

# **The Mesicopter: A Miniature Rotorcraft Concept**

## **Phase II Final Report**

Ilan Kroo, Fritz Prinz (P.I.'s)

Michael Shantz (Lecturer)

Peter Kunz, Gary Fay, Shelley Cheng, Tibor Fabian, Chad Partridge (Ph.D. students)

Stanford University

November 2001

### **Abstract**

This report summarizes work during Phase II of the NIAC-sponsored mesicopter program, which extended from November 1999 to October 2001. The research has dealt with an assessment of the feasibility and initial development of a very small-scale rotorcraft that flies on its own power and carries sensors for atmospheric research or planetary exploration. Initial devices are electrically powered, ranging in size from 2 to 15 cm, and have involved challenges in aerodynamics, control, and manufacturing. Many interesting scaling issues arise as one shrinks a flight vehicle down to this size. Certain scaling attributes are favorable, such as the increased strength and rigidity of structures at small scales, while others, such as aerodynamics, represent significant challenges. The report complements a series of progress reports over the past two years and describes the aerodynamic design of the rotor system, and approaches to fabrication, control, and power systems. Results of prototype tests suggest that the concept can be successfully produced and that the design methodology is appropriate, despite the insect-like scale of the rotors. Results of recent experiments with a larger versions of this concept intended to fly on Mars or for nearer-term terrestrial applications are also described.

## **1. Introduction and Summary**

### **1.1 Background**

The rapid development of microelectronics and microelectromechanical systems (MEMS) has made it possible to incorporate a variety of computing, communications, and sensing functions on air vehicles with masses of less than 100g. Micro air vehicles (MAVs) with 15 cm spans are now flying with realtime video, GPS, and sophisticated autopilot functions. With the idea that progress in these areas will continue, we have looked at what might be possible for future MAVs, focusing on airframe technologies that would enable flight at even smaller scales.

Such vehicles would have many unique capabilities including the ability to fly indoors or in swarms to provide sensor information over a wide area at a specific time. The very low mass of these devices might make them attractive for planetary exploration, of Mars or Titan for example, due to the high cost of transporting each gram [1]. Although sub-

gram imaging systems are not available, miniature aerial robots might be used in the near term for simple atmospheric sensing tasks.

Our investigation focuses on meso-scale systems – devices larger than microscopic, yet significantly smaller than conventional air vehicles. Such rotorcraft, termed mesicopters, are centimeter-scale devices with masses of 3 to 15g, powered by DC motors (figure 1). To achieve this goal, the program has begun with the development of somewhat larger vehicles with masses up to 60g. This provides near-term payload capabilities while reducing the cost of the development program.

#### *1.1.1 Hovering vs. Forward Flight*

That fixed wing aircraft dominate aeronautical development may be attributed to two factors. First, aircraft have generally been employed for transportation of people or goods. The goal is not just to remain airborne, but to get from point A to point B. Not all applications require the high speed capability of aircraft, but this has been an obvious feature to exploit. Although surveillance, communications, and imaging (information-related) applications have existed for some time, only recently has this started to become possible with very low mass systems. Second, and most importantly, hovering requires substantially more power than does forward flight – at least for conventionally-sized systems. Even if the rotor and wing dimensions are similar, the aircraft in forward flight requires a thrust of  $W / L/D$  to maintain level flight, while the rotor requires a thrust equal to the vehicle weight. However, the required power for a fixed wing airplane increases with speed when the thrust is given. So if the wing  $L/D$  is small and the speed that would be required to sustain flight is large, the discrepancy between the required power for forward flight and hover is less significant. In some cases, especially at very low Reynolds numbers, the two flight modes require very similar power inputs.

Rotorcraft may also be desirable for certain missions because of their compact form factor and ability to maintain their position in hover. In many imaging applications, the conventional aircraft's minimum speed limitations are problematic. Current designs for a Mars aircraft indicate that to avoid excessive vehicle dimensions, flight speeds of Mach 0.5 to 0.6 are required, limiting low altitude, high-resolution imaging options. Finally, with a rotorcraft design of this size we can provide sufficient control for a four-rotor vehicle using motor speed control, avoiding problems with control surface aerodynamics and actuation that plague small aircraft of conventional design.

