2.6 INS2d time-accurate analyses of an impulsively started NACA 4402 at Re=6000.

The results of these time-accurate computations at higher angles of attack indicate that much higher lift coefficients may be attainable within the unsteady range of operation. Although inviting, these operating points could be difficult to exploit in practice. The increase in lift comes at the cost of a very large increase in drag. Increasing the angle of attack from four to eight degrees more than doubles the drag, to an average value of over 1000 counts. The average lift to drag ratio is roughly maintained, and may even increase slightly, but the practical issue of providing sufficient power to overcome the drag rise, combined with the time variation of section performance, complicates the utilization of these unsteady operating points.

Under the described limitations, INS2d has proven to be a suitable tool for initial exploration of the ultra-low Reynolds number regime. All reasonable efforts have been made to validate the accuracy of the code and reduce the possible mechanisms of error, but a true validation of the code is difficult due to the dearth of relevant experimental and computational data in the literature. In many ways the data compiled here *is* the baseline.

Chapter 3

Analysis and Design of Airfoils for Use at Ultra-Low Reynolds Numbers

3.1 Introduction

In this chapter a series of airfoil sections are analyzed at chord Reynolds numbers below 10,000 using the INS2d [12, 13] two-dimensional, incompressible, Navier-Stokes solver. Two-dimensional analysis allows a broad spectrum of parameters to be considered and provides a baseline for future, more detailed, studies. The geometric parameters investigated include thickness, camber, and the effects of leading and trailing edge shapes. The current study represents the first detailed survey of two-dimensional airfoil performance at ultra-low Reynolds numbers and provides significant insight into the aerodynamic design space. The results address a fundamental question of whether section geometry is still important at ultra-low Reynolds numbers, and go further, revealing which factors are critical to performance and what metrics of performance are most useful for aerodynamic design.

Building on the results of this parametric study, an automated design optimization method has been developed that maximizes the two-dimensional lift-to-drag ratio at a given Reynolds number. Rather than a simple mimicry of nature or qualitative

extrapolation based on higher Reynolds number experience, the resulting geometries represent the first quantitatively designed airfoils for use under these conditions.

3.2 General Reynolds Number Effects

The most obvious effect of operation at ultra-low Reynolds numbers is a large increase in the section drag coefficients. Zero lift drag coefficients for airfoils range from 300 to 800 counts depending on the Reynolds number and geometry. The increase in drag is not reciprocated in lift. Lift coefficients remain of order one, resulting in a large reduction in the L/D. Flight at these Reynolds numbers is much less efficient than at higher Reynolds numbers and available power is a limiting technological factor at small scales. It is important to operate the airfoil at its maximum L/D operating point, but this requires operating close to the maximum steady-state lift coefficient. Even small increases in the maximum lift coefficient are significant and generally translate to higher L/D.

Flow at ultra-low Reynolds numbers is dominated by viscosity, and as the Reynolds number is reduced, the effects of increasing boundary layer thickness become more pronounced. The concept of a 'boundary layer' at such low Reynolds numbers in a fully viscous flow field is an inexact notion. Here it is generalized as the low velocity flow adjacent to the body over which the pressure gradient perpendicular to the surface is close to zero. Regions of constant pressure extend a significant distance away from the surface of the airfoil and the effective geometry is significantly altered by the presence of the boundary layer.

In this operating regime the boundary layer has a dramatic effect on surface pressures, closer to that of a separation streamline. This is demonstrated by considering the inviscid and viscous pressure distributions on the NACA 0008 at zero angle of attack, shown in Figure 3.1. Both the Re=6000 and Re=2000 cases are fully attached. As the Reynolds number is reduced, the value of the minimum pressure and the slope of the

adverse gradient in the pressure recovery are reduced. This weakened pressure recovery does impact the pressure drag, but at a positive angle of attack the largest effect is on lift. Figure 3.2 is a similar plot for the NACA 0002. Here again, all three cases are fully attached. For this 2% thick section at zero lift, there is essentially no recovery of pressure in the Re=1000 case. It behaves as if predominately separated. Also of note is the small increase in the magnitude of the minimum pressure, due to the effective thickening of the section.

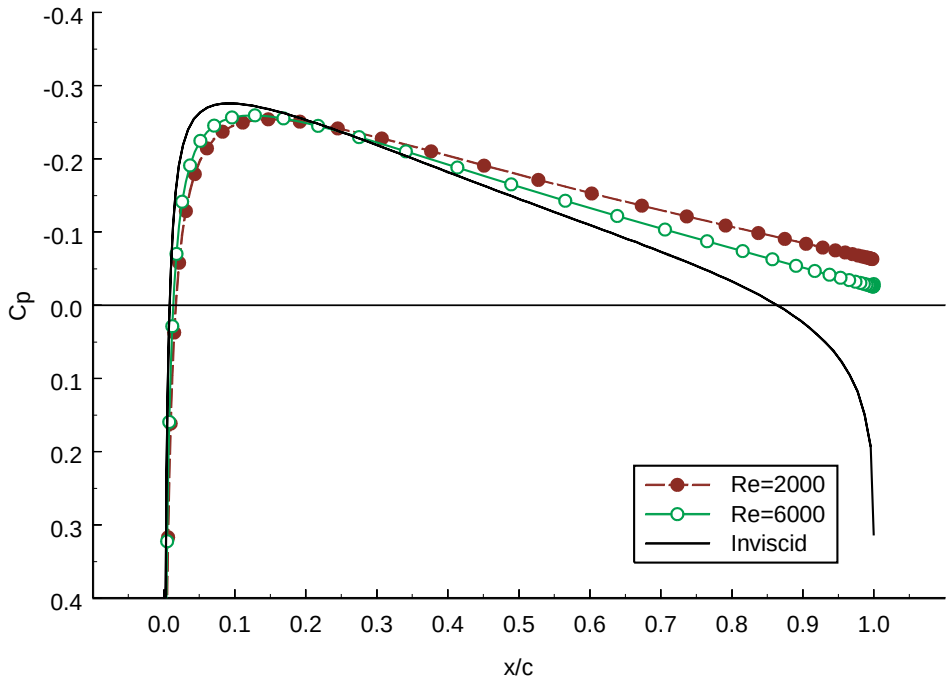

FIGURE 3.1 Zero lift C_p distributions for the NACA 0008.

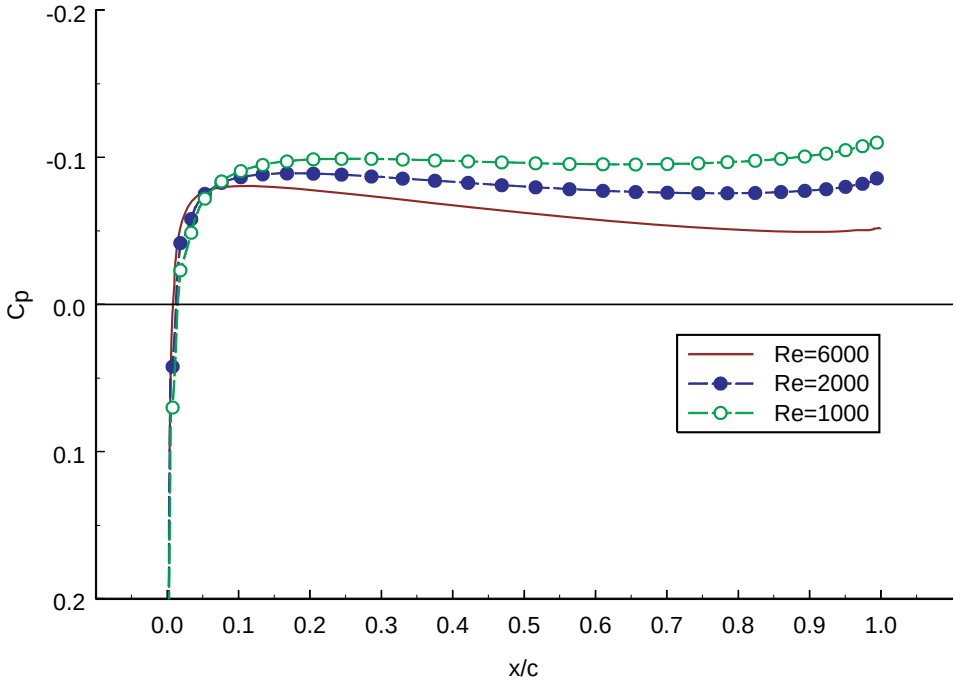

FIGURE 3.2 Zero lift C_p distributions for the NACA 0002.

Viscous effects alleviate and smooth the high gradients present in the nose region of an airfoil. The reduction in the height of the leading edge suction peak and the reduction in slope of the adverse pressure recovery gradient delay the onset of separation and stall. Leading edge separation is delayed in thin sections, with trailing edge separation delayed in thicker sections. The results are higher attainable angles of attack and higher maximum steady-state lift coefficients. Pressure distributions for the NACA 0008 airfoil at α=2.0 are presented in Figure 3.3. The Re=6000 case is on the verge of trailing edge separation, but the Re=2000 case does not separate until α=3.5. Lift coefficients for the two cases agree within 3.5%. The Re=2000 case achieves the same amount of lift with a much weaker suction peak, a less adverse recovery gradient, and an additional margin of separation-free operation.

The delay in trailing edge separation is supported by the form of several classical laminar separation criteria. Stratford [18] with later modifications from Curle and Skan [19] formulates the following laminar separation criteria:

$$x^2 C_p' \left(\frac{d}{dx} C_p' \right)^2 \approx 0.0104 \tag{3.1}$$

The canonical pressure coefficient (C_p') may be expressed as a function of the edge velocities (U_e) yielding an expression having no Reynolds number dependence and a dominant dependence on the pressure recovery gradient. In this instance, *x* represents the distance from the location of minimum pressure.

Thwaites [20] laminar separation criteria takes the form:

$$\frac{\theta^2}{\nu} \left(\frac{dU_e}{dx} \right) \approx -0.09 \tag{3.2}$$

This relation is also Reynolds number independent and dominated by the pressure recovery gradient. This only become clear upon the substitution of Thwaites's formulation for the momentum thickness:

$$\theta^2 = \frac{0.45\nu}{U_e^6} \int_0^x U_e^5 dx \tag{3.3}$$

Both of these formulations support the concept that the large growth in boundary layer thickness that accompanies a reduction in Reynolds number is potentially beneficial to lifting performance due to a reduction in the adverse pressure recovery gradient.

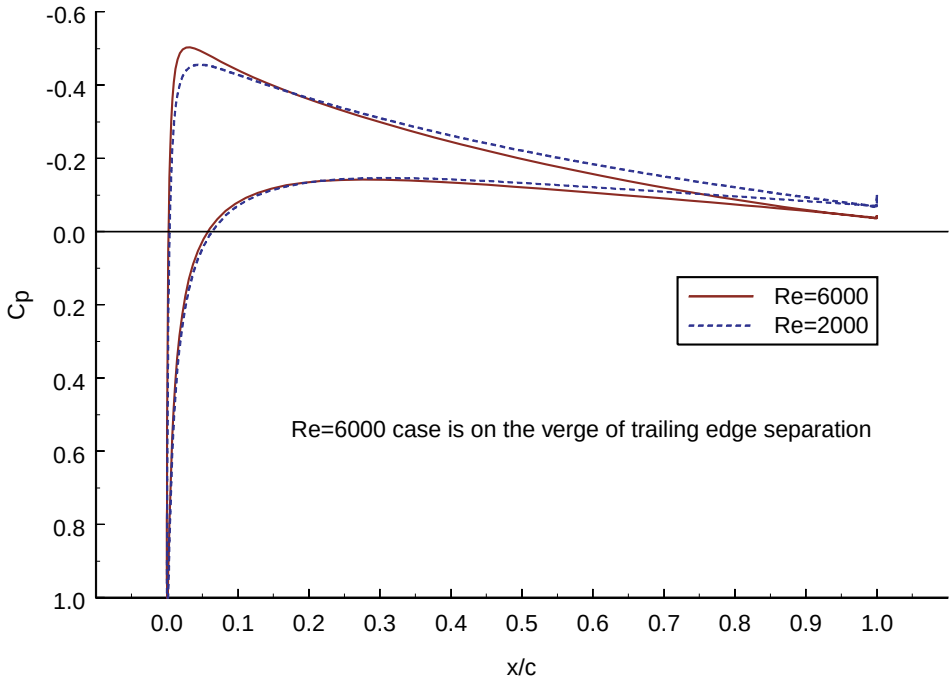

FIGURE 3.3 C_p distributions for the NACA 0008 at α=2.0. At Re=6000, trailing edge separation is imminent, but the Re=2000 case remains fully attached.

Reducing the Reynolds number affects the lift curve by reducing the slope in the quasi-linear range and extending the quasi-linear range to higher angles of attack. While operating within this range, the displacement effect of the boundary layer progressively reduces the effective camber of the section with increasing angle of attack. This change in the effective geometry increases as the Reynolds number is reduced. The delay in separation that accompanies reduced Reynolds numbers extends the quasi-linear range to higher angles of attack. Once the flow does separate, growth of the separated region is delayed by a reduction in the Reynolds number. The overall effect is a significant increase in both the maximum steady-state angle of attack and the lift coefficient.

Lift curves for the NACA 0002 and NACA 0008 are presented in Figure 3.4. The calculations are at Re=2000 and Re=6000. The reduction of slope is most apparent for the NACA 0002 airfoil, but both sections exhibit the extension of the quasi-linear lift range. The NACA 0002 lift curves remain linear across its entire operating range, with the flow field becoming unsteady due to leading edge separation. The Re=2000 case

reaches α=5.0 and a lift coefficient a full tenth greater than the Re=6000 case. Similar gains occur for the NACA 0008 at Re=2000.

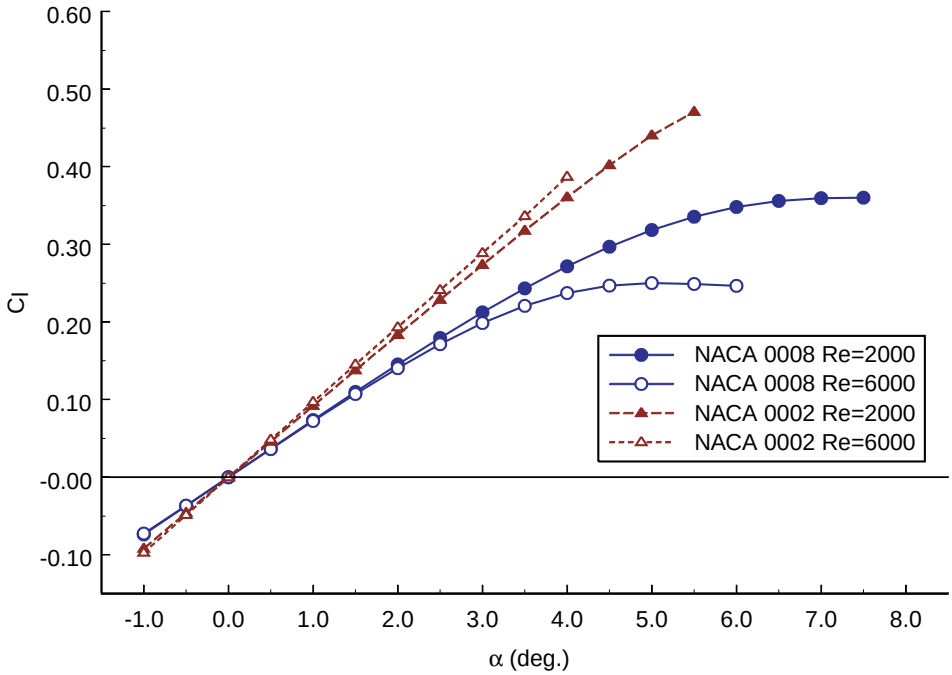

FIGURE 3.4 Lift curves for the NACA 0002 and NACA 0008.

Streamlines near the trailing edge of the NACA 0008 are displayed in Figures 3.5 and 3.6 for Re=2000 and Re=6000. The plots begin at the upper edge of the quasi-linear lift range for each Reynolds number. The streamlines originate from identical points in both figures. The onset of trailing edge separation is pushed from α=2.0 at Re=6000 to α=3.0 at Re=2000. As the angle of attack is increased, the lower Reynolds number case achieves almost two degrees higher angle of attack for similar amounts of trailing edge separation.

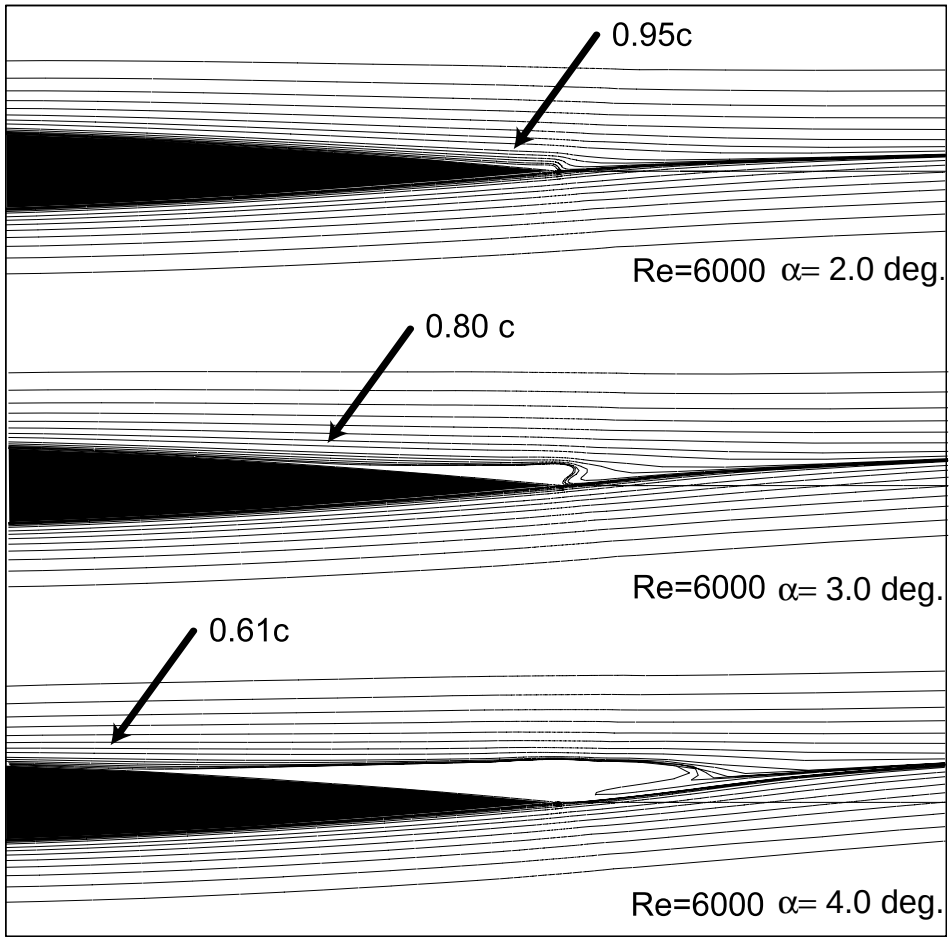

FIGURE 3.5 Streamlines for the NACA 0008 at Re=6000. The aft 45% of the airfoil is visible. The point of separation is indicated by an arrow.

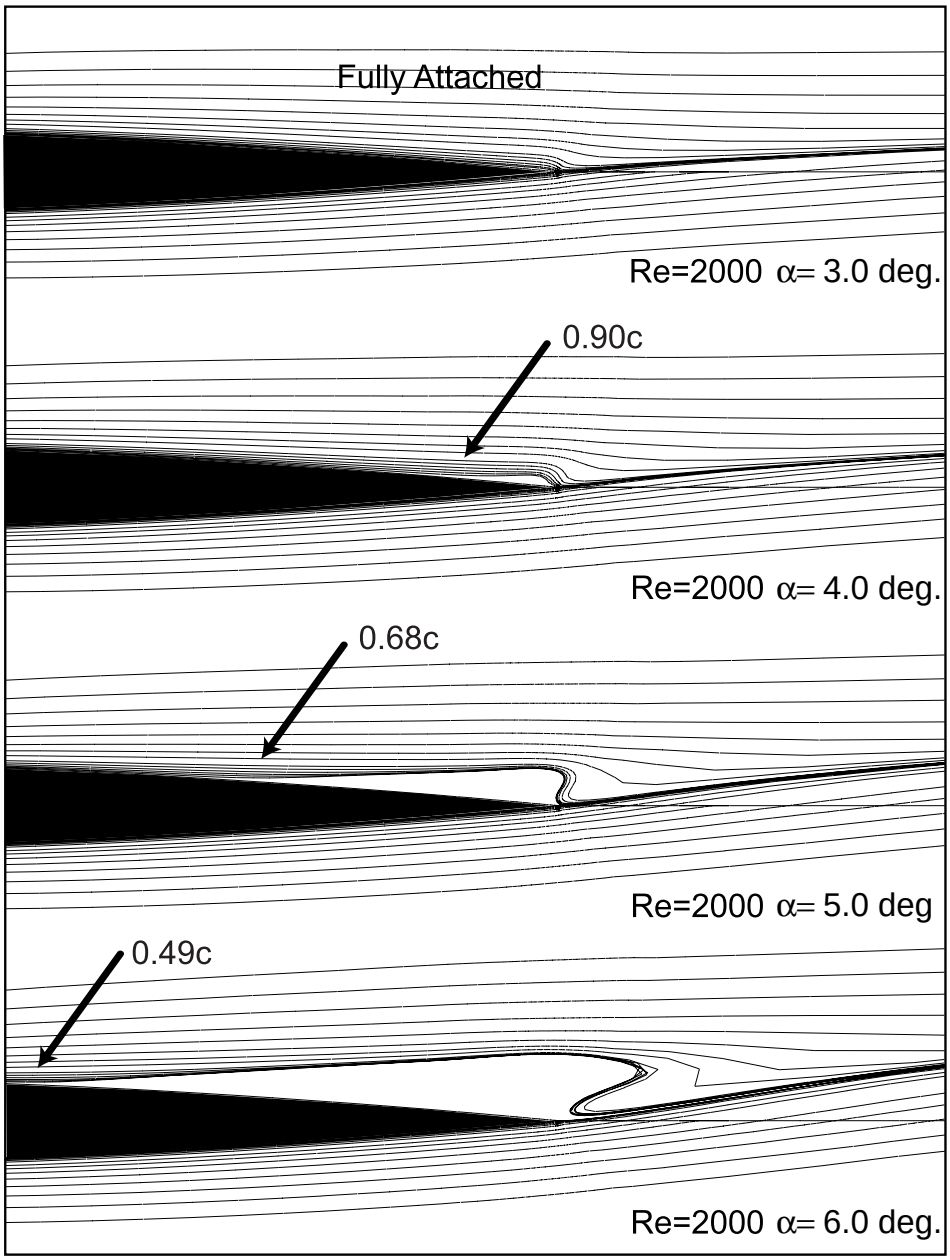

FIGURE 3.6 Streamlines for the NACA 0008 at Re=2000. The aft 45% of the airfoil is visible. The point of separation is indicated by an arrow.

3.3 Maximum Section Thickness Effects

The effects of airfoil thickness variations are investigated using uncambered NACA 4-digit airfoils ranging from 2% to 8% thick in 2% increments. Performance estimates for each section have been computed at Re=6000 and Re=2000. Airfoil thickness variations appear to have two principal performance effects: a drag penalty, due to the pressure recovery attributable to increased thickness, as expected, and a strong reduction in the lift curve slope, an unexpected result.

3.3.1 Effect of Thickness on Drag

The variations in drag with section thickness are illustrated by the airfoil drag polars in Figure 3.7. It is interesting to compare the calculated values with airfoils of the same family operating at higher Reynolds numbers and with the theoretical drag of a fully laminar flat plate. The results of this comparison are shown in Table 3.1. The experimental data for the NACA 4-digit family is taken from Abbot and Von Doenhoff [21] at Re=6 million. Over a practical range of section thickness and for a fixed Reynolds number, the relationship between maximum thickness and the zero lift drag is well approximated by a linear function for this family of airfoils. The effect of thickness variations on the zero lift drag may be expressed as a reference zero thickness drag and a slope. The theoretical laminar plate drag at comparable Reynolds numbers is also included for comparison.

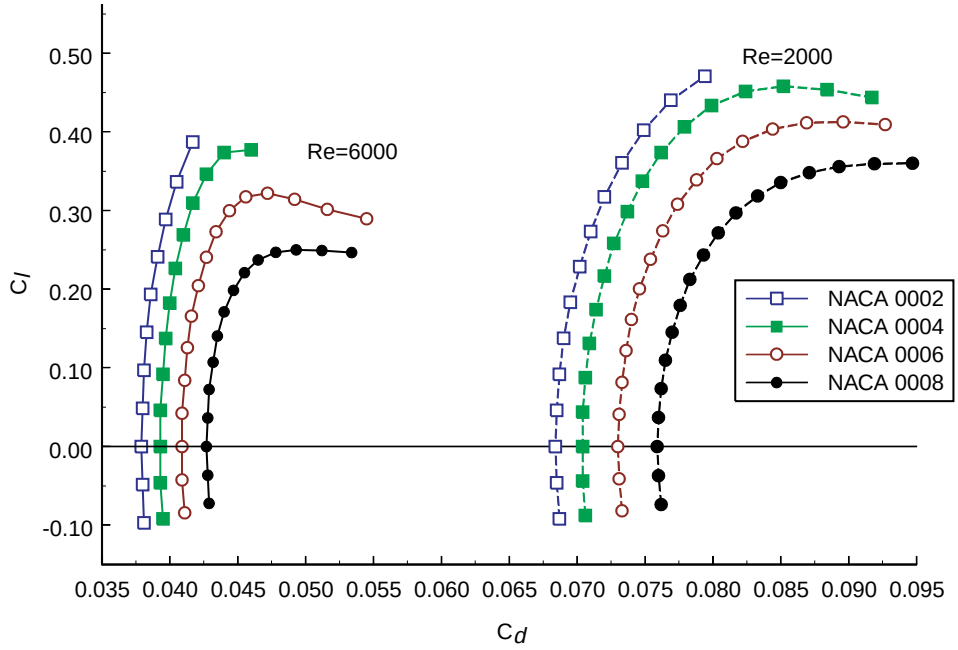

FIGURE 3.7 Drag polars for uncambered NACA 4-digit sections across a range of thickness.

TABLE 3.1 Effect of Airfoil Thickness Ratio and Reynolds Number on Zero Lift Drag.

	Zero Thickness, Zero Lift C_d	C_d Increase per % t/c	Laminar Plate C_d
Re=2000	0.0656	0.0013	0.0593
Re=6000	0.0362	0.0008	0.0342
Re=6.0e6	0.0040	0.0002	---

The most apparent results are a general consequence of operation at ultra-low Reynolds numbers. Lift coefficients are similar to those seen at much higher Reynolds numbers, but the drag coefficients increase by an order of magnitude. This results in section lift-to-drag ratios in the range of one to ten as opposed to 60 to 100. Also notable is the large increase in the drag coefficient, nearly doubled, between the Re=6000 and Re=2000 results. In this regime, small changes in the Reynolds number result in large variations in the drag coefficient. The theoretical laminar flat plate drag mimics these trends and provides a good estimate of the zero lift, zero thickness airfoil drag. The theoretical

laminar flat plate drag differs from the linear extrapolated zero thickness drag by only 5.3% at Re=6000 and 9.5% at Re=2000. This result is not surprising considering that the zero thickness reference is a flat plate. This supports the assumption of a linear relationship between thickness and zero lift drag for these airfoils.

The drag penalty associated with increasing thickness grows as the Reynolds number is decreased, but the rates of drag increase relative to the zero thickness drag are similar. Although the magnitudes of the drag coefficients increase dramatically, they are in line with simple laminar plate results. The variations with thickness also exhibit an order of magnitude increase, but the trends are consistent with those at higher Reynolds numbers.

3.3.2 Effect of Thickness on Lift

Within the quasi-linear range, the inviscid lift curve slope of an airfoil benefits from increased thickness, with thicker sections obtaining as much as a 10% increase in lift curve slope over the thin airfoil value of 2π per radian. At more conventional Reynolds numbers, viscous effects then degrade the lift curve slope, with the net result of lift curve slopes 5% to 10% below the inviscid thin airfoil value. The increased thickness of the upper surface boundary layer relative to the lower surface boundary layer at a positive angle of attack effectively reduces the camber of the airfoil or equivalently, the angle of attack. The inviscid gains and viscous losses typically cancel, resulting in lift curve slopes close to 2π per radian across a range of section thickness. This is not the case for Reynolds numbers below 10,000.

In this range of operation, the viscous boundary layer growth dominates and increasing thickness results in a significant decrease in lift curve slope in the quasi-linear range. The results in Figure 3.8 show as much as a 35% reduction in lift curve slope for the 8% thick section from the idealized value of 2π. The 2% thick sections come closest to the inviscid thin airfoil value, showing a 15% reduction. The effect of reducing the Reynolds number from 6000 to 2000 is a further reduction in lift curve slope. In Figure 3.9, the 2% section shows the greatest effect with a 5.2% reduction relative to the

Re=6000 result due to Reynolds number effects. The effect of Reynolds number variations decreases with increasing thickness.

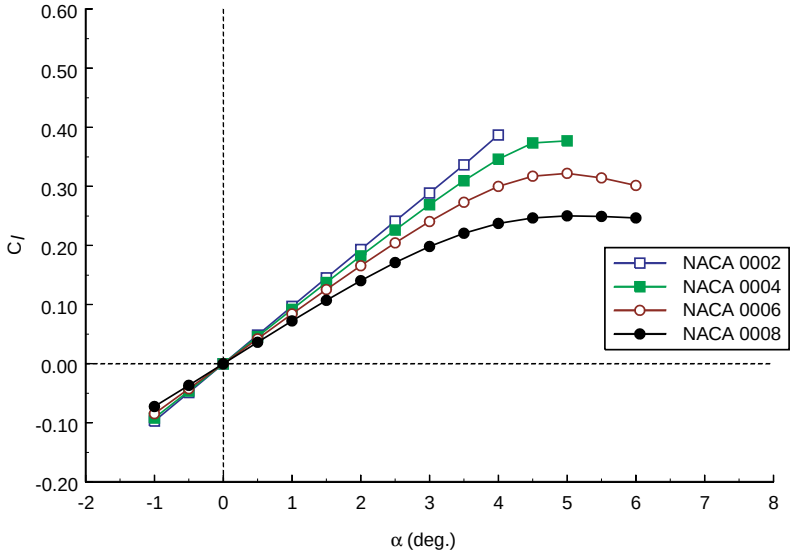

FIGURE 3.8 Lift curves for uncambered NACA 4-digit sections, Re=6000.

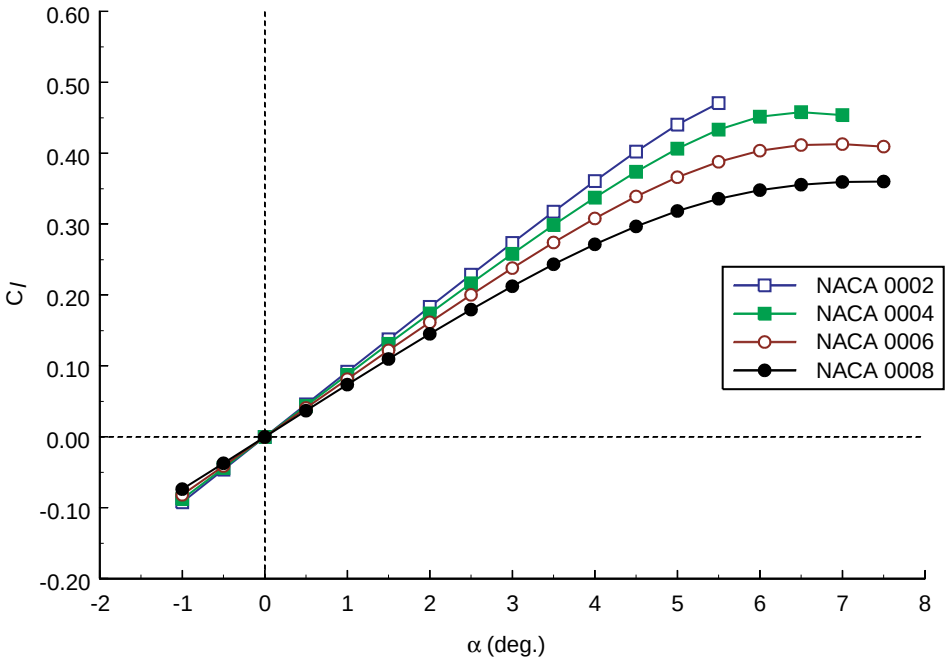

FIGURE 3.9 Lift curves for uncambered NACA 4-digit sections, Re=2000.

The decambering effect of the boundary layer is visualized by considering constant velocity contours in the flow field. The area of reduced flow velocity is large and the validity of defining a reference boundary-layer edge value is debatable. The contours are chosen at a fixed fraction of the freestream velocity, low enough to be considered within the boundary layer, providing a qualitative notion of the boundary layer geometry. Several $0.2V_\infty$ contours are drawn for the NACA 0002 and NACA 0008 sections at Re=6000 in Figure 3.10. Three angles of attack, 0.0, 2.0, and 4.0 degrees are shown. These respectively represent the zero lift condition, the upper limit of the quasi-linear lift range for the NACA 0008, and a point within the nonlinear range where trailing edge separation comes into play. The boundary layer has little effect on the effective geometry of the NACA 0002, but the thicker upper surface boundary layer of the NACA 0008 significantly decreases the effective camber of the airfoil.

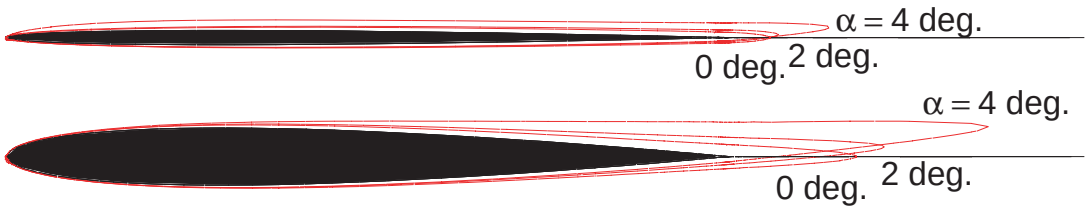

FIGURE 3.10 NACA 0002 and NACA 0008 boundary layer development at Re=6000.

A second effect of increasing thickness is a more rapid reduction in the lift curve slope once past the quasi-linear lift range, attributable to earlier and more severe trailing edge separation. The NACA 0002 is fully attached up to stall, however the NACA 0008 is not, as presented in Figure 3.5. Figure 3.5 begins at α=2.0, the edge of the quasi-linear range. The flow is almost fully attached at α=2.0, with visible trailing edge separation at 95% chord, but by α=3.0 there is significant separation at 75% chord. This moves to 60% at α=4.0. These separated regions result in a large displacement of the flow within the aft boundary layer, increasing the decambering effect and resulting in larger reductions in the lift curve slope compared to the fully attached NACA 0002.

3.4 Effect of Camber

The boundary-layer decambering effect just described suggests that the introduction of camber may offer the potential for significant performance gains over a simple flat plate. The effects of camber do not differ significantly from those at much higher Reynolds numbers, but the fact that the detailed geometry is still an effective driver of performance at such low Reynolds numbers is itself a useful conclusion.

A comparison of the NACA 0002 and NACA 4402 airfoils indicates the gross effects of camber on performance. Lift curves and drag polars are provided for Re=1000, Re=2000, and Re=6000 in Figures 3.11 and 3.12. As at higher Reynolds numbers, the first order effect on the lift curve is a translation towards lower zero lift angles of attack

with increasing camber. The addition of 2% camber results in a 2.0 to 2.5 degree shift in the zero lift angle of attack. The maximum steady-state lift coefficients also increase. In the case of these two sections, there is a 30% increase in the maximum steady-state lift coefficient. Although the drag also increases, in this regime the ability to attain higher lift coefficients generally results in a net gain in lift-to-drag ratio. Due to the introduction of camber, the maximum lift-to-drag ratio increases from 4.5 to 5.4 at Re=1000 and from 9.3 to 11.0 at Re=6000.

Within the quasi-linear range, the reduction in lift curve slope with a decrease in the Reynolds number is visible for the NACA 0002, but the effect of reducing the Reynolds number on the NACA 4402 is different. It appears as a drift towards higher zero lift angles of attack as the reduction in Reynolds number uniformly reduces the effective camber of the section across the entire quasi-linear lift range. All but the Re=1000 case eventually suffer from leading edge separation, but the onset is delayed slightly. This is most likely attributed to the increase in the ideal angle of attack that comes with the introduction of camber.

The delay in trailing edge separation attributed to reducing the Reynolds number is also visible in the NACA 4402 results. At α=5.0 and Re=6000, separation occurs at 55% chord while at Re=2000 separation is pushed back to 92% chord. The Re=1000 case, at α=5.0, remains fully attached, and at α=7.5, is still attached up to 65% chord. Consideration of the drag polars in Figure 3.12 reveals trends similar to those at higher Reynolds numbers. The addition of camber results in an increase in zero lift drag and an upward shift of the polars towards higher lift coefficients.

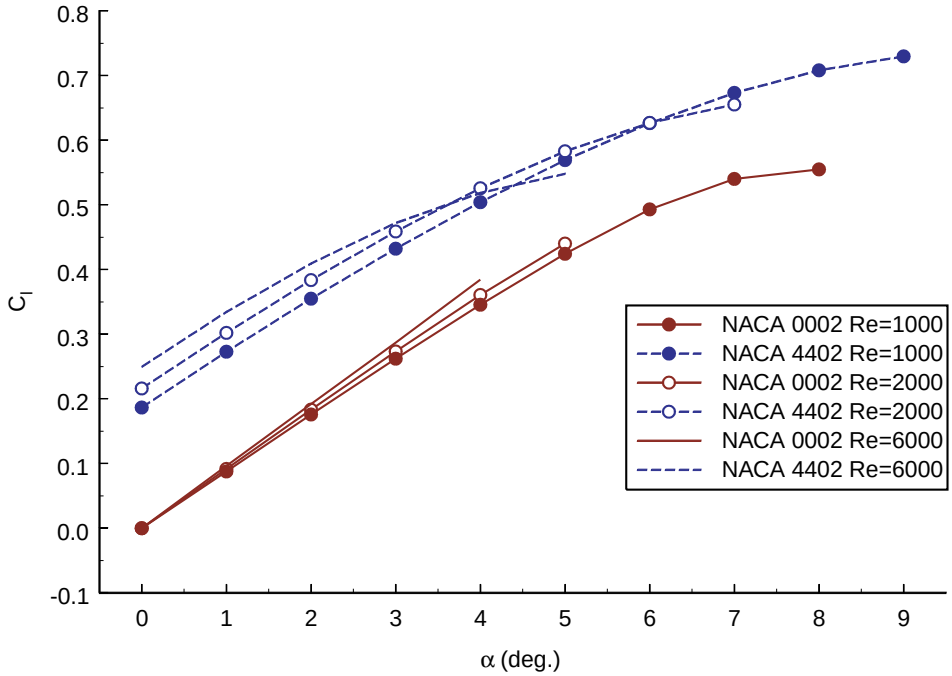

FIGURE 3.11 Lift curves for the NACA 4402 and NACA 0002.

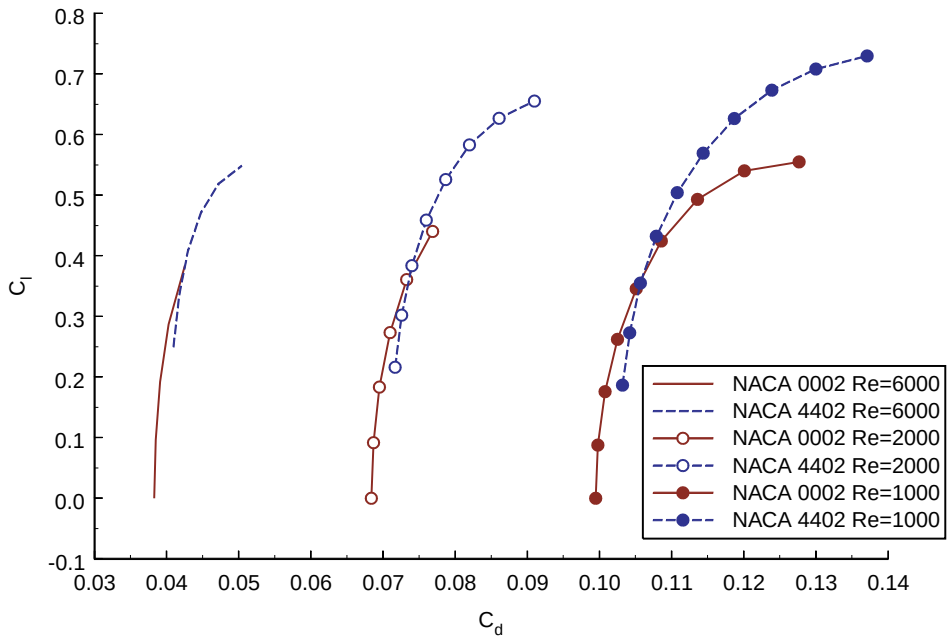

FIGURE 3.12 Drag polars for the NACA 4402 and NACA 0002.

Further analyses have investigated the possible benefits of varying the magnitude and distribution of camber. The design space has been explored using nine airfoils spanning 2% to 6% maximum camber located at 30%, 50%, and 70% chord. All of the sections are 2% thick NACA 4-digit profiles. All calculations have been completed at Re=12,000.

The lift curves provided in Figure 3.13 are for 2% and 4% camber at all three of the chord locations. In both plots, the aft shift of maximum camber results in a less severe reduction of lift past the quasi-linear range, higher attainable lift coefficients, and higher lift-to-drag ratios. This correlates with reduced trailing edge separation for a given angle of attack. The aft cambered sections exhibit separation at a lower angle of attack, but the growth of separation is retarded. As the angle of attack increases, the majority of the suction side experiences less adverse gradients than a similar section with forward camber. The region of separation contained aft of the maximum camber location by maintaining less adverse gradients ahead of it. The aft concentration of camber functions like a separation ramp in the pressure distribution [22].

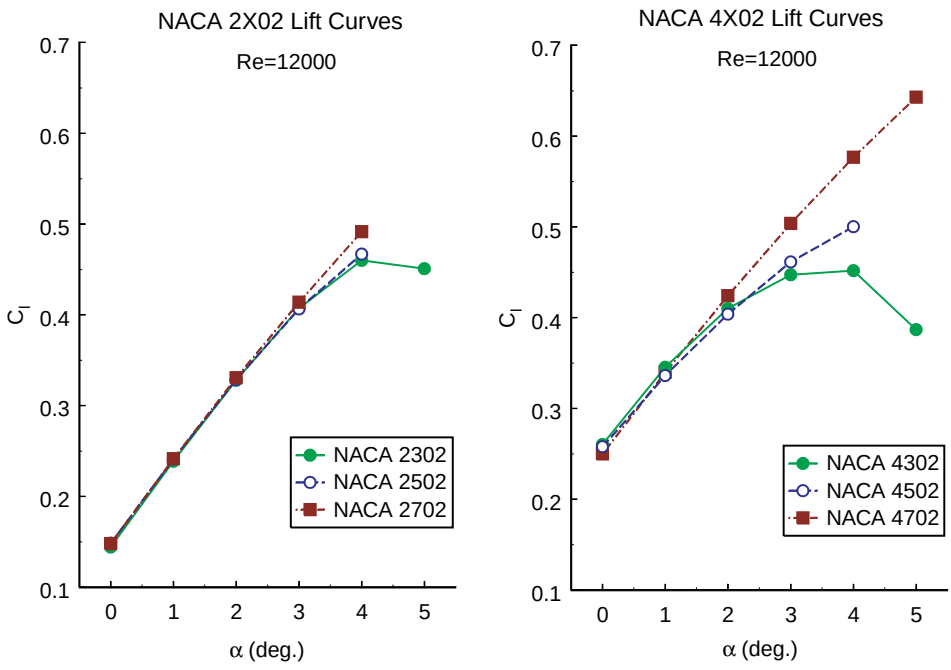

FIGURE 3.13 Lift curves for 2% and 4% cambered NACA 4-digit airfoils at Re=12000.