#### *1.1.2 Rotating vs. Flapping*

It has been noted that the only truly successful examples of flight at such small scales are insects, which achieve flight using wing flapping, rather than rotary motion. Indeed Ellington[2], Dickinson[3], and others have argued that unsteady aerodynamic effects may be significant features of insect flight, increasing the achievable maximum lift at very low Reynolds numbers. A number of successful powered ornithopters have been developed and, despite the fact that these have generally achieved poor efficiencies, there is little question that the approach can be used. It is most attractive at small scales where inertial loads do not dominate and where unsteady phenomena may be helpful. Still, the complexity of the flow field, the required wing motion, and the mechanism itself leads

one first to ask not whether a flapping device can be used, but whether it must be used. Just as automobiles depart from the paradigm suggested by walking animals, we have started by considering the simple steady motion of a rotor, eliminating mechanical complexity and simplifying stability and control issues. Rotary motion does not preclude the possibility of exploiting unsteady effects and we are currently considering the potential for increased maximum lift with rotor speed modulation.

### *1.1.3 Fundamental Scaling Issues*

The obvious success of small flying animals, in fact the absence of very large flying animals, suggests that flight at very small scales may be more easily accomplished than flight at larger scales. Some of the relevant scaling laws (e.g. increased strength and stiffness with smaller size) have a beneficial effect on the design of small flight vehicles, while others (e.g. aerodynamics) make the design of small-scale aircraft more difficult. Tennekes[4] shows that the wing loading ( $W/S$ ) of flying devices increases with size. While this result is likely related to various versions of a square-cubed law, it is surely more complicated than the simple constant density scaling suggested in this reference. One part of the explanation for this trend in nature is that smaller, lighter creatures can more feasibly contend with relatively larger wings because of improved structural properties (strength and stiffness / weight) at small scales. This is also a feature that can be exploited with a small rotorcraft. Rotor weight and aeroelasticity are not problems for tiny rotors with a disk loading of less than  $25 \text{ N/m}^2$  ( $0.5 \text{ lb/ft}^2$ ), while large helicopters have disk loadings of  $500 \text{ N/m}^2$  ( $10 \text{ lb/ft}^2$ ) or more and would become unmanageable with rotors 4 to 5 times as large. The lower attainable disk loading is one favorable attribute of small rotorcraft since the *required* power-to-weight ratio scales as the square root of the disk loading.

The *available* power-to-weight ratio depends on the motor and power source. While scaling of devices such as gas turbines may be quite different from permanent magnet motors, our initial devices consist of electrically driven rotors with energy stored in batteries. Available small DC motors show relatively little effect of size on their power-to-weight ratio, although the efficiencies of the very small motors are usually smaller than larger versions. The specific energy of batteries depends more on chemistry than size, although very small batteries are often dominated by case weight. A study of commercial DC motors and batteries suggested that scaling effects on available power-to-weight ratio were not significant over the 4-5 orders of magnitude considered. See figures 2, 3. If it were not for Reynolds number effects, a miniature electric rotorcraft would have some fundamental advantages over larger devices since the required power-to-weight ratio is reduced by low disk loading, while the available power-to-weight is little changed with size.

The reduction of Reynolds number with vehicle size poses one of the greatest challenges for meso-scale flight. While section lift-to-drag ratios of 200 and above are common for large airfoils, the increased skin friction at small scales leads to values in the range of 5 to 15 for Reynolds numbers in the thousands. (Figure 4.)

## 1.2 Approach

The development of the meso-scale rotorcraft is ongoing and has involved challenges in many disciplines. Some of the primary issues are summarized below.

*Insect-Scale Aerodynamics:* The Reynolds number of the mesicopter rotors lies in the range of 1,000 to 6,000 where aerodynamics are dominated by viscous considerations and few design tools are available. This is one of the areas in which scaling laws are unfavorable, with lower lift-to-drag ratios and limited rotor lift capabilities. Some of the aerodynamic features are poorly understood in this size regime and means by which improved performance may be realized have been little explored. Because the flow is viscous, some of the simpler tools used for propeller and rotor design are not applicable and basic design rules (e.g. nearly constant inflow) are not appropriate.

*3-D Micro-Manufacturing:* To achieve high lift-to-drag ratios smooth rotors with 3D surfaces at micro scale dimensions must be built. Traditional micro-fabrication techniques can generate features at and below the desired size scales. Yet the need to produce smooth 3D surface features requires rethinking processing steps commonly used for the building of IC and MEMS structures. Traditional 3D machining methods are not normally employed for the fabrication of parts and devices as thin as 50microns, yet their resolution of a few microns makes them attractive candidates for shaping surfaces within this size regime.

*Integration of Power and Control Systems:* Although many types of batteries with high specific energy are becoming available, identifying very small batteries suitable for the Mesicopter, with good specific energy and high current rates is not easy. The control of these small devices is also a problem. Because of their size, stability time constants are very short and the mass budget for motor/flight control sensors and processing is limited.

The basic approach has been to develop scalable design and fabrication methods and to start with devices that are larger than the eventual goal. The (super) scale model prototypes are sufficiently large that commercial motors, batteries, and electronics can be employed. The first such prototype is shown in figure 5 with a maximum takeoff weight of about 3g. This device was used to gather data on aerodynamic performance and required an external power supply since the planned Li-Ion batteries were not yet available. A second prototype with a maximum weight of 10-15g is currently being tested and can utilize existing batteries. Finally, a stability and control testbed with even larger dimensions and with a mass of 60g is also under development. As these systems are refined, the scale will be reduced to explore the limits of this technology.

### 1.3. Summary of Prototype and Testbed Development

Development of the mesicopter has relied on a variety of prototypes, of different scales, each intended to address particular problems that need to be solved before a true meso-scale vehicle is successful. The following sections summarize the characteristics and objectives of some of these vehicles. Additional detail is provided in subsequent chapters.

#### 1.3.1 Initial 3g Vehicle

The initial mesicopter work was undertaken to determine whether a very small scale vehicle would be at all feasible. The device was constructed using 3mm Smoovy motors with 1.5cm diameter rotors and is shown in figure 1-1. With power supplied from an external source the device was used to measure thrust, power, and efficiency of the motors and rotors.




Figure 1-1. Initial 3g mesicopter and rotor tests.

#### 1.3.2 Nearer-Term Mesicopter (15g)

Since the motor speed controllers and batteries that are currently available commercially were too heavy to provide energy and control for the 3g mesicopter, a larger device was constructed. Figure 1-2a shows the assembled mesicopter prototype. The total weight of the prototype is 17 grams. Figure 1-2b shows the comparison of single propeller lift versus mesicopter lift as a function of controller voltage. The shrouds produce a negligible effect on the lift performance due to the current large gap between rotor tips and the shroud. Using this lift information from it can be estimated that at least 15 V at the controller input are needed to take off. Tests with external power supply confirmed this estimation. The mesicopter lifts off at 13 V and is able to hover out of ground effect with 16 V draining almost 1 A. Even with weight reduction to 15g, the prototype cannot maintain flight for a useful period with the power available from the onboard supercapacitors and alternative power sources would be needed to make the device useful.

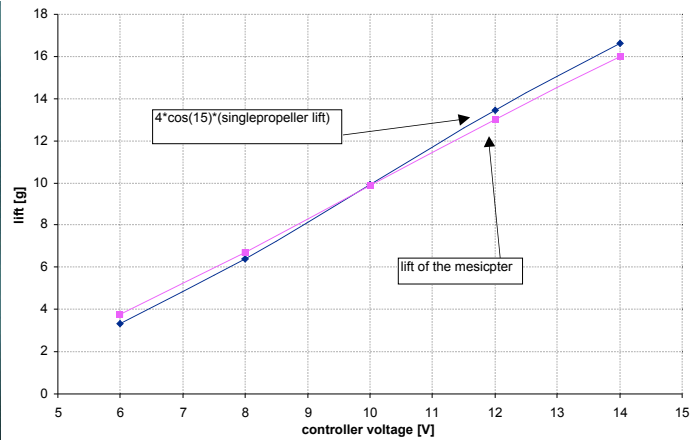



Figure 1-2: (from left) a) Assembled mesicopter, b) Comparison of single propeller lift versus mesicopter lift as a function of controller voltage

### 1.3.3 Stability and Control Testbed

The small mesicopters have been developed primarily to study aerodynamics and power issues. For development of stability and control strategies, it is more effective to start with larger devices with larger lifting capacity and slower dynamics. A series of stability and control testbeds have been designed and constructed. Photos of these prototypes are included below.




Figure 1-3. Six-inch (rotor tip to rotor tip) stability and control testbed.




Figure 1-4. Side View Illustrating Rotor Cant Angle

The device weighs approximately 60g with component weights as follows:

Frame:	3.35 gm
Motor:	5.65 gm x 4
Rotor:	0.6 gm x 4
Speed Controller:	1.08 gm x 4
Micro-Receiver:	2.34 gm
Batteries:	26.18 gm (Seven cell NiCd pack)
TOTAL:	61.19 gm

The testbeds were built using mostly off-the-shelf components available from hobby retailers. Control is sent to the mesicopter via a standard programmable transmitter with channel mixing described previously. A five channel micro-receiver then sends the command signals to the four micro-scale speed controllers. The speed controllers convert the PWM signal generated by the receiver into a variable voltage to the motors. Power to the motors is supplied using either lithium ion or NiCd cells. The motors are brushed DC motors that are commercially available. Finally, the rotors are manufactured at Stanford since clockwise and counter-clockwise versions are needed. The current frame is balsa.