The maximum L/D for all nine airfoils are provided in Table 3.2. Comparison of the L/D of the nine sections reveals the complexity that is common in airfoil design. Only the 4% aft camber airfoil manages to out perform the 2% camber sections. The other two 4% camber cases are lower in L/D and roughly equivalent in lift coefficient.

TABLE 3.2 Effect of Camber Variations on L/D for a 2% Thick Airfoil.

Max. Camber	L/D with Max. Camber at: 0.3c	0.5c	0.7c
2%	14.7	14.8	15.4
4%	13.7	13.9	15.7
6%	10.4	11.5	14.8

The effects of varying the amount of camber, while fixing the maximum camber location at 70% chord, are depicted in Figure 3.14. The increase in camber causes a nonlinear penalty in drag for a given lift coefficient, but this is tempered by the ability to attain higher lift coefficients. For these three cases, the maximum L/D is attained by the 4% camber section. The effects of varying the location of maximum camber for a fixed 4% camber is depicted in Figure 3.15. The aft movement of camber results in significantly higher drag below C_l=0.45, but this geometry is able to achieve higher lift coefficients within the steady-state operating limitation.

This simple nine-point test matrix indicates that aft camber is beneficial. Selection of the amount of camber is less clear, but for this particular camber definition, it should lie in the midrange of the values considered. This study is not meant to be a detailed indicator for design; the camber-line is rather rigidly defined and the sampling is sparse. It is however, indicative of the large variations in performance that exist within the design space and some of the physical trends responsible.

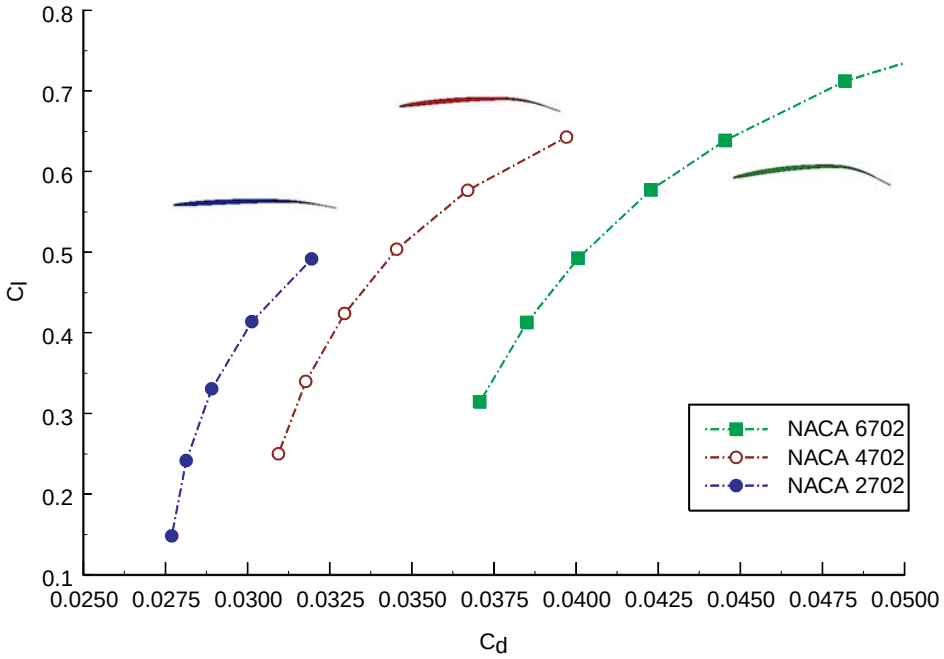

FIGURE 3.14 Drag polars for NACA 4-digit airfoils with varying amounts of camber. The maximum camber location is fixed at 70% chord.

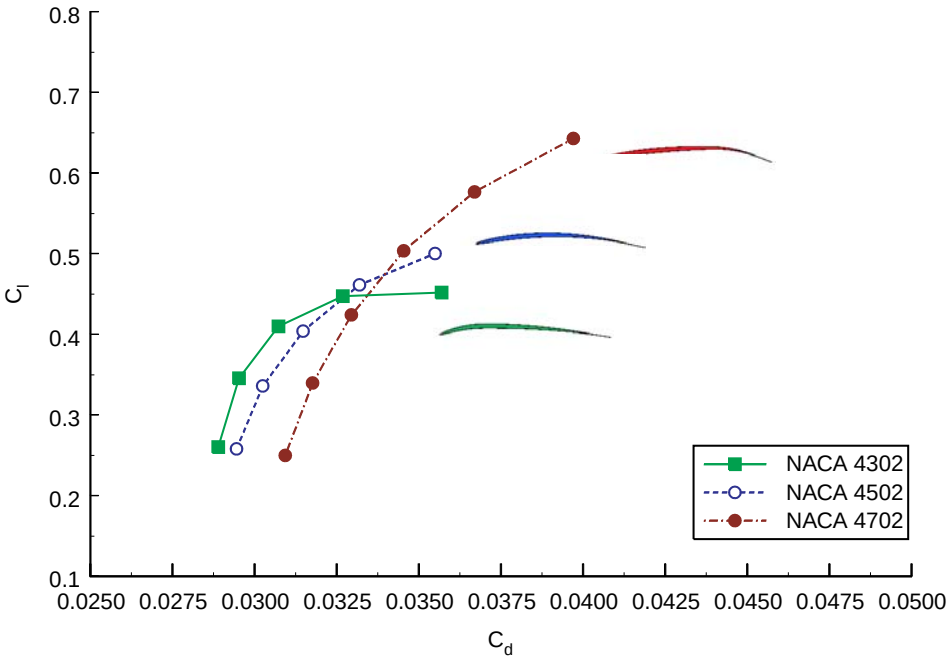

FIGURE 3.15 Drag polars for NACA 4-digit airfoils with variations in the maximum camber location and a fixed 4% camber.

3.5 Effect of Leading Edge Shape and Constant Thickness Profiles

Manufacturing limitations at small scales place minimum gage constraints on the design of highly detailed cross-sections. For very small chord-lengths, the choice must be made between a traditional airfoil profile, at the expense of greater maximum thickness, or a constant thickness distribution at the minimum gage. This trade has been explored by comparing the performance of the NACA 0002 airfoil with 2% constant thickness plates. Two leading edge shapes are considered for the plates: a leading edge radius and a nearly blunt leading edge. The trailing edge begins as a radius at 99% chord, but is brought to a point to simplify grid generation. The three edge shapes are shown in Figure 3.16.

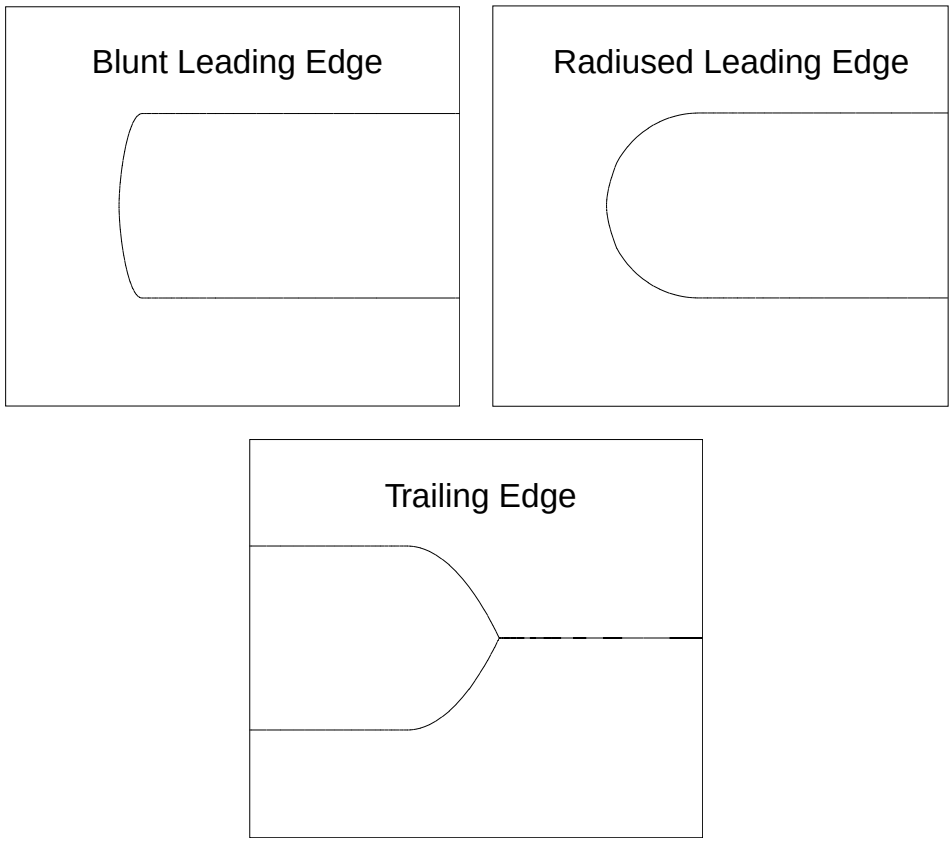

FIGURE 3.16 Leading and trailing edge shapes used in the generation of constant thickness airfoils.

The primary effect of substituting a constant thickness profile is a uniform increase in drag across the operating range of the section. Drag polars for the NACA 0002 and both plate sections are shown in Figure 3.17. The drag increase ranges from 9% at Re=6000, down to 5% at Re=1000. The increase is comparable to changing from a 2% thick airfoil to a 5% thick airfoil, but the plates suffer penalties in L/D that are minimal compared to an actual increase in thickness. An increase in section thickness results in a significant reduction in the lift curve slope and lower maximum steady-state lift coefficients, but the constant 2% thickness plates exhibit no reduction in lift curve slope relative to the 2% thick airfoil. An actual increase in thickness would result in a 20% to 25% reduction in maximum L/D. The penalty for the constant thickness plates is only 5% at Re=1000 and Re=2000, increasing to 18% at Re=6000.

The leading edge shape effects the formation of the leading edge separation bubble. At Re=1000, both plates are fully attached past α=4.0. Once leading edge separation occurs, the radiused plate gains less than one-half degree of angle of attack for equivalent bubble lengths. Leading edge separation bubbles appear on the plates earlier than on the airfoil, but the leading edge stall on the airfoil occurs very quickly. The net effect is a minimal penalty in lift for the plates. At Re=6000, the leading edge separation bubble forms almost immediately on the blunt section, but does not form on the radiused plate until 1.5 degrees later. Avoiding a blunt profile is advisable, but as the Reynolds number and maximum section thickness are reduced, the details of the thickness distribution become less relevant and the camber-line becomes the dominant factor in performance.

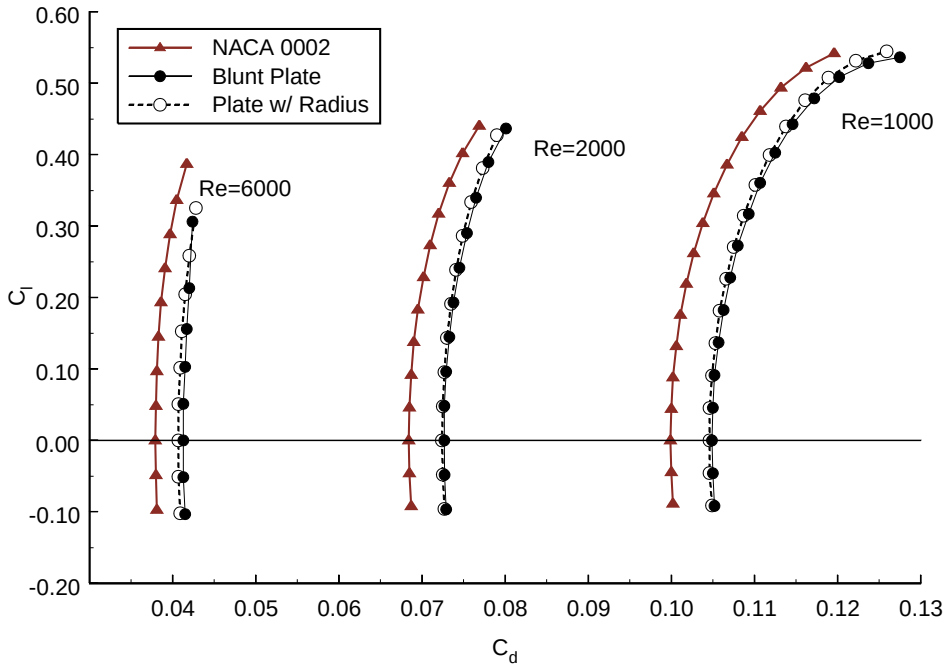

FIGURE 3.17 Drag polars comparing the radiused and blunt constant thickness airfoils with the NACA 0002.

3.6 Design of Optimal Camberlines for Ultra-Low Reynolds Numbers

It is not difficult to select the section with the highest L/D from the analyses of camber variations from Section 3.6, but the 9-point matrix, a very rigid and sparse sampling of the design space, only scratches at the surface. Detailed exploration of the design space requires additional degrees of freedom for the geometry and calls for some form of automated optimization.

The optimization study makes use of the previous results to simplify the problem to its most essential elements. The maximum section thickness is fixed at 2% and the NACA 4-digit thickness distribution is used. This should not affect the utility of the results and greatly facilitates automated grid generation. A specified thickness distribution reduces

the number of variables considered, but also simplifies the problem by removing minimum thickness constraints.

The free design element is the camber-line. A wide range of variation is achieved by modeling the camber-line with an Akima spline [23] anchored at the leading and trailing edges. Akima splines provide a curve that has the benefits of a tensioned spline, avoiding the undulations that can occur with simple cubic splines, but are generally smoother across the control points and require no tuning to achieve a satisfactory interpolation. Four interior control points, or knots, are used to define the camber-line. These are evenly distributed and their chordwise locations are fixed. The knots move perpendicular to the chord-line bounded by upper and lower camber limits.

The optimization study utilizes a constrained simplex optimizer, a modified Nelder-Mead simplex [24], coupled with the INS2d code and a grid generator to maximize the section lift to drag ratio. With only four design variables, this simple optimization method is sufficient. In addition, each 2-D steady-state solution of the flow solver is relatively inexpensive. The simplex method is simple to implement and does not require (possibly noisy) gradient calculations. It is likely not the most efficient option, but the small problem size and inexpensive flow calculations make it a good solution.

Function evaluations are also designed to increase the robustness of the method. A coarse performance polar is generated for each function evaluation. Calculations begin at a low angle of attack where convergence is likely, regardless of geometry. The angle of attack is then increased until the L/D passes its maximum or the flow solver fails to converge. In either case, an objective function value is generated for that geometry and the process continues. This is computationally expensive and not particularly efficient, but it is robust and effective.

Two airfoils have been developed using this approach for Re=6000 (R6) and Re=2000 (R2). Both sections are shown in Figure 3.18. The optimization runs were initialized with a flat plate airfoil, but the converged solutions have been checked by restarting with a geometry near the upper camber limits. Both airfoils exhibit similar features with a

prominent droop near the nose, well-defined aft camber, and distinct hump in the camber distribution that begins near 65% chord and reaches its maximum height at 80% chord. Since the spline knots are located between 20% and 80% chord, the regions between the edges of the airfoil and the outermost control points are constrained to be nearly linear. The camber distributions are plotted in Figure 3.19. The NACA 4402 and NACA 4702 are also plotted for reference.

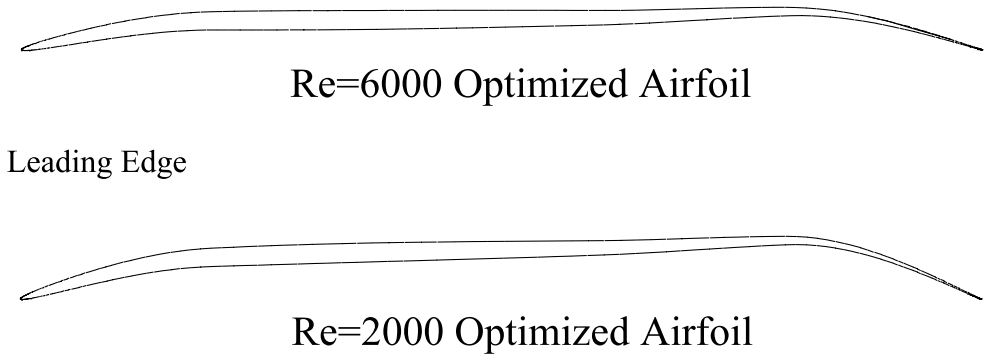

FIGURE 3.18 Optimized airfoils for Re=6000 and Re=2000.

Optimization at Re=6000 resulted in a maximum camber close to 4%, but the R2 solution increases to 6% camber. This increase compensates for the larger reductions in effective camber at lower Reynolds numbers. The R2 airfoil achieves a maximum L/D of 8.2, 5% higher than best 4-digit section analyzed at this Reynolds number, the NACA 4702.

The 4% camber of the R6 airfoil is closer to the 4-digit airfoils examined earlier and provides a better point for comparisons. This airfoil achieves an L/D of 12.9, 4% better than the NACA 4702 and a 16% improvement over the NACA 4402. Figure 3.20 shows the L/D versus geometric angle of attack for these three airfoils. The optimized section begins to show gains past α=3.0, increasing until the maximum L/D is reached at α=5.0.

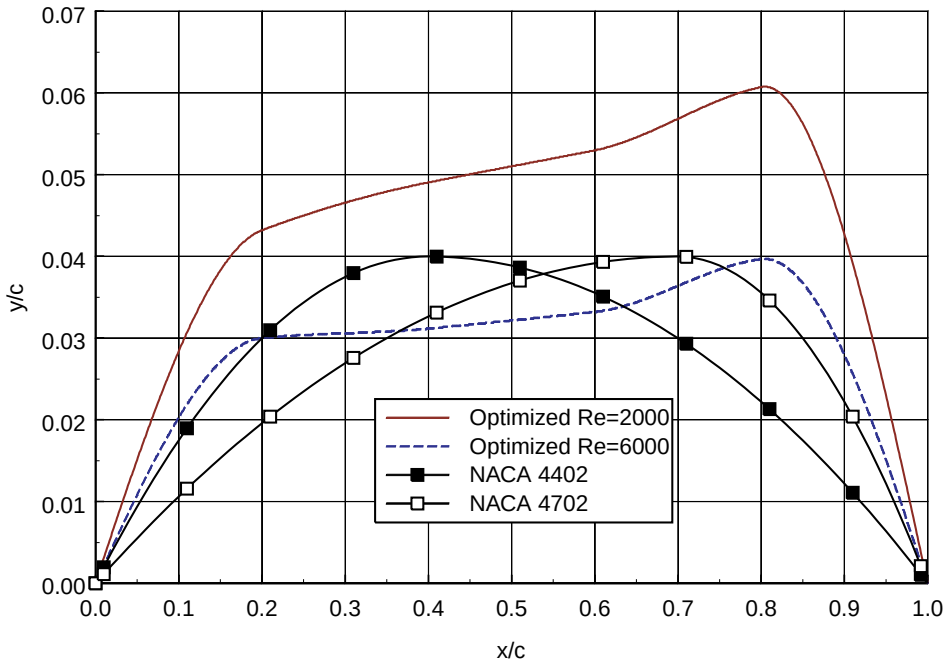

FIGURE 3.19 Camber distributions for the Re=6000 and Re=2000 optimized airfoils.

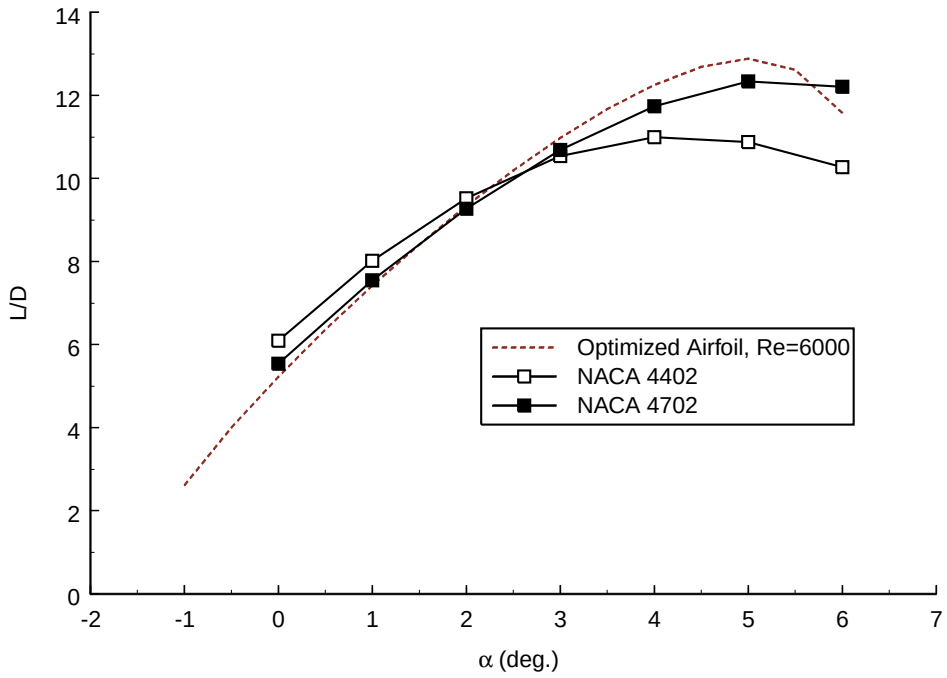

FIGURE 3.20 L/D for the optimized Re=6000 airfoil and two NACA 4-digit airfoils.

Small improvements in lift and drag allow the R6 airfoil to outperform the NACA 4702. The majority of the gains in lift and drag are connected to 5% less trailing edge separation on the R6 compared to the NACA 4702. The optimizer is attempting to exploit the benefit of limiting trailing edge separation. The maximum camber is moved to the aft control point and this region is once again operating similarly to what has been described as a separation ramp [22].

For the R6 airfoil at α=4.0, trailing edge separation occurs at 88% chord, growing to 86% at α=5.0. Beyond α=5.0, the separation point almost immediately moves forward to 30% chord. Figure 3.21 shows the aft 60% pressure distributions at α=4.0 and α=5.0. At α=4.0, there is a distinct inflection in the pressure distribution at 82% chord. The reduced adverse gradients ahead of this location do not allow separation to move forward, but the steeper gradients following the inflection point promote separation on the aft portion of the airfoil. This is the primary source of the drag penalty seen with aft camber at low angles of attack. At α=5.0, the inflection is no longer present. Beyond α=5.0, separation moves far forward and performance rapidly degrades.

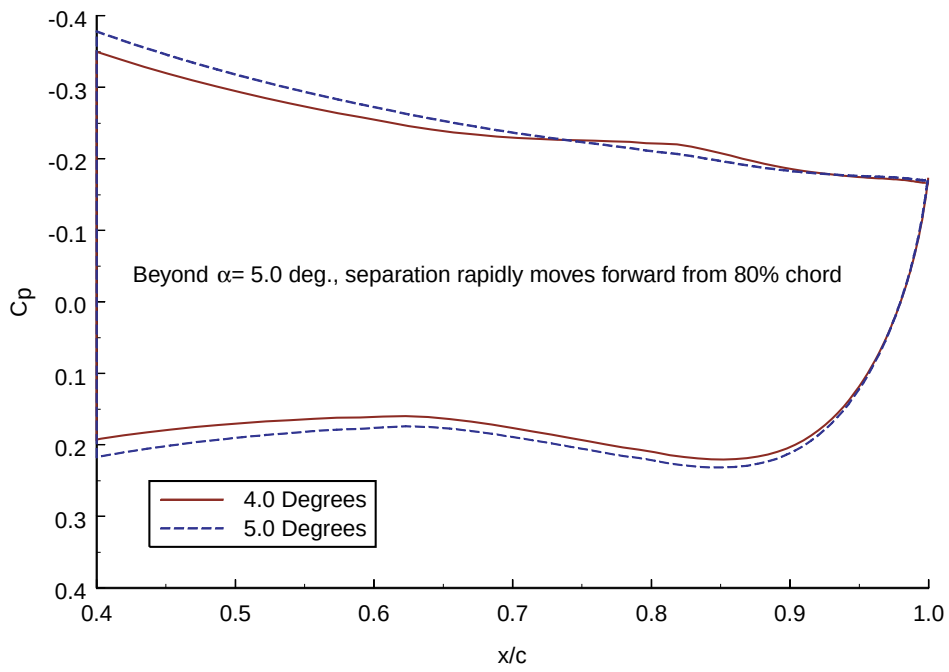

FIGURE 3.21 C_p distributions for the Re=6000 optimized airfoil.

The optimized design of these two airfoils highlights the ability of small modifications in geometry to be very effective in altering section performance. Additional degrees of freedom may be easily added to the problem by introducing more spline knots, but this simple four-variable problem succeeds in achieving significant performance gains over smooth formula-based camber-lines. The lack of experience at ultra-low Reynolds numbers makes optimizers an effective and important tool, not only for design, but also for enhancing our understanding of this flight regime.

Chapter 4

Hybrid Method for Rotor Design and Analysis

4.1 Introduction

A rapid rotor analysis code and a closely related design code have been created for the development of ultra-low Reynolds number rotors. These tools build upon the exploratory two-dimensional aerodynamics analyses of Chapter 3 towards the goal of designing, manufacturing, and testing micro-rotors for practical application. The toolset consists of a rotor performance package coupled with a nonlinear optimizer and section data from two-dimensional Navier-Stokes analyses. Key to the design capability is the integration of theoretical or experimental motor data within the optimization framework. The limited number of suitable small scale motors is a critical practical constraint and the inclusion of motor curves assures that the design is developed at an obtainable operating point. This code is currently applicable for the static thrust (hover) case only, but the development of relations for steady climbing rotors is also described.

4.2 Derivation of the Rotor Thrust and Torque Equations for Hover and Vertical Climb

The process of solving the rotor analysis and design problem for vertical flight can appear rather convoluted, but one form of the solution may be reduced to the following problem statement:

Given the following parameters:

Rotor speed (ω), radius (*R*), ascent rate (*U)*, blade count (*B*), and the chord and lift distributions.

Find:

Thrust (*T*), torque (*Q*), and the required blade pitch distribution (θ(r)).

As part of the process, it is necessary to solve the intermediate problem determining the velocities induced by the rotor. This section develops the necessary relations, beginning with the simplest and most fundamental forms and progressively adding fidelity (and complexity) as necessary for this application.

The formulation of the design and analysis tools utilizes the combination of actuator ring theory and the blade element approach. Thrust and torque expressions are developed for a differential blade element and are then integrated along the length of the blade. Two equations are developed for thrust and two for torque. One pair is based on actuator ring theory; the other is based on blade-element theory. A typical section is presented in Figure 4.1 and depicts the relevant angles and velocity components.

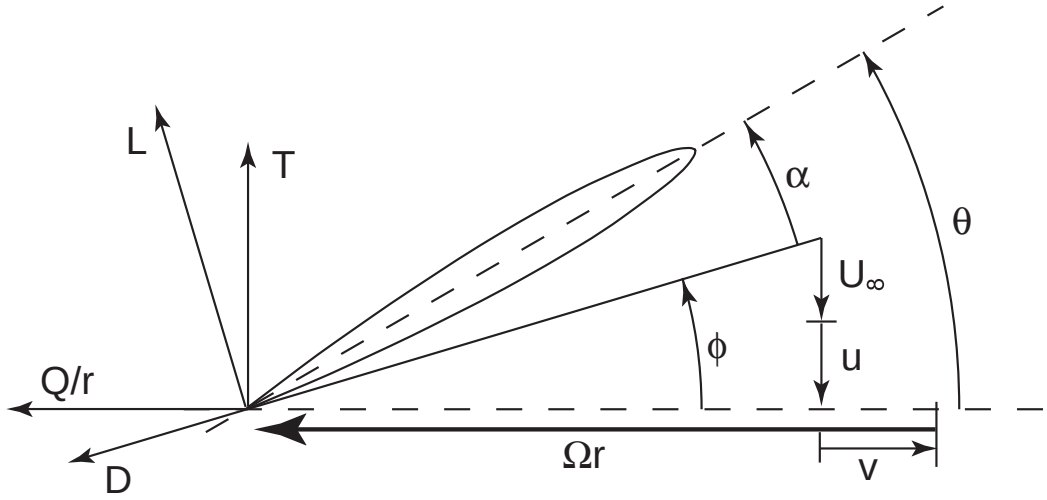

FIGURE 4.1 Typical section diagram for the rotor equations. All variables are positive as depicted.

4.2.1 Fundamental Actuator Ring Equations

The actuator ring relations for thrust and torque are based on the conservation of linear and angular momentum across a differential annulus of the rotor disk and its associated streamtube combined with the incompressible Bernoulli equation to relate pressure and velocity. The action of the rotor is modelled by continuous velocity through the rotor disk, but a discontinuous jump in pressure across it. The flow is assumed to be incompressible and initially, the flow is treated as inviscid and uniform azimuthally across the entire annulus, equivalent to a rotor with an infinite number of blades. For a concise description and derivation of the initial relations, refer to the text by Glauert [25]. Taking the most basic form as a stating point, the thrust and torque per unit radius may be expressed as:

$$dT = 2\rho u(u + U_\infty)(2\pi r)dr \tag{4.1}$$

$$dQ = 2\rho v(u + U_\infty)(2\pi r)r\,dr \tag{4.2}$$

Neglecting for now any corrections to the swirl velocity, the viscous and pressure drag effects are partially accounted for by adding the two-dimensional section drag:

$$dT = (2\rho u(u + U_\infty)(2\pi r)dr) - (Bq_{local}cC_d \sin(\phi)dr) \quad (4.3)$$

$$dQ = (2\rho v(u + U_\infty)(2\pi r)rdr) + (Bq_{local}cC_d \cos(\phi)rdr) \quad (4.4)$$

4.2.2 Blade Element Equations

The blade element formulation assumes that at any given differential blade element the generated forces can be determined from two-dimensional section properties and the local bound circulation. Given the geometry of the blade and the local inflow vector, thrust and torque can be determined. These forces are then integrated over the length of the blade and the total number of blades. The lift force is often expressed in terms of the aerodynamic angle of attack and a constant lift curve slope. In this formulation, however, the lift is represented by the Kutta-Joukowski relation: $d\vec{F} = \rho(\vec{V} \times \vec{\Gamma})dr$. At higher Reynolds numbers, the lift curve slope is nearly constant over the operating range, but operation at ultra-low Reynolds numbers results in highly nonlinear lift curves and a need to operate close to the steady $C_{l\,max}$ as demonstrated in Chapter 3. Resolving the circulatory force into axial and azimuthal components and adding viscous terms as shown in Eqns.4.3 and 4.4, yields two blade element relations for thrust and torque:

$$dT = (B\rho(\Omega r - v)\Gamma dr) - (Bq_{local}cC_d \sin(\phi)dr) \quad (4.5)$$

$$dQ = (B\rho(u + U_\infty)\Gamma rdr) + (Bq_{local}cC_d \cos(\phi)rdr) \quad (4.6)$$

4.2.3 Elimination of Trigonometric Terms

The trigonometric terms are eliminated in all four relation by utilizing the relationship between the sectional lift and the Kutta-Joukowski relation:

$$dL\sin(\phi) = \rho(u + U_\infty)\Gamma dr \tag{4.7}$$

$$dL\cos(\phi) = \rho(\Omega r - v)\Gamma dr \tag{4.8}$$

Substitution yields:

Actuator Ring:

$$dT = (2\rho u(u + U_\infty)(2\pi r)dr) - (B(C_d/C_l)(u + U_\infty)\rho\Gamma dr) \tag{4.9}$$

$$dQ = (2\rho v(u + U_\infty)(2\pi r)rdr) + (B(C_d/C_l)(\Omega r - v)\rho\Gamma rdr) \tag{4.10}$$

Blade Element:

$$dT = (B\rho(\Omega r - v)\Gamma dr) - (B(C_d/C_l)(u + U_\infty)\rho\Gamma dr) \tag{4.11}$$

$$dQ = (B\rho(u + U_\infty)\Gamma rdr) + (B(C_d/C_l)(\Omega r - v)\rho\Gamma rdr) \tag{4.12}$$

For a given lift distribution, rotor speed, and chord distribution, the thrust and torque relations of Eqns. 4.9 - 4.12 are a closed set of equations for (u,v) and subsequently (T,Q). However, in their present form, the two resulting equations for (u,v) are both tightly coupled and highly idealized. Additional effort is required to obtain a more usable form and to incorporate potentially important inviscid and viscous effects currently neglected in the rotor model.

4.2.4 Prandtl Tip Loss Correction

Up to this point, the actuator ring equations have assumed constant induced velocities across any particular annulus. In reality, the presence of a finite number of blades can result in a circulation distribution that is markedly different from the infinite blade limit. A simple correction to this assumption is obtained by applying a form of the Prandtl tip loss factor. This correction is based on a cylindrical vortex helices in the wake. Contraction of the wake is not considered in this model, but a modification for contraction effects will be introduced later in this chapter. For details of the development, the reader is referred to the text by McCormick [26]. Defining the Prandtl tip loss factor as κ:

$$B\Gamma_{actual} = \kappa\Gamma_{\infty blades} \tag{4.13}$$

$$\kappa = (2/\pi)\,\mathrm{acos}(e^{f}) \tag{4.14}$$

$$f = (B/2)\left(1 - \frac{r}{R}\right)\left(\frac{1}{\sin\phi_{tip}}\right) \tag{4.15}$$

This correction is applied to the local bound circulation as modeled by the inviscid portions of the actuator ring equations (Eqns. 4.9 and 4.10):

$$dT = (2\kappa\rho u(u + U_\infty)(2\pi r)dr) - (B(C_d/C_l)(u + U_\infty)\rho\Gamma dr) \tag{4.16}$$

$$dQ = (2\kappa\rho v(u + U_\infty)(2\pi r)rdr) + (B(C_d/C_l)(\Omega r - v)\rho\Gamma r dr) \tag{4.17}$$

4.2.5 Swirl Velocity Considerations

The formulation, at this stage, incorporates only the inviscid induced tangential velocity, also referred to as inviscid swirl. In most conventional large scale, high Reynolds number applications, this is sufficient. For the small scale, very low Reynolds number

applications of interest here, the viscous flow entrainment is an important consideration. The thick wake regions generated by each blade produce a significant 'viscous swirl' effect. The models used for this phenomenon are described later in this chapter, but for the purpose of formulating the equations this term is incorporated by separating the tangential velocity into $v_{inviscid}$ (v_i) and $v_{viscous}$ (v_v). No viscous correction is applied to the vertical induced velocity. The viscous swirl is proportional to Cos(ϕ) while any viscous downwash term would be proportional to Sin(ϕ) and roughly an order of magnitude smaller than the swirl correction.

The viscous swirl is incorporated into all terms of the formulation except in the inviscid portion of the actuator ring equation for torque. The viscous losses are already accounted for in the viscous drag portion of that equation. These substitutions result in the final versions of the four basic relations:

Actuator Ring:

$$dT = (2\kappa\rho u(u + U_\infty)(2\pi r)dr) - (B(C_d/C_l)(u + U_\infty)\rho\Gamma dr) \tag{4.18}$$

$$dQ = (2\kappa\rho v_i(u + U_\infty)(2\pi r)rdr) + (B(C_d/C_l)(\Omega r - v_v - v_i)\rho\Gamma r dr) \tag{4.19}$$

Blade Element:

$$dT = (B\rho(\Omega r - v_v - v_i)\Gamma dr) - (B(C_d/C_l)(u + U_\infty)\rho\Gamma dr) \tag{4.20}$$

$$dQ = (B\rho(u + U_\infty)\Gamma r dr) + (B(C_d/C_l)(\Omega r - v_v - v_i)\rho\Gamma r dr) \tag{4.21}$$

4.2.6 Uncoupled Equations for the Rotor Induced Velocities

Algebraic manipulation of Eqns. 4.18 - 4.21 result in an uncoupled pair of equations for the inviscid rotor induced velocities (u and v_i). The general vertical flight equation for v_i can be expressed as the following quartic:

$$v_i^4 + (2A_1^2A_2)v_i^3 + \left(A_1^4(A_2^2 + U_\infty^2) - 2A_1^2\left(A_2^2 + \frac{U_\infty^2}{2}\right)\right)v_i^2 \qquad (4.22)$$
$$- (2A_1^4A_2(A_2^2 + U_\infty^2))v_i + A_1^4\left(A_2^4 + (A_2^2U_\infty^2) - \frac{U_\infty^4}{4}\right) = 0$$

where:

$$A_1 = \frac{BcC_l}{8\pi r\kappa} \qquad (4.23)$$

$$A_2 = \Omega r - v_v \qquad (4.24)$$

Note that if the Prandtl tip loss correction is applied as defined in Section 5.2.4, this equation is still loosely coupled to u by the presence of ϕ_{tip} in the definition of κ. If this model is used the value of u_{tip} is found iteratively, but the modified contracted wake model presented later in the chapter eliminates this coupling.

The vertical induced velocity (u) can then be expressed as a function of v_i:

$$u = \frac{-U_\infty + \sqrt{U_\infty^2 - 4(v_i^2 - A_2v_i)}}{2} \qquad (4.25)$$

For the hovering cases considered in this work, the restriction that $U_\infty = 0$ allows the expressions for v_i and u to be simplified to:

$$v_i^2 + (A_1^2A_2)v_i - A_1^2A_2^2 = 0 \qquad (4.26)$$

$$u = \sqrt{A_2v_i - v_i^2} \qquad (4.27)$$

4.2.7 The Distinction Between the Analysis and Design Problems

The four relations for thrust and torque (Eqns. 4.18 - 4.21) yield two equations for two unknowns (u, v_i) for each differential blade element. The other three unknown quantities (Γ, κ, and v_v) are treated as dependent functions of the input parameters: the lift distribution, rotor speed, ascent rate, number of blades, and chord distribution. Details of these models are presented in the following sections of this chapter. With values for u and v_i, the required blade pitch distribution, θ(r) may be found as:

$$\theta(r) = \alpha_{geo} + \operatorname{atan}\left(\frac{U_\infty + u}{\Omega r - v_i - v_v}\right) \tag{4.28}$$

The determination of θ(r) is the last step in what can be described as the design case. The geometry of the rotor is only partially defined with the remaining aspects of the geometry revealed as part of the solution. The problem can be solved directly without iteration, but has limited applicability.

Single point design is useful, but solving analysis problems, such as assessing a new design at multiple operating points, or the performance of an existing design, is also an essential capability for developing rotors for practical application. Unfortunately, this requires some form of iteration with this rotor model. Rather than develop a separate method for analysis, a simple modification in the definition of the subsequent optimization problem is described in Section 4.6. This allows a single unified method to be used for analysis and design.

4.3 Viscous Swirl Modeling

The extreme operating conditions being considered precipitates a more detailed consideration of viscous swirl effects, where it might otherwise be neglected at higher Reynolds numbers. Typical section lift to drag ratios range from 50 to 100 at chord Reynolds numbers above 100,000, but, as shown in Chapter 3, section lift to drag ratios drop below ten in the ultra-low Reynolds number regime. In addition, the rotors developed for ultra-low Reynolds number applications typically exhibit high solidity, decreasing the separation between adjacent blades and increasing the likelihood of strong leading/trailing blade viscous wake interactions. The fact that commonly used swirl models have been implemented only on large scale, high Reynolds number rotor blades, if at all, led to the development of several alternative models that directly utilize the 2-D CFD analyses of Chapter 3. The goal is to determine both the need and effectiveness of enhanced viscous swirl models.

4.3.1 Average Wake Deficit Viscous Swirl Model

The first viscous swirl model is based on computed airfoil wake profiles at ultra-low Reynolds numbers. Based on this data, the value of v_v is taken as a predicted average wake deficit velocity. The current data is from INS2d calculations at Re 1000, 2000, and 6000. The section is a 2% constant thickness airfoil with a NACA 4402 camberline, a blunt leading edge, and a radiused trailing edge. The airfoil is operating at 4.0 degrees geometric angle of attack placing it close to conditions for maximum lift to drag ratio. The 2% constant thickness is representative of manufacturing minimum gage constraints and moderate variations to the camberline would have a small effect on the wake profiles compared to the much larger variations caused by Reynolds number and distance from the trailing edge. Therefore, analysis of this single airfoil provides a sufficient basis for this model.

Two input parameters determine the viscous swirl correction for any individual blade element, the chord Reynolds number of the blade element and the local arc length between the trailing edge of one blade and the leading edge of the next. The model applies the computed two-dimensional wake along this arc. Figure 4.2 illustrates the effects of varying the distance from the trailing edge at a fixed Reynolds number. The pronounced translation of the profiles with increasing distance occurs because the grid is aligned with the chordline and the airfoil is at a positive angle of attack. The mean deficit velocity is calculated over the region of the profile where $(u/U_\infty) < 1.0$. As expected, the wake diffuses and dissipates as it moves down stream, decreasing the mean deficit velocity for a given Reynolds number. The effects of reducing the Reynolds number are seen in Figure 4.3 as an increase in the width and intensity of the wake, subsequently raising the mean deficit velocity at a given distance. The distilled INS2d results and the final model are displayed in Figure 4.4.

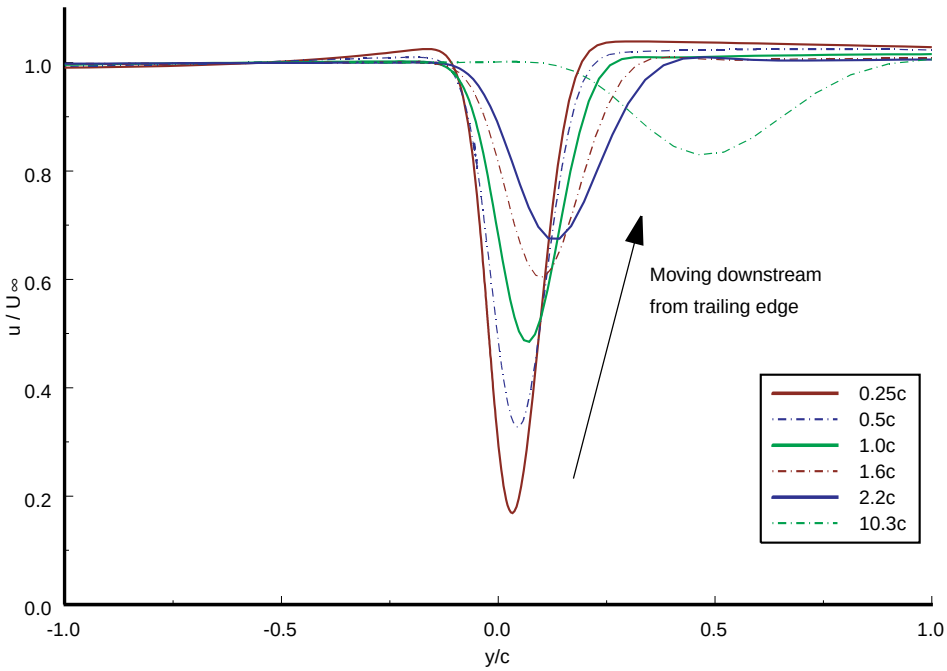

FIGURE 4.2 Effect of downstream distance on wake velocity profiles. INS2d calculation of a 2% thick NACA 4402 camberline, Re=1000, α=4.0 degrees.