The rotors on this testbed are capable of producing more than 20 gm of thrust each, meaning that the vehicle can hover using only 75% of its maximum thrust. This provides a sufficient margin for maneuvers and limited payload or additional sensors. The maximum thrust may be increased further using optimized rotors as described in the aerodynamics section. It has been used successfully for the vision-based closed-loop control tests.



#### *1.3.4 Payload-Carrying Testbed*

A larger testbed vehicle was also constructed using larger, coreless DC motors of much higher efficiency. This testbed weighs approximately 100g but the rotors provide approximately 160g of lift making it possible to carry significant payloads. The total dimension is about twice that of the smaller testbed. Rotors were constructed from carbon-fiber.




Figure 1-5. 100g testbed

#### *1.3.5 Systems Integration Testbed*

A systems integration testbed has been constructed for exploring various approaches to communication, sensing, stabilization, and control. Exploration of algorithms and techniques for flight stability and semi-autonomous navigation requires a flexible and programmable prototyping platform. Such a platform is needed to validate physical simulations of both the flight characteristics and sensor functions. In order to move rapidly on prototypes, the testbed should be built from readily available chips and components to avoid the time and expense of custom ASICs. Completely custom silicon could, of course, produce a much smaller system but would not be a good investment of time and money at this stage of the project. We chose to design with off the shelf surface mount chips giving us a good tradeoff between effort and small size. The components were chosen for their support of digital communications, sensor support, and reconfigurability. The system consists of two main parts, a **custom transmitter** on the ground for radio communications and control, and a **flyer** with onboard micro-controller, radio, and other peripherals. Without motors or battery the flyer weighs 9 grams.



The custom transmitter shown below uses a Great Planes Real Flight Futaba controller with a joystick port output plugged into a board containing a Microchip PIC17C756A micro-controller and a Linx Technologies 418MHz RF digital transmitter. This radio transmits serial data at 4800 baud with a 300 ft range, and can easily be upgraded to a higher bandwidth two-way digital communications link. The PIC micro-controller samples the joystick port, converts the joystick values for roll, pitch, yaw, and throttle, to motor controls for front, left, back, and side motors. It then transmits this data plus button settings to the flyer.




Figure 1-6. Controller and transmitter for systems integration testbed.

The following equation expresses the conversion.

$$\begin{bmatrix} \text{roll} \\ \text{pitch} \\ \text{yaw} \\ \text{throttle} \end{bmatrix} = \begin{bmatrix} 0 & d & 0 & -d \\ d & 0 & -d & 0 \\ k & k & k & k \\ 1/4 & 1/4 & 1/4 & 1/4 \end{bmatrix} \begin{bmatrix} \text{front} \\ \text{left} \\ \text{back} \\ \text{right} \end{bmatrix}$$

where  $d$  represents the distance from the center of mass to each motor and  $k$  expresses a linear (approximate) relationship between lift and drag. The matrix is inverted in order to solve for the power to the front, left, back, and right motors given the joystick settings for roll, pitch, yaw, and throttle.

The flyer, shown below, is essentially a flying printed circuit board. This is in line with an emphasis on multipurpose components. The thin, light, 20mil circuit board also serves as the frame to which the motors are attached with CA glue. The connector for the

battery is also the battery support post. The PIC17 micro-controller can easily be extended with additional sensory and communication components as well as additional control software to incorporate their functionality. The surface mount PIC17 reads the incoming motor controls from the Linx Technologies 418MHz receiver via a serial digital interface and one of the PIC17's two UARTs. The PIC17 uses its 3 on-chip PWM outputs plus a fourth PWM, implemented in software using timers, to drive the motors. L293DD power chips input the PWM and deliver 1.2A (max) of current to each motor. The vehicle with motors and battery weighs 65 grams and generates 80 g of total thrust. The PIC17 software is written in C giving the capability for expansion through use of the 12 10bit A/D inputs, I<sup>2</sup>C interface, dual UARTs, pulse width modulators, and many robust I/O pins.





Figure 1-7. PCB forms the structure of this prototype.

The flyer has been tested and it lifts off and flies. However, it is very difficult to control without some form of sensor based stabilization. An investigation was conducted into the use of a CMOS image sensor and fixed lens as input into a vision stabilization system. The Photobit PB0100 CMOS camera chip was selected for its flexible register control, dynamic on-chip exposure control, and convenient I<sup>2</sup>C interface to the PIC17 micro-controller. The PIC17 and 4800 baud radio are not fast enough to communicate video to an image processor on the ground but initial calculations indicate that the PIC17 is fast enough to sample small regions of the camera image and to compute the optical flow vectors for these regions. These vectors may be used to compute an estimate of the six degree-of-freedom spatial vector describing the vehicle motion and may be used to stabilize the vehicle.