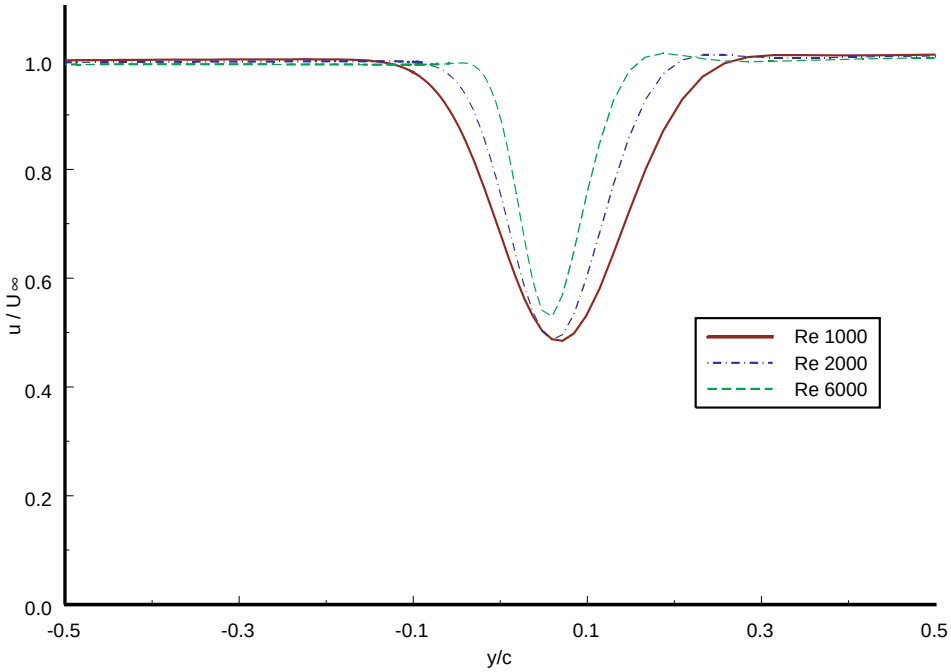

FIGURE 4.3 Effect of Reynolds Number on wake velocity profiles one chordlength aft of trailing edge. INS2d calculation of a 2% thick NACA 4402 camberline, α=4.0 degrees.

This average wake deficit model for v_v is based on a power law fitting across the blade separation arclength and quadratic fitting of the coefficients across Reynolds number:

$$v_v = E_1(arclength)^{E_2} \tag{4.29}$$

where:

$$E_1 = -(3.0e^{-10}(Re_{\Omega r}{}^2)) - (3.0e^{-6}(Re_{\Omega r})) + 0.241 \tag{4.30}$$

$$E_2 = (3.0e^{-9}(Re_{\Omega r}{}^2)) - (7.0e^{-5}(Re_{\Omega r})) - 0.372 \tag{4.31}$$

The most conspicuous simplification in this model is that the CFD computations are for a single airfoil in a constant free-stream flow. Modeling the viscous swirl effect as a two-element system, the forward blade and the trailing blade, neglects the fact the forward blade itself is operating in a non-uniform flow field. This is a large

simplification, particularly at low advance ratios, but attempting to include the coupled effects of all the blades would be difficult with this simple model. An iterative approach to solving the system would steadily drive the total velocity to zero as the viscous terms accumulated. For this reason, the local rotational velocity is taken as the normalization velocity for v_v and as the velocity for the chord Reynolds number. This model should capture the first order viscous wake effects for the specified range of Reynolds numbers, discouraging designs with extremely high local solidity, particularly in the hub region. Unfortunately, as will be discussed in Chapter 6, this model appears to generally over-estimate the viscous swirl effect.

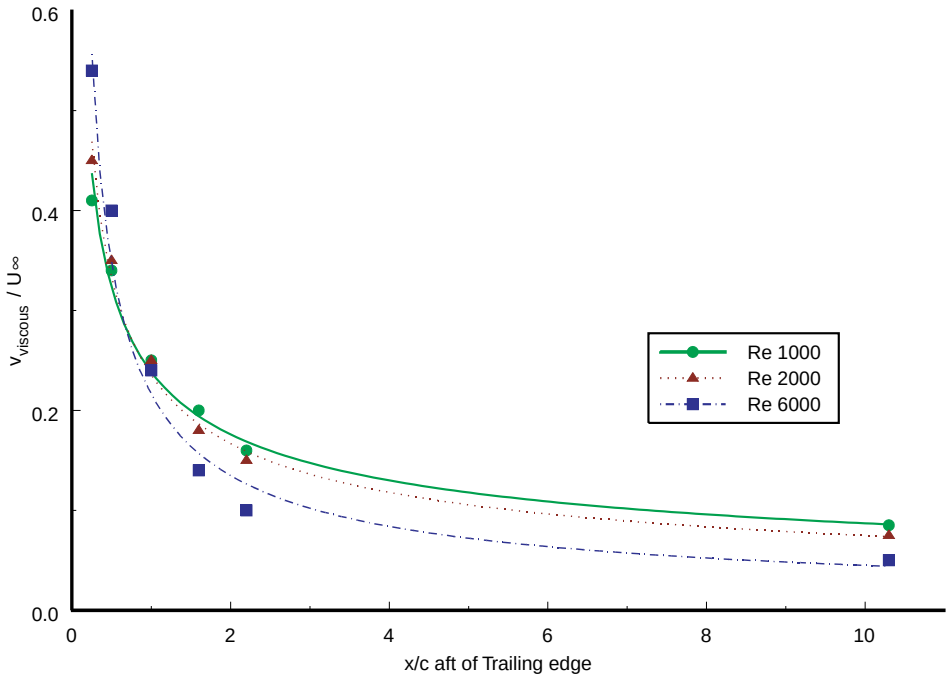

FIGURE 4.4 Average wake deficit model and INS2d data points at three Reynolds numbers.

4.3.2 Gaussian Wake Viscous Swirl Model

The average wake deficit model neglects the effects of rotor downwash and the detail of the wake velocity distribution. That model has been modified in an attempt to account for these two factors. Based on the same INS2d data, the wake deficit velocity profile is

modeled as a Gaussian distribution that varies with Reynolds number, distance aft from the trailing edge, and distance above the trailing edge. This distribution is translated downward based upon the local helix angle and blade spacing. The viscous swirl velocity is then taken as the value of the Gaussian distribution at the intersection of the translated profile and the next blade's leading edge. The velocity distribution takes the form:

$$\frac{u}{U_\infty} = \frac{u_{mid}}{U_\infty}\exp\left(-\frac{(\bar{y}-\mu)^2}{2\sigma}\right) \tag{4.32}$$

where σ is the standard deviation of the wake deficit and $\mu = mean\left(\frac{u}{U_\infty}\right)$

There are several common problems with both of these models. Both treat each set of leading and trailing blade sections as if isolated from the rest of the rotor and operating in a uniform free-stream. Applying either model only once cannot account for the coupled effect of each section on the total rotor system, but the models are unstable if applied iteratively in an attempt to account for the other blades. The additive effects continuously reduce the Reynolds number, increasing the viscous swirl component. Neither model can account for the combined effects of viscous entrainment and rotor downwash. The first model assumes no lift on the rotor, emulating the viscous properties of a spinning solid disk. The Gaussian distribution model incorporates the downwash, but since each pair of sections is treated in isolation there is no rotational flow entrainment permitted ahead of the leading blade (above the rotor). The sharp roll-off of the wake deficit velocity also makes the Gaussian model highly sensitive to the prescribed induced velocities. Once again, this model has proven to be unsatisfactory for reasons to be discussed in Chapter 6.

4.3.3 Conservation of Angular Momentum Viscous Swirl Model

The third viscous swirl model is consistent with the blade element / actuator disk theory used in the rotor performance model. The final form of this model is presented in many texts [27], but the derivation is typically neglected. It is presented here for clarity. This provides a reasonable mechanism for translating the effects of individual blade sections into a uniform viscous swirl velocity. The basis for the model is the conservation of angular momentum within the rotor/wake system. The sum of the moments in the rotor plane applied to the annular wake by the drag of the corresponding blade elements may be expressed as:

$$dM_{rotor} = Brq_{local}cC_d\cos(\phi)dr \tag{4.33}$$

The change in the angular momentum of the wake annulus is:

$$\frac{dH}{dt} = \rho v_v r\frac{d}{dt}(AnnularVolume) = \rho v_v r(2\pi r(u + U_\infty)) \tag{4.34}$$

Solving for v_v yields:

$$v_v = \frac{Brq_{local}cC_d\cos(\phi)dr}{\rho r(2\pi r(u + U_\infty))} \tag{4.35}$$

Substituting into the denominator of Eqn.4.35 from Eqn. 4.3 and a modified form of Eqn.4.5 from blade element theory:

$$dT = Bq_{local}c(C_l\cos(\phi) - C_d\sin(\phi))dr \tag{4.36}$$

yields the simple relation:

$$v_v = 2u\left(\frac{C_d}{C_l}\right) \tag{4.37}$$

Beyond the assumptions intrinsic in blade element / actuator ring theory, the only additional assumptions are that lift is inviscid and plays no direct role in the viscous swirl and any tip loss/wake corrections are neglected in the substitution of the thrust equation. There are no small angle approximations and the simplicity of the final form is due to exact cancellation. Viscous swirl as defined here also incorporates the pressure drag of the section since both are included in the section drag coefficient. The versatility of this model appears to extend reasonably well to the ultra-low Reynolds number regime.

4.4 Development of Stream-Function-Based Vortex Ring Wake Model

Blade-element theory and actuator ring theory alone provide a simple model for the wake and its effects on the rotor. They do not account for any effects of discrete vorticity in the wake due to a finite blade count, instead assuming the wake is composed of continuously shed stream-wise vorticity along stream-tubes. The physical analog is the presence of an infinite number of blades. This model typically overestimates the lift generated near the blade tips.

The Prandtl tip loss correction described earlier is a significant improvement. It is based upon helical vortices of constant strength and diameter emanating from each blade tip. The vertical component of the shed vorticity is neglected and the wake model reduces to a semi-infinite column of vortex rings. The spacing of the rings is determined from the blade spacing and the wake tip helix angle, assuming uniform down-wash. From this potential flow model, the vorticity distribution on the blade is determined and expressed as a correction (κ) to the infinite blade solution. This model is well suited for lightly loaded rotors and rotors with large advance ratios. In these two cases, the assumption of a cylindrical wake is reasonable. The tendency of the helical wake to contract as it

moves downstream is less important in the near field either due to lower vorticity in the lightly loaded case or highly pitched helices due to large advance ratios.

The meso-scale rotor designs typically have a disk loading two to three times lower than a full-scale helicopter, but high rotor solidity, as much as 30%, coupled with a primary interest in hover performance, increases the potential importance of wake contraction. The first-order effects of wake contraction are captured using a model based on an axis-symmetric streamline solution for a vortex ring. Vortex rings are initially stacked as in the Prandtl method with a constant radius equal to the rotor radius, but the rings are then iteratively resized to obtain a wake stream-tube with constant mass flow and no leakage.

The vortex ring stream-function [28] may be expressed in terms of the complete elliptic integrals F_1 and E_1 as:

$$\psi(x, r) = \frac{\Gamma}{2\pi}\sqrt{rr'}\left(\left(\frac{2}{k} - k\right)F_1(k) - \frac{2}{k}E_1(k)\right) \tag{4.38}$$

$$k^2 = \frac{4rr'}{(x - x')^2 + (r + r')^2} \tag{4.39}$$

(x,r) = center-line coordinate and radius of the point of interest

(x',r') = center-line coordinate and radius of the vortex ring

The initial ring strength (Γ) and fixed spacing are determined from the inviscid constant downwash rotor having equivalent thrust to an inviscid rotor with a given C_l distribution:

$$\Gamma = \frac{2\pi T}{\rho\Omega B\pi R^2} \tag{4.40}$$

$$dx = \frac{2\pi u_{ideal}}{B\Omega} = \frac{2\pi\sqrt{\frac{T}{2\rho\pi R^2}}}{B\Omega} \tag{4.41}$$

This equation assumes that the helical wake is convected downstream at the idealized constant downwash velocity based on actuator disk theory. The ring separation distance is equivalent to the resulting pitch of the helix. This is consistent with the model implemented for the Prandtl tip loss correction due to a cylindrical wake as described by McCormick [26].

The Prandtl tip loss model depends upon the choice of a tip helix angle. This is based on the radius, rotation rate, and induced velocities at the tip. This parameter could be determined rigorously from iterative solution of the rotor and wake models, but for an optimal hovering rotor it does not vary considerable from the actuator disk value. The iteration required would also result in a significant increase in the computational cost. Given the approximate nature of the governing wake model, the worth of this improved fidelity is questionable. The additional computational cost would likely be better spent on a more complex wake model such as a helical filament model.

The contracted wake is obtained iteratively by calculating the mass flux through each ring due to the entire wake structure and resizing the ring radius in proportion to the flux ratio of the rotor disk and the wake ring. A great advantage of the streamline formulation is that the mass flux through any axisymmetric circle due to a vortex ring can be directly calculated:

$$S = -2\pi\psi(x, r) \tag{4.42}$$

After calculating the flux through all rings including the rotor disk, the rings are resized according to:

$$r_{new} = r_{old}\left[1 + \left(\frac{S_{rotor} - S_{ring}}{S_{rotor}}\right)\right] \tag{4.43}$$

The resized ring models the horizontal component of a helical filament with a modified pitch. The helical filament must have a constant strength, so the strength of the horizontal ring is modified by the ratio of the cosine of the local pitch angle to the cosine

of the pitch angle at the rotor with the vertical ring spacing held constant. Other models are possible, such as varying the separation at fixed strength, but this model is straight forward and allows the wake length to be rigidly defined as an input.

This procedure is repeated until the wake structure reaches equilibrium, typically five to ten cycles. The ring model extends downstream for five rotor radii. The variations in the computed induced velocities appear to be negligible for larger wake lengths while any increase in length also increases the computational cost due to additional rings. The equilibrium wake form is axisymmetric and symmetric from end to end, with a 'bell-mouth' at the stream-tube exit equal to the rotor disk area. Additional wake length moves the exit 'bell-mouth' further downstream but does not significantly affect the solution and adds to the computational expense. The initial and converged ring configuration for a candidate rotor are displayed in Figures 4.5 and 4.6. The rotor plane is at $x=0$ and the streamtube exit is below the lower extent of the figures.

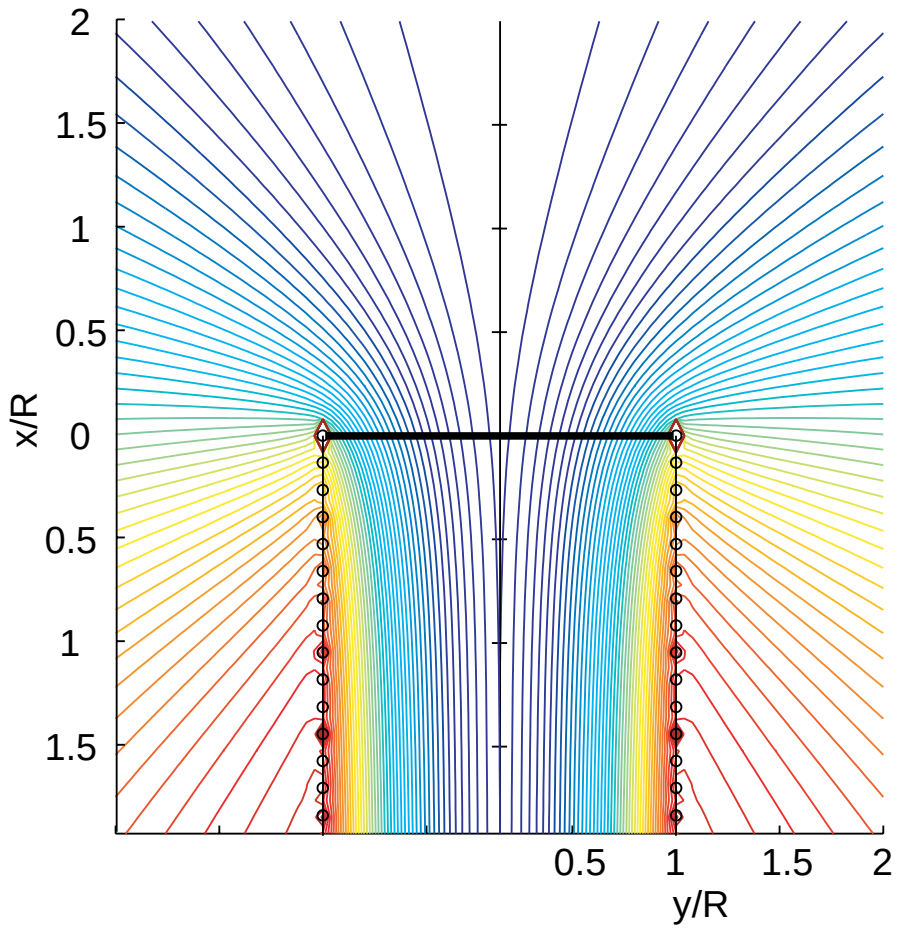

FIGURE 4.5 Initial streamlines for the contracted wake model of a candidate rotor.

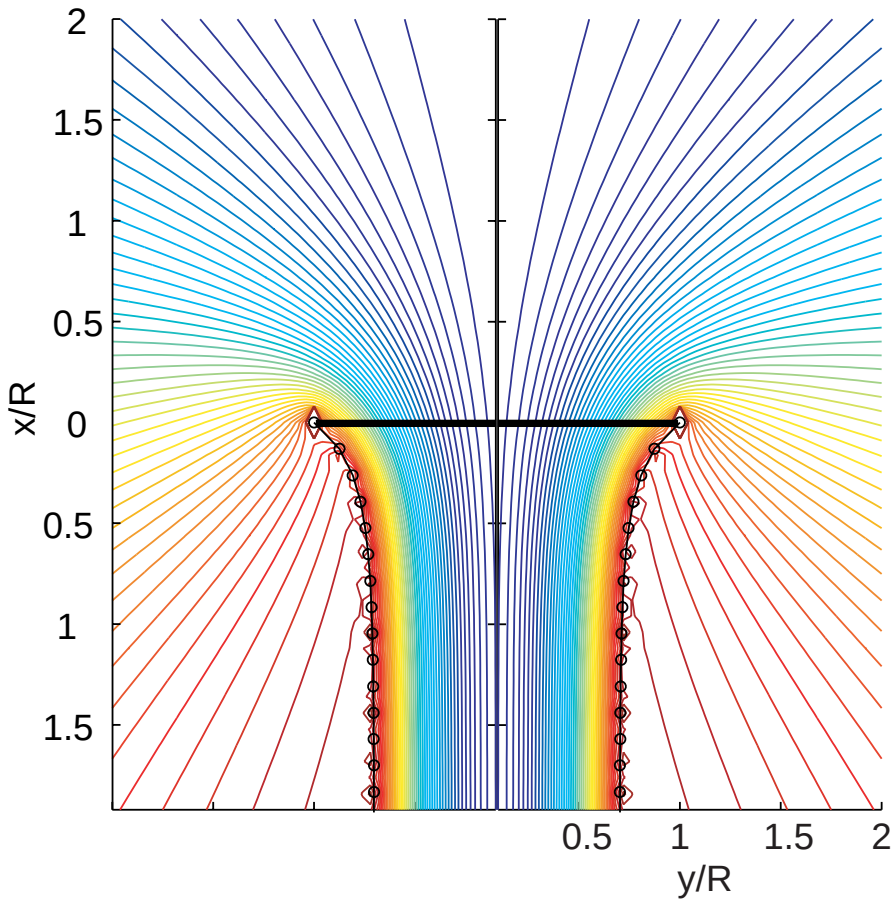

FIGURE 4.6 Converged wake streamlines for a candidate rotor after 6 iterations.

The inviscid induced rotor inflow velocities are calculated using second order central-differencing of the stream-function along the blade. These velocities are utilized to derive a modified κ distribution. Eliminating κ and directly using the results to modify the inflow velocities is inconsistent with the initial separation of viscous and inviscid components of the thrust and torque equations as derived earlier. Utilizing a corrected κ allows the classically derived equations to be used in an unmodified form. The inviscid momentum and blade element thrust equations may be solved for $B\Gamma$ and equivalently $\Gamma_{\infty \text{ blades}}$:

$$B\Gamma = \Gamma_{\infty\, blades} = \frac{4\pi r u(u + U_\infty)}{(\Omega r - v)} \tag{4.44}$$

Here, u is the constant downwash velocity predicted by actuator disk theory:

$$u(u + U_\infty) = \frac{T}{4\rho\pi r} \tag{4.45}$$

Using the downwash velocities (u') calculated from the contracted wake model the corrected value of $B\Gamma$ may be expressed as:

$$B\Gamma = \kappa\Gamma_{\infty\ blades} = \frac{4\pi r u'(u' + U_\infty)}{(\Omega r - v)} \tag{4.46}$$

The final expression for κ simplifies to:

$$\kappa = \frac{u'(u' + U_\infty)}{u(u + U_\infty)} \tag{4.47}$$

The effectiveness of this model and its impact on both performance estimation and design are assessed later in Chapter 6.

4.5 Higher-Order Modeling of 2-D Viscous Effects

The commonly used simplifications of linear lift curve slopes and parabolic drag polars become increasingly inaccurate as the Reynolds number drops into the region of interest. The most attractive operating point, around the sectional maximum lift to drag ratio, is also the area of greatest non-linearity in the lift curves; it is also typically very close to the steady-state stall point. The necessary fidelity is attained by utilizing an database of 2-D section characteristics. The method, first implemented by Kunz [29], sequentially generates spline fits across flap deflection and Reynolds number, resulting in final spline curves for the geometric angle of attack and C_d as a function of C_l. This method was developed for natural laminar flow airfoils used on high-performance sailplanes. These sections also generally exhibit performance curves that differ significantly from the

idealized models. The method is purely interpolative and includes a stall model which returns a mark-up drag value and a stall indictor flag to the calling routine. The only modification to the original method was the replacement of the original bi-cubic spline routines with the Akima spline formulation previously mentioned in Chapter 3.

The cost associated with generating the initial airfoil databases is not insignificant. In this case the data is comprised of a large number of INS2d runs for each section, resulting in a collection polars across a range of Reynolds numbers. Typically five polars were used per airfoil at Re (1000, 2000, 6000, 8000, 10,000) with no more than 15 points per polar. Fifty 2-D steady state calculations are still a small fraction of the computational cost of a single 3-D rotor calculation. An additional and major benefit is that once the up front cost the airfoil database is incurred, it can be added to a library of sections to be reused for further analysis and design with no additional cost.

4.6 Rotor Design and Analysis via Gradient-Based Optimization

The rotor analysis method is only capable of estimating the performance of a given geometry and operating condition. Alone, it is incapable of autonomously developing or improving a rotor design for a particular application. The method also does not incorporate the chosen power plant into the analysis. A particular power plant cannot arbitrarily provide any amount of power at any RPM. The rotor operating condition must be matched to the capabilities of the power plant to have a physically realizable system.

A complete rotor analysis and design tool has been developed by coupling the rotor performance program with a non-linear optimization package. The optimization code being used is the SNOPT package developed by Gill, Murray, and Saunders [30]. In the design mode, the goal is to maximize thrust for a given radius and for a particular motor. The spanwise discretization of the rotor blade is also specified. The chord distribution,

lift distribution, and rotational speed are treated as the free design variables with the primary constraint that the power required matches the motor's power available. The motor model is flexible in that the code requires a user supplied subroutine that returns the power available for a given RPM. The overall structure is pictured in Figure 4.7.

As mentioned earlier, the analysis mode is a different problem and requires the ability to input only the geometry of the rotor, chord distribution, incidence, and the rotor speed. Rather than create a second iterative rotor analysis code for this problem, this case is also solved using the same rotor analysis code and non-linear optimizer, but the problem posed to the optimizer is modified. The specified incidence angle at each station is treated as an equality constraint. This drives the geometry to the specified incidence angles. The optimal solution that satisfies the incidence constraints provides the correct lift distribution for the analysis of rotor performance. With RPM as an input, the constraint on the power required is removed. It is assumed in the analysis mode that the case represents a realizable operating point. The design code produces an optimal solution at a single operating point. The analysis code allows these solutions to be evaluated over a range of operating conditions. It also provides a means of validating the method by comparing predicted performance with experimental results for existing rotors and propellers. The next chapter discusses the effectiveness of the various enhancement models and the validation of the overall method.

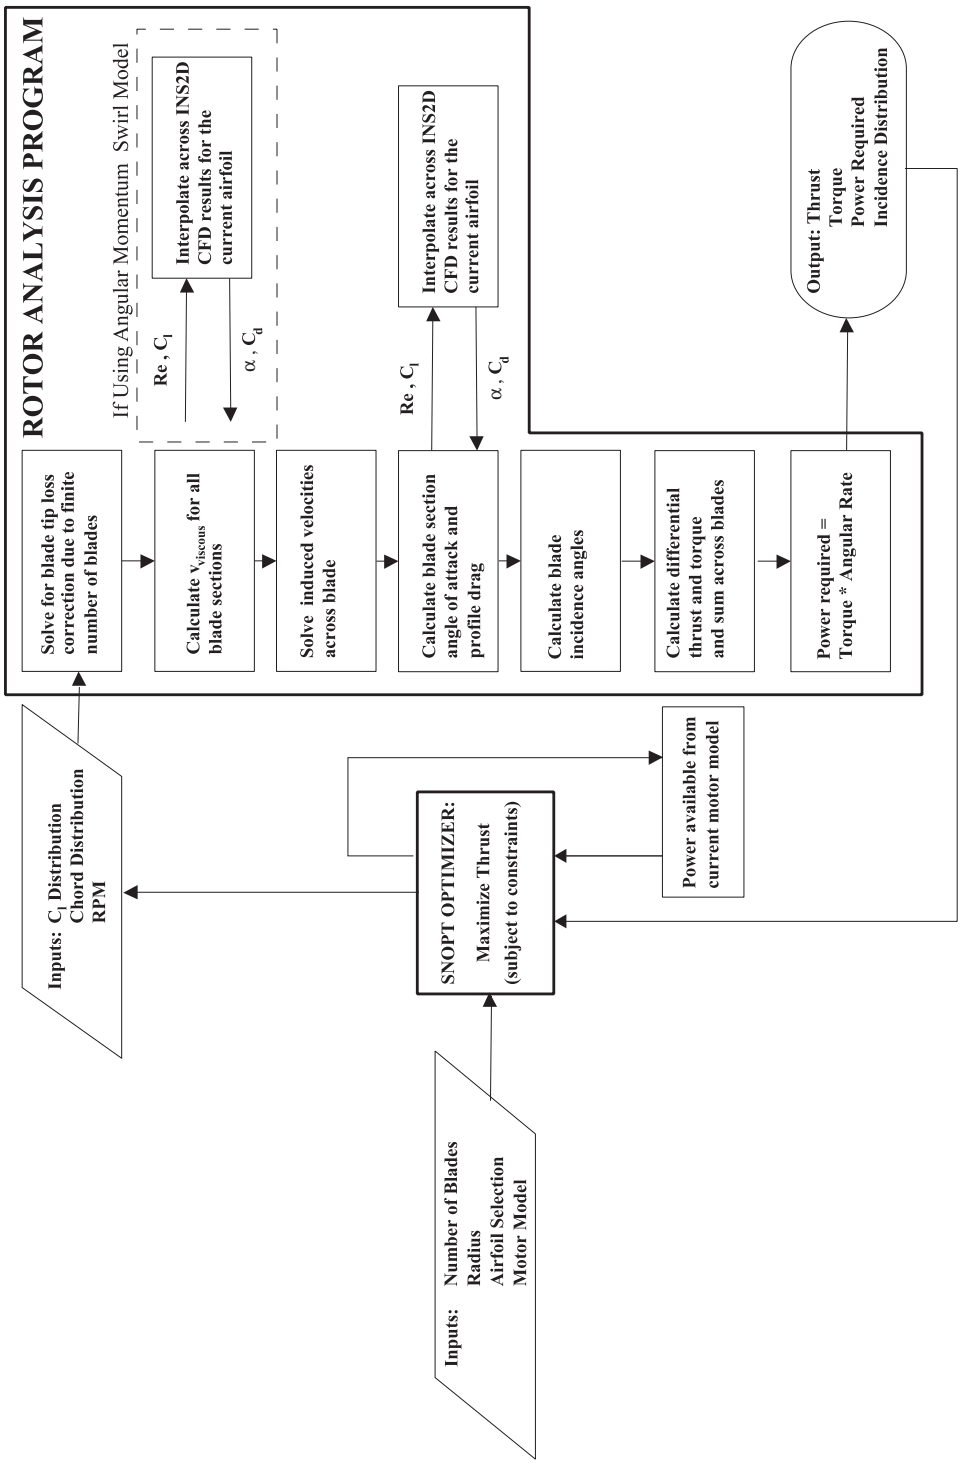

FIGURE 4.7 Flowchart of the rotor analysis and design process.

Chapter 5

Overview of Experimental and Computational Validation Methods

5.1 Introduction

Fabrication of centimeter-scale rotors represents a significant achievement due to their small size and required geometric precision. This work was accomplished using a sophisticated technique for precision micro-manufacturing developed in part at Stanford. Testing poses its own challenges due to the size and fragility of the test articles and the need to resolve small forces and moments. This has required the fabrication of several test fixtures specifically tailored to the problem. The experimental work provides data on gross performance parameters such as total thrust and torque, but the computational validation provides the most in depth look at the flow physics. The numerical calculation is also guaranteed to have the specified rotor geometry, thereby serving as an indicator of the geometric accuracy of the test articles. The following sections describe the methods and tools used in both the experimental and computational validation studies.

5.2 Rotor Manufacturing

5.2.1 Shape Deposition Manufacturing

Three different methods have been used in the manufacturing of the rotor prototypes. The smaller rotors are fabricated by the Rapid Prototyping Laboratory at Stanford University. The rotors constructed of epoxy resin have been built using the Shape Deposition Manufacturing (SDM) process [31]. A pictorial summary of this method is shown in Figure 5.1. This method can yield very accurate geometries during manufacturing within certain limitations, but in the case of these micro-rotors, the accuracy of the specified geometry after handling and use has proven to be a problem.

Each blade has a minimum thickness of only 20 mils, but the desired camberline is accurately reproduced as seen in Figure 5.2. The blade airfoil geometry is the result of the two-dimensional airfoil optimization at Re = 6000. The NACA four-digit thickness distribution is clearly not reproduced in this case. The desire to minimize section thickness is tempered by manufacturing minimum gage constraints. If the design exceeds these minimums the tooling forces may deform or tear away the part. The trade is then between a thicker airfoil section with a more conventional thickness distribution, or a constant thickness section at minimum gage. The results of Chapter 3 indicate that the latter is preferred. A second cross-section, in this case a NACA 4402 camberline, is shown in Figure 5.3. In both cases the constant thickness regions are approximately 2% chord, representing this minimum gage constraint. In this case, some shaping of the leading and trailing edge regions is visible. This is achieved by reversing the manufacturing orientation such that the kerf of the tool at the bounding tool paths provides beneficial shaping. A comprehensive discussion of the manufacturing aspects of the mesicopter project is provided by Cheng [32]. It is important to recall that, as shown in Chapter 3, the camberline is the primary determinant of performance for very thin sections.

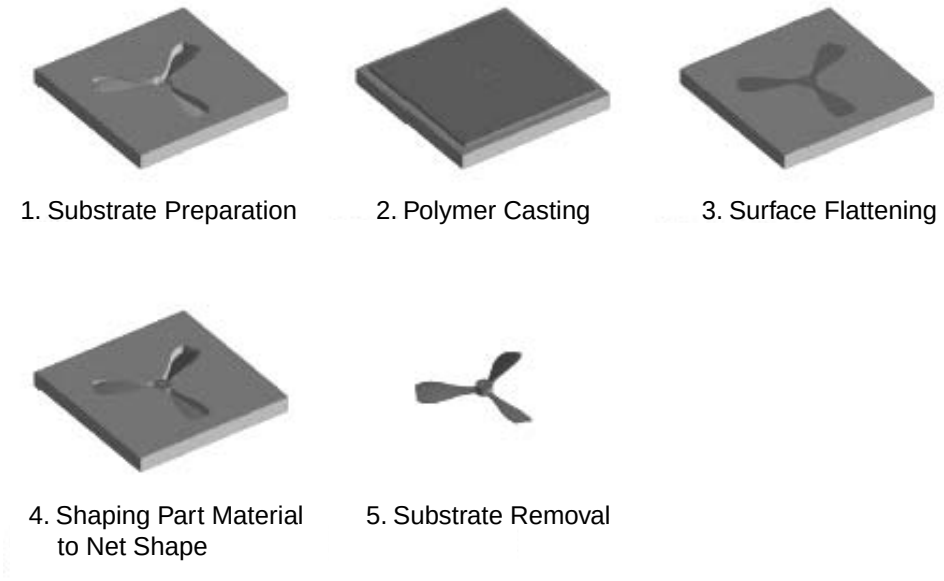

FIGURE 5.1 Summary of rotor SDM process (from Ref. 32)

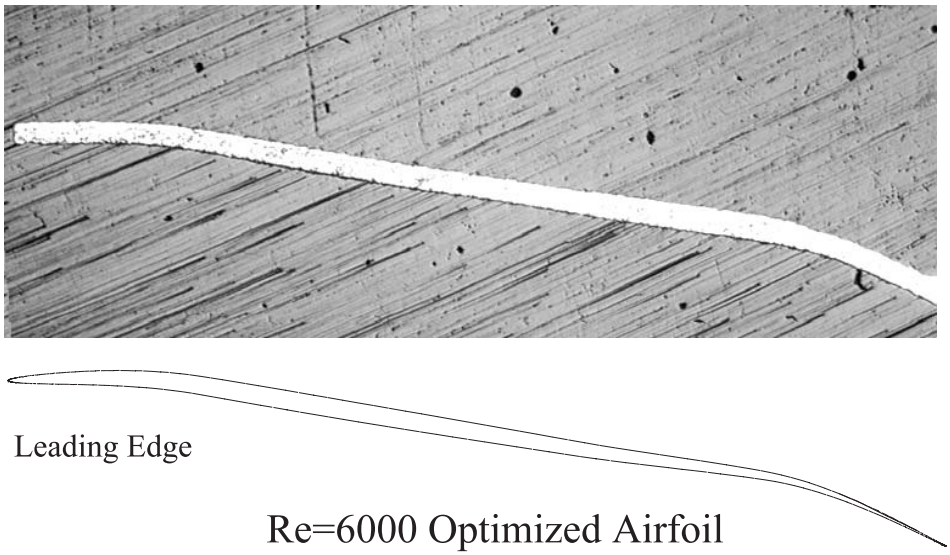

FIGURE 5.2 Specified airfoil and blade section photomicrograph.

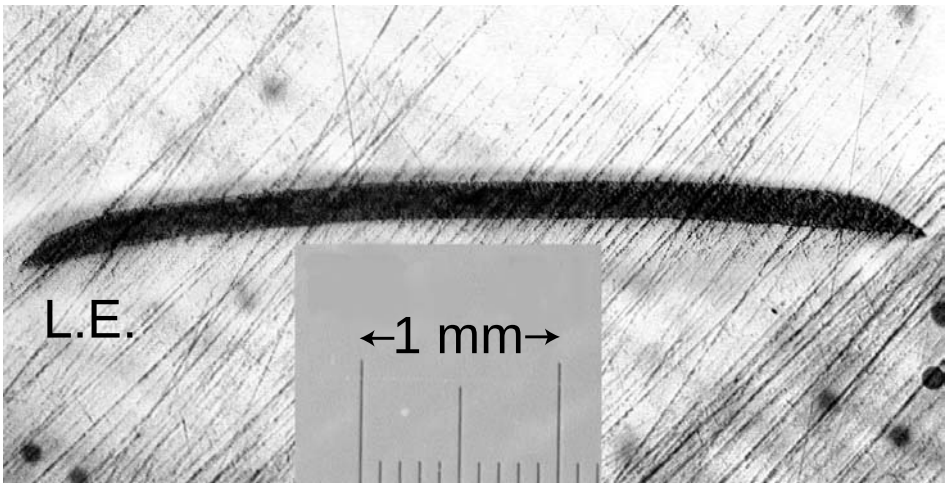

FIGURE 5.3 Photomicrograph of an SDM wing cross-section, based on the NACA 4402 camberline.

5.2.2 Dual Surface Machining

In addition to the epoxy SDM micro-rotors, samples of the same design have been manufactured in aluminum in order to evaluate the benefits of a stiffer material. The material selected was 7075-T6 aluminum alloy. This process involves precision machining of both the rotor itself and a fully supporting fixture which allows machining of the second side of the rotor. For a detailed description of the manufacturing process the reader is once again directed to the thesis by Cheng [32].

5.2.3 Composite Press Molding

The largest rotor being considered here, a five inch radius two-blade configuration, is manufactured using a two-piece press molding technique. One of the two-piece molds is pictured in Figure 5.4. Two sets were required for left and right handed rotors. The fabrication of the molds was completed by the Stanford Rapid Prototyping Laboratory; rotor fabrication has been completed by the author. The rotors consist of a five-ply wet lay-up of fiberglass, carbon fiber, and Kevlar with epoxy laminating resin. The carbon fiber provides the majority of the structural stiffness, the fiberglass provides a smooth

surface finish and aids in resin retention during the room temperature cure cycle, and the Kevlar is incorporated into the hub region primarily as a safety measure to prevent complete blade separation in case of a tip strike or some other incident leading to a structural failure at high RPM. A center hub and set-screw attachment, machined from aluminum, is bonded to the rotor post-cure. Finally, each rotor is trimmed and balanced.

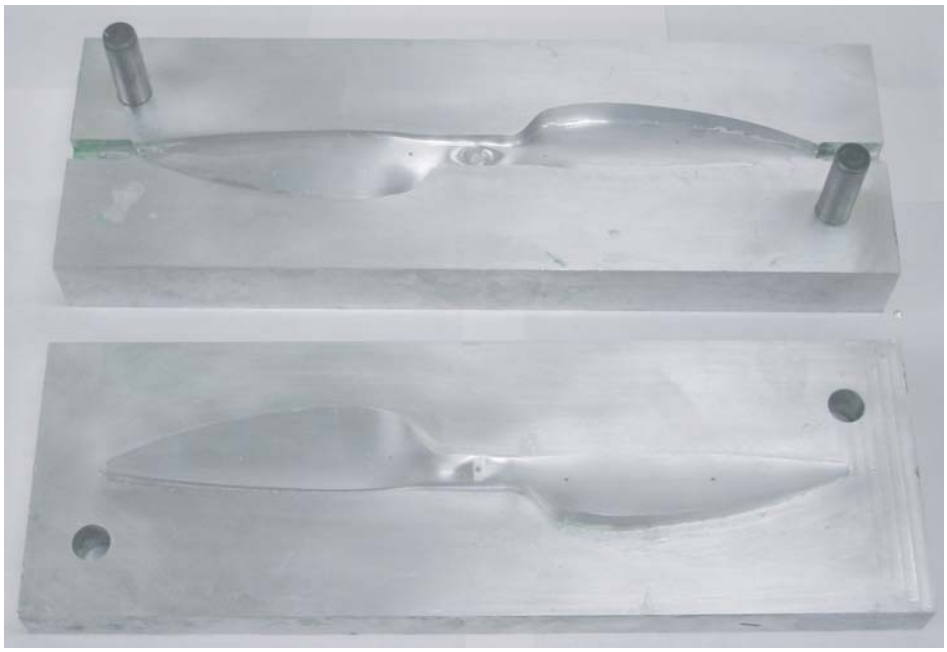

FIGURE 5.4 Two-piece aluminum press molds for the five-inch radius, two-blade rotor.

5.3 Experimental Methods

Testing methods and apparatus have been developed for the experimental evaluation of candidate rotor designs and validation of the design analysis tools. The same apparatus can also be configured to quantify the performance of candidate electric motors. The detailed aerodynamic design of the rotors is highly dependent on the performance characteristics of a given motor. Manufacturers commonly provide some form of performance data, but it is often either insufficient, or based on theoretical models rather than detailed experimental data.

The test fixtures for measuring thrust consist of an aluminum bar with a motor-mount on one end and a contact point for an electronic two milligram-accurate balance on the other. It is supported at near mid-span by a shaft and ball-bearings, providing a single axis of rotation. Upward thrust is transferred to a downward force on the balance that can be read directly, then converted to the thrust value by the ratio of the moment arms. Two fixtures have been fabricated to cover a wider range of rotor sizes.

The smaller set of test equipment, configured for thrust measurement, is pictured in Figure 5.5. Rotor torque for larger rotors may be measured by mounting the rotor and motor to the shaft of the balance bar with the rotor axis of rotation parallel to the shaft. The registered force on the balance is converted to a torque across the moment arm from the balance contact point to the aluminum bar's axis of rotation. This configuration is shown in Figure 5.6.

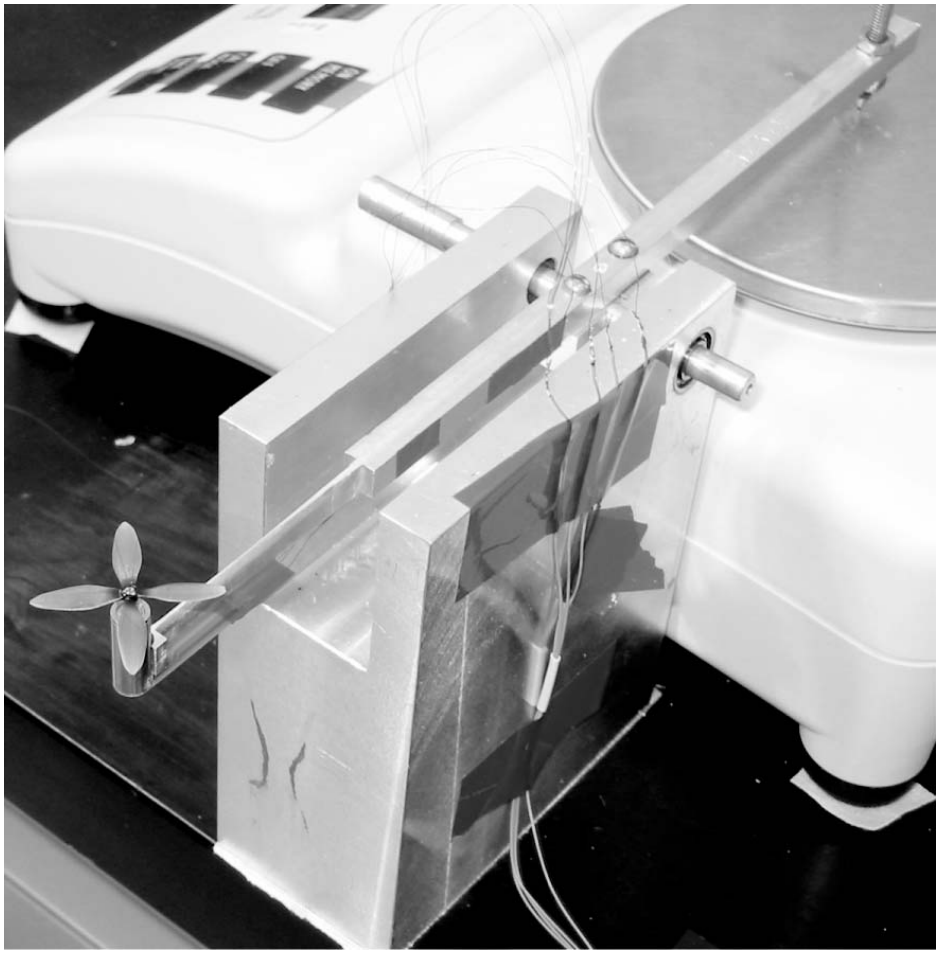

FIGURE 5.5 Small test fixture in thrust testing configuration.