The PB0100 chip is shown below on the PCB for the new vision stabilized flyer. The lens in its black plastic lens holder and the PCB for the camera are also shown.



*Figure 1-8. PC board for vision-stabilized prototype. Note lens and CMOS camera chip.*

## 1.4 Future Work

Work to date has illustrated both the feasibility of meso-scale rotorcraft and the challenges still ahead before practical implementations are possible. Continued work in each of the areas reported here is still required, with special emphasis on stability and control.

Many of researchers with whom we have talked, have suggested that a device that is capable of carrying 10-50g of payload would be useful today. While we expect that advances in sensor technologies will provide interesting missions for vehicles carrying less than 1g of payload in 10-40 years, such systems are not presently available. Thus, in addition to efforts directed at the smallest mesicopters, future work should include development of the larger vehicles which serve as useful testbeds while providing greater near-term payload capabilities.

## 1.5. References

1. Kroo, I., Kunz, P., "Development of the Mesicopter: A Miniature Autonomous Rotorcraft," presented at the American Helicopter Society International Vertical Lift Aircraft Design Specialists' Meeting, San Francisco, CA, Jan. 2000.
2. Kroo, I., Kunz, P., "Meso-scale Flight and Miniature Rotorcraft Development," Proceedings of a Workshop on Fixed and Flapping Flight at Low Reynolds Numbers, Notre Dame, June 2000.
3. Kunz, P., Kroo, I., "Analysis, Design, and Testing of Airfoils for Use at Ultra-Low Reynolds Numbers," Proceedings of a Workshop on Fixed and Flapping Flight at Low Reynolds Numbers, Notre Dame, June 2000.
4. Henry Bortmann, "Whirlybugs," New Scientist, June 5, 1999.
5. Neil Gross, ed., "Developments to Watch: Tiny Helicopters that Go Boldly Where No Man...", Business Week, June 21, 1999.

# Chapter 2

## Aerodynamic Research and Design Development

### 1.0 Introduction

The research and development of meso-scale rotor-craft requires both a greater understanding of the relevant aerodynamics and the development of specialized design tools for this unique operating environment. The sectional Reynolds numbers for these devices are generally below 10,000. The authors describe this as the ultra-low Reynolds number regime, differentiating it from low Reynolds number aerodynamics which generally consider Reynolds numbers one to two orders of magnitude greater. The work in the area of aerodynamics undertaken in this NIAC Phase II study can be separated into several focus areas including two-dimensional aerodynamics at ultra-low Reynolds numbers, airfoil design at ultra-low Reynolds numbers, development of an efficient rotor design methodology, 3-D Navier-Stokes modeling of candidate rotors, and experimental testing.

### 2.0 Two-Dimensional Airfoil Analysis and Design at Ultra-Low Reynolds Numbers

The operating regime of the meso-scale helicopter poses certain difficulties for aerodynamic analysis and design. It has been unclear to what extent classical airfoil and finite wing analysis and design methods are applicable in this flow regime. The highly viscous nature of the flow field results in large increases in the boundary layer thickness and the potential for large regions of separated flow. These factors result in large discrepancies in performance from what might be expected from experience at higher Reynolds numbers. A key advantage of operating at ultra-low Reynolds numbers is that the boundary layer can safely be assumed to be fully laminar, eliminating the complexities and inaccuracies of transition and turbulence modeling.

Very little experimental or computation work exists to date for aerodynamic lifting surfaces operating at such low Reynolds numbers. Under this phase II study, significant effort has gone towards a comparison of the analysis tools currently available, investigation of the unique characteristics of two-dimensional airfoils operating in this regime, and the assessment of the performance benefits of detailed airfoil design. The majority of this work is detailed in the phase II interim report but the key elements are revisited here in addition to subsequent material. A more thorough discussion may also be found in the paper by Kunz and Kroo<sup>[1]</sup>, a paper based on work completed as part of this NIAC study.



## 2.1 Computational Methods

A range of computational tools has been explored for application to airfoil analysis at ultra-low Reynolds numbers. Low section Mach numbers occur for a large number of possible vehicle applications, but Navier-Stokes solvers for the compressible flow equations generally require some form of preconditioning at very low Mach numbers. Operation small rotor-craft at high altitudes or in unique environments such as the Martian atmosphere would result in ultra-low section Reynolds numbers at section Mach numbers greater than 0.5. Significant compressibility increases the complexity of the flow-field, but it also would allow use of more conventional RANS solvers.