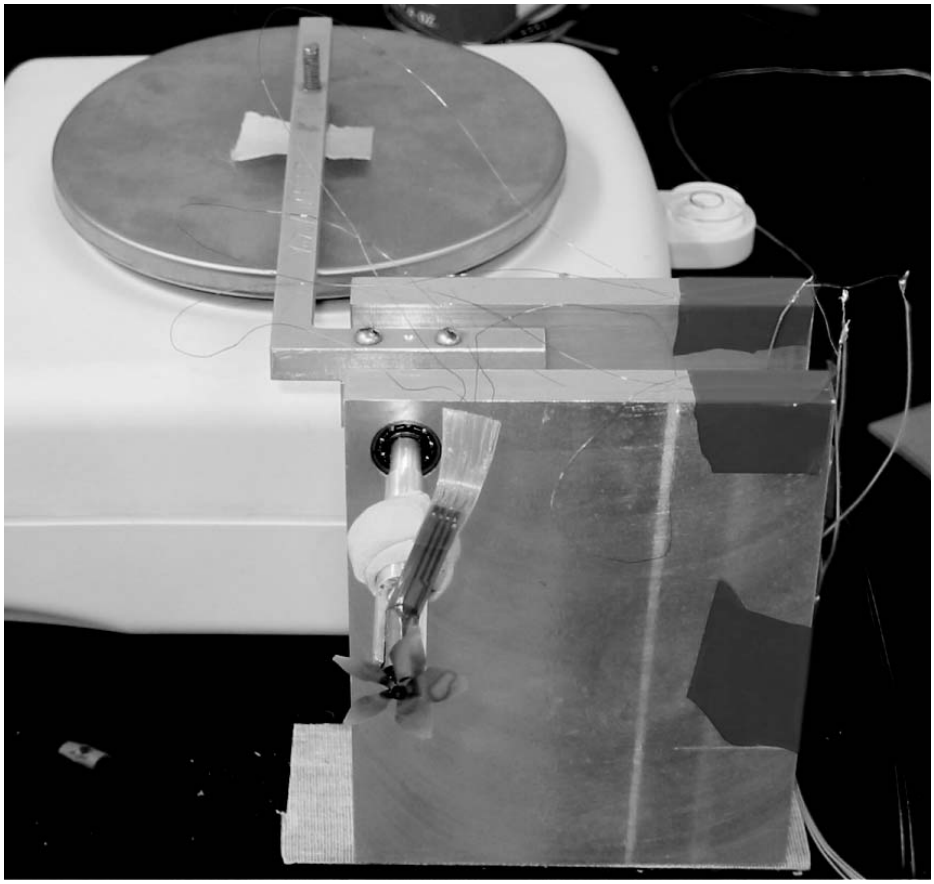

FIGURE 5.6 Small test fixture in the large torque testing configuration.

The ability to measure rotor torque with this setup is limited by the maximum precision of the electronic balance and physical constraints that limit the minimum length of the moment arm. The micro-rotors are under three centimeters in diameter and generate roughly 100 μNm of torque. At these low magnitudes, factors such as bearing stiction and the dynamics of the load cell come into play, making accurate and repeatable measurement difficult.

Another test configuration has been developed to address the majority of these issues. This approach uses the mounting bracket of the balance arm as a simple fulcrum balance. After leveling the bracket, mass is added to one side and the rotor is spun up until it supports the added mass. The mount is machined with precisely located threaded holes providing both a known offset and a secure means of attachment for the added mass.

The direct influence of the scale is removed since it is only used to measure the added mass prior to the experiment. Bearing stiction errors are reduced by lightly perturbing the system once balanced to assure it returns to equilibrium. Finally, the resolution is greatly increased by reducing the moment arm from 0.73” to 0.25”. This revised test fixture is pictured in Figure 5.7.

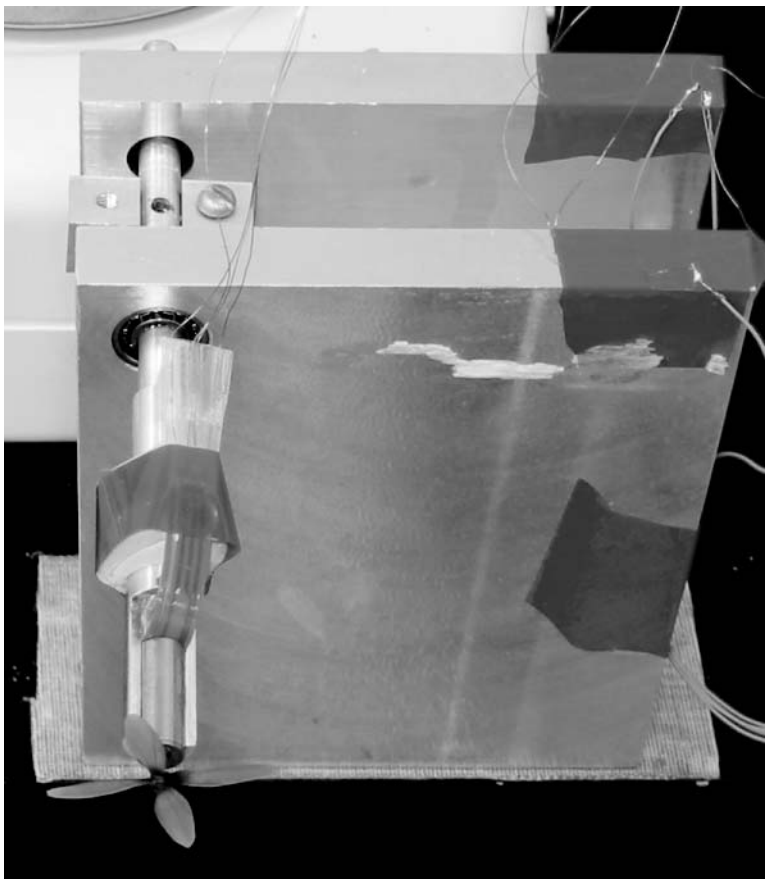

FIGURE 5.7 Small test fixture configured for small torque measurements.

The evaluation of motor performance is accomplished by coupling the shafts of two motors. The test motor is mounted to the stand in the same manner as for large magnitude torque testing. A second motor acts as a torque load and is mounted in a second fixture such that the drive shafts are co-linear. The test motor is powered by a constant voltage power supply. Torque, RPM, and current data are recorded as the load torque is gradually increased until the rated stall current is achieved. From this data,

power and efficiency curves can then be calculated. This configuration is shown in Figure 5.8.

FIGURE 5.8 Motor testing experimental configuration.

The larger thrust stand is pictured in Figures 5.9 and 5.10. This stand can be configured similarly to the small stand for direct thrust measurement of the larger rotors with the increase in length keeping the rotor wash clear of the pivot point and balance. This fixture can also be configured as a pendulum balance. In this configuration thrust is measured by observing the deflection from vertical. This capability was developed for qualitative testing in closed environments where use of the electronic balance is not possible.

FIGURE 5.9 Large test fixture in thrust measurement configuration.

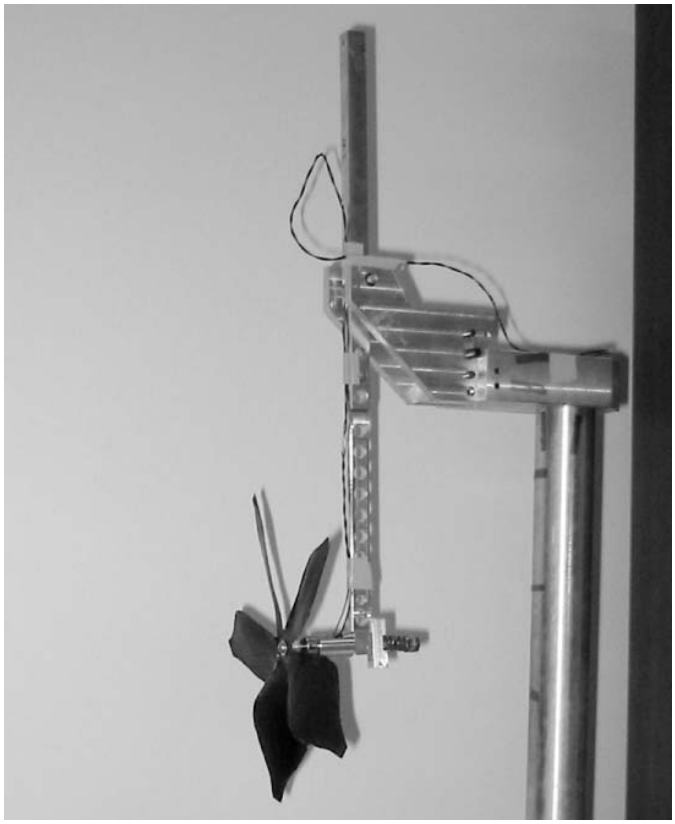

FIGURE 5.10 Large test fixture in thrust pendulum configuration.

5.4 Discussion of Experimental Error

The various test fixtures and methods used for thrust and power required testing have been developed to be as simple and straightforward as possible, but there are still potential sources or error that should be described and their impact on the measurements assessed. The sources considered are the motor power supply, determination of rotor RPM, the electronic balance, errors due to the geometry or construction of the fixtures, and aerodynamic blockage effects. The geometric accuracy of the rotors themselves is considered separately.

5.4.1 Motor Power Supply

All tests have utilized a *Global Specialties Model 1335* DC power supply for motor power. Typical voltages range between 6V and 13V for all tests and current ranges from 100mA to 250mA for all micro-rotor tests and 100mA to 400mA for the larger carbon-fiber rotors. The integrated digital displays of the power supply have been used for current and voltage measurement. The significant error incurred is the round-off error of the least significant digit of the display. The readability limits are 0.1V and 0.01A resulting in a possible error of (+/-)0.05V and (+/-)5mA. The manufacturer's stated display accuracy is much worse, 0.5% of reading + two counts, but voltage and current have periodically been monitored using a Protek 506 DMM in addition to the supply's displays, this DMM is accurate beyond the round-off error of the supply's displays and has agreed with the supply's values within this error bound.

Errors in recorded current and voltage will only manifest themselves as errors in the calculated electro-mechanical efficiencies derived from experimental power required for two micro-rotors to be discussed in Chapter 6. These efficiencies are then used to estimate the power required for several other micro-rotors. The typical input power for this system ranges from 0.5W to 3W, rendering any effect of voltage and current measurement error negligible.

5.4.2 Determination of Rotor RPM

Rotor RPM is measured using a *TNC Electronics Model PT10S* optical tachometer. The manufacturer's stated accuracy is (+/-) one RPM up to 100,000 RPM, but in practice the total error is somewhat higher since this is a handheld device and positioning and stability during reading can effect the output depending on the application. The user is left to interpolate across mildly fluctuating readings. For the large, 10" diameter, 2-blade carbon-fibre rotors, the observed fluctuations are roughly (+/-) 10 RPM, with RPM ranging from 1,000 to 2,500, representing a maximum error of (+/-) 1%. The micro-rotors are roughly one inch in diameter spinning at 20,000 to 50,000 RPM. Positioning

is much more challenging and an auxiliary light source directed through the rotor at the tachometer sensor is required for steady readings. Here the observed fluctuations during measurement increase to roughly (+/-) 200 RPM, but still represent a maximum (+/-) 1% error in RPM.

This measurement has also been validated by periodically using the frequency counting function of a *Hewlett-Packard model 54501A* digitizing oscilloscope attached to one phase of the 5mm Smoovy stepper motor described in Chapter 7. The results consistently agree within the stated (+/-) 200 RPM.

5.4.3 Electronic Force Measurement

Other than the pendulum test stand, all test configurations use an *Acculab VI-1mg* electronic balance. This balance is calibrated prior to each testing session and has 1mg readability with (+/-) 2mg accuracy. This does not constitute a relevant source of error for these experiments since the minimum thrust is on the order of one gram.

5.4.4 Test Fixture Errors

The effects of vibration and bearing stiction are significant sources of uncertainty and error in the test results. Vibration causes fluctuating readings on the balance forcing the user to interpolate. Bearing stiction can result in either preload or residual forces on the balance, but is unpredictable and also affected by vibration. The uncertainty due to both these issues has been assessed by observation on a test by test basis, by making note of residual forces during the test cycle and annotating measurements with the range of variation due to vibration.

For the small rotor thrust tests, vibration has not been an issue, but stiction effects account for an error of up to (+/-) 25mg on the balance reading which translates to (+/-) 20.5mg of thrust. This results in a thrust uncertainty of (+/-) 2.1% at 1.0g of thrust dropping to 0.5% at 4.0g of thrust. The equivalent torque is approximately (+/-) 21μNm,

or roughly (+/-) 25% of the expected torque of the rotors. This is the primary reason for developing the fulcrum balance method for small rotor torque measurement.

For large rotor testing, typical forces and torques increase by a factor of 10 to 20 removing bearing stiction as a source of significant error. Vibration is the bigger issue resulting in an uncertainty of (+/-) 0.5g at the balance and (+/-) 0.3g of thrust. This translates to a (+/-) 1.5% uncertainty at 20g of thrust, decreasing to (+/-) 0.5% at 60g.

5.4.5 Aerodynamic Blockage

An aerodynamic blockage correction has been applied to all experimental thrust results for the 5mm Smoovy motor. This motor utilizes a relatively large flex-print connector as a wiring harness attachment point. This connector is visible at the base of the motor in Figure 5.5. The exposed area of the connector directly under the rotor is large, 13% of the disk area in the case of the 2.5cm diameter rotors, and must be taken into account. The correction has derived experimentally, but two theoretical arguments placed upper and lower bounds on the correction.

The worst case scenario is a complete blockage of the rotors momentum flux over the area of the connector. This scenario would model the effect of the connector if located at the rotor plane. This would result in the measured values being 15% below the actual thrust. The connector is not at the rotor plane however, and actually sits 6mm below it. The low end correction is approximated by considering the flat plate drag, C_d=1.2, of the connector in a uniform flow field determined by actuator disk theory. This results in a 4% deficit in thrust. The actual value should lie between these two extremes.

The experimental determination consisted of two thrust tests of a single rotor. The first test used the standard test configuration, the second doubled the area of the connector by placing a second obstruction of the same dimensions and vertical location, 180 degrees off of the flex print connector. The result were compared and the difference taken as the correction. This resulted in an average deficit of 10.8% between 20,000 and 50,000

RPM with a standard deviation of only 0.4%. The value and the independence from RPM are reasonable considering the theoretical bounds just described.

5.5 Three-Dimensional Analysis using OVERFLOW-D

The Reynolds-averaged Navier-Stokes flow solver used to validate the results from the rapid design method is based on a version of the OVERFLOW code developed by Buning *et al* [33]. OVERFLOW is a general-purpose Navier-Stokes code designed for overset-grid computations on static grids. Meakin [34, 35] has adapted this code to accommodate arbitrary relative motion between vehicle components and to facilitate off-body solution adaption. The modified code automatically organizes grid components into groups of approximately equal size, facilitating efficient parallel computations of multi-body problems on scalable computer platforms. On parallel machines, each processor is assigned a group of grids for computation, with inter-group communications performed using the Message Passing Interface (MPI) protocol. This code is known as OVERFLOW-D. The solution of hovering-rotor problems requires a number of modifications to the OVERFLOW-D flow solver. Strawn and Djomehri [36] describes these modifications in detail. Post-processing of the sectional and global rotor forces uses the FOMOCO force integration code [37].

This modified version of OVERFLOW-D has been validated with experimental data for a model UH-60A rotor by Strawn and Djomehri. Figure 5.11 compares the experimental and computational global thrust and torque coefficients. Figure 5.12 compares the computed and sectional thrust distribution for the same model. The over prediction of thrust visible near the tip is to some extent an artifact of the OVERFLOW-D analysis and grid density issues [36]. Grid related vortex diffusion exacerbates the blade vortex interaction (BVI), resulting in a local exchange in circulation between the BVI and the

newly forming blade tip vortex. In spite of this possible anomaly, the overall results do show good overall agreement with the best experimental data currently available.

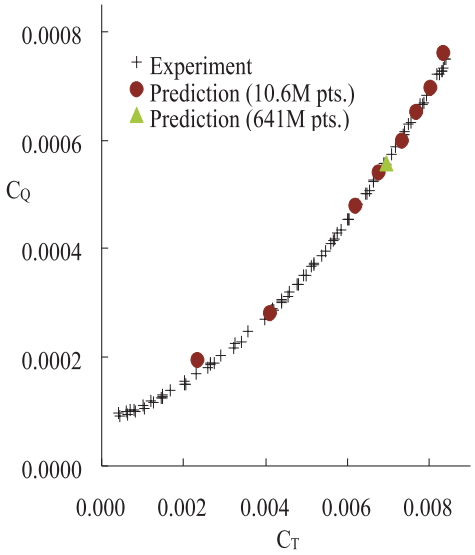

FIGURE 5.11 Thrust and power results for a model UH-60A rotor (from Ref. 36).

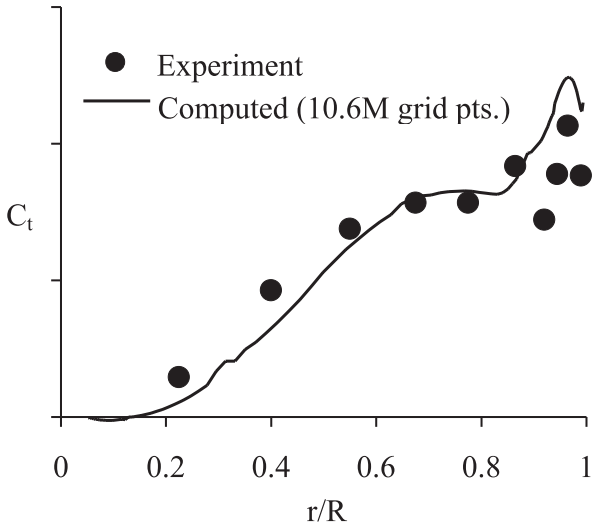

FIGURE 5.12 Sectional thrust distribution for a model UH-60A rotor (from Ref.36).

The micro-rotor case that has been analyzed is a four-blade 2.5cm diameter rotor. The computational cost is reduced by assuming a steady-state flow-field and by implementing periodic boundary conditions on the blade entrance and exit planes of a single rotor segment. In spite of reducing the computational volume by one quarter, the computational grid still contains over five million grid points. The near-body grids comprise 40% of the cell count.

The computational domain utilizes a Chimera overset grid topology with two cell minimum overlaps. One quadrant of the rotor is modeled in a domain extending five radii above and below the rotor plane and five radii out along the vertical inflow and outflow planes. The near body grids consists of 250,000 grid points about the hub, 180,000 grid points about the blade tip-cap, 520,000 grid points about the blade-hub intersection collar, and one million grid points about the blade itself, with 189 chordwise panels. The initial grid point spacing on the blade places at least 25 grid points within a laminar boundary layer at x/c=0.1. The resulting rotor model is pictured in Figure 5.13. The near-body grid geometry is depicted in Figure 5.14.

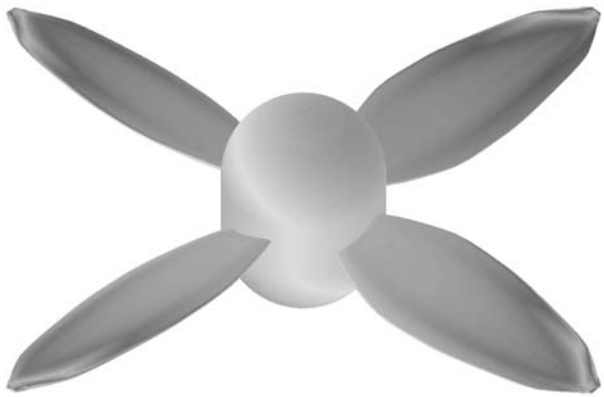

FIGURE 5.13 Computational Model of the four-blade 2.5cm diameter rotor.

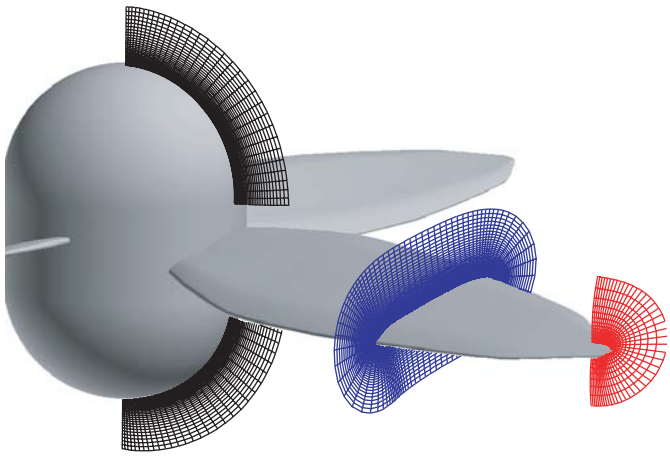

FIGURE 5.14 Near-body grid geometry.

The remaining cells are distributed across three levels of Cartesian off-body grids, surrounding the near-body grids, with the innermost domain matching the outer cell spacing of the near-body volume grids. The total grid size is approximately 5.5 million cells and required approximately 24 hours to run a single operating point using 16 processors of a SGI Origin 2000.

In order to reasonably assume steady-state flow, the hub model had to differ significantly from the actual part. The as-built rotors employ a 1.0mm high, 2.5mm diameter disk as a hub. In use, the underside of the hub is effectively faired by the presence of the motor. The flat upper and lower surfaces of the hub, combined with the absence of the motor body in the computation, would result in large amounts of unsteady, separated flow. This problem is dealt with by increasing the height of the modeled hub and capping both ends with hemispheres.

Separation is less likely, but the modified hub adds significant wetted area. The skin friction drag of the hub is not incorporated into the calculation of the global rotor forces,

but the no slip condition generates axial vorticity via a thick boundary layer on the hub. The primary effects are reduction in dynamic pressure and Reynolds number, resulting in a reduction in the inboard lift and Reynolds number. This is a problem unique to these rotors since proportionally the hub diameter is much larger than those of typical helicopter rotors, but even with the larger hub radius, the local dynamic pressure is low enough that the effect on total blade forces is negligible.

Chapter 6

Design Examples and Comparisons with Experiment

6.1 Introduction

Several small rotors have been designed using the rapid rotor analysis and design method developed in Chapter 4. The results of experimental testing, in conjunction with three-dimensional Navier-Stokes analysis of one of the designs, provides significant insight into the value of the various viscous swirl and wake enhancements and validation of the overall method as a means of preliminary rotor design at ultra-low Reynolds numbers.

Three different rotors are the focus of this chapter. The first is a five-blade 2.2cm diameter rotor. This rotor is the result of design using a predecessor of the current rotor code, incorporating the classical Prandtl tip loss correction and angular momentum viscous swirl models. In addition to providing another test case for the current analysis method, this design provides some insight into the effects of both viscous swirl modeling and aero-structural deformations on rotor design and performance. The follow-on design is a four-blade 2.5cm diameter rotor. This rotor design incorporates the classical

Prandtl tip loss model and the average wake deficit viscous swirl model. Both of these designs experience chord Reynolds numbers ranging from 1,000 to 9,000.

The final design is a much larger 5 inch radius two-blade rotor developed for a large scale flight test model. The Reynolds numbers seen by this rotor are higher than the two micro-rotor designs, ranging from 8,000 to 25,000. This rotor is a re-design, taking the chord distribution of an existing commercial product, applying a different airfoil, and optimizing the incidence distribution. The same version of the code used for the four-blade rotor is used in this case. It is provided a reference point for the analysis method at a larger physical scale, where the confidence in the tested rotor geometry is higher, rotational speed lower, and the measured forces much larger, reducing any measurement error.

These rotors and others have been designed at various points during the evolution of the rotor code of Chapter 4. These three cases do not represent all of the designs manufactured and tested, but provide an experimental reference by which others can be evaluated. Further iterations and analysis of these designs and others would be desirable, but the constraints imposed by the rapid pace and structure of the overall mesicopter development program limited the ability to complete detailed in-depth development of any single design.

6.2 Motor Selection and Characterization

6.2.1 Myonic 5mm Smoovy Motor

The 2.5cm and 2.2cm diameter rotors have been designed in conjunction with a 15 gram prototype vehicle to be described in Chapter 7. The motor selected for this application is the 5mm Smoovy bi-directional brushless DC motor, Model #59002, manufactured by Myonic Inc. (formerly RMB Inc.) [38]. The motor has a total mass of 1.4 grams and is

capable of speeds up to 100,000 RPM. This motor is currently out of production. A closed-loop motor controller, also manufactured by RMB, was used for all tests. This motor is pictured in Figure 6.1.

FIGURE 6.1 Myonic (RMB) 5mm Smoovy motor, inch scale.

Limited performance data is provided by the manufacturer, but this has been augmented by in-house testing. The required motor data for the rotor design method consists of motor output power as a function of the motor RPM. The simplest data-set to generate experimentally are curves based on a constant input voltage. Multiple curves may be generated across the range of permissible motor voltages or a single composite curve may be constructed based on a desired input current or input power. The best method depends on the limiting factor, current or voltage, of the energy storage and motor systems.

The output power versus RPM curves used for design are based on manufacturer's data for a maximum 75 second run time. These motors are rated as high as 40% efficient under nominal operating conditions of 6 Volts and 130mA draw, but the specified efficiency drops to approximately 28% for the 75 second conditions. Manufacturer's data is unavailable for the ten second conditions. It is important to note that these

efficiencies are only a part of the total electo-mechanical efficiency of the power and propulsion system, battery, wiring harness, and controller efficiencies further reduced the total system efficiency.

For their small size, these motors provide very good performance both in terms of torque production and high rotational speeds. One limitation of this three-pole stepper motor is the need for additional circuitry to cycle power through the motor coils. This adds mass to the system; in the final configuration of the 15g vehicle, the controllers represent a mass fraction of 20% to 30%. Simple circuitry is required to run the motor without feedback, but the motor efficiency is greatly reduced and the maximum RPM under load drops considerably, in this case from well over 50,000 RPM to below 25,000 RPM. Feedback is required to account for phase lags in the coil/shaft position induced by the rotational inertia of the shaft and rotor, friction, and aerodynamic drag. The feedback circuitry does not add considerably to the mass, but the complexity of the circuity increases. Small traditional or coreless DC motors, such a pager motors, are not a viable alternative. Although they require no additional circuity, these motors provide insufficient torque and rotational speed at his scale. Their efficiencies are also typically much lower, often as low as 5% to 10%. It may be possible that a design-specific traditional DC motor could be built with performance closer to the level of the 5mm Smoovy, but this is beyond the scope of this work.

6.2.2 Astroflight Firefly Motor

The ten inch diameter rotor is powered by the Astroflight Firefly coreless DC motor with an integrated 16:1 gearbox, Model #800 [39]. This motor and gearbox have a combined mass of 12g and a manufacturer's indicated efficiency of 83%. The Firefly and gearbox are pictured in Figure 6.2. This, once again, is not the total system electro-mechanical efficiency, simply that of the motor and gear box. This factor of two to four increase in efficiency relative to the 5mm Smoovy illustrates some of the practical difficulties associated with developing very small flight vehicles. In each case these motors represent the upper levels of performance at their respective scales.

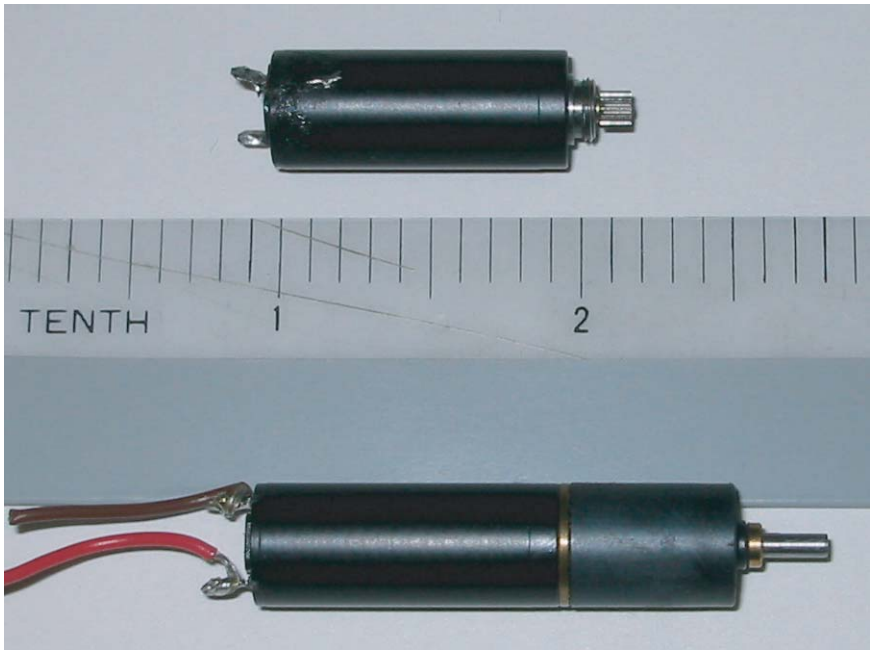

FIGURE 6.2 The Astro Flight Firefly 800 with 16:1 gearbox, inch scale.

The output power / RPM curve used for design with the Firefly motor has been theoretically derived from motor characteristics provided by the manufacturer using a model of an idealized DC motor [40].

6.3 Rotor Design Specifications

6.3.1 Five-Blade 2.2cm Diameter Rotor

This rotor is the first manufactured and tested for the 5mm Smoovy three-phase stepper motor. It utilizes a NACA 4402 camberline with a nominal 2% constant thickness profile. It was developed using a predecessor to the current design tool and incorporates the angular momentum swirl model, Prandtl tip loss correction, and 2-D drag polars based on INS2d analyses. The diameter and blade count have been chosen by a parametric study completed with the same code. This rotor is pictured in Figure 6.3.

with its chord and incidence distributions presented in Figure 6.4. The mid-chord line is unswept. Two versions are discussed here, the original design and a modified version having an additional three degrees of incidence and a modified spanwise thickness distribution. The differences in thickness distribution are shown in Figure 6.5. Both versions use a 2mm diameter hub.

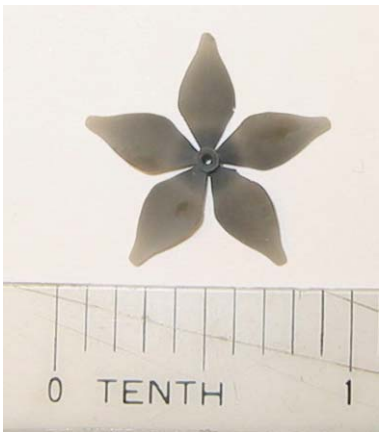

FIGURE 6.3 Epoxy five-blade 2.2cm diameter rotor, inch scale.

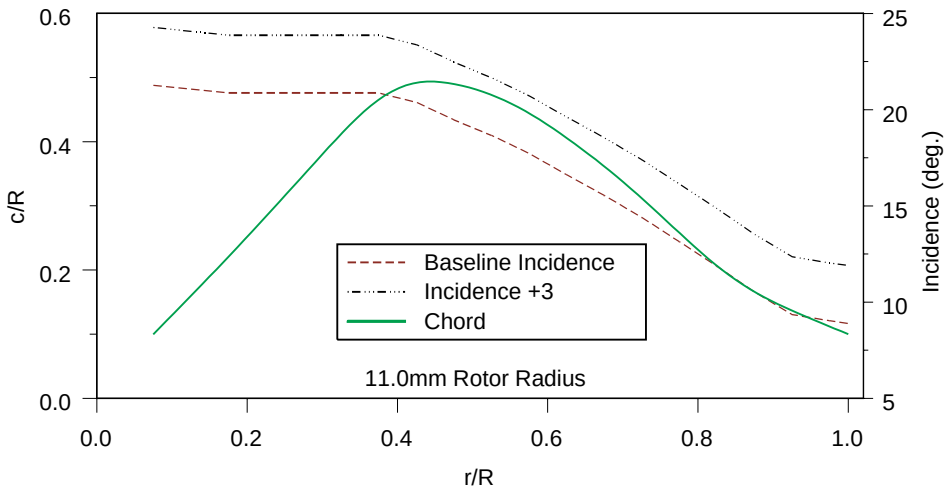

FIGURE 6.4 Chord and incidence distributions for the five-blade 2.2cm diameter rotor.

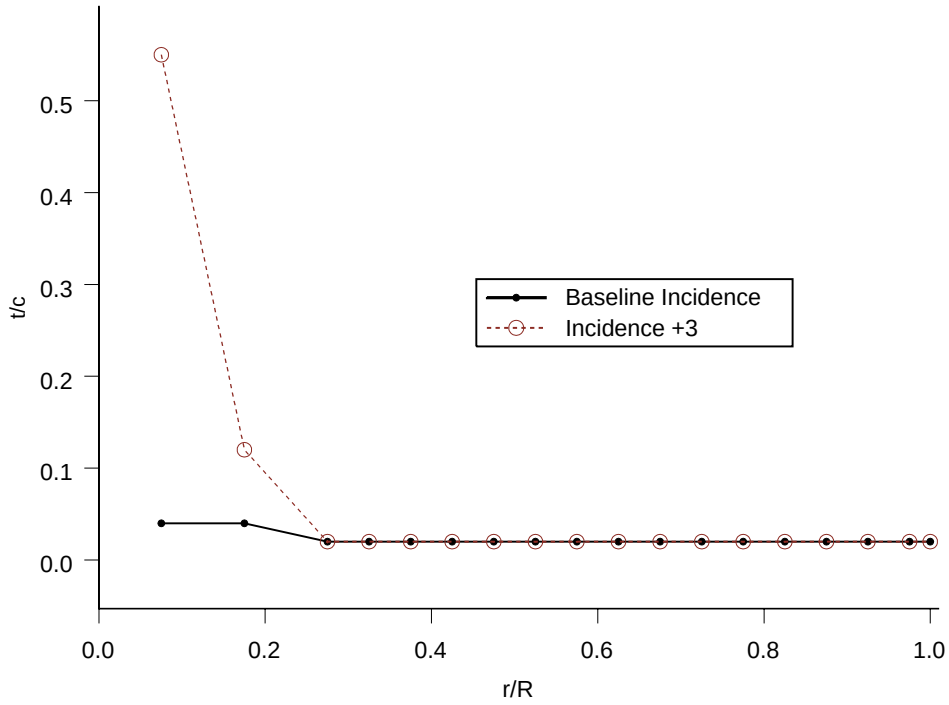

FIGURE 6.5 Thickness distributions for the five-blade 2.2cm diameter rotor.

6.3.2 Four-Blade 2.5cm Diameter Rotor

The four-blade 2.5cm diameter rotor was designed using the method of Chapter 4 and incorporates the wake-deficit viscous swirl model and the Prandtl tip-loss correction. The goal was to design the smallest rotor capable of generating 4g of thrust using the 5mm Smoovy motor and this resulted in a design operating point of 47,000 RPM. Two versions have been produced and are pictured in Figure 6.6. The first has a hub diameter of 2mm, which was later increased to 5mm to improve the torsional stiffness of the blades. The rotor diameter and blade count are based on a parametric study completed with the same code. In this case a four-blade design prevailed. The reduction in blade count from the 2.2cm diameter rotor can be largely attributed to the change in viscous swirl models. The average wake deficit model will be shown later to discourage high local solidity, particularly at the inboard stations. This rotor is the last micro-rotor design to be manufactured during the mesicopter program. The subsequent 3-D CFD analyses, which provide additional insight into the relevant flow physics and certain trouble spots

in the rotor geometry, occurred later and sufficient time was not available to manufacture and test another design iteration.

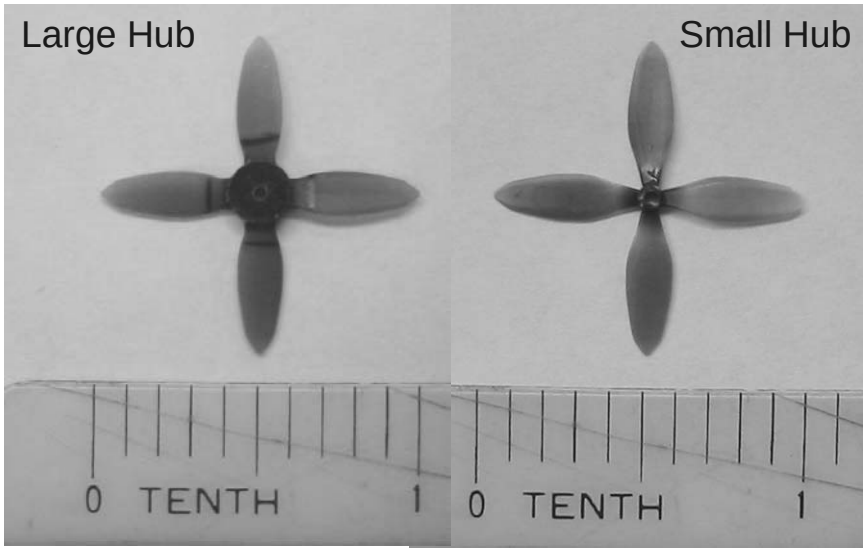

FIGURE 6.6 Four-blade 2.5cm diameter rotors, large and small hub versions, inch scale.

The airfoil is the 2-D optimized design for Re=6000. The maximum lift coefficient from INS2d calculations at Re=10,000 is 0.62, increasing to 0.66 at Re=6,000. The operating point lift coefficient for this rotor is constrained to 0.5 to provide stall margin and to account for uncertainties in the analysis. The design uses 30 spanwise stations, resulting in a total of 61 design variables (30 chord values, 30 lift coefficients, and RPM). Chords are limited by an upper bound of 80% local solidity, taken as the ratio of the blade chords to the local circumference at any given spanwise station. Power required is constrained to equal power available from the 75 second motor performance curve. The resulting chord, incidence, and thickness distributions are shown in Figures 6.7 and 6.8. The mid-chord line is unswept. Thickness ratio variations near the tip are due to minimum thickness manufacturing constraints. Variations at the root for the large blade hub were incorporated to further increase the torsional stiffness.

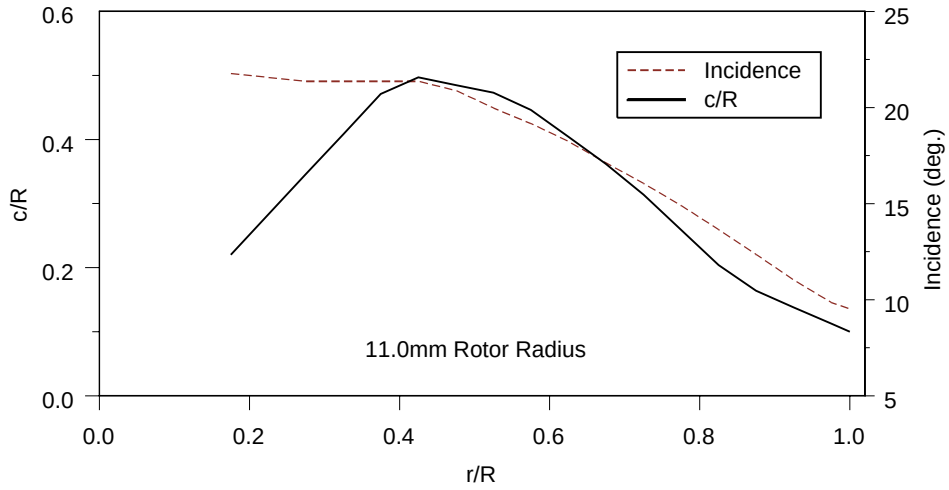

FIGURE 6.7 Chord and incidence distributions for the four-blade 2.5cm diameter rotor.

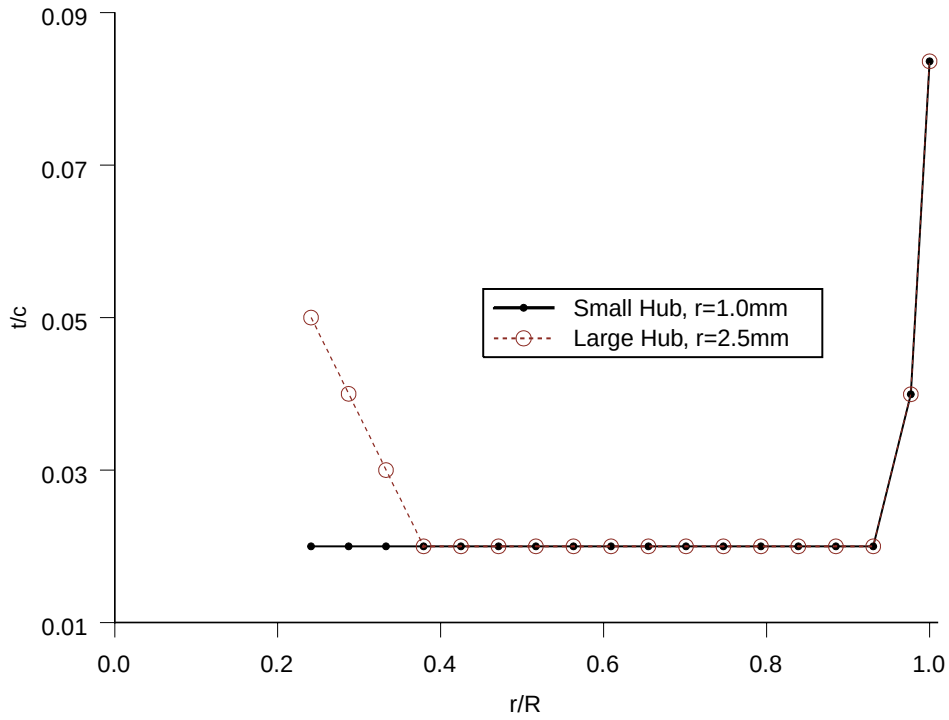

FIGURE 6.8 Thickness distributions for the four-blade 2.5cm diameter rotor.

6.3.3 Two-Blade Ten Inch Diameter Rotor

The two-blade ten inch diameter rotor is a design based on the WES-Technik Carbon 25/12 propeller [41]. This carbon fiber prop is intended for slow-flight remote control aircraft but exhibits performance that makes it marginally suitable for use in a prototype rotor craft. This vehicle was intended originally as a systems integration and control test-bed, making the selection of an off-the-shelf rotor expedient. Initially it had been intended to simply manufacture identical mirror image props necessary for reverse rotation, but analysis of the rotor showed that a re-design could achieve significant gains for the static thrust hover case.

Based on the results of Chapter 3, the airfoil section is specified as a NACA 4703. The two-blade geometry was maintained to simplify the composite tooling. The resulting chord distribution from a full re-design proved to be very close to the WES-Technik rotor, so the planform of the original rotor was maintained. This allowed for a more direct evaluation of any performance gains. The incidence-only design with the new airfoil is shown in Figure 6.9. For this rotor, the quarter-chord line is unswept. The manufactured carbon-fiber composite blade is presented as Figure 6.10. The design maximum thrust operating point is 1850 RPM.

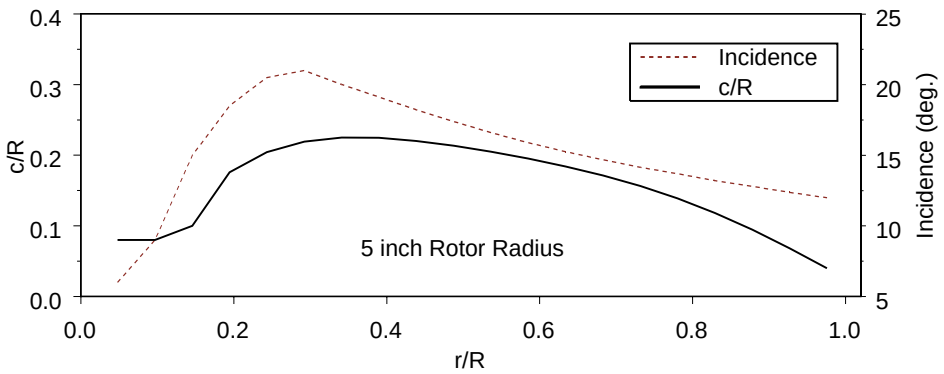

FIGURE 6.9 Chord and incidence distributions for the two-blade ten inch diameter rotor.

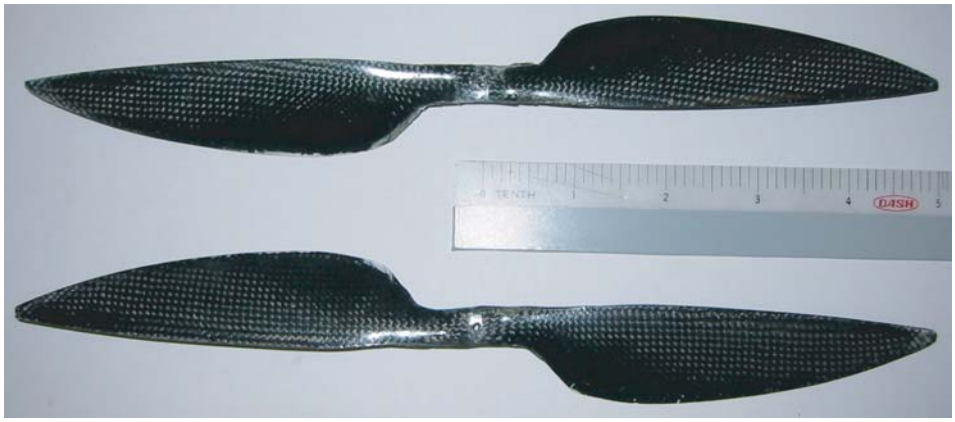

FIGURE 6.10 Carbon-fiber two-blade rotors, inch scale.

6.4 Comparison of Total Thrust and Power

6.4.1 Two-Blade Ten Inch Diameter Rotor

Global metrics, the total rotor thrust and power required, are both the principal points of interest in design and the simplest to obtain experimentally. Whether or not they form a concrete and comprehensive measure of accuracy will be a topic of later discussion. The first comparison of analysis and experiment is for the ten inch two-blade rotor. The larger scale greatly increases the confidence in the as-built geometry and the lower rotational speeds, typically around 2000 RPM, diminish the likelihood of structural deflection. The measured forces and torques increase by a factor of ten, reducing the effect of approximately constant magnitude measurement errors.

Close agreement is obtained in both thrust and rotor power required for this design. The maximum difference in thrust is approximately 10% at the highest tested RPM, corresponding to a less than 5% difference in power required. One point of note to be expanded on later is the indication of stalled tips when analyzed with the contracting wake model. The rotor was designed with the Prandtl tip loss model and at that point had no indication of stall. The thrust comparison for three rotor samples and the rapid

analysis method are presented in Figure 6.11. The corresponding power required for two samples is shown in Figure 6.12. All samples were not available for power required testing.

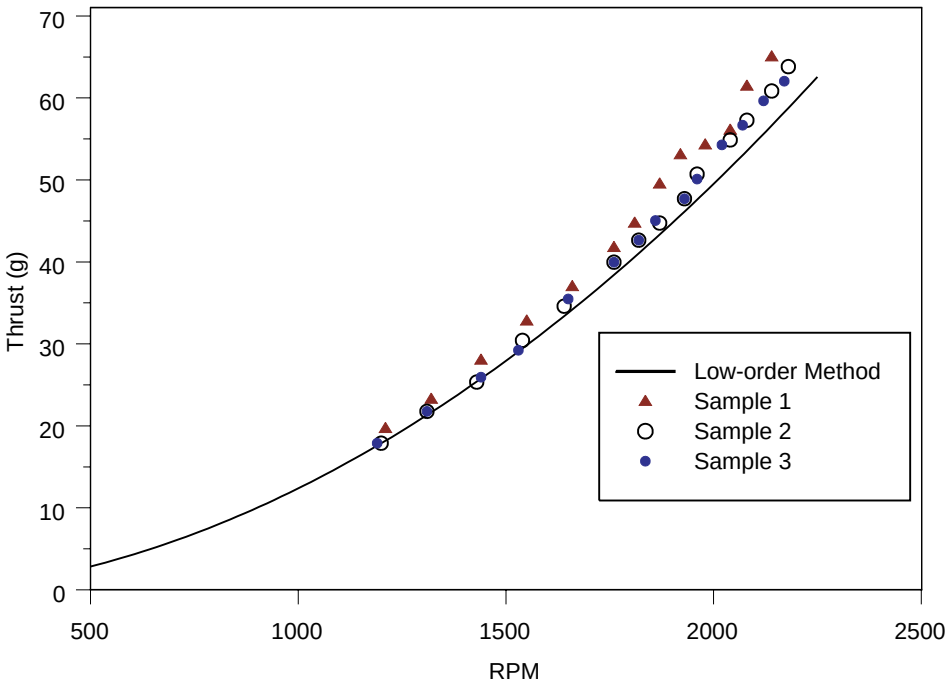

FIGURE 6.11 Experimental and predicted thrust versus RPM for the two-blade ten inch rotor.

For this case, the absolute and percentage difference between experimental and predicted thrust increases with RPM, but the power required is also seen to follow this trend. The most probable explanation is the presence of the partially stalled tips mentioned previously and the inability of the analysis method to properly account for the local section behavior in this regime. The rapid analysis method currently places a fixed cap on $C_{\ell\,max}$, and a prescribed drag markup. In reality, a section will continue to provide increases in lift with large increases in drag slightly beyond the onset of stall.

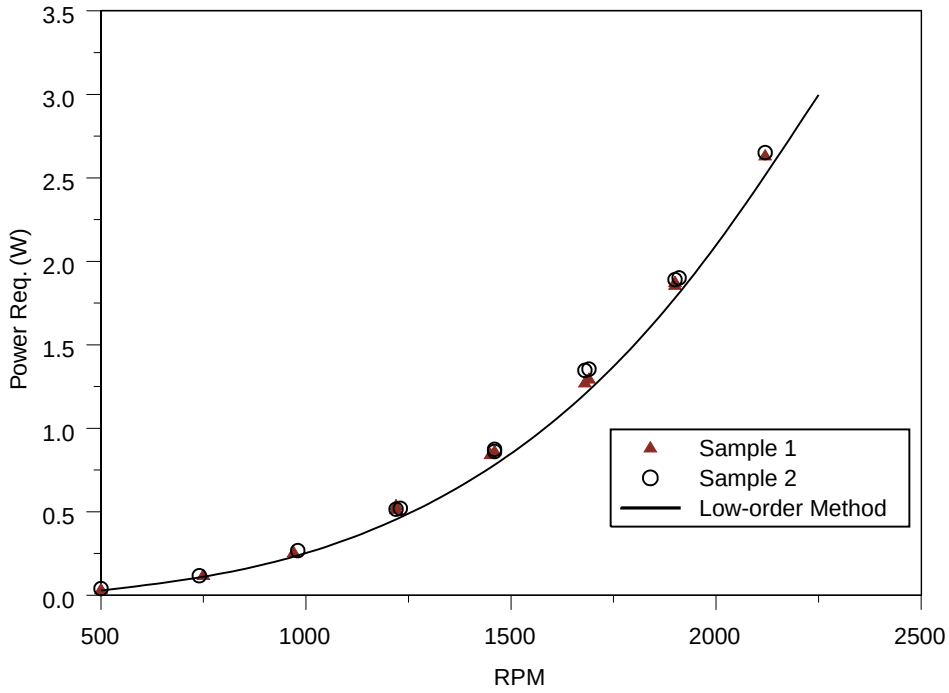

FIGURE 6.12 Experimental and predicted power required for the two-blade ten inch rotor.

6.4.2 Four-Blade 2.5cm Diameter Rotor

The agreement in both thrust and power required for the largest and slowest rotating of the rotor designs provides confidence in the method and allows exploration of the performance of the smaller design test cases. The four-blade 2.5cm rotor is the most thoroughly investigated design. It has been evaluated with the rapid analysis method, OVERFLOW-D analyses at three rotational speeds, and experimental testing of both thrust and power required. Several distinct versions of the rotor have been tested including the small hub epoxy rotor, the large hub epoxy rotor, and a large hub aluminum rotor. The experimental and predicted thrust performance for this rotor design is presented in Figure 6.13. The experimental data has been adjusted for blockage as described previously.

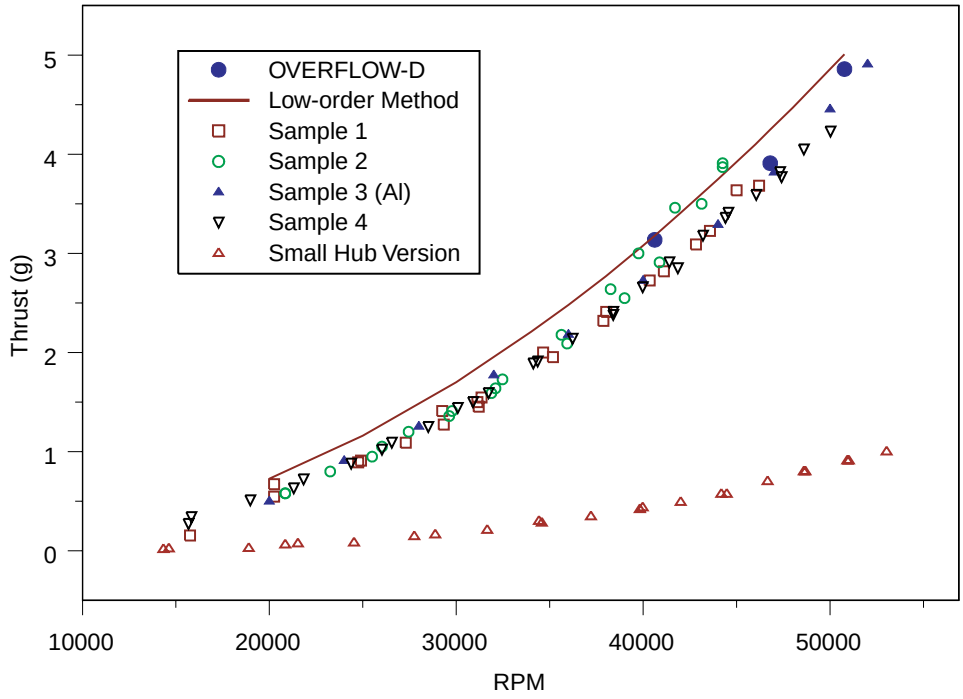

FIGURE 6.13 Thrust versus RPM for the four-blade 2.5cm diameter rotor.

The most obvious trend is a tremendous loss of performance for the small-hub rotor. The designed chord and incidence distributions are identical for all these rotors excluding the hub region, but clearly the performance is dramatically different. This is caused by large reductions in the blade incidence from structural deformations due to rotation. Other than this one anomalous case, the rapid analysis method predicts thrust values consistently 5% to 10% above the experimental values. The agreement with the three OVERFLOW-D analyses is considerably better and the comparative results for these two methods are presented in Table 6.1.

On average, OVERFLOW-D predicts 4.4% lower thrust and 7.1% higher power. Given the tremendous reduction in computational intensity required by the rapid analysis method these results are encouraging. A primary reason for limiting the number of OVERFLOW-D analyses to a small number is this computational expense, roughly 24 hours of computation on 16 SGI Origin 2000 processors, compared with a one to two minute analysis cycle on a typical desktop PC for the rapid analysis method.

TABLE 6.1 Comparison of predicted thrust and power required at three operating points.

@ 40,620 RPM	Thrust	Power Required
Overflow-D:	3.14 g	0.337 W
Current Method:	3.18 g	0.305 W
@ 46790 RPM		
Overflow-D:	3.91 g	0.473 W
Current Method:	4.25 g	0.456 W
@ 50,760 RPM		
Overflow-D:	4.86 g	0.627 W
Current Method:	5.01 g	0.576 W

Measurements of the required rotor power have not been completed on the identical samples used for thrust testing with the exception of the small-hub version. The time between the initial thrust testing and the development of a workable solution for rotor torque measurement resulted in rotor deformations due to aging. However, the total electro-mechanical efficiencies, encompassing the 5mm Smoovy motor, wiring harness, and closed-loop controller, derived from testing two large-hub samples, provides a model that has been applied to the other large-hub rotor thrust test data to derive the power required from the known electrical power input. For a fixed test setup, this efficiency is purely dependent on RPM, current, and voltage, independent of what is loading the motor.

These determining factors correspond reasonably well between the two large-hub epoxy rotors that have been tested for power required and the rotors used for thrust testing. The determination of electromechanical efficiency is subsequently reduced to a single parameter model based on RPM. The experimental data for RPM versus input current and voltage for all of the large-hub rotors are presented in Figures 6.14 and 6.15. The electro-mechanical efficiency curves for the two large-hub power required test cases are shown in Figure 6.16 along with the quadratic fit used for the other large-hub samples.

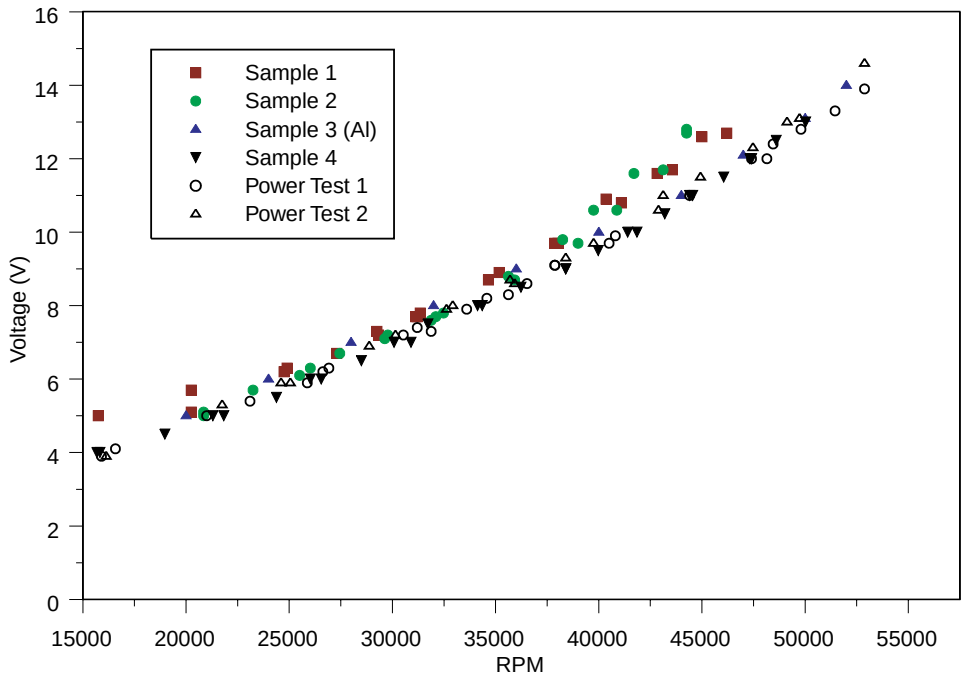

FIGURE 6.14 Experimental input voltage for the four-blade 2.5cm rotors.

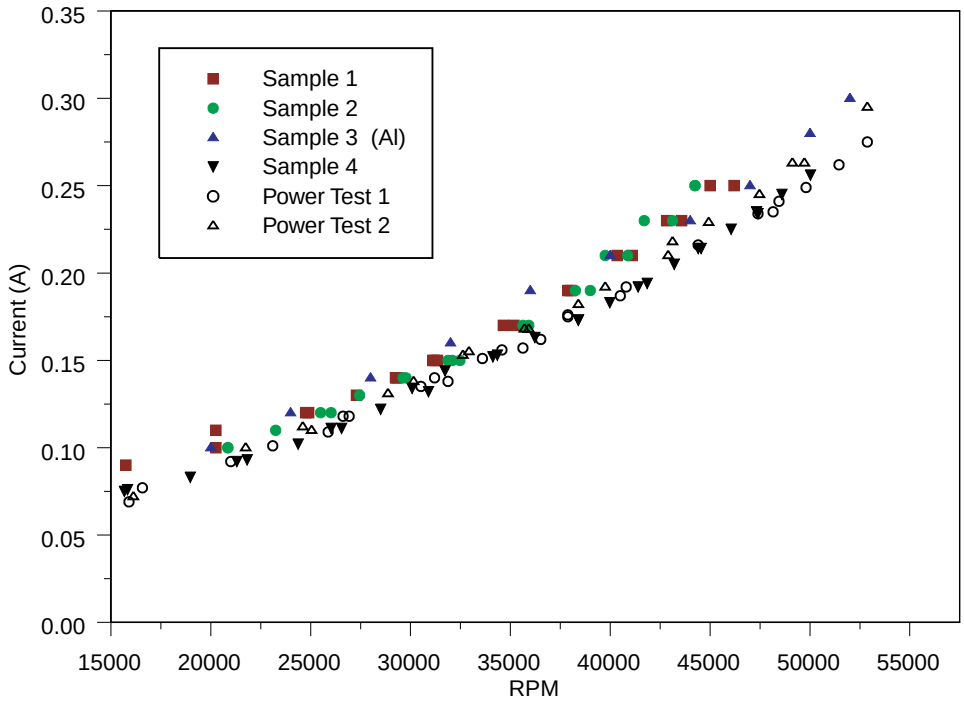

FIGURE 6.15 Experimental input current for the four-blade 2.5cm rotors.

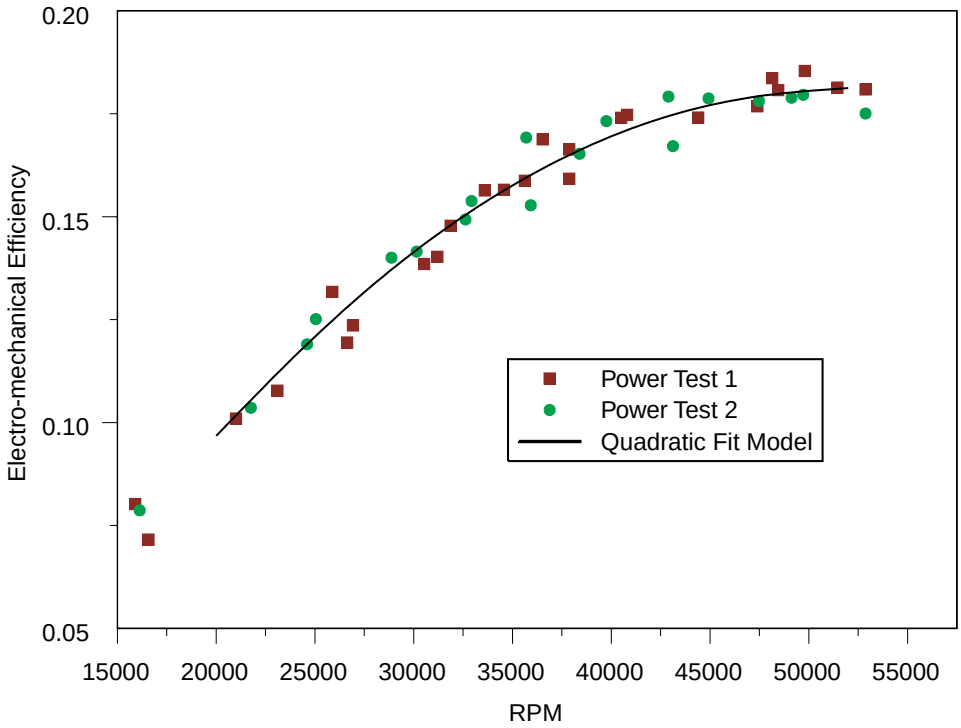

FIGURE 6.16 Experimental and modeled efficiency of the 5mm Smoovy system.

For a rotor with radically different thrust/power versus RPM characteristics, this single parameter model fails because it in effect is assuming that the torque-loading versus RPM is the same for all samples. The test cases considered are not identical in this respect but are close enough that the total deviation in efficiency at any RPM is approximately 1%.

The resulting power required for the four-blade 2.5cm diameter rotors is presented in Figure 6.17. The data for the small-hub version is directly obtained by experiment. Once again the performance of the small-hub version stands out. The deflection of the rotor due to rotation unloads the rotor to such and extent that the power required is reduced almost to the zero-thrust base drag value. For the remaining cases, we see trends similar to those for the two-blade ten inch diameter rotor. The rapid analysis method predicted power compares well at lower RPM with increasing discrepancies as RPM is increased, reaching 15% at 50,000 RPM for the aluminum rotor, but closer to

10% for the other samples. This coincides with predictions of increasing regions of stall near the tip. This is reinforced by similar behavior in the OVERFLOW-D results.

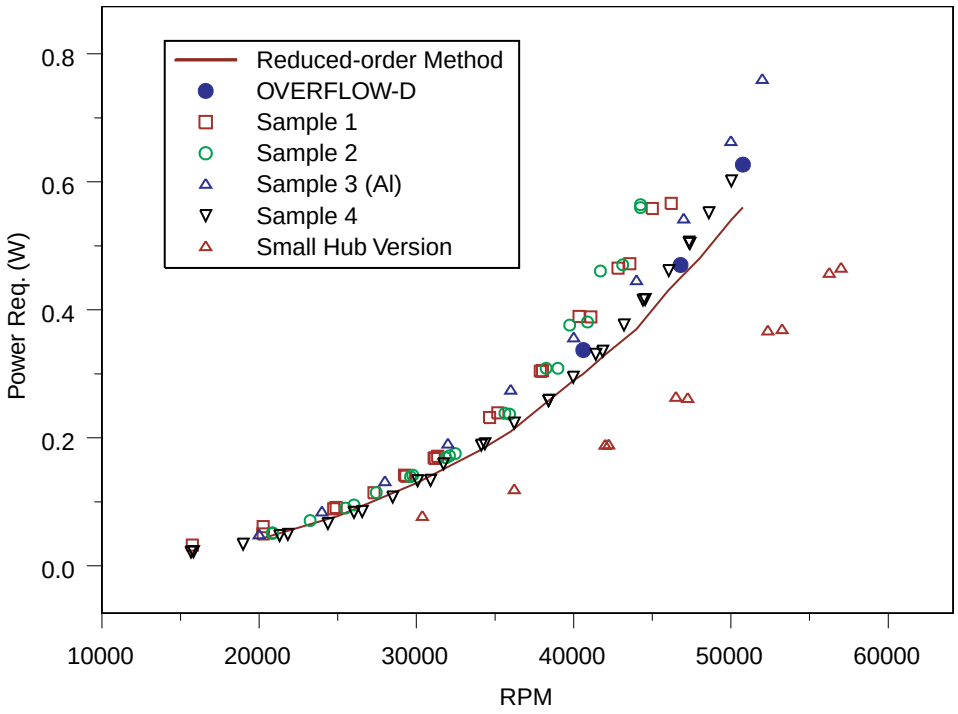

FIGURE 6.17 Predicted and experimental power required for the four-blade 2.5cm diameter rotor.

Another key point of interest is the large amount of power required by the Sample-1 and Sample-2 rotors, particularly apparent in the 40,000 to 50,000 RPM range, while the aluminum and Sample-4 rotors are very consistent with each other and the two analysis predictions. Sample-2 does exhibit additional thrust in this region, but Sample-1 does not and simply requires more power for a given amount of thrust.

It is important to stress that variations between the rapid analysis prediction and experiment discussed here are less than 10% to 15% in power required, typically less than 5% to 10% in thrust. This is adequate for preliminary design and feasibility studies and demonstrates the value of this tool. The following sections investigate the sources of

these variations and possible solutions, but even the current error magnitude only marginally diminishes the utility of the rapid analysis method. The one glaring exception is the small-hub rotor. In this case performance radically differs from both the predictions and the other experiments. Similar problems are seen in the last of the three designs.

6.4.3 Five-Blade 2.2cm Diameter Rotor

This rotor has the most non-traditional planform of all three designs. Key features are the small hub radius and root chords, large mid span chords, and high local solidity. This rotor is predicted to perform similarly to the four-blade rotor, generating three to four grams of thrust between 35,000 and 45,000 RPM, but the experimental results have been surprising and quite disappointing. The predicted and measured thrust for the original version of this rotor is presented in Figure 6.18. The behavior is very similar to that of the small hub four-blade rotor and is explained in the next section.

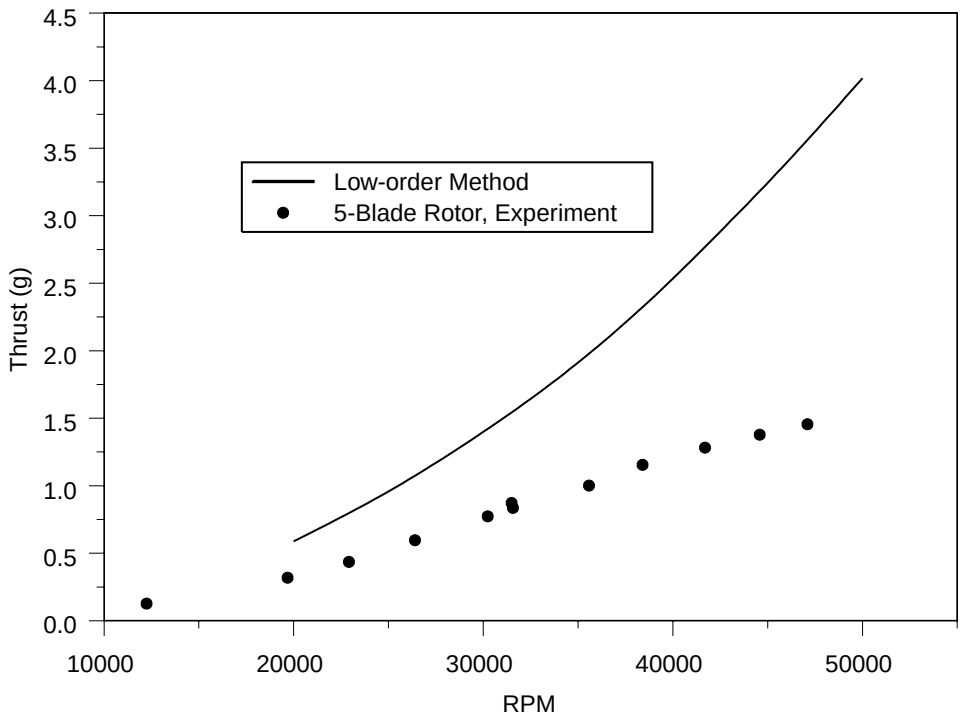

FIGURE 6.18 Predicted and experimental thrust of the five-blade 2.2cm rotor.

6.5 Effects of Structural Deformation

Several observed characteristics in the experimental results can be attributed to deviations of the flying geometry from the prescribed design. One component is due to static deformations incurred at some point prior to testing. A second, and much more universal source, is torsional deflection due to rotational mass effects augmented by the applied aerodynamic forces and moments. The former accounts for the anomalous performance of the four-blade Sample-1 and Sample-2 rotors. The later accounts for the dramatic loss of performance for the four-blade and five-blade small hub rotors, and has larger ramifications on the overall design process and choice of swirl models.

6.5.1 Static Structural Deformations

Although the exact cause is unknown at this time, both the Sample-1 and Sample-2 four-blade rotors have been shown to have static deformations that change the incidence of all four blades from the original design. The aluminum version also has deformations, but most likely due to the much higher stiffness of the material, the deflections are much smaller. All three rotors have been laser scanned as documented by Cheng [32], but further treatment and analysis of the raw data is provided here.

The initial data as provided by the scanning service consisted of a map of X-Y-Z surface points. This data was then reduced by Cheng to local incidence angles along each blade. At this point the scatter in the data along any individual blade appeared to be severe, making it difficult to extract a reasonable trend in incidence. This scatter can be attributed to a known error band consisting of plus or minus the angle represented by the thickness of the airfoil relative to the chord as depicted in Figure 6.19. In determining incidence from the raw X-Y-Z data, the highest and lowest points in the chord-wise direction were sought out, representing the local leading and trailing edges. The thin 2% sections combined with high incidence angles make it difficult to discern between the

upper and lower surface corners of each edge, introducing a possible error in the calculated incidence equal to +/- ArcTan(t/c).

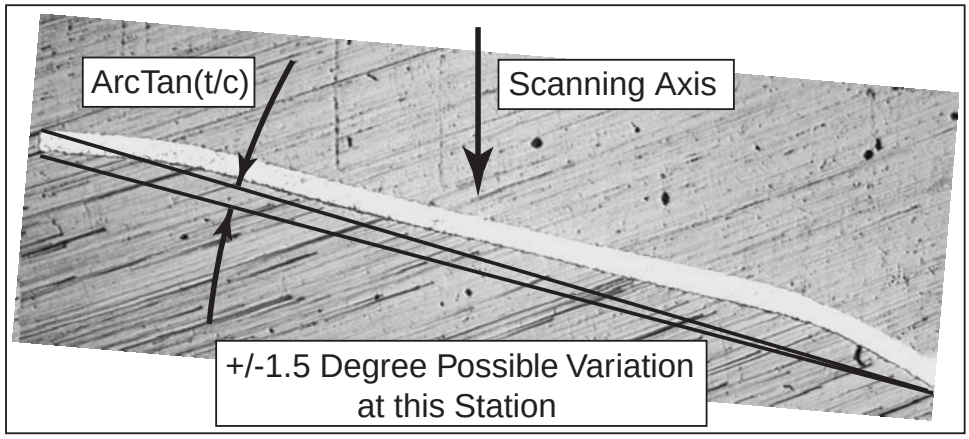

FIGURE 6.19 Photomicrograph of a blade cross-section from an aluminum four-blade rotor, demonstrating the potential for error in incidence determination.

The incidence values determined from scanning for each blade have been modeled by least squares fits of the data. Error bounds based on this assumed source contain the scatter, even capturing the effect of local variations in the thickness ratio. The results shown in Figure 6.20 for a single blade of the four-blade Sample-1 rotor are typical.

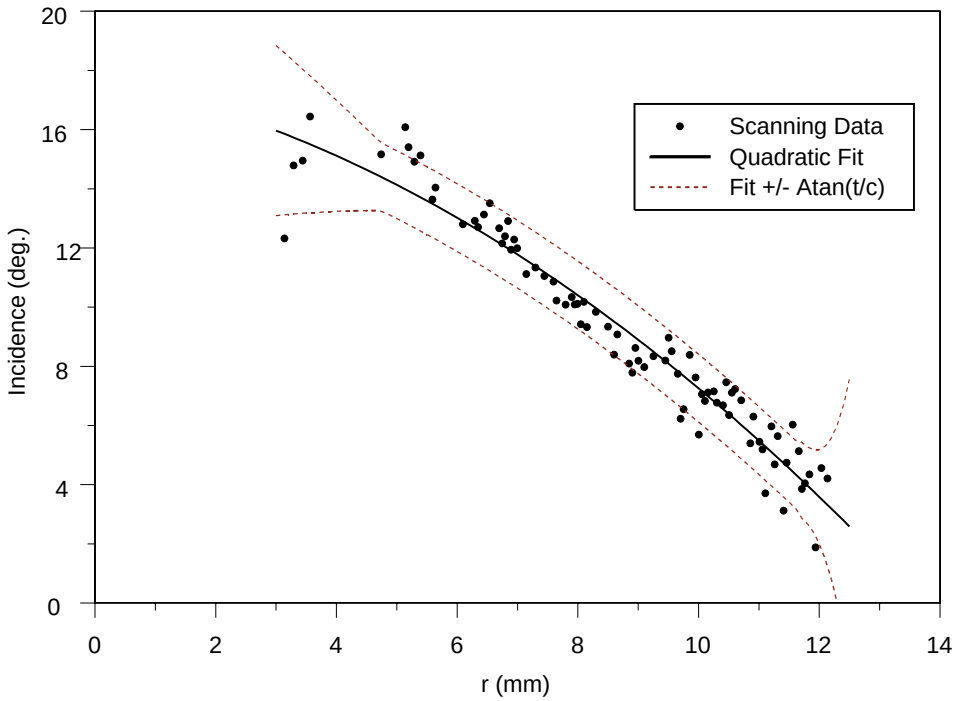

FIGURE 6.20 Comparison of laser scanning incidence data and quadratic fit with error bounds for one blade of the Sample-1 four-blade 2.5cm rotor.

The incidence distributions for the Sample-1, Sample-2, and Sample-3 rotors based on quadratic fitting of the scan data are shown in Figures 6.21, 6.22 and 6.23. The cause of these variations is not clear at this time. The SDM process results in the correct geometry prior to the part being removed from the substrate, so the deformation must occur either during the extraction of the finished part or after production due to material aging or environmental factors. This is one area for further study, but is outside the focus of this work.

Knowledge of the as-tested rotor geometries does permit further insight into some of the variations seen in the thrust and power required data presented earlier for rotors that are ostensibly the same design. The rotor performance with dissimilar blades is difficult to estimate quantitatively, the rapid analysis method assumes identical blades, as does the current OVERFLOW-D calculations using a periodic domain, but reasonable qualitative arguments can be made using this information.

The Sample-1 rotor exhibits similar thrust to the Sample-3 aluminum rotor but requires additional power. Comparing Figures 6.21 and 6.22, the aluminum rotor is closest to the specified geometry, while the Sample-1 rotor has two blades at greatly reduced incidence and one blade at increased incidence. The design code limits the section lift coefficients to a prescribed value, typically one to two tenths below the maximum steady-state lift coefficient, meaning that additional lift is possible.

Considering the drag polars of Chapter 3, near stall this lift would come at an ever increasing cost in drag. The single blade of Sample-1 with increased incidence compensates for the loss of lift from the two low incidence blades, but at a total increase in power required as compared to the design and Sample-3. A similar argument may be made for the Sample-2 rotor. In this case two blades have excessive incidence with another slightly low. This creates another situation requiring excess power, but in this case additional thrust is also generated. This is particularly evident between 35,000 and 45,000 RPM.

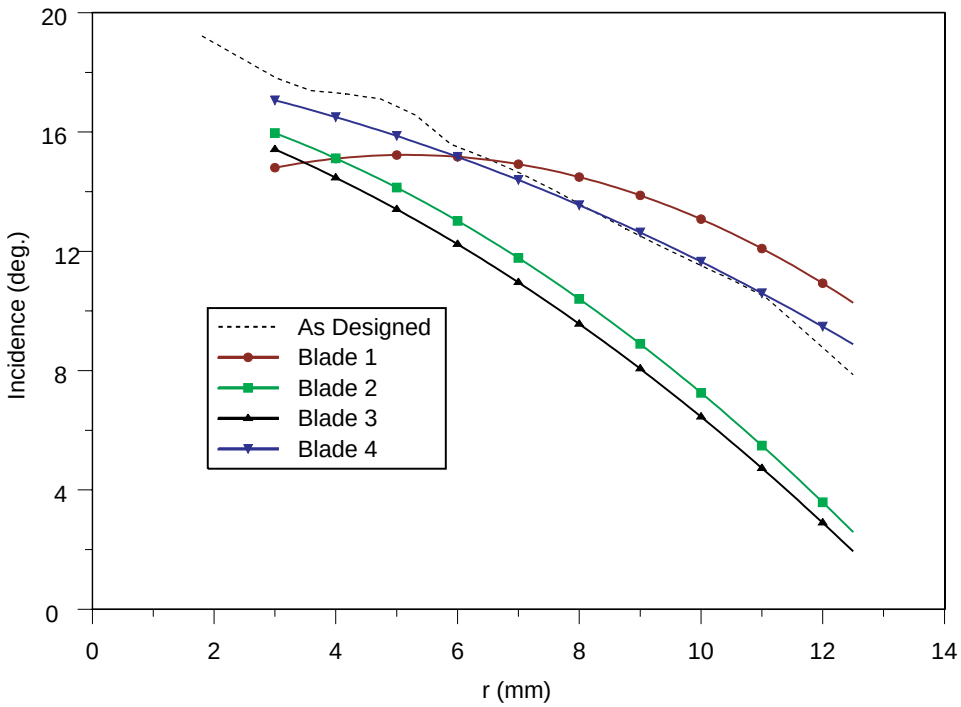

FIGURE 6.21 Sample-1 blade incidence distributions based on quadratic fitting of laser-scan data.

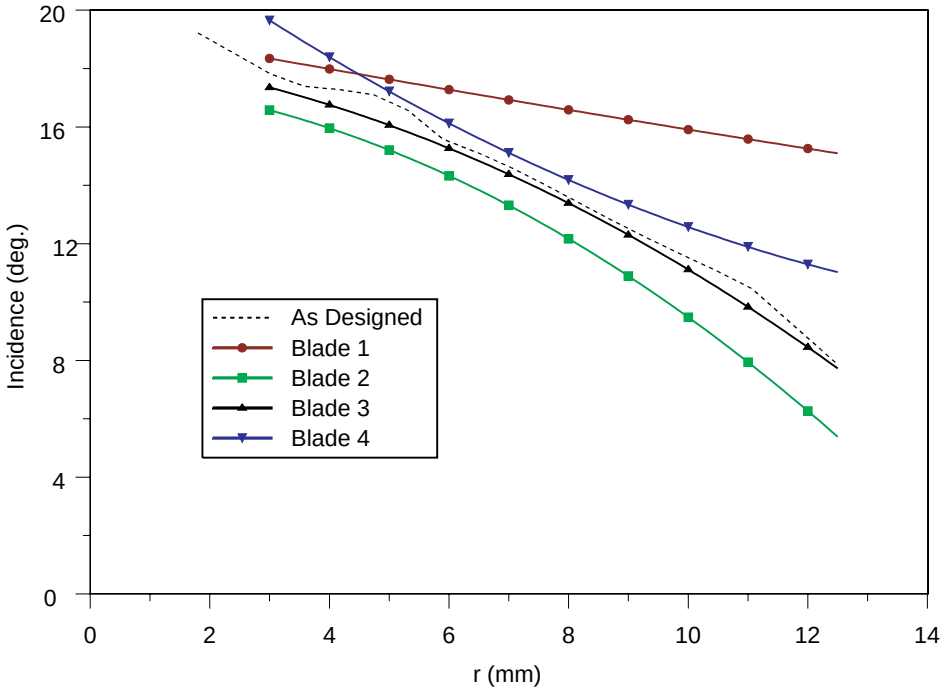

FIGURE 6.22 Sample-2 blade incidence distributions based on quadratic fitting of laser-scan data.

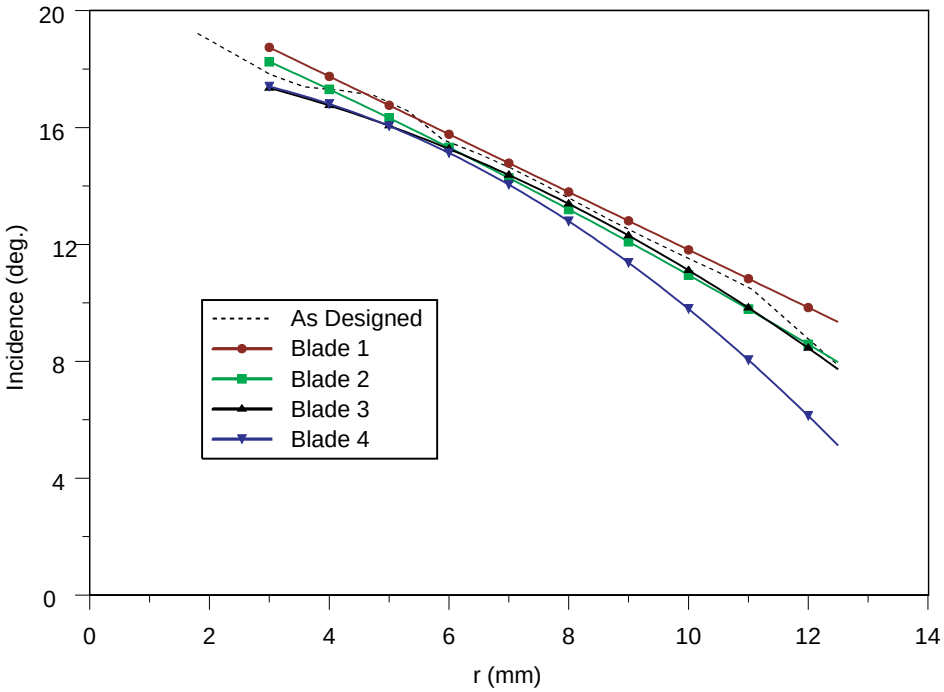

FIGURE 6.23 Sample-3 blade incidence distributions based on quadratic fitting of laser-scan data.

The static deformations described in this section account for some of the detailed behavior of particular rotor samples and emphasize the potential problems associated with precision manufacturing at this scale. This is particularly true of any aerodynamic surface where structural variations of only a few degrees can have a distinct effect on performance. While in this section deformations were examined on a case by case basis in a necessarily qualitative manner, the next section considers, in detail, deformations that are systemic in nature and broadly, and somewhat uniquely, affect rotor design at these small physical scales.

6.5.2 Operational Deformations

Rotors spinning from 30,000 to over 50,000 RPM experience tremendous inertial forces due to centripetal acceleration. At the design speed of 47,000 RPM the tip of the four-blade 2.5cm rotor experiences centripetal accelerations 30,000 times the force of gravity. These forces should be beneficial axially, stiffening the blade against longitudinal bending, but the thin structures coupled with the incidence of blades has the potential for strong torsional effects.

Thin Plate Beam Torsion Model

The magnitude of torsional deflections due to rotation have been modeled using beam torsion theory. The thin blade sections and relatively high aspect ratio of the blades make use of this simple model acceptable. The local blade cross-section is reduced to a flat plate with the structural axis at mid-chord. This simplifies the derivation of a closed form for the local torsional moment and is reasonable considering the thin and lightly cambered, constant thickness airfoils being used, and the fact that both the four-blade and five-blade rotors have zero sweep at mid-chord. Other than these simplifications, the full rotor geometry is modeled, with chord, thickness, and design incidence permitted to vary along the span.

The governing equation relates the rate of twist to the applied torque subject to the boundary conditions that Q(tip) = 0 and θ(0)=0:

$$\frac{d\theta}{dR}(r) = \frac{1}{GJ(r)} \int_{r}^{tip} Q(R)dR \tag{6.1}$$

The polar moment of inertia for a thin flat plate can be expressed as:

$$J(r) = \frac{c(r)t(r)^3}{3} = \frac{1}{3}c(r)^4\left(\frac{t}{c}(r)\right)^3 \tag{6.2}$$

What remains is to derive an expression for the torque applied to spanwise differential element. The model used in deriving the toque expression is depicted in Figure 6.24. A differential element of an inclined blade section is displaced above or below the rotational plane by some distance *h*. This element having mass *dm* experiences the centripetal acceleration $\omega^2\hat{r}$. The resulting force vector $dm\omega^2\hat{r}$ has no component in the y direction and a component in the x direction proportional to *Sin(β).* The moment about the structural axis due to *dm* can be expressed as:

$$dQ(r) = -dm\omega^2\|\hat{r}\|h\sin(\beta) \tag{6.3}$$

This is then expressed in terms of c, t/c, $\rho_{material}$, R, and ζ, and integrated across the chord to yield:

$$Q(r) = -\frac{1}{12}\rho_{mat}\frac{t}{c}c^4\omega^2\cos(\zeta)\sin(\zeta)dR \tag{6.4}$$

This would be the torque without any torsional deflection, and would provide a first order estimate, but the actual torque will be a function of the built incidence and the torsional deflection, taking the final form:

$$Q(r) = -\frac{1}{12}\rho_{mat}\frac{t}{c}c^4\omega^2\cos(\zeta+\theta)\sin(\zeta+\theta)dR \qquad (6.5)$$

This problem is solved as a least squares vector minimization of Eqn.6.1 posed as a homogeneous problem with $\theta(r)$ as the unknowns. The left-hand side of Eqn.6.1 is computed using second-order central differencing; the integration of $Q(r)$ uses a trapezoidal approximation. In addition to the torque expressed by Eqn.6.5, the aerodynamic torque is modeled by including a nominal section pitching moment and the lift acting at the quarter chord.