One approach to the preconditioning issue is the use of the artificial compressibility method, first introduced by Chorin<sup>[2]</sup>, to deal with incompressible flows. 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,

$$\frac{\delta u}{\delta x} + \frac{\delta v}{\delta y} = 0$$

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

$$\frac{\delta \rho}{\delta \tau} + \frac{\delta u}{\delta x} + \frac{\delta v}{\delta y} = 0$$


Density is related to pressure via an artificial equation of state, where  $\delta$  is the artificial compressibility.

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

This introduces an artificial and finite acoustic speed governed by the selection of the  $\delta$  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 is seen to act similarly to a relaxation parameter.

The computational studies in this phase II study make extensive use of the INS2d two-dimensional incompressible Navier-Stokes solver developed by Rogers<sup>[3],[4]</sup>. All of the INS-2d calculations in this study use a C-grid topology with either 256 by 64 cells or 512 by 128 cells. The airfoil is paneled using 70% of the stream-wise cells, with 10% 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 1. Error values are based on analysis with a 1024 by 256 grid. The 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 1.**

The INS-2d code has been validated for this application using the experimental results of Thom and Swart<sup>[5]</sup> for a RAF-6a airfoil at ultra-low Reynolds numbers. In general, 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, using 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. 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  $\alpha = 10.0$ . The computational result matches the experimental value of  $C_l = 0.52$  within 3.0%.



### Experimental and Computed Drag for the R.A.F. 6a



**FIGURE 2.**

Navier-Stokes solvers offer high-fidelity, but at the cost of time and computational expense. Using integral boundary layer formulations in conjunction with inviscid flow field solutions offers the potential for significant computational savings over viscous flow solvers. The MSES program developed by Drela<sup>[6]</sup> has been applied in this study with limited success. This is a two-dimensional Euler solver, coupled with an integral boundary layer formulation. It gives 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 coupled inviscid/boundary layer methods of this type.

A comparison of results from MSES and INS2d for the NACA 4402 and NACA 4404 airfoils at  $Re=1000$  are presented in Figures 3 and 4. The upper end of each curve represents the maximum angle of attack for which a steady-state solution was attainable. Over the range in which MSES does converge to a solution, the trends in the results agree with INS2d, and in both figures the effects of increasing thickness are the same. Both analyses indicate similar reductions in the lift curve slope and increases in drag. The MSES solutions predict a lower lift curve slope and a slightly higher zero lift angle of attack, resulting in a deviation in predicted lift, and approximately 5% lower drag than the equivalent INS2d cases. That this method works at all is surprising, given the limitations of the boundary layer formulation, but under the restriction of low angles of attack, the much faster inviscid/integral boundary layer codes can provide a functional alternative to

full viscous flow field solutions.




FIGURE 3.

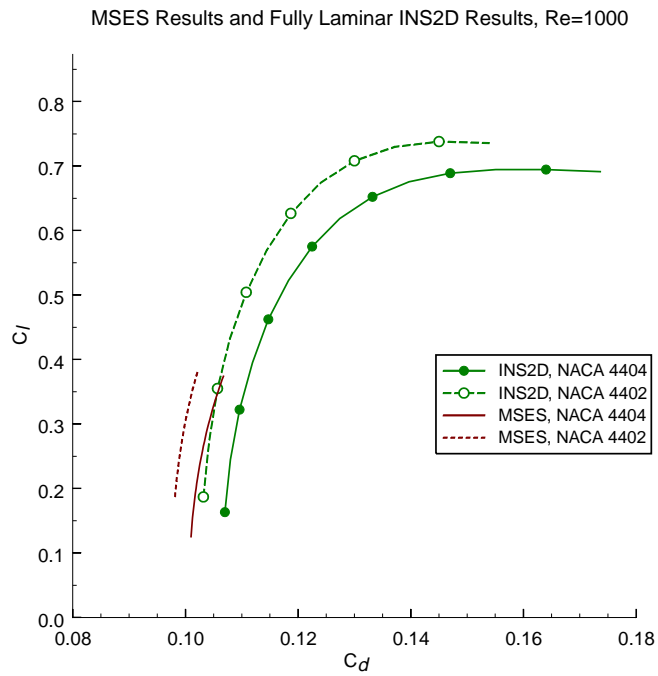


FIGURE 4.