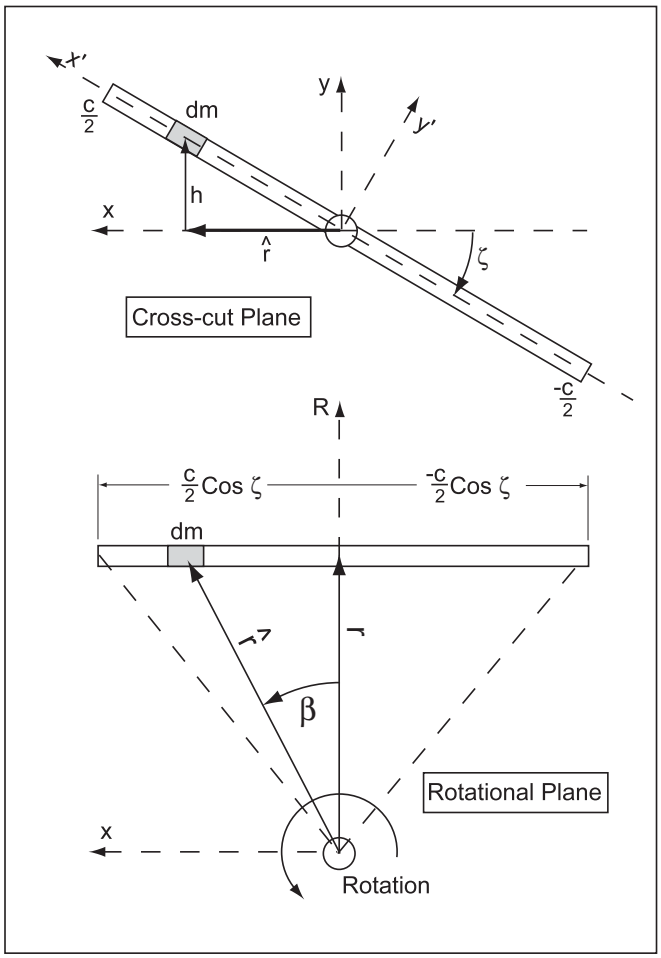

FIGURE 6.24 Depiction of a chord-wise blade element for a rotational torsional deflection model.

The importance of the operational deflections for micro-rotors becomes apparent if Eqn.6.1 is considered with the restrictions of constant chord and thickness along a blade:

$$\frac{d\theta}{dR}(r) = -\frac{3}{12}\rho_{mat}\omega^2\frac{1}{G}\frac{1}{(t/c)^2}\int_r^{tip}\cos(\zeta+\theta)\sin(\zeta+\theta)dR \qquad (6.6)$$

This may be further manipulated to show a predominant dependence of θ on ωR (the tip speed) and the inverse of the thickness ratio squared. The tip speeds of the micro-rotors presented here are roughly one third to one fourth those of full scale helicopter rotors, but t/c is reduced by a factor of three to four and dominates. This effect is also seen on large scale rotors, but for the development of small rotors it should be considered an essential component of design. Large rotors typically have some form of collective control that can be used to partially compensate for the loss of incidence, but the only control available to the micro-rotors is RPM.

The effects of the aerodynamic forces are not negligible and are strongly coupled to the structural deformations. The pitching moment coefficient is typically insensitive to angle of attack and provides a relatively small nose down moment that works in unison with the structural moment due to rotation. However, the local lift coefficient is obviously very sensitive to the incidence angle and, for a positive lift coefficient, opposes the inertial effects of rotation. The correct deflected rotor solution would be at structural equilibrium and operating at lift coefficients representative of the equilibrium geometry. This requires some form of iteration. Ideally this would be integrated into the design process, but at this time that has not been implemented. For the purpose of the analyses presented here, several iterations have been carried out manually at each rotor speed, but the solutions are not fully converged, meaning that the assumed lift coefficients used for the structural analyses do not precisely match the predicted value from the final performance analyses.

Structural Analyses

The method described above has been applied to five cases consisting of the aluminum four-blade rotor, the epoxy large-hub and small-hub four-blade rotors, and the two previously described versions of the five-blade 2.2cm rotor. The initial geometries are taken as the designed geometries, no static deformations are included. The aluminum rotor is constructed of 7075-T6 aluminum. The material properties used are E=72 GPa, G=27Gpa, and ρ=2796 kg/m^3. The epoxy rotors are constructed from Adtec EE 501/530 epoxy. The shear properties of this resin have been approximated from the results of tensile testing reported by Kietzman [42]. The UTS and elongation to failure have been used do obtain an approximate Young's modulus of 4.2GPa. Typically this should be based upon yield characteristics, but this data was not available. Since typical values of Poisson's ratio are between 0.2 and 0.3 for many materials, a value of 0.2 has been assumed, yielding an approximate shear modulus of G=1.75GPa. The use of UTS and the lower Poisson's ratio results in a conservative stiffness value.

The aluminum four-blade rotor was found to have the smallest static deformations of the samples tested and is expected to be significantly stiffer. For these reasons it is included essentially as a baseline analysis, since the effects of rotation should be minimal in this case. The predicted torsional deflections for the aluminum and epoxy large-hub four-blade rotors at five different operating speeds is presented in Figure 6.25. The aluminum rotor performs as expected with minimal deflection, typically under 0.25 degrees, but the epoxy rotor fares worse both in magnitude and the amount of variation over the speed range.

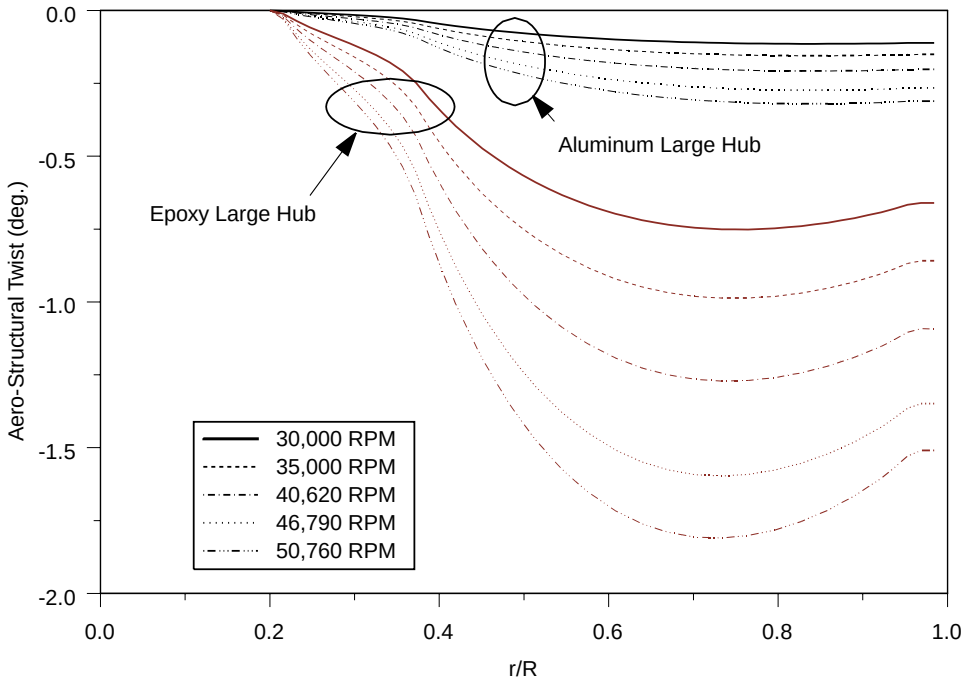

FIGURE 6.25 Predicted torsional deflections for large-hub four-blade rotors.

The change in slope visible near the tips is the result of aerodynamic loading. In this case the lift coefficient is assumed to be 0.5 and the quarter chord pitching moment is set at -0.08. The effect of neglecting the aerodynamic forces and moments is displayed in Figure 6.26 for the large hub epoxy rotor. At both speeds, the tip deflection almost doubles. This demonstrates the importance not only of including these effects, but also the need to accurately predict, or converge to, the final static equilibrium loading.

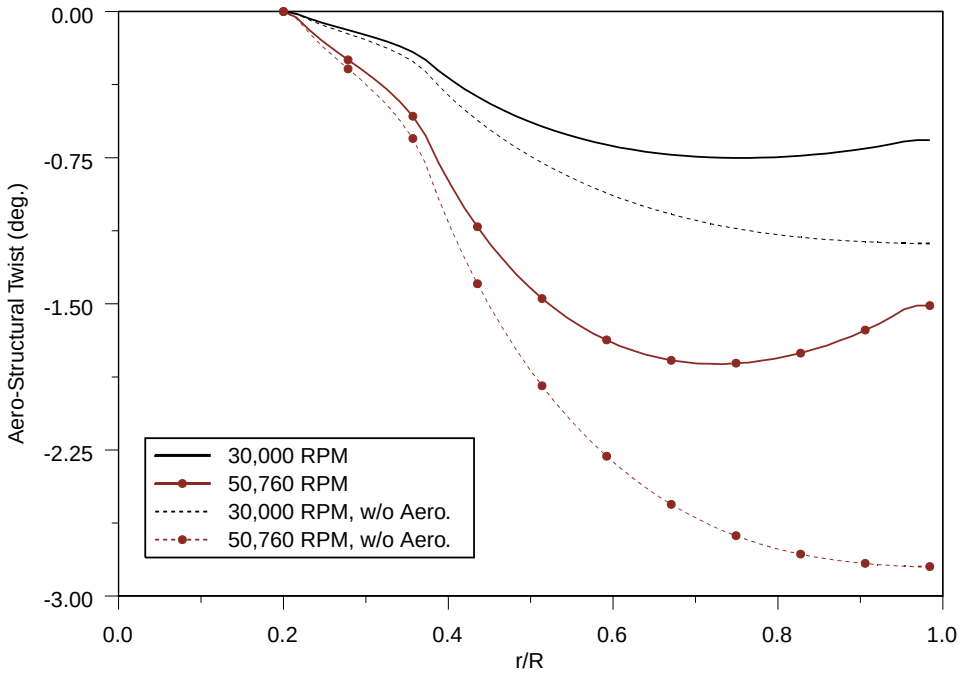

FIGURE 6.26 Predicted torsional deflections with and without aerodynamic effects.

Moving on to compare the large hub and small hub four-blade rotors, the dominant reason for the dramatic difference in performance is revealed. Figure 6.27 compares the predicted torsional deflections for both rotors. For the small hub version, the lift coefficient specified in the structural analysis has been reduced to 0.15. This is still slightly above the values predicted by the performance analysis of the deformed blade. The smaller chords and reduced t/c cause roughly a nine degree deflection over the first 20% of the blade and larger variations across the speed range. This is the explanation for both the reduced thrust and power required seen in Figures 6.13 and 6.17.

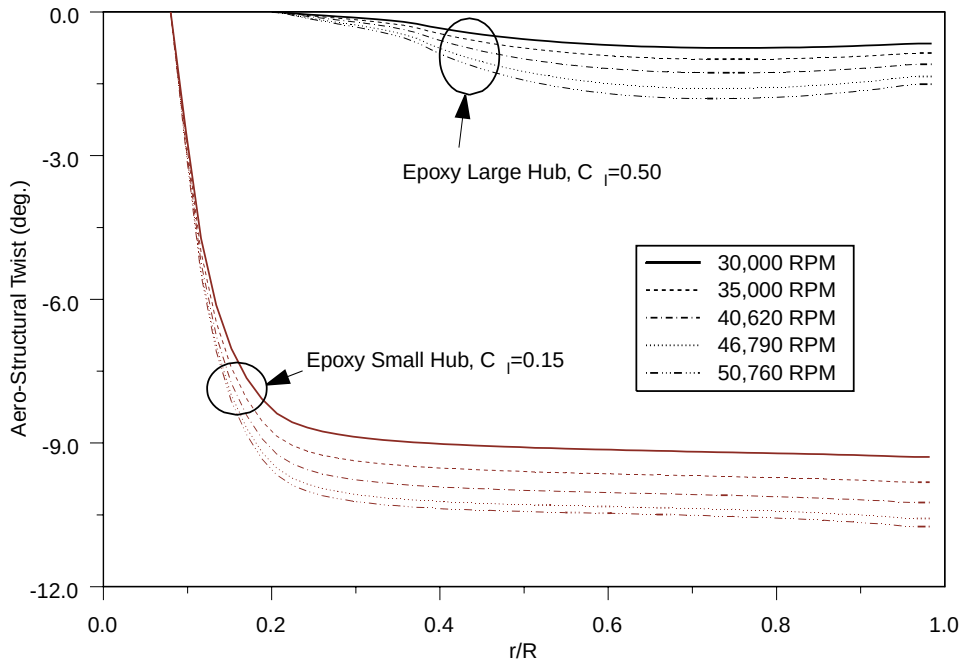

FIGURE 6.27 Predicted torsional deflections for large-hub and small-hub four-blade rotors.

The deflected geometry for each of five operating points have been analyzed using the rapid analysis method of Chapter 4. The results for the aluminum large hub and epoxy small hub rotors are presented in Figures 6.28 and 6.29, along with the experimental results, the original OVERFLOW-D results, and the original rapid analysis method performance predictions. There is minimal movement in the aluminum rotor case indicating that this is not the sole source of the typical 10% variation seen between the rapid analysis method and experiment. However, the torsional model clearly captures the large-scale performance loss of the small hub rotor, both in power and thrust.

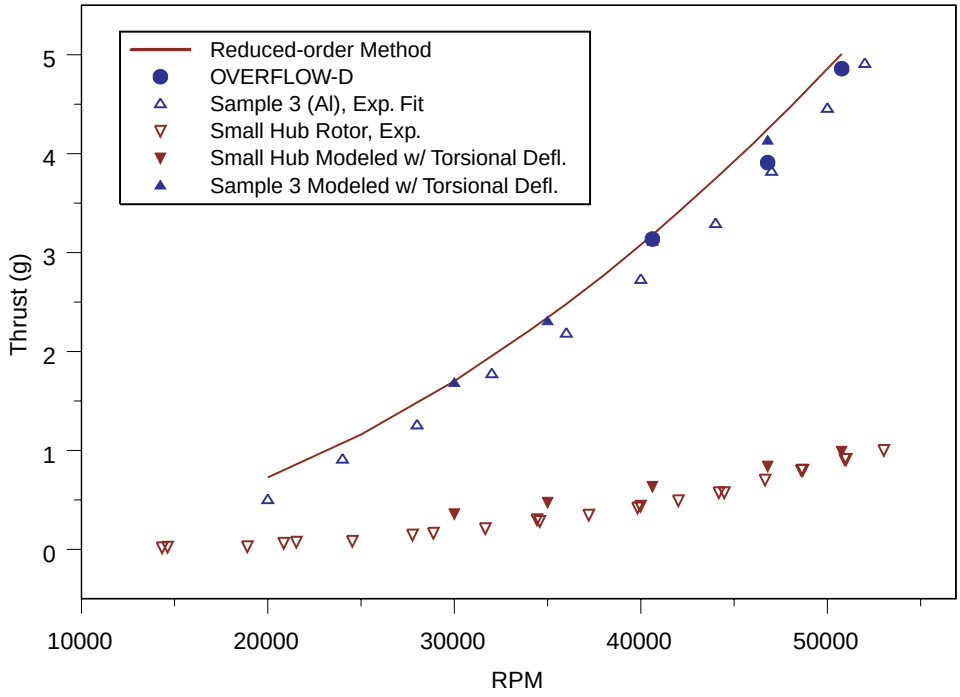

FIGURE 6.28 Comparison of experimental and predicted thrust with torsional deflections.

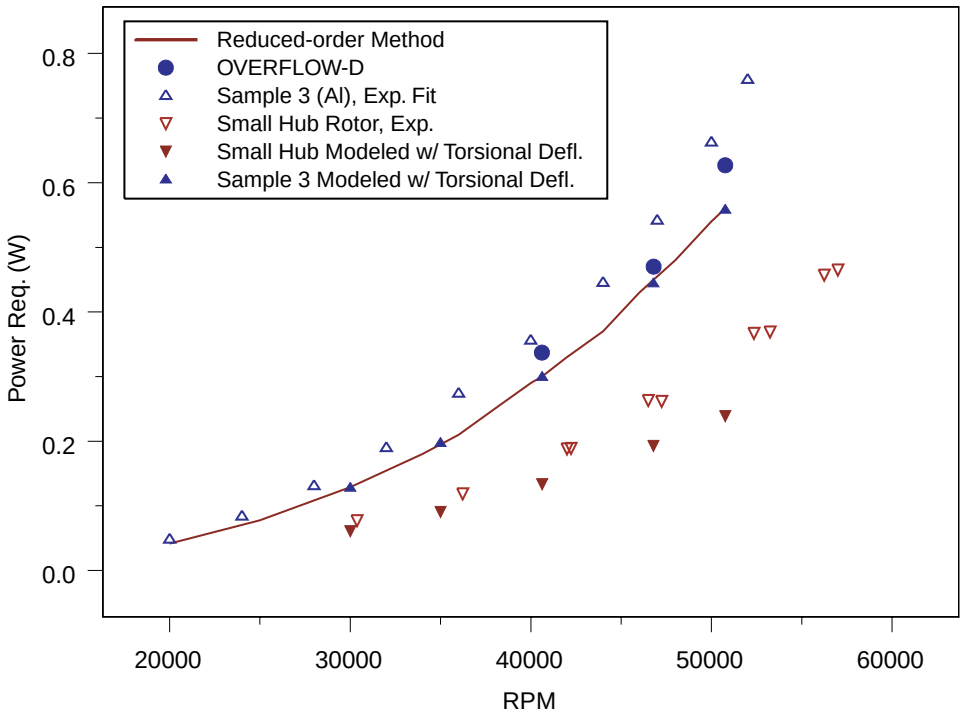

FIGURE 6.29 Comparison of experimental and predicted power required with torsional deflections.

Similar results are found for the two versions of the five-blade 2.2cm diameter rotors. The predicted deflections are shown in Figure 6.30. The large differences result from the variations in the thickness distribution over the first 25% of the blade. The slight change in incidence has negligible effect and in this case the planform and hub diameter is identical for both rotors. Compared to the four-blade rotor this design exhibits roughly twice as much variation across the sampled operating speeds, even though the roots chords and tip speeds are similar. The large mid-span chords of this design contribute to this difference.

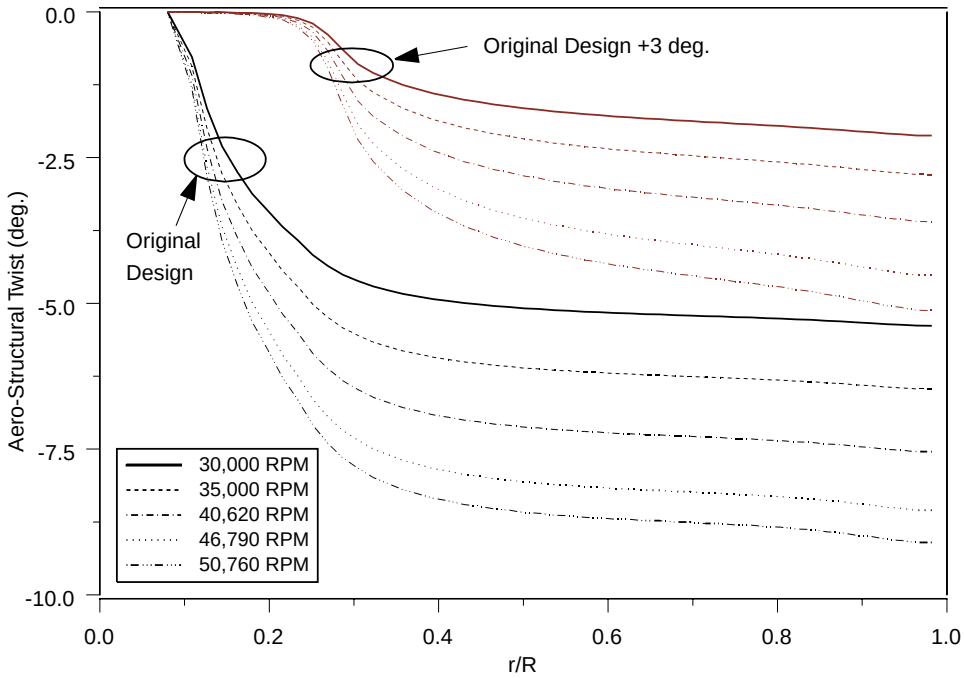

FIGURE 6.30 Predicted torsional deflections for two versions of the five-blade 2.2cm rotor.

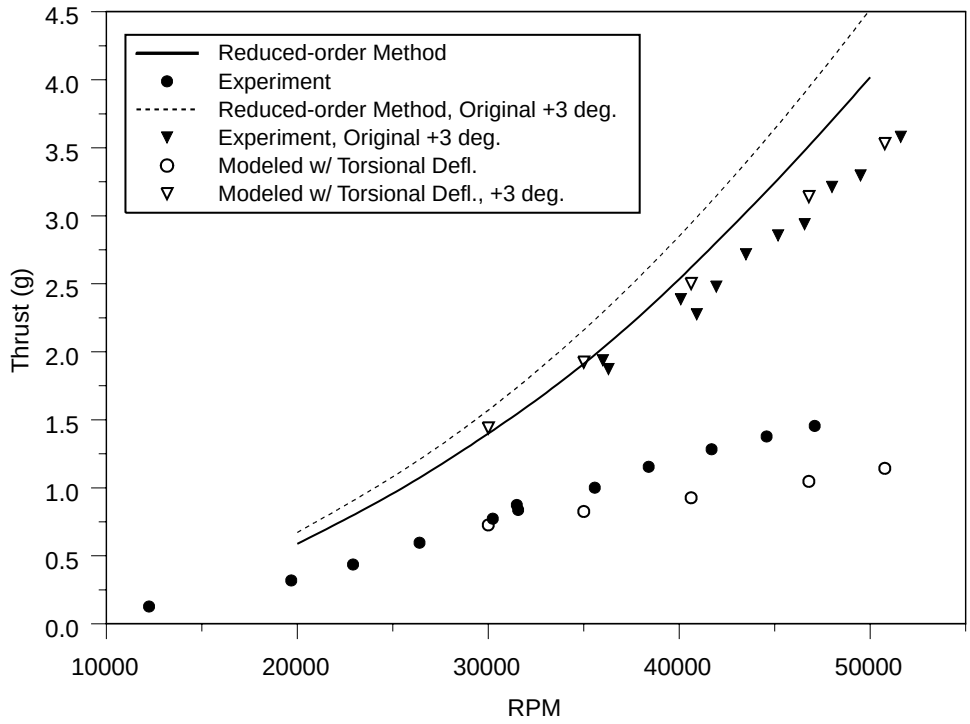

FIGURE 6.31 Comparison of experimental and predicted thrust with torsional deflections for two versions of the five-blade 2.2cm rotor.

The rapid analysis method thrust predictions for the deflected forms are presented with the experimental results and the undeflected predicted performance in Figure 6.31. Here the agreement is not as good as in the four-blade case, but the overall trend is well captured considering that only two iterations on the aerodynamic forces were performed. Performance analysis of these deflected forms reinforces the importance of torsional effects and the ability of the current tools to capture them. It also provides closure on the most anomalous experimental behavior observed over the course of this research.

6.6 Spanwise Thrust and Torque Distributions

The experimental results have so far limited the discussion of performance to the gross performance parameters of total thrust and total power required. Three-dimensional

computational fluid dynamics, however, brings with it a wealth of information and the ability to probe further into the performance of these micro-rotors. The remaining sections are concerned with a detailed comparison of the rapid analysis method predicted performance for the four-blade 2.5cm diameter rotor and the OVERFLOW-D calculations of the same geometry.

Predictions for single blade sectional thrust and torque distributions, using the contracting wake model and angular momentum swirl model, are provided in Figures 6.32 and 6.33. The two predictions show distinct differences in the spanwise thrust distributions. The OVERFLOW-D analyses predict higher thrust outboard of 83% span and lower thrust inboard. The over prediction of thrust is similar to the effect seen previously in the OVERFLOW-D validation case, but, unlike that experiment, which included direct measurement of section pressure distributions, a direct comparison with experimental section properties is not possible in this instance, making it difficult to assess the presence or magnitude of the same effect.

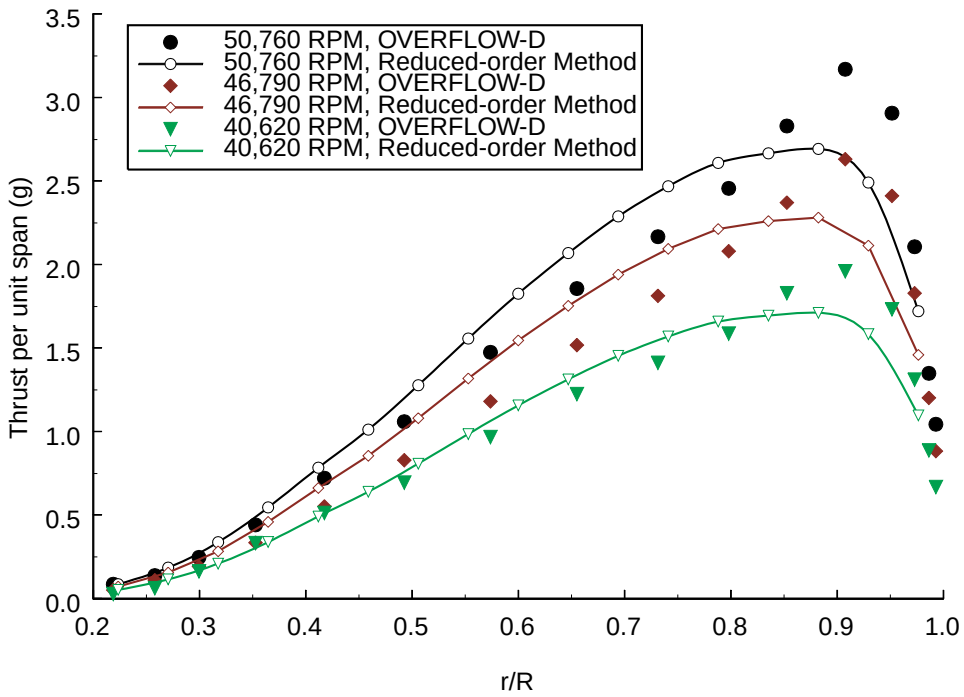

FIGURE 6.32 OVERFLOW-D and rapid analysis method spanwise thrust distributions for the four-blade 2.5cm rotor.

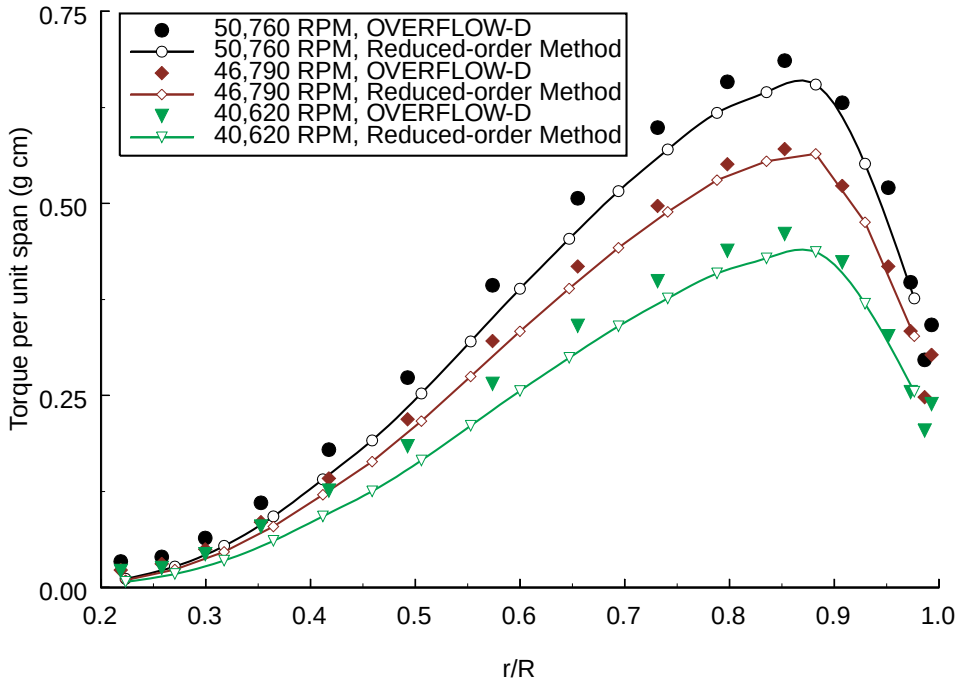

FIGURE 6.33 OVERFLOW-D and rapid analysis method spanwise torque distributions for the four-blade 2.5cm rotor.

The assumption that the discrepancy at the tip region is an artifact of the OVERFLOW-D analysis does not explain the local differences, as much as 15%, in the prediction of thrust between the rapid analysis method and OVERFLOW-D. The shape of the distributions are also consistently different, concave in the case of the OVERLFOW-D analyses and convex for the blade-element method. This is indicative of distinctly different downwash distributions.

6.7 Modeling Effects on Performance Estimation

6.7.1 Effect of Swirl Modeling on Performance Estimation

Application of the three viscous swirl models described in Chapter 3 results in significant variations in both the global and spanwise distributed thrust and torque for a given rotor geometry. The four-blade 2.5cm diameter rotor operating at 48,000 RPM is used as a test case to investigate these differences. The predicted thrust and torque distributions using each of the three viscous swirl models are presented in Figures 6.34 and 6.35. For reference, the classical angular momentum correction has been used for all previous comparisons to OVERFLOW-D results.

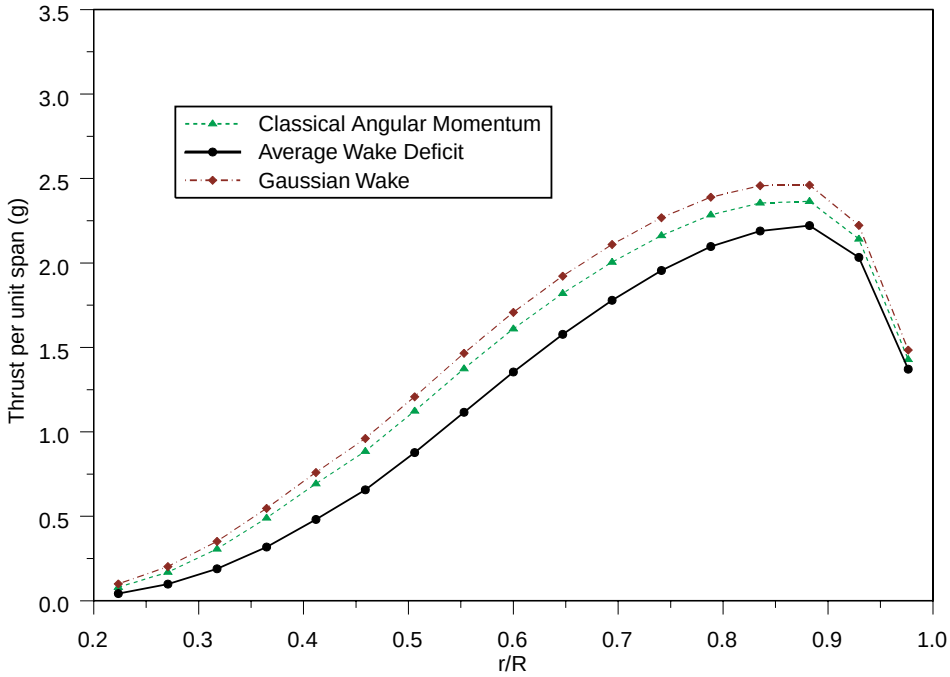

FIGURE 6.34 Predicted spanwise thrust distributions for the four-blade 2.5cm rotor using three different viscous swirl models.

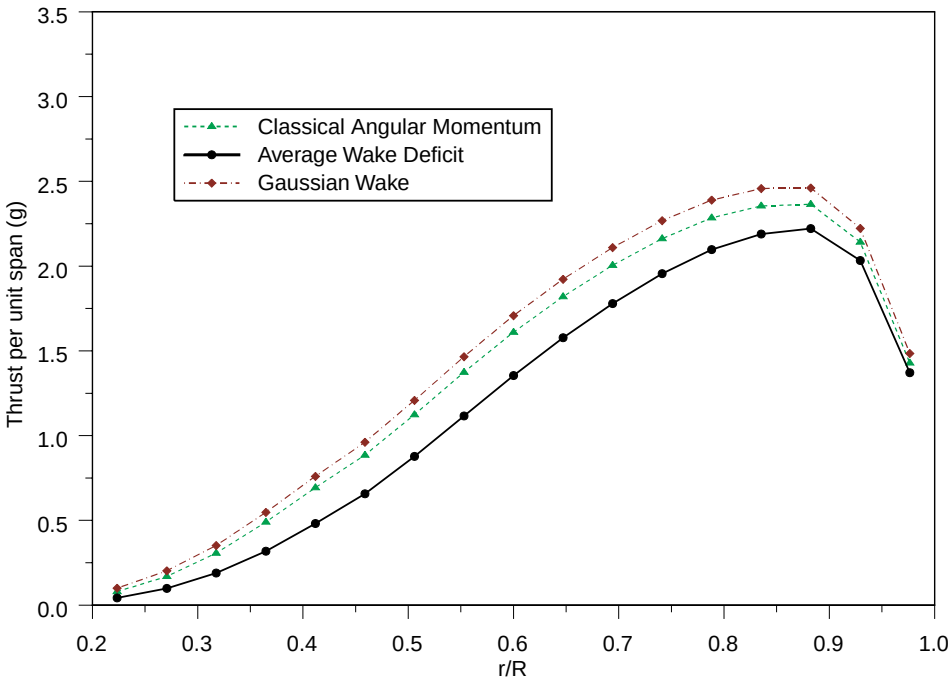

FIGURE 6.35 Predicted spanwise torque distributions for the 4-blade 2.5cm rotor using three different viscous swirl models.

The trends are consistent across the span and globally for both metrics. The Gaussian wake model predicts the highest thrust and required power, 0.0455N and 0.485W. For this example, the localized nature of the Gaussian velocity deficit distribution results in zero viscous swirl effect all along the blade except at the inner-most stations. The thrust and power values drop 5.5% and 4.7% respectively for the angular momentum model. The average wake deficit model results in the lowest values, 17.4% and 15.3% lower than the global thrust and power required predicted with the average wake deficit model.

It was initially thought that the choice of viscous swirl model might have a significant effect on both the inflow angle and the local flow velocity, but this has proven not to be the case. Due to the coupling between the vertical and tangential induced velocities in the rotor relations, the effect on inflow angle is minimal. The predominant effect of the viscous swirl corrections is a reduction in the flow velocity, resulting in a reduction in Reynolds number and dynamic pressure.

The inflow angle variations across all three viscous swirl models are shown in Figure 6.36. The only visible effects occur inboard of half span and increase towards the root as the local solidity increases and the local Reynolds number decreases towards zero. This inboard region contributes only a small portion of the total thrust and torque, consequently inflow variations due to viscous swirl modelling are not a primary factor in performance estimation.

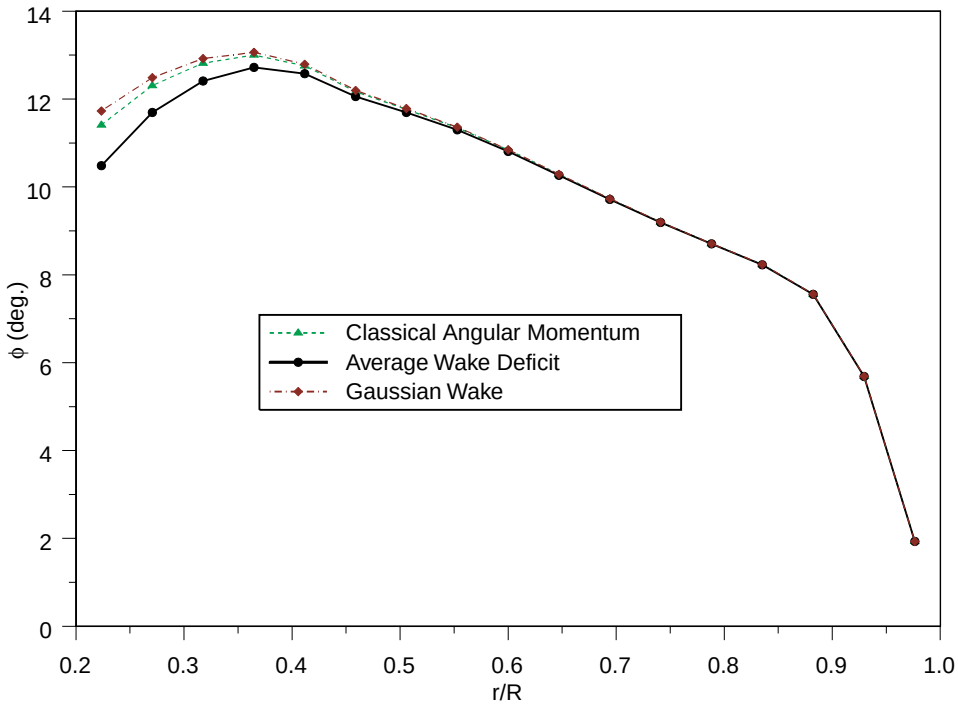

FIGURE 6.36 Predicted inflow angles using three different viscous swirl models.

The reduction in the local relative flow velocity is depicted in Figure 6.37 with the relative flow velocity non-dimensionalized by the local rotational velocity. Once again, the Gaussian wake model is roughly equivalent to having no viscous swirl correction. The reduction in velocity and Reynolds number causes an increase in C_d and a reduction in section L/D, but these effect are relatively small. The maximum variation in C_d across the three models ranges from 2% at the tip up to 5% at mid span. The larger effect is the reduction in the dynamic pressure. This reduces the both the thrust and torque in spite of the increases in the section drag coefficients.

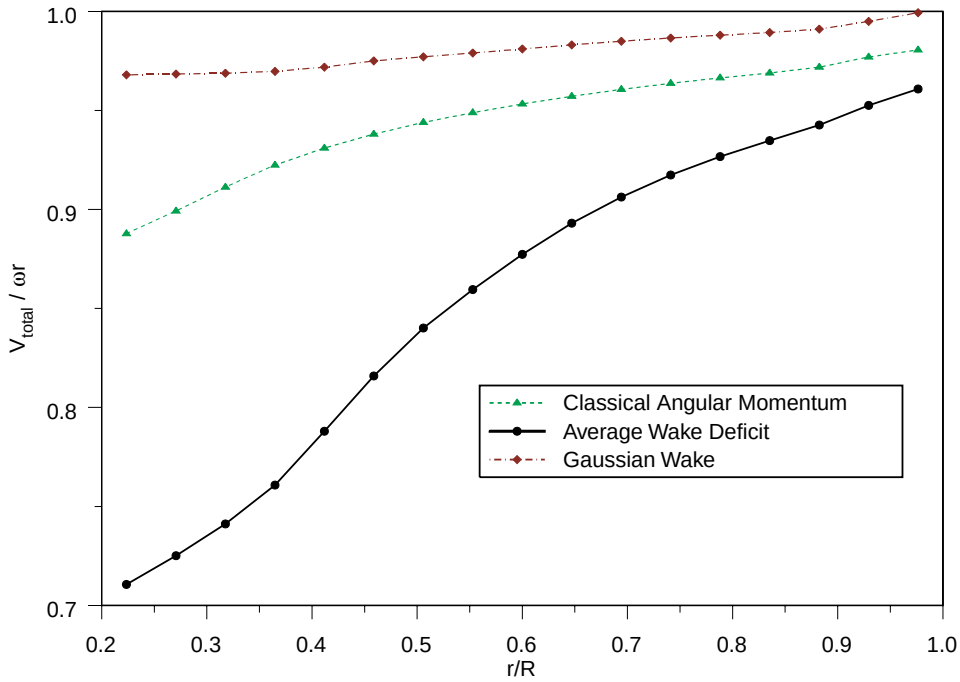

FIGURE 6.37 Local relative flow velocities using three different viscous swirl models.

Of the three models considered, the classical angular momentum provides the best agreement with the OVERFLOW-D results presented earlier. The Gaussian wake model as currently implemented has no effect and the average wake deficit model grossly under-predicts the sectional thrust and torque. None of these models provides an ideal solution. The angular momentum model is appropriate in the outer blade region where the Reynolds numbers are higher and the local annular solidity much lower, but inboard, the models based on actual viscous wake profiles and a direct accounting of local solidity should more accurately represent the flowfield. In spite of these issues, the classical angular momentum correction is satisfactory overall for analysis since higher fidelity is not needed inboard due to the triangular loadings typical of rotors.

6.7.2 Effect of Wake Modeling on Performance Estimation

The impact on performance estimation of the two wake models, the classical Prandtl tip loss correction and the contracted vortex ring model, has also been explored using the

four-blade 2.5cm diameter rotor operating at 48,000 RPM. All analyses incorporate the angular momentum viscous swirl model. The results indicate that the impact on predicted performance is small, but the possible impact on rotor design is significant. Predicted global thrust and power required agree within 1% of each other and the thrust and torque distributions displayed in Figures 6.38 and 6.39 compare well overall. The only significant discrepancies can be seen near the tip.

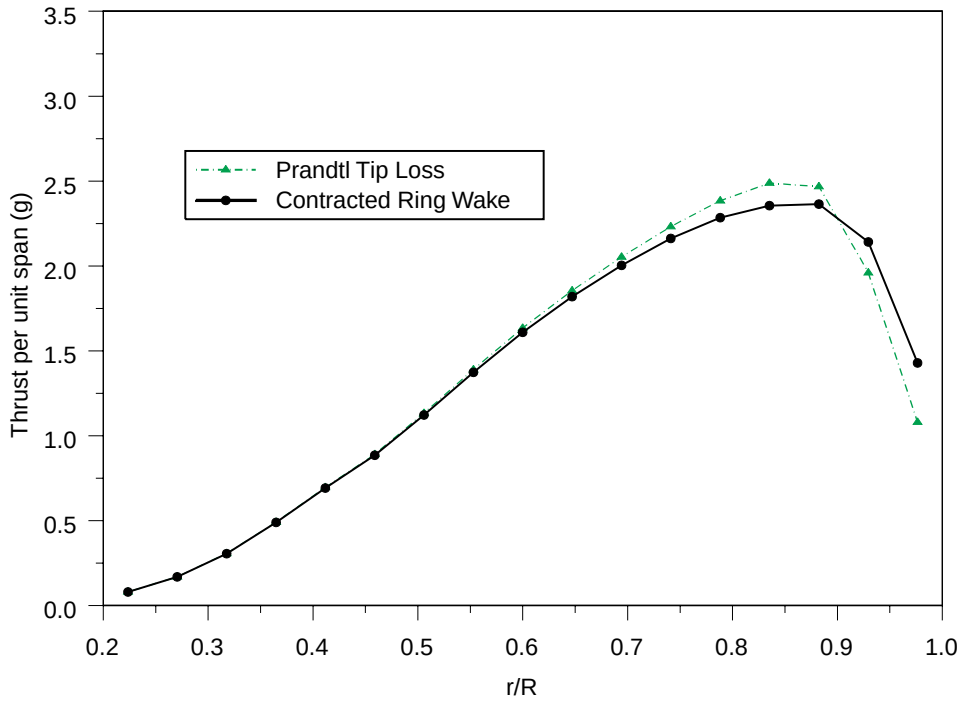

FIGURE 6.38 Predicted spanwise thrust distributions for the four-blade 2.5cm rotor using two different wake models.