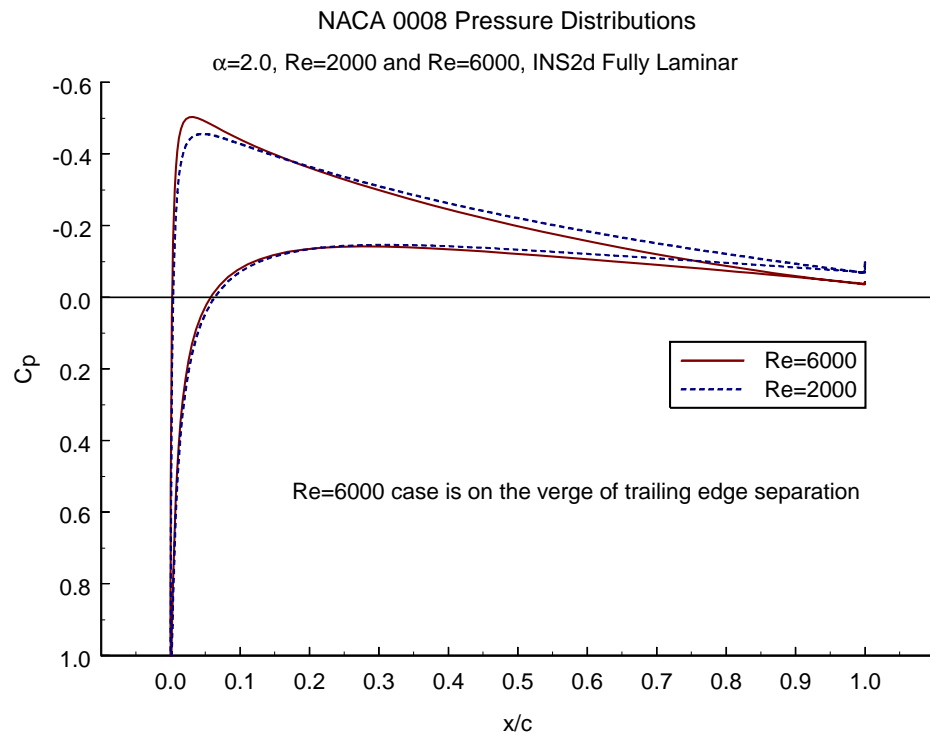
## 2.2 Effects of Reynolds Number on Airfoil Performance

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.

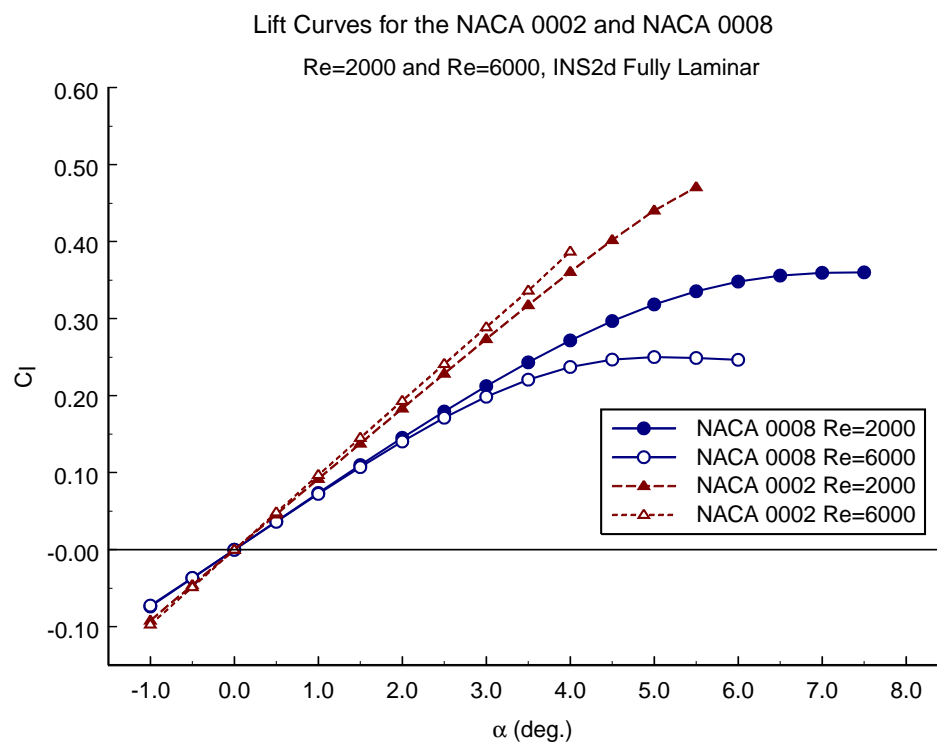
Flow at ultra-low Reynolds numbers is viscously dominated, and as the Reynolds number is reduced, the effects of increasing boundary layer thickness become more pronounced, significantly altering the effective geometry of the airfoil. A reduction in the height of the leading edge suction peak and the 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  $\alpha=2.0$  are presented in Figure 5. The  $Re=6000$  case is on the verge of trailing edge separation, but the  $Re=2000$  case does not separate until  $\alpha=3.5$ . 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.