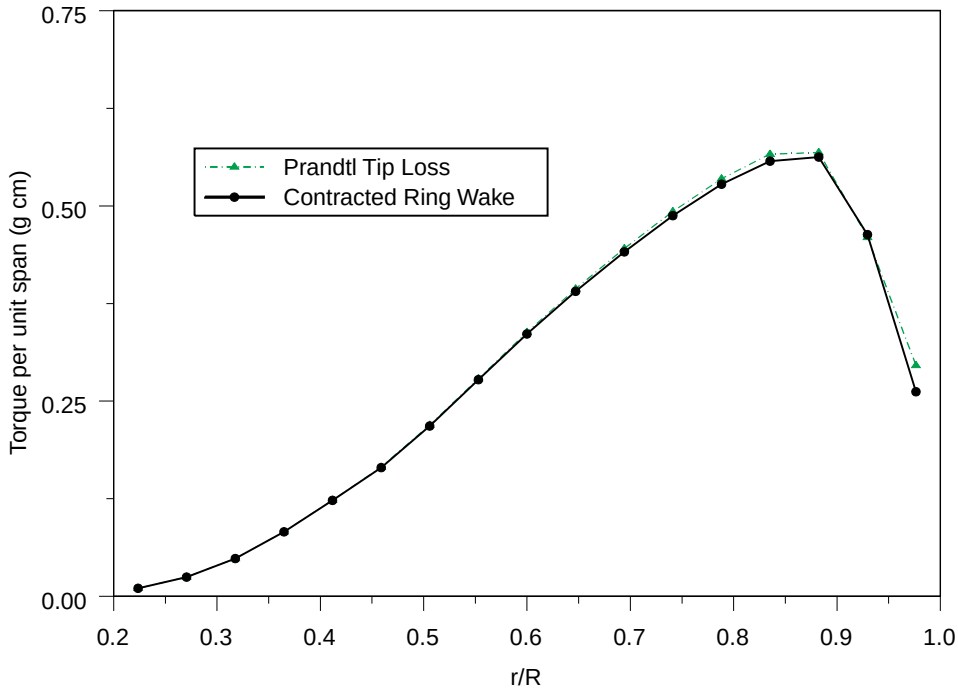

FIGURE 6.39 Predicted spanwise torque distributions for the four-blade 2.5cm rotor using two different wake models.

Unlike the viscous swirl models, which have been shown to primarily effect the local relative velocity, the effect of the wake models is primarily a modification of the rotor inflow angles on the outboard portions of the blade. The relative velocity distributions shown in Figure 6.40 are essentially identical, but in Figure 6.41 significant variations in the inflow angle are apparent outboard of 70% span with a maximum variation of three degrees at the tip. The reduction in inflow angle at the tip for the contracted ring wake model results in higher predicted section lift coefficients, increasing at the tip from $C_\ell = 0.50$ for the Prandtl tip loss model to $C_\ell = 0.66$ for the contracted ring wake.

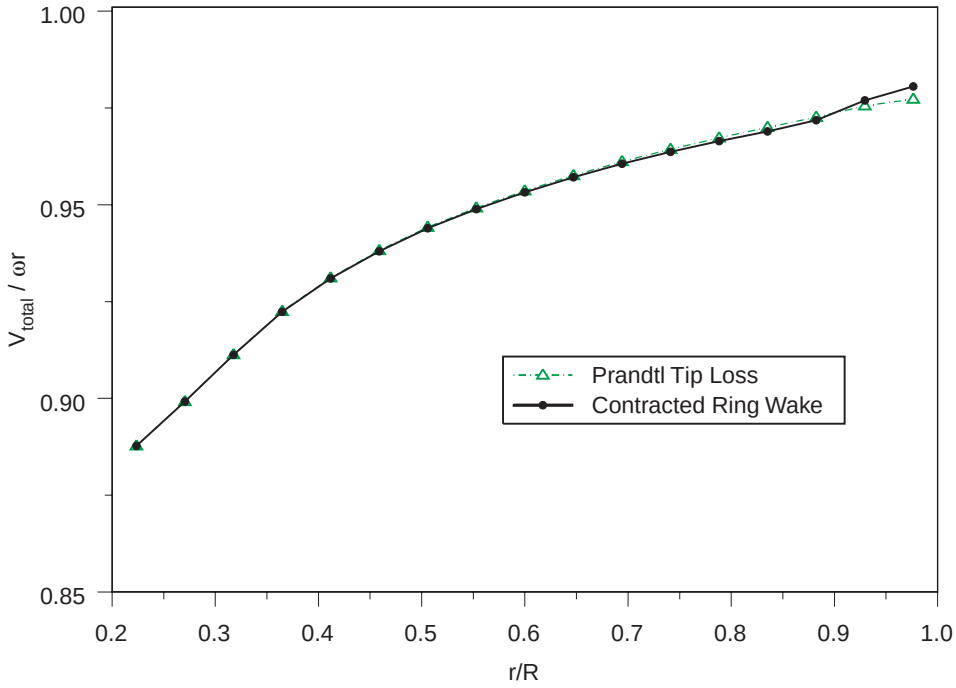

FIGURE 6.40 Local relative flow velocities using two different wake models.

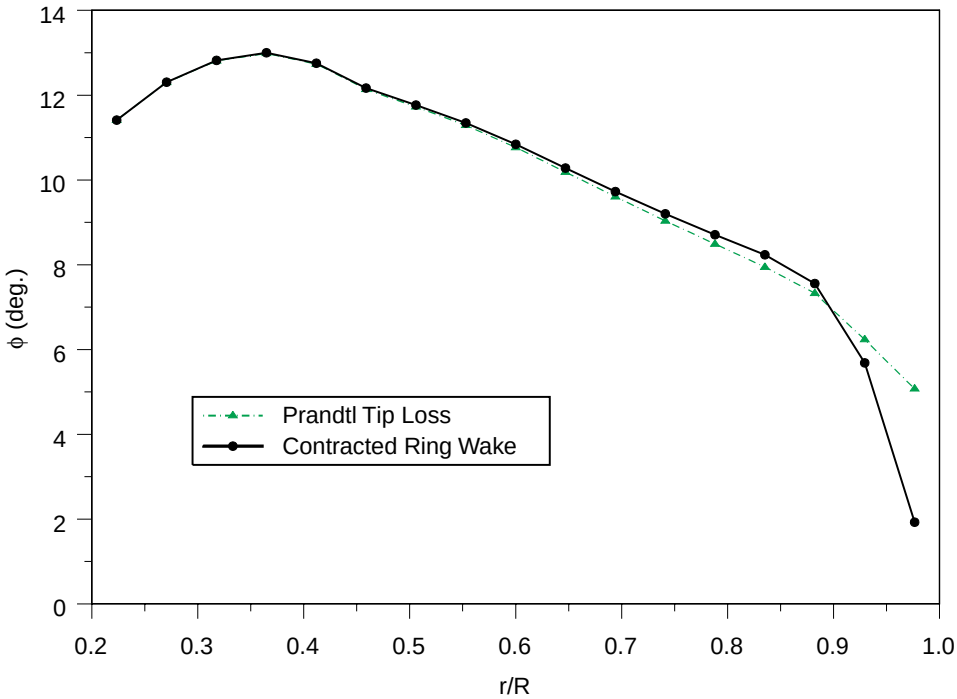

FIGURE 6.41 Predicted inflow angles using two different wake models.

Recall that this rotor was designed using the Prandtl tip loss model with a maximum section lift coefficient of 0.5 based on steady-state INS2d results. Although it does not reduce the utility of the results or introduce any global metric discrepancies, analysis with the contracted wake model indicates that outermost sections are operating significantly past the prescribed design maximum lift coefficient, the veracity of this prediction is supported by the OVERFLOW-D thrust distributions presented in Figure 6.32.

OVERFLOW-D predicts significantly higher sectional thrust near the tip, but even with the contracted wake model active, the rapid analysis tool cannot model sections operating beyond the steady-state limit imposed by construct of the two-dimensional section database. Above this limit, a heuristic stall model is implemented, capping the lift coefficient and increasing the section drag. In the case of rotor design, this successfully bounds the geometry within the steady-state region, but the best that the rapid analysis tool can do is indicate a potential problem in the tip region.

6.8 Modeling Effects on Design Configuration

6.8.1 Effect of Swirl Modeling on Design

The impact of the three different viscous swirl models on rotor design has been explored by developing three different configurations using the rapid analysis and design method described in Chapter 4. All the rotors are constrained to have four blades, a 5mm diameter hub and a 2.5cm total diameter and all use the classical Prandtl tip loss correction. This restricts the design variables to chord, incidence, and operating RPM. The 5mm Smoovy 15 second duration motor data is used for all cases. These conditions and constraints have been implemented to allow the as-built and tested four-blade rotor to be used as a point of reference. Following the design process, the rapid analysis

method with the contracted ring wake and angular momentum swirl model have been used as a common method of comparison.

The resulting chord and incidence distributions are provided in Figures 6.42 and 6.43. The largest differences are seen in the inboard half of each blade. The average wake deficit case is identical to the as-built and tested four-blade 2.5cm diameter rotor. The other two viscous swirl models predict much lower viscous swirl velocities and reduce any penalty associated with blockage due to high local solidity, this permits and makes desirable an increase in chord closer to the hub. With the reduced blockage and solidity penalty, the optimizer increases the total lifting area and increases the Reynolds number at the inboard stations, lowering the section drag coefficient.

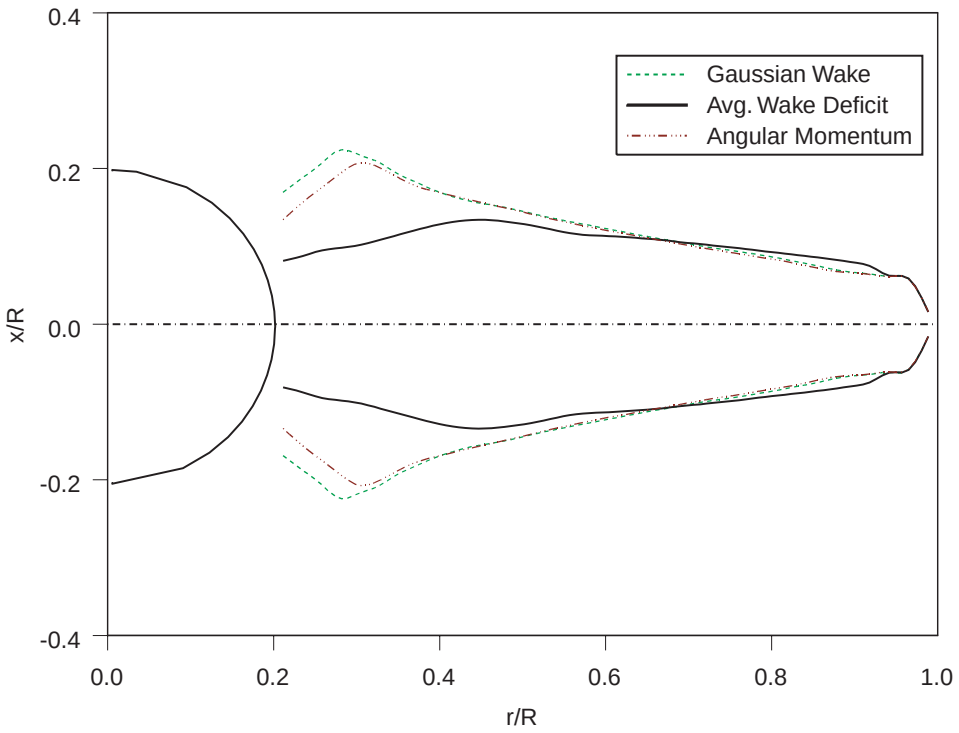

FIGURE 6.42 Blade planforms obtained by applying the rapid design tool with three different viscous swirl models in conjunction with the classical Prandtl tip loss correction.

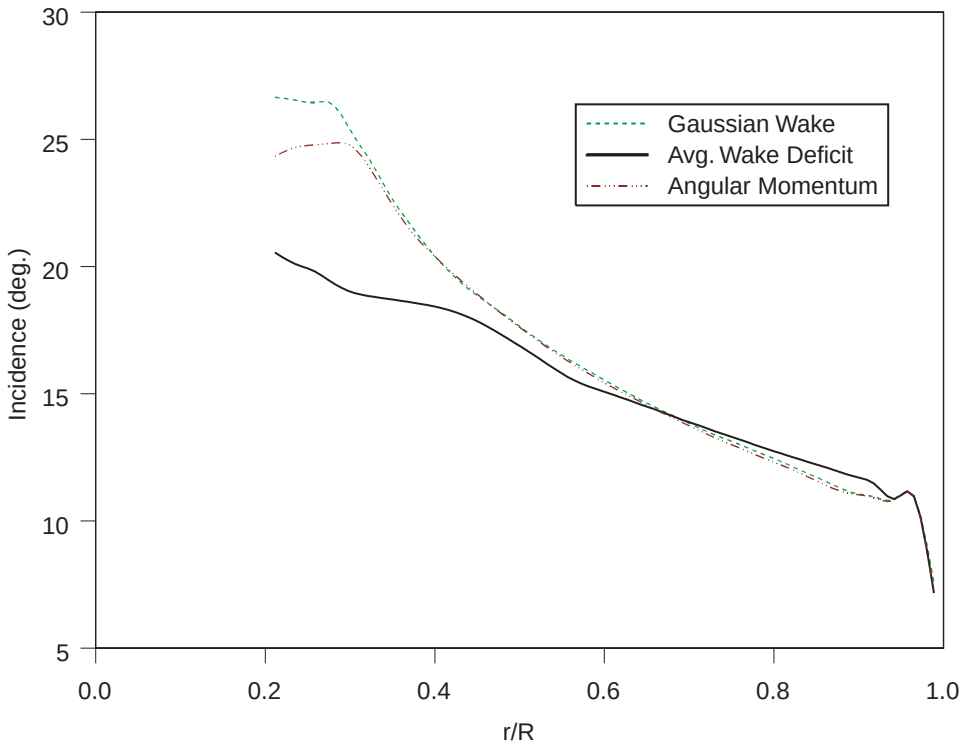

FIGURE 6.43 Blade incidence distributions obtained by applying the rapid design tool with three different viscous swirl models in conjunction with the classical Prandtl tip loss correction.

The inboard incidence variations result primarily from the effect of the viscous swirl velocity on the inflow angle, as seen previously in Figure 6.36. The reduced viscous swirl effects of the Gaussian wake and angular momentum model result in increased inflow angles as defined from the rotor plane. This necessitates increased incidence on the inboard stations relative to the average wake deficit model in order to maintain the similar lift coefficients.

The predicted thrust and power required for all three cases are shown in Figures 6.44 and 6.45. Results are shown for both a common analysis method, using the conservation of angular momentum swirl model and the contracted ring wake, and each individual design method. The operating points from the original design analyses exhibit nearly identical power requirements, but this is expected given a total spread of only 2500 RPM and the power constraint provided by the motor model. This small range in RPM relative to the average operating speed of 47,000 RPM has essentially no effect on the

output power of the motor. The higher operating speed of the average wake deficit design is due to the higher viscous swirl velocities predicted by this model. This both forces the rotor to operate faster and allows it to operate faster. The lower dynamic pressure, decreased blade area, and reduced Reynolds numbers at a given RPM all force a higher operating speed to maintain thrust, but these same factors mitigate the drag rise associated with higher speed, maintaining roughly the same power required as the other two designs.

The effect of these geometry variations appears negligible when all three cases are evaluated with a common version of the rapid analysis method incorporating the contracted ring wake and the angular momentum swirl model. The results for the three designs are bounded by the two lower lines in each figure. Common analysis indicates only very small variations in both thrust and power required. This is not surprising considering the similarities in the chord distribution and incidence distribution over the outer half of each blade. Of the three viscous swirl models, the angular momentum model is preferred for design simply because of its previously discussed superior performance in the analysis case and the desire to use a common model for both types of work.

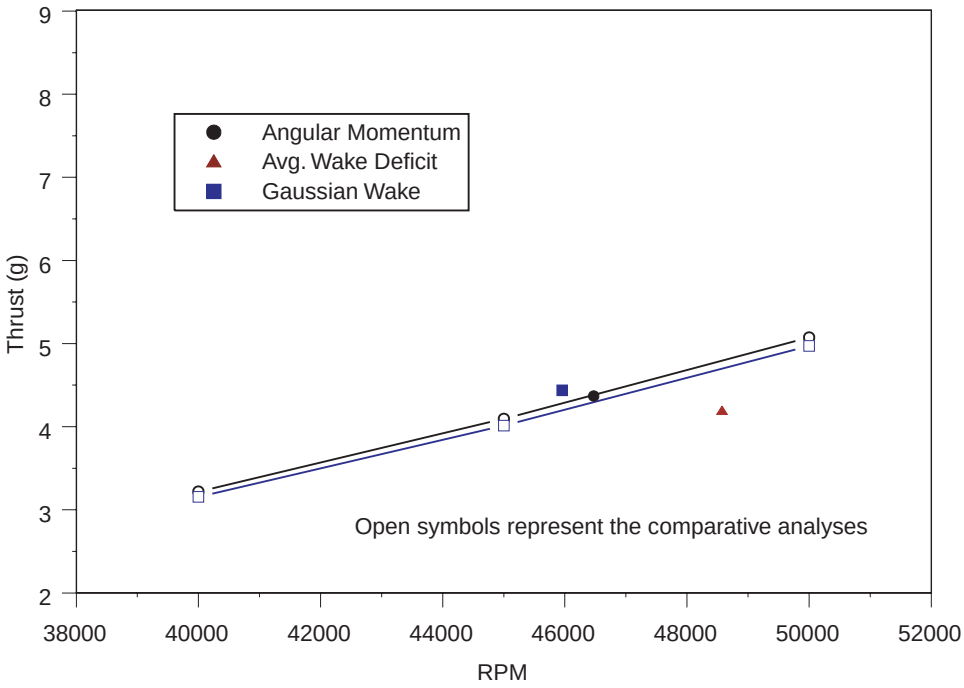

FIGURE 6.44 Predicted thrust for three different 2.5cm diameter rotor designs utilizing various viscous swirl models.

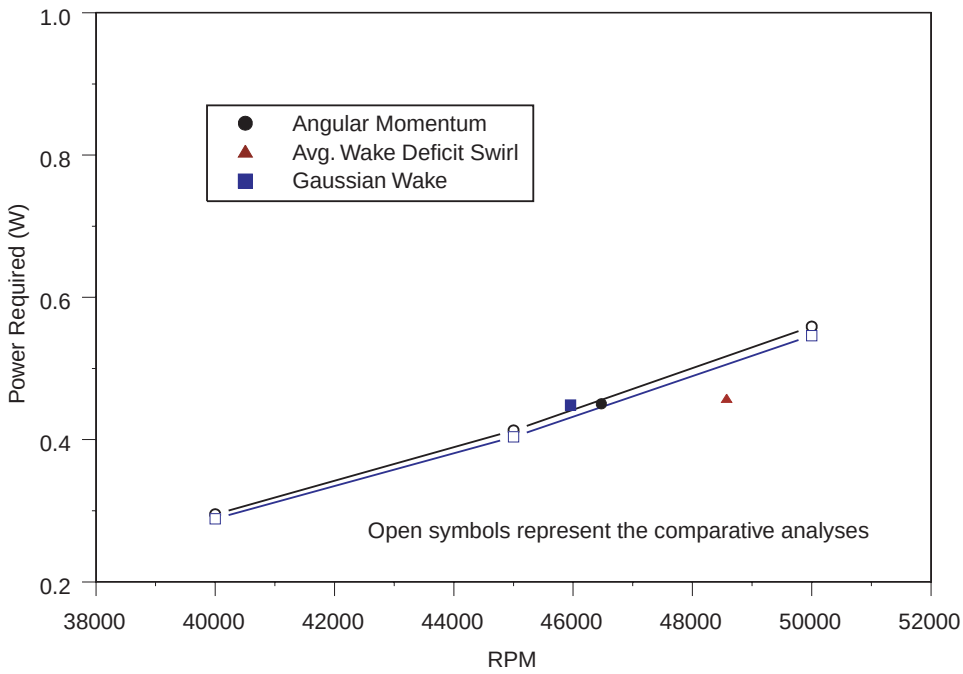

FIGURE 6.45 Predicted power required for three different 2.5cm diameter rotor designs utilizing various viscous swirl models.

6.8.2 Effect of Wake Modeling on Design

The choice of rotor wake model has a more significant impact on the rotor design output. While the effects of the viscous swirl models are localized near the root of the blade where lower rotational speeds minimize the impact on the performance, the wake models primarily effect the outermost regions of the rotor blade where the dynamic pressure is greatest. All analyses use the angular momentum swirl model and the same conditions and constraints used in the viscous swirl model study of the previous section. Three cases are considered:

- Design with the classical Prandtl tip loss correction.
- Design with the contracting ring wake model.
- A two-step process where the initial design is completed using Prandtl tip loss, followed by an secondary optimization of only the incidence distribution and RPM, using the contracting ring wake.

The design incorporating the contracted ring wake exhibits a dramatically different planform from all other cases and from rotor and propeller designs in general. The chord distributions from this design study are shown in Figure 6.46. The large growth in chord at the tip is most likely a by-product of wake modelling. In reality, the wake is continuously shed from the entire length of the blade based on the local circulation. Here it is modelled as discrete horizontal vortex rings representing the total vorticity of the blade. The ring wake model has no knowledge of the lift distribution of the blade, only the total thrust, but the ring wake has a strong effect on the local circulation via the modified tip loss correction factor.

This decoupling combined with the interaction between the first and second rings creates a region where a large growth in chord is beneficial. Qualitatively, there is some logic to this approach, as it attempts to maintain a uniform blade loading in the presence of a rapidly contracting wake, but its potential value from an aerodynamic standpoint has not been assessed. Due to practical considerations such as manufacturing and aero-structural twist this solution has been set aside.

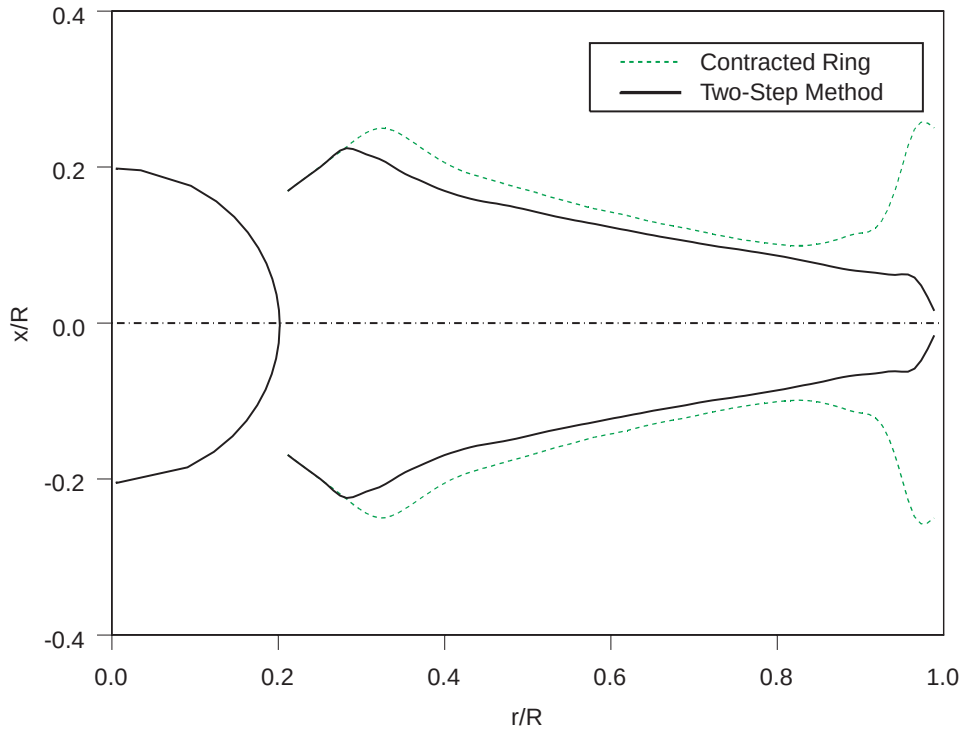

FIGURE 6.46 Blade planforms obtained by applying the rapid design tool with two different wake models in conjunction with the angular momentum swirl correction.

The design using the Prandtl tip loss correction has already been described in the previous section on viscous swirl model effects. As noted earlier, a primary issue with this wake model is an increased risk of tip stall. The primary observed benefit of the contracted ring wake model is an increase in tip loading relative to the predicted loading using the Prandtl tip loss correction. For the design of rotors, this results in a reduction in the tip incidence angle and consequently a reduction in the risk of tip stall. The two-step method works around the planform issues of a pure contracted ring design, but still benefits from the change in tip loading. The reduction in incidence is clearly visible in Figure 6.47.

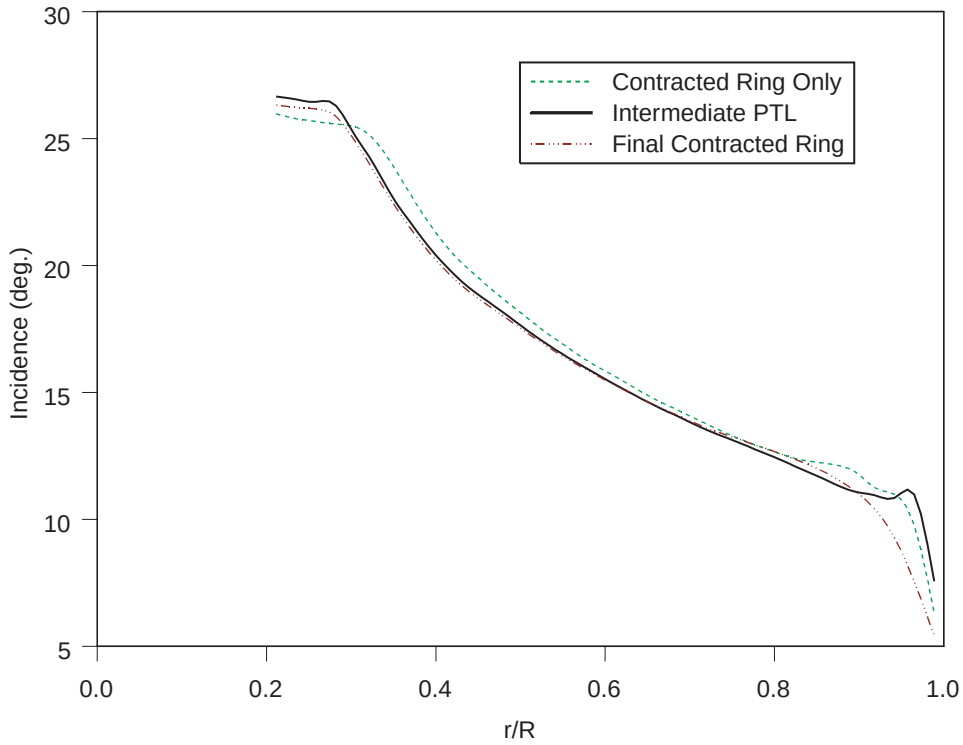

FIGURE 6.47 Blade incidence distributions obtained by applying the rapid design tool with two different wake models in conjunction with the angular momentum swirl correction.

The predicted lift coefficient distributions, all utilizing the contracted ring wake and angular momentum swirl model, are presented in Figure 6.48. For all three cases the design maximum lift coefficient was set at 0.600. This was increased to 0.627 for the analysis to demonstrate the variations present between designs and the value of the contracted ring model in mitigating risk. The Prandtl tip loss design is predicted to be at risk of stall over the outer ten percent of the blade. This is qualitatively supported by the OVERFLOW-D results for the four blade 2.5cm diameter rotor. The OVERFLOW-D results do not indicate complete stall but do show a rapid increase in blade loading near the tip and primary vortex shedding slightly inboard of the tip, exacerbating blade-vortex interactions.

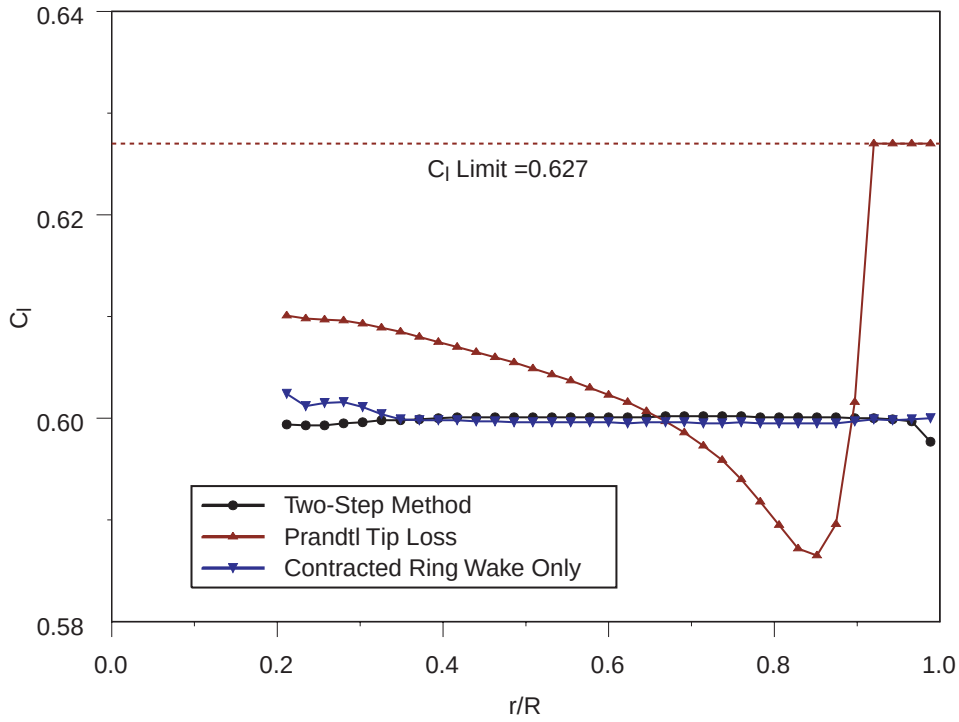

FIGURE 6.48 Lift coefficient distributions predicted by the rapid analysis tool for three different rotor designs emphasizing the effect different wake models.

6.9 Three-Dimensional Boundary Layer Effects

Sectional torque is consistently underestimated by the rapid analysis method compared with OVERFLOW-D, even where the thrust is over predicted. Given the high dependence of the sectional torque on the sectional lift, this discrepancy can only be attributed to errors in the section properties when compared to those predicted by the three-dimensional Navier-Stokes results.

This problem poses an impediment to using blade-element methods for detailed design at ultra-low Reynolds numbers. An investigation of section pressure and skin friction distributions indicate large scale variations are most likely due to three-dimensional effects, primarily Coriolis and centripetal accelerations, on the boundary layer development. The largest consequences are an increase in skin friction and a reduction

in boundary layer thickness. The reduction in boundary layer thickness increases the effectiveness of the airfoil geometry, particularly aft where the cumulative effects on boundary layer development are greatest. For the optimal airfoil designs of Chapter 3, this means that the kink in the camberline at 80% chord will have much more effect than intended.

Detailed comparisons of blade sectional aerodynamics are presented for a single chord-wise station at r/R = 0.48 and 50,760 RPM. This case was selected to minimize tip effects. Reynolds number variations are minor since the rotational velocity term dominates any induced velocities. Cross-flow is minimal up to separation, which occurs aft of 90% chord. The blade-element method predicts that at this station C_l = 0.45 corresponding to α=2.5 degrees. Pressure and skin friction distributions for this section calculated with INS2d and those taken from the OVERFLOW-D analysis are presented in Figures 6.49 and 6.50. The total pressure thrust agrees within 5%, but the pressure distributions are dramatically different. The OVERFLOW-D section is operating at a lower local angle of attack, but achieves the same pressure thrust with a large increase in the aft loading. The section torque due to skin friction is 30% higher in the three-dimensional case.

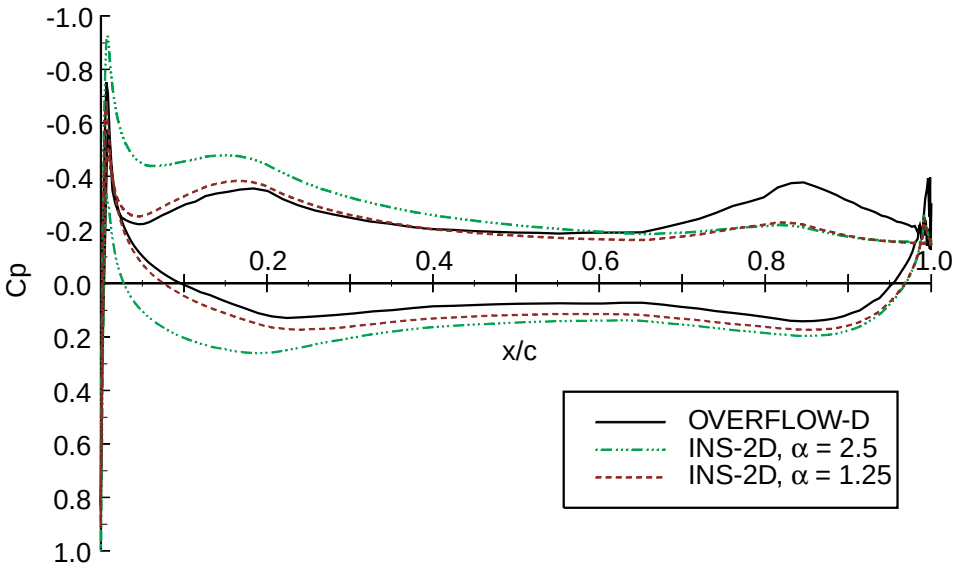

FIGURE 6.49 Chordline pressure distribution at r/R=0.48, 50k RPM.

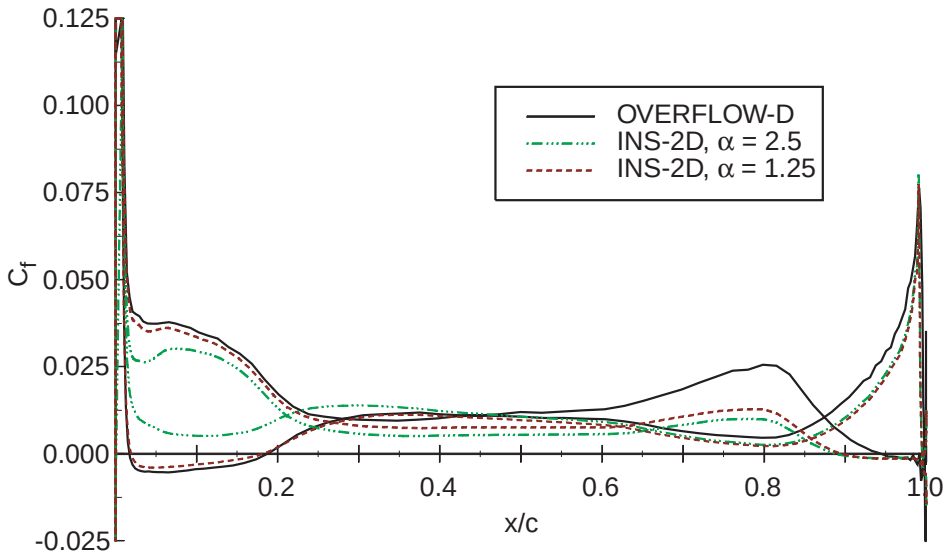

FIGURE 6.50 Distribution of the chord-wise component of skin friction at r/R=0.48, 50k RPM.

Also included in these plots is an INS2d calculation at α=1.25 degrees. At this angle of attack the stagnation points in the two-dimensional and three-dimensional cases agree. The pressure and skin friction distributions also agree over the first 50% of the airfoil. This is the effective angle of attack in the OVERFLOW-D result. This close agreement over the front half of the airfoil also serves as a mild validation of the INS2d calculations. The three-dimensional case, however, has 20% greater pressure thrust and 33% higher viscous torque. These discrepancies in sectional details are representative of the entire span, but in spite of this, the predicted global thrust and power agreed on average within 5% with maximum errors under 10%.

These large variations in section performance necessitate some way of modeling the rotational effects in order for a blade-element method to be viable beyond preliminary design. Even then, the method may not be useful if, as in this case, the section properties in 3-D mandate a redesign of the airfoil. This particular airfoil is very sensitive to off-design conditions, so it is possible that a different section may exhibit less variation, particularly in the pressure distribution.

Unfortunately there does not appear to be a straight-forward solution to this problem. In the case of the Coriolis force, the rotational accelerations are dependent on the local velocity in the fluid and cannot be applied as a simple body force in a two-dimensional calculation. For analysis, heuristic corrections may be applied to the airfoil polar database, but this would not be possible for design. An approximate method such as a three-dimensional integral boundary layer formulation might be used to correct the section properties, but such a solution may be too complicated and unreliable for practical use.

Chapter 7

Micro-Rotorcraft Prototypes

7.1 Introduction

The penultimate goal of the preceding chapters is to develop the understanding and the facility to achieve controlled powered flight at unprecedented small physical scales. Moving further towards that end, several rotorcraft of varying physical size and capability have been designed, fabricated, and tested. The goal of this closing chapter is to introduce the vehicles that have been developed in conjunction with this body of research and affirm the practical value of the preceding work. These prototypes have demonstrated both the high potential for such vehicles, and some of the limitations imposed by current technology and the relevant aerodynamics.

7.2 The 15g Prototype

The smallest rotorcraft developed using the rapid rotor analysis and design method is the 15g prototype pictured in Figure 7.1. Like all of the vehicles associated with this research program, it is referenced simply by its target gross take-off mass. This vehicle incorporates four of the four-blade 2.5cm diameter epoxy SDM rotors described in

Chapter 6. Four Myonic 5mm Smoovy motors power the rotors, controlled on-board by four Philips TDA 5145 SMT closed-loop motor controller integrated circuits. The rotor ducts and core structure of the vehicle are an integrated SDM manufactured part and is described in detail by Cheng [32]. The ducts have not been rigorously designed for aerodynamic benefit and serve primarily as protection for the fragile rotors. The applicability of aerodynamic shrouds to micro-rotors is one area of future research that could offer a significant performance benefit, but at this time the issues involved in developing an effective design are not fully understood. The complete mass allocation for this vehicle is provided in Table 7.1.

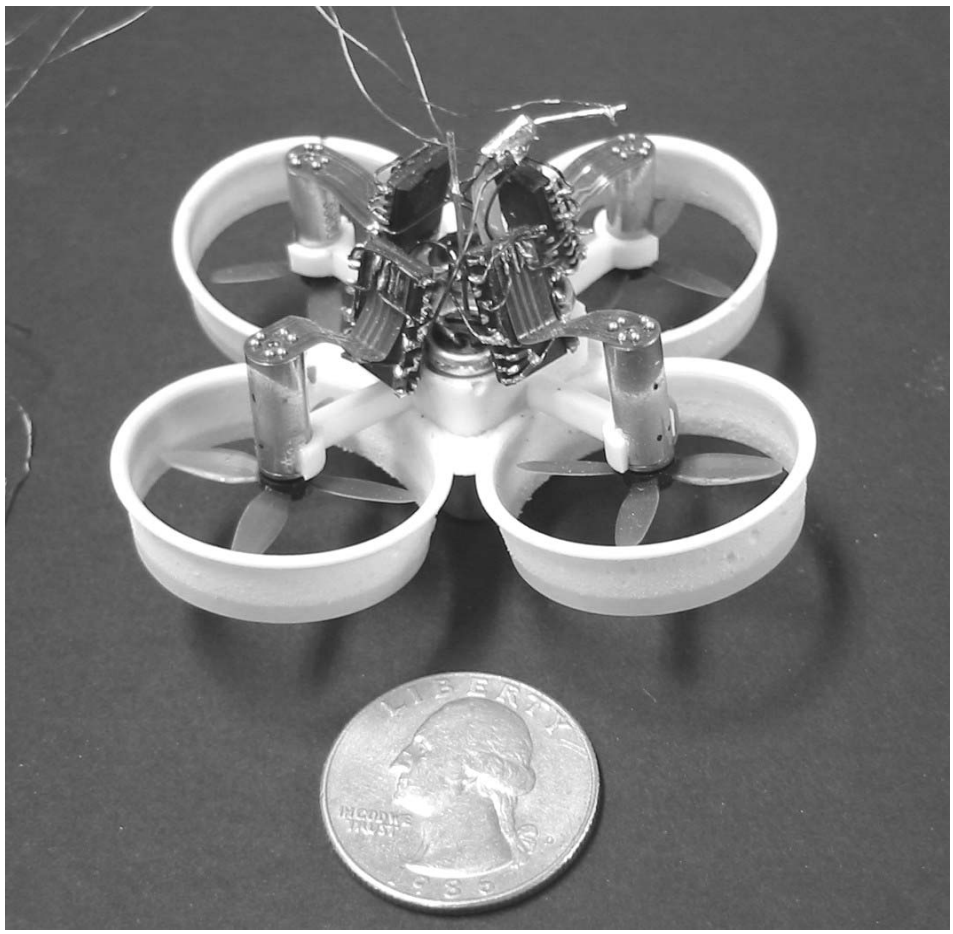

FIGURE 7.1 The 15g prototype electric rotorcraft.

This vehicle does not currently incorporate active control, but the planned control scheme is common to all of the prototypes. No tail rotor is used. Instead, torque balance in the yaw axis is maintained by having one pair of opposing rotors spinning clock-wise and the other opposing pair rotating counter-clockwise. Pitch and roll are controlled by using pairs of rotors along a perpendicular to the desired axis, accelerating one rotor while decelerating the other to generate the necessary moment. The stability and control of these vehicles, including the development of linearized dynamic models has been explored by Kroo [43].

TABLE 7.1 Mass allocation for the 15g prototype electric rotorcraft.

		Measured Mass
Super Capacitors:		6.4g
Motor/Gearing:	4 RMB 5mm Smoovy Motors	6.0g
Rotors:	4 Epoxy SDM Rotors	0.3g
Electronics:	4 TDA 5154 SMD Controllers	2.8g
Structure:		1.5g
Total mass of primary flight systems:		**17.1g**

Electrical energy storage is a key issue for micro-rotorcraft. For a given chemistry, battery cell capacity is closely coupled to the volume, and volume is quickly lost as the linear dimensions are reduced. The mass of casings and other supporting elements relative to the mass of the actual electro-chemical elements of a battery cell also increases as cell size is diminished. For this particular vehicle, current commercially available batteries do not offer a feasible solution. Instead, the use of four one-Farad super capacitors were investigated, but the capacitors were incapable of maintaining sufficient voltage over the course of their discharge cycle [32]. The absence of a suitable on-board power source has limited testing of this prototype to externally powered hover tests with reduced degrees of dynamic freedom.

The fully assembled prototype, including super capacitors, has demonstrated the ability to generate sufficient thrust for hover while constrained to only yaw and vertical translation. The rotorcraft is powered by an external power supply and constrained by a vertical rod via a guide tube passing through the center of the vehicle. The voltage

required for hover is 16V at a current draw of one ampere. Sixteen Watts is considerably more power than predicted for the rotors alone. Each rotor has been shown to require roughly 0.6W of power when generating the necessary thrust, a value that has been verified by OVERFLOW-D computations and experiment. This tremendous loss of power between source and rotor represents a practical challenge for developing micro-rotorcraft and likely places this scale of vehicle somewhat beyond the current state-of-the-art, but efficient electrical power management and motor control are beyond the scope of this thesis.

7.3 The 65g Prototype

The intermediate size rotorcraft, with a target gross mass of 65g, has been created primarily as a testbed for electronic systems integration, and implementation of augmented stability and control. Two versions have been built, one using remote radio control without any stability augmentation, and the second incorporating on-board solid state gyroscopes, a microprocessor, and small transceiver. Photographs of each version provided in Figures 7.2 and 7.3. The mass allocation for the two versions is provided in Table 7.2.

TABLE 7.2 Mass allocation for the 65g prototype electric rotorcraft.

		R/C Version	M.Proc. Version
Batteries:	7 Cell 110mAh NiCad Pack	26.2g	26.2g
Motor/Gearing:	4 Mabuchi Watt-Age Motors	22.6g	22.6g
Rotors:	4 Polyester SDM Rotors	2.4g	2.4g
Radio/Control Electronics:		6.7g	25.0g
Structure:		3.4g	Integral
Total mass of primary flight systems:		**61.3g**	**76.2g**

FIGURE 7.2 The 65g prototype electric rotorcraft, remote control version.

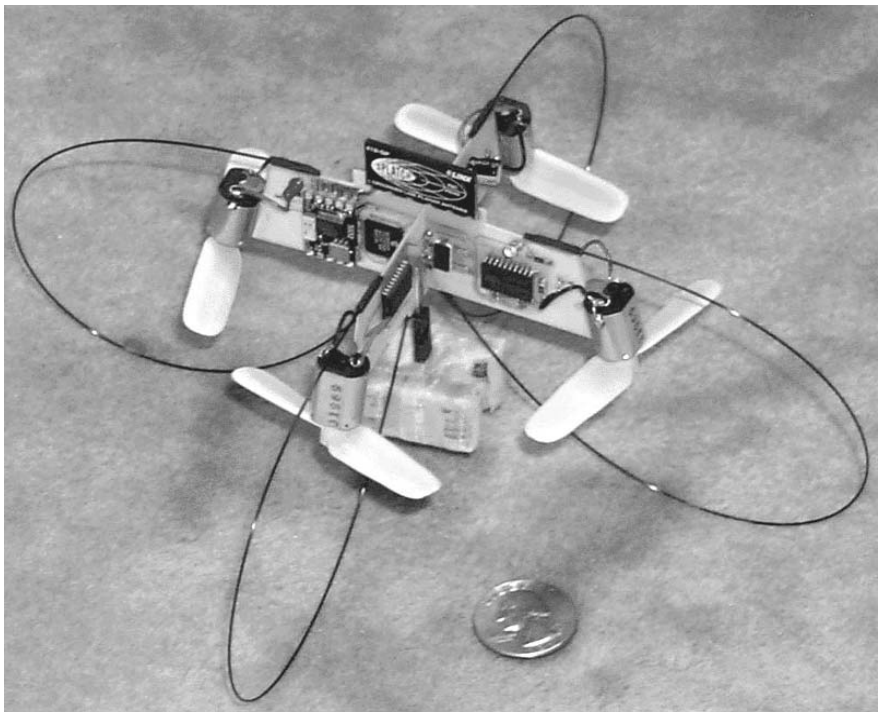

FIGURE 7.3 The 65g prototype electric rotorcraft, microprocessor version.

Since this prototype was not primarily an aerodynamics test bed and since minimizing total development time was a priority due to contractual constraints, the rotor design is based entirely on a commercially available flying model airplane propeller. As such, no work has been done to optimize the performance, but it is likely that considerable improvement could be achieved by using the rapid design method to carry out a rigorous rotor design for this target application. This two-blade 6.35cm diameter rotor has a constant chord planform with a section thickness to chord ratio of roughly 7%. Comparison with the 2.5cm diameter rotor designs indicate that considerable gains could be had by reducing the blade section thickness and increasing the rotor solidity.

Both prototypes have undergone tethered testing powered on-board by an 8.4V package of seven 110mAh NiCd cells. This is sufficient power for up to one minute of operation, powering four brushed direct-current motors. The power required for hover at 65g has been measured at 15 Watts. These motors were obtained at a local hobby shop and to date only their manufacturer, Mabuchi, has been identified. They are sold under the Watt-Age brand, by Global Hobby Distributors [44] as replacement motors for their smallest line of flying model aircraft.

The short flight duration combined with hardware problems with the stability augmentation system limited the free flight potential for these models, but they both did demonstrate sufficient thrust for hover under their own power and demonstrated more than adequate control authority in all axes. The remote controlled version saw further use as a constrained flight testbed for off-board stability and control via on-board infrared optical emitters and ground-based cameras [45].

7.4 The 150g Prototype

The largest of the three vehicles is representative of a scale that is attainable within the current state of the art in terms of electronics, energy storage, and high efficiency electric motors. The 150g prototype is pictured in Figure 7.4. The mass allocation is provided in Table 7.3. Over the four year course of this research, this vehicle has gone from being truly unique to being one of a number of small remote controlled free-flying rotorcraft. The majority are commercially available and marketed to hobbyists. The key characteristic that still sets this vehicle apart is the capability for useful work.

FIGURE 7.4 The 150g prototype electric rotorcraft.

TABLE 7.3 Mass allocation and payload estimate for the 150g prototype electric rotorcraft.

			Initial Estimate	Measured
Batteries:	4 Tadiran 430mAh Cells		46g	49g
Motor/Gearing:	4 Astroflight Firefly Motors		52g	50g
Rotors:	4 Carbon 10” dia. Rotors		15g	19g
Radio/Control Electronics:	In-house System Design		25g	25g
Structure:			10g	10g
Total mass of primary flight systems:			**148g**	**153g**
WES-Technik:				
Total thrust at 8.3V motor input:		49.5g X 4 =	198g	
Total motor power required:		12.9 W		
Stanford 10” Rotor:				
Total thrust at 8.5V motor input:		50.1g X 4 =		200g
Total motor power required:		10.8 W		
16% reduction in power required				
Available payload mass:			**24g**	**21g**
(Maintains max. thrust to weight ratio of 1.15)				

Improved aerodynamic efficiency and novel structural and systems integration translates to increased payload and/or efficiency relative to the commercially available examples. Payload mass can effectively be traded with battery mass, trading endurance for additional payload. The end result is a vehicle capable of carrying up to 20 grams of payload with an endurance of five to 20 minutes depending on the battery size and chemistry. The integrated sensors, microprocessor, and transceiver currently allow augmented stability control with a near term goal of achieving autonomous flight. It also creates a flexible system, reprogrammable for varying conditions and missions.

To date, the vehicle has been successfully remotely piloted with augmented stability, both tethered and in free flight. Power is currently supplied by a 12.0 Volt Tadiran [46] lithium / manganese-dioxide power pack consisting of four cells in series of either 430mAh or 780mAh capacity. Thrust is provided by four of the two-blade 10cm diameter rotors and the Astroflight Firefly motor with 16:1 gearing, both described in Chapter 6. High frequency PWM motor controllers provide the connection between the receiver or microprocessor and each powerplant.

At hover, with a mass of 153g, the entire system consumes ten Watts of power. The predicted and experimental results indicate 1.3 Watts of power required per rotor at this thrust level indicating a total electro-mechanical efficiency of roughly 52%. This is a tremendous increase over the performance of the 15g prototype and once again demonstrates a key issue with developing electrically powered micro-rotorcraft: the rapid degradation of electro-mechanical efficiencies at reduced physical scales.

7.5 Insights Gained, Limiting Technologies, and the Potential for Future Development

With the goal of self-powered autonomous micro rotorcraft, the three prototypes span the design space from 15g (considered infeasible with current technology), up to 65g (that with further design refinement would represent the state-of the-art), finally increasing to 150g. The largest is a design that poses challenges in system integration and automated flight control, but otherwise represents what can be accomplished with current consumer level technology and hardware.

Experience with these three prototypes indicates that the cost of reduced scale on the overall power requirements is severe. A summary of key sizing parameters and hover power requirements for these rotorcraft is provided in Table 7.4. The data without battery mass is included to indicate the lower bound on hover power required.

TABLE 7.4 Summary of physical and performance data for three prototype electric rotorcraft.

As built and designed w/ batteries/caps

Vehicle Mass (g)	P Req. @ Hover (W)	Footprint (sq. cm)	Rotor Dia. (cm)	Disk Loading (g/cm^2)	Batt. Mass (g)
17	16	25	2.5	0.865	6
65	15	160	6.35	0.513	26
153	10	1300	25.4	0.075	49

w/o batteries/caps, based on static rotor test data (experimental motor/rotor power with additional 70% efficiency factor)

Vehicle Mass (g)	P Req. @ Hover (W)	Footprint (sq. cm)	Rotor Dia. (cm)	Disk Loading (g/cm^2)
11	12	25	2.5	0.560
39	8	160	6.35	0.307
104	6.51	1300	25.4	0.051

Three areas of technology play the largest roles in determining the feasibility of a design: aerodynamic efficiency, electro-mechanical efficiency, and the energy storage density at the necessary current levels to drive the rotors.

7.5.1 Emerging Battery Technologies

A detailed discussion of battery technologies is beyond the scope of this work but comparisons of several widely available chemistries and cell sizes is provided by Kroo [43]. As a footnote to their discussion of power storage issues for micro-air-vehicles, lithium polymer cells have now become widely available in the consumer market. Due to the current physical dimensions of these cells, this development does not immediately affect the 15g vehicle, but for the larger vehicles, this technology offers significant performance gains over the previously considered chemistries. At approximately 50mAh/g, current consumer lithium polymer cells offer a higher energy density than the 43mAh/g 780mAh Tadiran lithium / manganese-dioxide cells. In addition to the benefits of increased energy density, further gains are possible due to the increased base voltage of 3.7V and a 4C to 6C discharge rate, versus 3.0V and a 3C rate for the Tadiran cells. This results in a power density of 185mWh for the lithium polymer cells versus 129mWh for the Tadiran cells, a 43% increase. The term ‘C-rate’ refers to the one hour discharge current, equal to the cells mAh capacity.

Using the 150g prototype as an example, conversion to lithium polymer technology would permit a reduction from four to three cells resulting in roughly a 50% increase in endurance at the same total battery mass. There is a one volt drop in the no-load voltage, but the operational voltages of the two cell types would be much closer due to the additional headroom in discharge rate and associated reduction in internal losses provided by the lithium polymer cells. The higher discharge rate also permits application of these cells to higher current draw situations such as the 65g prototype, where previously NiCd and NiMh chemistries, with significantly lower energy densities, provided the best solution.

7.5.2 Electro-mechanical Efficiency

While the primary focus of this work has been the aerodynamics of ultra-low Reynolds number flight, the design and development of a complete vehicle cannot be undertaken without consideration of the electro-mechanical efficiencies of the supporting systems. This issue, in conjunction with energy storage issues, poses a significant impediment to the development of extreme micro-rotorcraft such as the 15g prototype. This area encompasses not only the motors and speed controllers, but additional equipment required for voltage conversion and regulation, communications, and control. As discussed here, this does not include the battery system, which is considered separately, or the power required for mission-based payload such as cameras or other sensors.

The 150g vehicle exhibits an overall power efficiency of 52%, this is considerably lower than the experimentally determined motor efficiency of approximately 65% to 70%, but in addition to the motors, this system includes a radio receiver, four motor controllers, and four solid-state piezo gyros. Although the 15g prototype's rotor system requires less than half of the power of the 150g vehicle's rotors, the total system power required for hover increases by 60%. This results in a total system efficiency of only 15%. This is once again lower than the experimentally determined electro-mechanical efficiency for the motor, in this case combined with the closed-loop controller, of approximately 17%, but there are no additional system components in this case. The 40% drop in power plant

efficiency is devastating to the concept's feasibility, but it also appears to suffer an additional 2% loss, most likely due to the addition of the four rotor-ducts without consideration being given to the rotor design. As was seen in Chapter 6, this amount of variation could also be attributable to geometric variations in the four epoxy SDM rotors. In both cases, the motors and controllers are representative of the state-of-the-art in small, high speed electric motors.

For the 5mm Smoovy motors, a significant portion, perhaps as much as half, of the losses are due to the high speed switching controller required for brushless motors. It might be possible to improve the overall efficiency and achieve a significant mass reduction by developing a suitable brushed or coreless direct-current motor, but this is only speculative and beyond the focus of this work.

It is clear that significant gains must be made in the supporting technologies of electronics, electric motors, and energy storage before even optimal aerodynamic design would permit success at the scale of the 15g prototype. Storage-based electric power has been the focus of this effort, but there are other approaches that have not been explored. The absence of concepts such as beamed energy, combustion, and other novel concepts is not meant to imply that they should be discounted; the current focus has been chosen because the existing technology was thought to be scalable without significant development effort.

7.5.3 Rotor Aerodynamic Efficiency

Given the large non-aerodynamic handicaps placed on micro-rotorcraft, maximizing rotor performance should be considered an even more critical issue than in larger applications. Inefficiencies in the rotor design are effectively multiplied through the inefficiencies of the motors, controllers, and energy storage system. As an example, any increase in rotor power is seen by the batteries as doubled when passed through a 50% efficient system such as the 150g prototype. Unfortunately, rotors at small scales face a fundamental reduction in performance relative to their large scale brethren.

Basic rotor theory has been applied to develop a better understanding of the performance costs associated with micro-rotors and key factors responsible for these penalties. From this effort, reasonable expectations on performance can be established. Rotor figure of merit (*M*) is a common criteria for the comparison of rotor designs and represents the hovering efficiency of a given rotor relative to an idealized reference value. Figure of merit is defined as:

$$M = \frac{P_{ideal}}{P_{actual}} \tag{7.1}$$

The ideal power is taken as the induced power required for hover from momentum theory, absent any viscous effects. The ideal power can then be expressed as:

$$P_{ideal} = T u_{ideal} = \sqrt{\frac{T^3}{2\rho A}} \tag{7.2}$$

Assuming than an optimized design will have an induced power close to the minimum, the largest factor in the variation of P_{actual} from P_{ideal} is the profile power, $P_{profile}$. This replaces Eqn.7.1 with the approximate form:

$$M \approx \frac{P_{ideal}}{P_{ideal} + P_{profile}} \tag{7.3}$$

This idealized form makes no attempt to account for hub effects on induced power. Under the limitations of a constant chord blade, small inflow angles, and the assumption that the section drag and lift coefficients remain essentially constant, the profile power may be approximated by:

$$P_{profile} \approx \frac{\rho\sigma C_d (\omega R)^3 A}{8} \tag{7.4}$$

Similarly the thrust may be approximated by:

$$T \approx \frac{\rho \sigma \pi C_l (\omega R)^2 R^2}{6} \tag{7.5}$$

For an optimal rotor designed for hover, Gessow and Myers [47] indicate that the constant C_l approximation is reasonable, as the mean lift coefficient across the blade should be close to the maximum lift coefficient in order to maximize the figure of merit. The assumptions made here neglect variations in C_{lmax} and C_d that may occur with Reynolds number variations across the blade, but still capture average Reynolds number effects when comparing different rotors.

Finally, upon substitution of Eqn.7.2, Eqn.7.4, and Eqn.7.5 into Eqn.7.3, the approximate expression for the rotor figure of merit becomes:

$$M \approx \frac{1}{1 + \left(\frac{3}{2} \frac{C_d}{C_l} \sqrt{\frac{3}{\sigma C_l}} \right)} \tag{7.6}$$

Equivalently, this may be expressed in terms of the thrust coefficient, C_T:

$$M \approx \frac{1}{1 + \left(\frac{3}{2} \frac{C_d}{C_l} \sqrt{\frac{1}{2 C_T}} \right)} \tag{7.7}$$

The figure of merit can be related to the thrust specific power required for hover and the disk loading by manipulating Eqn.7.1 and Eqn.7.2 to yield:

$$\frac{P_{req}}{T} = \frac{1}{M} \sqrt{\frac{1}{2\rho} \left(\frac{T}{A} \right)} \tag{7.8}$$

From the discussions of Chapter 3, the expected maximum steady state lift coefficients at the Reynolds numbers seen by centimeter scale rotors are approximately 0.6 to 0.7 over

most of the blade. These values are significantly lower than reasonable hover lift coefficients of larger rotors, ranging from 0.9 to 1.2. As seen in Eqn.7.6, this only serves to increase the impact of the dramatic reduction in the lift to drag ratio that occurs at ultra-low Reynolds numbers.

Some relief is provided by an increase in solidity (σ), but the solidity of a single rotor has an upper bound of 1.0 and in reality will be bounded at a considerably lower value due to structural considerations, the wetted area increase at a fixed radius or the increase in disk loading at reduced radius, powerplant limitations, and other factors.

Comparison with typical helicopter operating conditions reveals that there must be a fundamental reduction in the figure of merit for micro-rotors operating at ultra-low Reynolds numbers. Given a reduction in section lift to drag ratio by conservatively a factor of five, and a reduction in maximum lift coefficient of roughly 30%, performance will undoubtedly suffer. Even if the solidity of the micro-rotor was increased from typical helicopter values, ranging from 6% to 10%, to the maximum, but unobtainable, value of 100%, there would still be approximately a 9% reduction in figure of merit. Once again, since it is highly coupled to other parameters and constraints, reasonable upper limits on solidity will be considerably lower than 1.0, resulting in a much larger performance penalty.

In order to establish rational expectations for micro-rotor performance relative to typical helicopter rotor performance, Eqn.7.6 has been applied at two different section lift coefficients, representative of operation at the two different physical scales, across a range of rotor solidity and section lift to drag ratios. Figure 7.5 displays the results at C_l=1.0, reasonable for a large-scale rotor in hover while Figure 7.6 displays the results at C_l=0.6, representative of centimeter-scale rotors.

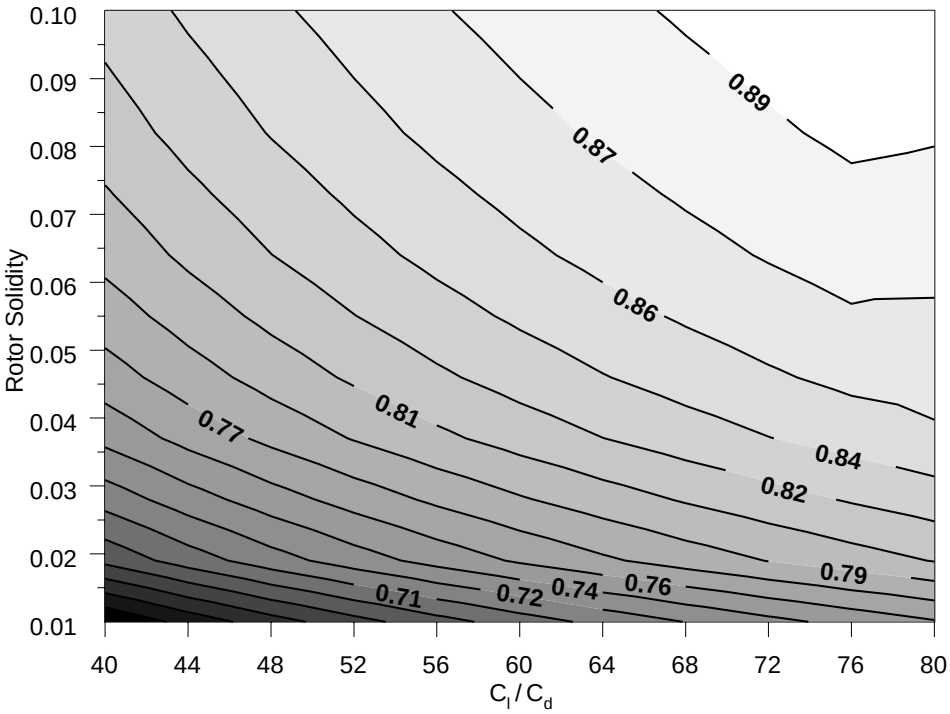

FIGURE 7.5 Figure of merit as a function of mean section performance and solidity at C_l=1.0.

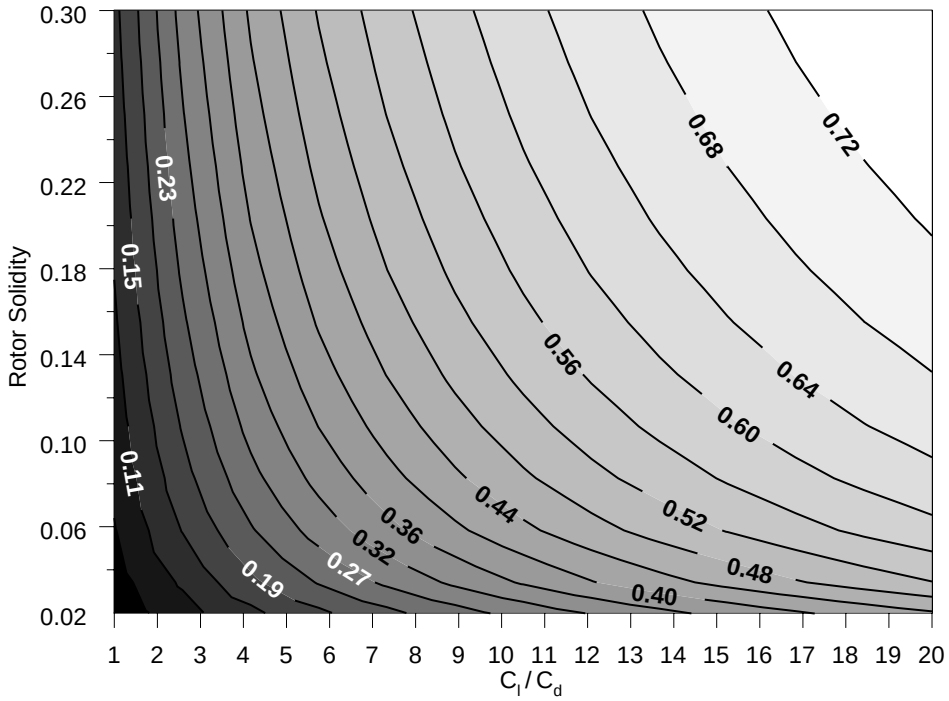

FIGURE 7.6 Figure of merit as a function of mean section performance and solidity at C_l=0.6.

For large scale helicopters, Figure 7.5 predicts figure of merit values between 0.80 and 0.87. This agrees reasonably well with the actual values for many full-scale helicopters. The results are seen to be relatively insensitive to airfoil performance. For micro-rotors, Figure 7.6 paints a very different picture. In this regime the figure of merit is highly sensitive to section performance. The two micro-rotor designs considered here, the four-blade 2.5cm diameter rotor and the and the five-blade 2.2cm diameter rotor, have comparatively high solidity values of 24% and 35% respectively, but the upper bound on figure of merit for these rotors appears to be between 0.60 and 0.65 based on the predicted optimal section performance of Chapter 3. Under the best case scenario then, centimeter scale micro-rotors can expect a 20% to 30% reduction in figure of merit relative to large scale helicopter rotors.

With the upper limits of performance reasonably well defined, there is still room for improvement and motivation for research and design for centimeter scale rotors. The rotor design examples described here still fall short of this target. The computed and measured values of the figure of merit for the four-blade 2.5cm diameter rotor are provided in Figure 7.7 as a function of the rotor RPM. The rapid analysis method provides the most optimistic estimate of approximately M=0.55, the measured values for this single sample are lower, at approximately M=0.39 for the higher RPM of primary interest.

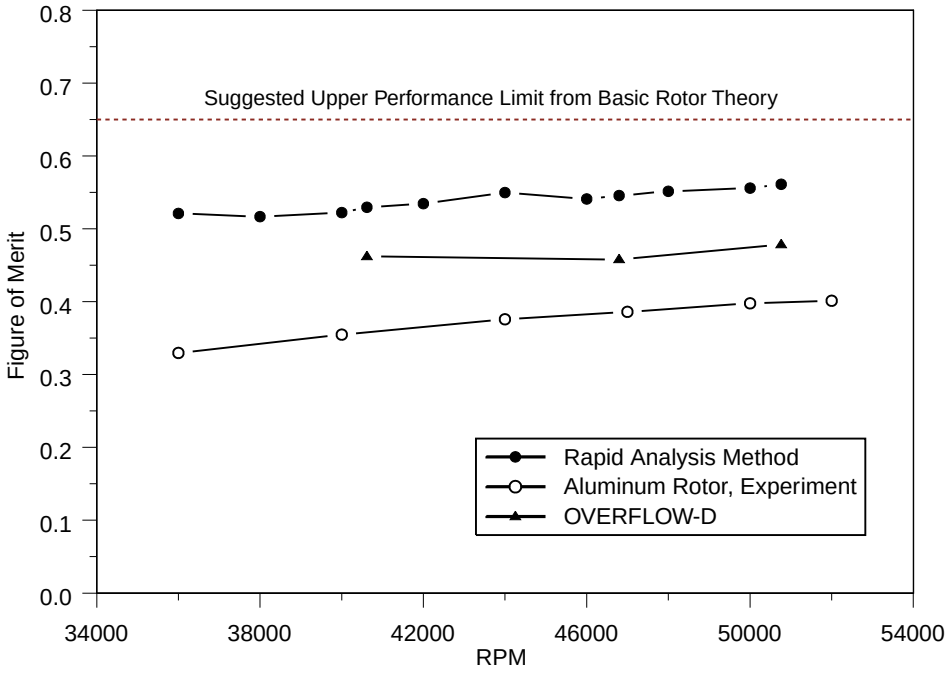

FIGURE 7.7 Figure of merit for the four-blade 2.5cm diameter rotor.

The potential benefits of improved performance on the power required for hover can be demonstrated by utilizing Eqn.7.8. Improving the measured performance from M=0.39 to M=0.6 at a vehicle mass of 17.1g would represent a decrease in the thrust specific power from 15.1W/N to 9.8W/N. Once translated through the system efficiency of 15%, this represents a 35% drop, or almost six Watts, in the system power required. It is clear that improving upon the measured aerodynamic performance of these rotors would partially mitigate the system efficiency issues and move further toward achieving powered flight at the smallest physical scales.

Chapter 8

Conclusions and Recommendations

Achieving powered flight at ultra-low Reynolds numbers and true micro-aircraft scales requires a greater understanding of aerodynamics in a flow regime that, until now, has only been minimally explored. The goals of this research have been to broadly explore the aerodynamics of this environment and then to use the insight gained to analyze, design, manufacture, and test micro-rotors of an unprecedented small physical scale. In the culmination of this effort, micro-rotorcraft have been proposed and developed that represent both the current state-of-the art and what lies just beyond it.

8.1 Summary of Results and Contributions in the Area of Two-Dimensional Aerodynamics

Airfoils in this regime face unique operating conditions and the two-dimensional computational study represents the first comprehensive and quantitative exploration of the airfoil design space at ultra-low Reynolds numbers. Under these conditions, the flow is dominated by viscous effects and the growth of thick boundary layers causes significant modifications to the effective geometry of the section. This effective loss of camber is exacerbated by the fact that flow separation occurs at low angles of attack. However, the onset of separation does not have to have the same catastrophic

performance ramifications seen in more traditional applications. Properly configured airfoils will continue to operate in a steady-state manner while significant portions of the section are separated, albeit with reduced performance.

At more conventional Reynolds numbers, separation is usually a rapid precursor to stall, but at ultra-low Reynolds numbers it is a prevalent feature of the normal operating range of an airfoil. This leads to a fundamental difference between airfoil design for more conventional Reynolds numbers and design for ultra-low Reynolds numbers. Typical high Reynolds number airfoil design considers separation only as a factor at the extremes of airfoil performance, such as near maximum lift or in designs emphasizing long runs of laminar flow; much more effort is spent on transition and other sources of drag that come into play long before separation is an issue. At ultra-low Reynolds numbers separation *is* the issue.

Small changes in the Reynolds number cause large changes in drag, and as the Reynolds number is reduced, section L/D quickly falls to single digits. Maximizing performance requires operation at or near maximum lift, causing separation and its impact on the effective camber to be dominant performance factors. Within the scope of this study, as the Reynolds number is reduced, the maximum steady-state lift coefficients are seen to generally *increase*. Since laminar separation is essentially independent of Reynolds number, the reduction in the steepness of the recovery gradient due to boundary layer growth dominates, delaying the onset of separation and permitting higher maximum lift coefficients.

These results are part of an even broader conclusion which provides a prime motivation for continued research. As the Reynolds number drops, the dramatic increase in viscous effects, such as the rapid growth of boundary layers and the prevalence of separation, do not simply washout the sensitivity to the detailed shape; the study of NACA 4-digit sections demonstrates that at ultra-low Reynolds numbers, geometry variations *still* have a tremendous effect on the aerodynamic performance of an airfoil. The impetus for two-dimensional research and design at these Reynolds numbers is further strengthened by

the significant performance gains achieved by the simple airfoil optimization study, resulting in the first quantitatively designed airfoils for ultra-low Reynolds number applications.

8.2 Summary of Results and Contributions in the Area of Ultra-Low Reynolds Number Rotors

The knowledge gained in the first segment has directly contributed to the development of a rapid hybrid method for rotor design and analysis. The effectiveness of this code in the analysis and preliminary design of centimeter-scale rotors has demonstrated the applicability of classical rotor analysis methods when properly augmented. Manipulation of basic rotor theory has shown that small rotors invariably suffer from a significant reduction in the figure of merit, meaning that once again, the unique operational parameters strongly call for some form of optimal design methodology. The combined implementation of classical rotor theory, 2-D viscous flow solutions, contracting wake models, and a rigorous treatment of viscous swirl, has culminated in an automated, non-linear optimizer based, design and analysis tool capable rapid of design cycles on modest computing platforms.

With an eye towards the goal of flight vehicle development, this tool has been used to design the first rotors specifically tailored to this regime. Three-dimensional viscous calculations of one of these geometries have demonstrated that the rapid analysis and design method provides sufficient accuracy, both in the global metrics of thrust and power and the detailed distributions of thrust and torque, to be a useful tool for preliminary design, but it also reveals two shortcomings. The presence of three-dimensional boundary layer effects result in behavior similar to a modification of the airfoil geometry and complex wake behavior, including blade-vortex interactions, result in an inflow distribution that differs from the reduced-order model. These issues deserve

further investigation, but do not significantly diminish the utility of the rapid analysis and design method.

Experimental testing has also served to both validate the method and expose another potential problem that strongly influences micro-rotor performance. Certain designs, having small hub diameters or high local solidities, exhibited thrust and power levels significantly below predicted values. The two-dimensional aerodynamic studies have demonstrated the strong need to minimize airfoil thickness at ultra-low Reynolds numbers, but the derivation of an aero-structural model for torsional deflections reveals that one consequence of reduced airfoil thickness is a disproportionate increase in torsional deflections due to rotational effects. Failure to account for this in a design can result in reduced incidence and aerodynamic unloading of the rotor.

Finally, this progression of work has led to the design and fabrication of prototype micro-rotorcraft. The largest design brings together current existing technologies into a flight-vehicle that could provide a valuable test-bed for further research. Full realization of the smallest of these designs appears to be beyond the state of the art, but it sets the bar for the future. Any complex device requires a convergence of necessary technologies for it to come into existence. Motors, batteries, microelectronics, and many other necessary components are progressing at accelerated rates, driven by the wants and needs of society. The aerodynamics are still in their infancy and progressing slowly. Even if the end product is not viable now, the research must still continue in order to be ready at the point that it all comes together - to do anything else is simply short-sighted.

8.3 Consideration of Future Work

This work was not envisioned as all-encompassing, and in that regard it has been an overwhelming success. There are many possible avenues for continued effort. Much can be done to improve and expand upon the methods and analysis completed here, but there

are also opportunities to use this work as a starting point for research in new directions and as a source of supporting technology for work in related fields.

In the area of airfoil performance and design there is a fundamental need for more data, whether experimental or computational, to increase the confidence in the results provided not only by INS2d, but by any number of other codes, currently available or in development, that are capable of these calculations. Since its initial publication in 2001, this data has already been used by others outside of the Stanford community as a point of comparison, and these comparisons are welcomed and encouraged, but this work is still very limited and without exception computational.

Adding new and relevant experimental data to the almost non-existent current offerings would provide an invaluable counter-point to computational results. High-quality airfoil testing at these Reynolds numbers is a technical challenge due to the possible requirement for a fluid medium other than air. This complicates the acquisition of forces, moments, and pressure distributions, but this should not be a deterrent. Technology has improved significantly since Thom and Swart tested their hand-shaped airfoil in 1940 and it would be exciting to see what modern experimental methods could bring to this problem.

Within the scope of airfoil optimization, there is room for improvement over the simple geometry model imposed in Chapter 3. The promising results shown here involve only four design variables. This low number is dictated by the optimization method employed, but application of a more sophisticated optimizer would permit much more freedom in defining the airfoil. One challenge is how to accomplish this efficiently given that the operating points of interest are so close to the steady-state operating limits.

The rapid rotor analysis method has been demonstrated and proven as a tool for preliminary design, but the issues of three-dimensional boundary layer effects and wake effects on the prescribed inflow velocities warrant further work. One area of investigation is the effects of vortex dissipation on the wake structure. The OVERFLOW-D results depict fairly rapid dissipation of the wake structure, but it has

not been determined whether this is an issue of gridding or numerics, or if this is supported by the flow physics. If it is a real phenomena of ultra-low Reynolds number operation, it suggests a simple modification to the current semi-infinite wake model that would strongly influence inflow velocities.

Additional improvements in the performance of the rapid analysis and design method could be readily achieved by integrating the current torsional deflection model directly into the design and analysis process. In the very least, this would eliminate the tedious manual iteration on deflection and aerodynamic loading carried out in Chapter 6.

If performance constraints warranted it, considerably more effort could be applied to move beyond this level of refinement and to fully account for three-dimensional effects in design. The rapid analysis and design method does not presume to achieve the fidelity of three-dimensional computational Navier-Stokes solutions, and shape optimization within 3-D Navier-Stokes computations has become more common. Such an approach would be currently be computationally expensive, but advances in computational technology and reductions in the financial cost of computing resources are rapidly making this a viable solution. However, it is not yet clear if the costs and complexity are warranted.

A novel approach to experimental testing of ultra-low Reynolds number rotors would be to make scale models, but in the opposite direction than is typical, larger not smaller. Reynolds number scaling would simply dictate slower rotational speeds. Increased size would simplify manufacturing, and facilitate maintaining and assessing geometric accuracy, and potentially improve the experimental accuracy by increasing the magnitude of the global metrics. Larger rotors would also open the possibility of experimental visualization, providing a comparison point to the details of the OVERFLOW-D solutions.

In a more supporting role, development of the rotorcraft prototypes would also facilitate new avenues of research. They potentially represent small, inexpensive, and highly maneuverable test-beds for research into distributed systems, collective intelligence,

formation flight, and many other interesting topics. There is no shortage of opportunities and I have been fortunate to participate in research on a topic in its infancy.

References

[1] Spedding, G.R., Lissaman, P.B.S., “Technical Aspects of Microscale Flight Systems,” *Journal of Avian Biology,* Vol. 29, No. 4, pp.458-468, 1998.

[2] Schmitz, F.W., “Aerodynamics of the Model Airplane,” (Translated from the German), RSIC-721, Redstone Scientific Information Center, Redstone Arsenal, Alabama, November, 1967.

[3] Althaus, Dieter, *Profilpolaren für den Modellflug, Band 2,* Neckar-Verlag, GmbH.

[4] Kunz, P.J., Kroo, I., “Analysis and Design of Airfoils for Use at Ultra-Low Reynolds Numbers,” *Fixed and Flapping Wing Aerodynamics for Micro Air Vehicle Applications, edited by T. J. Mueller, Vol. 195, Progress in Aeronautics and Astronautics,* AIAA, Reston, VA, 2001, Chap. 3.

[5] Kroo, I., Kunz, P.J., “Mesoscale Flight and Miniature Rotorcraft Development,” *Fixed and Flapping Wing Aerodynamics for Micro Air Vehicle Applications, edited by T. J. Mueller, Vol. 195, Progress in Aeronautics and Astronautics,* AIAA, Reston, VA, 2001, Chap. 23.

[6] Sunada, S., Kawachi, K., “Comparison of Wing Characteristics at an Ultralow Reynolds Number,” *Journal of Aircraft,* Vol. 39, No. 2, 2002.

[7] Azuma, A., Okamato, M., Yasuda, K., “Aerodynamics Characteristics of Wings at Low Reynolds Number,” *Fixed and Flapping Wing Aerodynamics for Micro Air Vehicle Applications, edited by T. J. Mueller, Vol. 195, Progress in Aeronautics and Astronautics,* AIAA, Reston, VA, 2001, Chap. 17.

[8] Ellington, C.P., “The Aerodynamics of Hovering Insect Flight,” *Philosophical Transactions of the Royal Society of London,* Series B, Vol. 305, pp.1-181, 1984.

[9] Weis-Fogh, T., “The Energetics of Hovering Flight in Hummingbirds and Drosophilia,” *Journal of Experimental Biology,* Vol. 56, pp.79-104.

[10] Weis-Fogh, T., “Quick Estimates of Flight Fitness in Hovering Animals, Including Novel Mechanisms for Lift Production,” *Journal of Experimental Biology,* Vol. 59, pp.169-230.

[11] Rayner, J.M.V., “A Vortex Theory of Animal Flight. I. The Vortex Wake of a Hovering Animal,” *Journal of Fluid Mechanics,* Vol. 91, pp.697-730.

[12] Rogers, S. E. and Kwak, D., “An Upwind Differencing Scheme for the Steady-state Incompressible Navier-Stokes Equations,” NASA TM 101051, November 1988.

[13] Rogers, S. E. and Kwak, D., “An Upwind Differencing Scheme for the Time Accurate Incompressible Navier-Stokes Equations,” *AIAA Journal*, Vol. 28, No. 2, 1990, pp. 253-262.

[14] Chorin, Alexandre Joel, “A Numerical Method for Solving Incompressible Viscous Flow Problems,” *Journal of Computational Physics*, Vol. 2, No. 1, 1967.

[15] Thom, A., Swart, P., “The Forces on an Aerofoil at Very Low Speeds,” *Journal of the Royal Aeronautical Society*, Vol. 44, pp. 761-770, 1940.

[16] Drela, M., Giles, M.B., “ISES - A Two-Dimensional Viscous Aerodynamics Design and Analysis Code,” AIAA Paper 87-0424, January 1987.

[17] Schlichting, Hermann, *Boundary-Layer Theory*, 7th ed., McGraw-Hill, New York, 1979.

[18] Stratford, B. S., Aeronautical Research Council R. & M. 3002, 1954.

[19] Curle, N. and Skan, S. W., *Aeronautical Quarterly*, No. 8, 1957, pp. 257.

[20] Thwaites, B., *Aeronautical Quarterly*, No. 1, 1949, pp. 245.

[21] Abbot, Ira H. and Von Doenhoff, Albert E., *Theory of Wing Sections*, Dover Publications Inc., New York, 1959.

[22] Maughmer, Mark D., Somers, Dan M., “Design and Experimental Results for a High-Altitude, Long Endurance Airfoil,” *Journal of Aircraft*, Vol. 26, No. 2, 1989, pp. 148-153.

[23] Akima, Hiroshi, “A Method of Univariate Interpolation that Has the Accuracy of a Third-Degree Polynomial,” *ACM Transactions on Mathematical Software*, Vol. 17, No. 3, 1991, pp. 341-366.

[24] Nelder, J. A., Mead, R., “A Simplex Method for Function Minimization,” *Computer Journal*, Vol. 7, 1965, pp. 308-313.

[25] Glauert, H., *The Elements of Aerofoil and Airscrew Theory*, Cambridge University Press, London, 1948.

[26] McCormick, Barnes W. Jr., *Aerodynamics of V/STOL Flight*, Dover Publications Inc., New York, 1999.

[27] Johnson, W., *Helicopter Theory*, Dover Publications Inc., New York, 1994.

[28] Lamb, Sir Horace, *Hydrodynamics*, Cambridge University Press, Cambridge, United Kingdom, 1997.

[29] Kunz, Peter J., "Development of a Software Package for the Assessment of High Performance Sailplanes," Master of Science Thesis, Department of Aerospace Engineering, Pennsylvania State University, University Park, Pennsylvania, August 1997, pp. 29-30, 43-51.

[30] Gill, Murray, and Saunders, "User's Guide for SNOPT 5.3: A Fortran Package for Large-scale Nonlinear Programming," Technical Report SOL 98-1, Department of Engineering Economic Systems and Operations Research, Stanford University, Stanford, CA, May 1998.

[31] A.G. Cooper, S. Kang, J. W. Kietzman, F. B. Prinz, J. L. Lombardi and L. Weiss, "Automated Fabrication of Complex Molded Parts Using Mold SDM," *Proceedings of the Solid Freeform Fabrication Symposium, University of Texas at Austin,* Austin, Texas, August 1998.

[32] Y. Cheng, "Fabrication Methods for a Mesoscopic Flying Vehicle," Ph.D. thesis, Stanford University, Stanford, CA, June 2001.

[33] Buning, P. G., Jesperson, D. C., Pulliam, T. H., Chan, W. M., Slotnick, J. P., Krist, S. E., and Renze, K. J., "OVERFLOW User's Manual, Version 1.8g," NASA Langley Research Center, March 1999.

[34] Meakin, R., "Composite Overset Structured Grids," Handbook of Grid Generation, edited by Thompson, J. F., Soni, B. K., and Weatherill, N. P., CRC Press, Washington, DC, 1999, pp. 11-1 to 11-19.

[35] Meakin, R. L., and Wissink, A. M., "Unsteady Aerodynamic Simulation of Static and Moving Bodies Using Scalable Computers," AIAA-99-3302, Proc. 14th AIAA Computational Fluid Dynamics Conf., Norfolk VA, July 1999 pp. 469-483.

[36] Strawn, R. C. and Djomehri, M. J., "Computational Modeling of Hovering Rotors and Wake Aerodynamics," presented at the AHS 57th Annual Forum, Washington, DC, May 9-11, 2001.

[37] Chan, William M. and Buning, Pieter G., "User's Manual for FOMOCO Utilities - Force and Moment Computation Tools for Overset Grids," NASA TM-110408, 1996.

[38] Myonic AG, Eckweg 8, PO Box 6121, CH – 2500 Biel-Bienne 6.

[39] Astro Flight Incorporated, 13311 Beach Ave., Marina Del Rey, CA 9029*2.*

[40] Internet Technical Library, MicroMo Electronics Incorporated, 14881 Evergreen Ave.,Clearwater, FL 33762-3008.

[41] WES-Technik GmbH, Klosterstr. 12, D-72644 Oberboihingen, Germany.

[42] Kietzman, John, “Rapid Prototyping Polymer Parts via Shape Deposition Manufacturing,” Ph.D. thesis, Stanford University, Stanford, CA, 1999.

[43] Kroo, Ilan, “Meso-Scale Flight Vehicle for Atmospheric Sensing,” NASA Institute for Advanced Concepts, Phase 1 Final Report, CP 98-01, 1999.

[44] Global Hobby Distributors, 18480 Bandilier Circle, Fountain Valley, CA 92708.

[45] Partridge, Chad B.,”A Distributed Dynamics and Control Simulator,” Eng. Thesis, Stanford University, Stanford, CA, 2001.

[46] Tadiran U.S. Battery Division, 2 Seaview Blvd., Port Washington, NY 11050.

[47] Gessow, A. and Myers, G. C., Jr., *Aerodynamics of the Helicopter,* Macmillan, New York, 1952.