Reducing the Reynolds number affects the lift curve by reducing the slope in the linear range and extending the 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 overall effect is a significant increase in both the maximum steady-state angle of attack and lift coefficient.

Lift curves for the NACA 0002 and NACA 0008 are presented in Figure 6. The reduction of slope is most apparent for the NACA 0002 airfoil, but both sections exhibit the extension of the linear lift range. The  $Re=2000$  case reaches  $\alpha=5.0$  and a lift coefficient a full tenth greater than the  $Re=6000$  case. Similar gains occur for the NACA 0008 at  $Re=2000$ .



**FIGURE 5.**



**FIGURE 6.**

## 2.3 Effects of Section Thickness on Airfoil Performance

Airfoil thickness variations appear to have two principal performance effects. A drag penalty, due to the pressure recovery attributable to increased thickness, is to be expected, but a strong reduction in the lift curve slope is also apparent. The variations in drag with section thickness are illustrated by the airfoil drag polars in Figure 7. The drag penalty associated with increasing thickness grows as the Reynolds number is decreased.

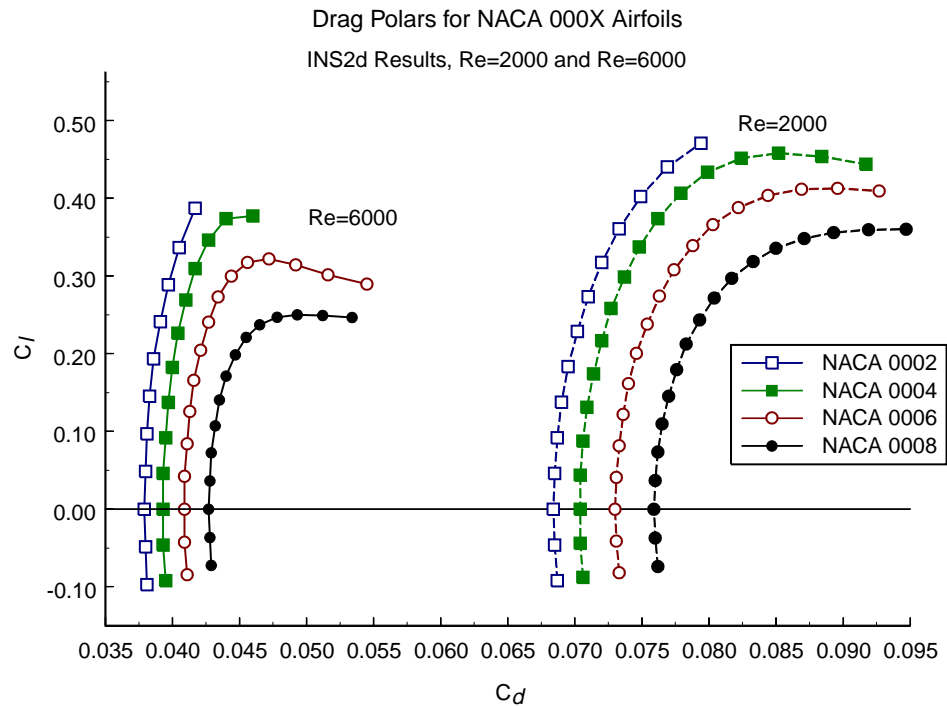


FIGURE 7.

The reduction in lift curve slope is an extension of behavior at higher Reynolds numbers. Within the linear range, the inviscid lift curve slope of an airfoil benefits from increased thickness. Viscous effects degrade the lift curve slope. 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. At more conventional Reynolds numbers, the net result is a 5% to 10% reduction in lift curve slope below the inviscid thin airfoil value of  $2\pi$ . This is not the case for Reynolds numbers below 10,000. The viscous boundary layer growth dominates and increasing thickness results in a significant decrease in lift curve slope. The results in Figure 8 show as much as a 35% reduction in lift curve slope for an 8% thick section. 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.

The decambering effect of the boundary layer is visualized by considering constant

velocity contours in the flow field. Several  $0.2V_\infty$  contours are drawn for the NACA 0002 and NACA 0008 sections at  $Re=6000$  in Figure 9. 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.

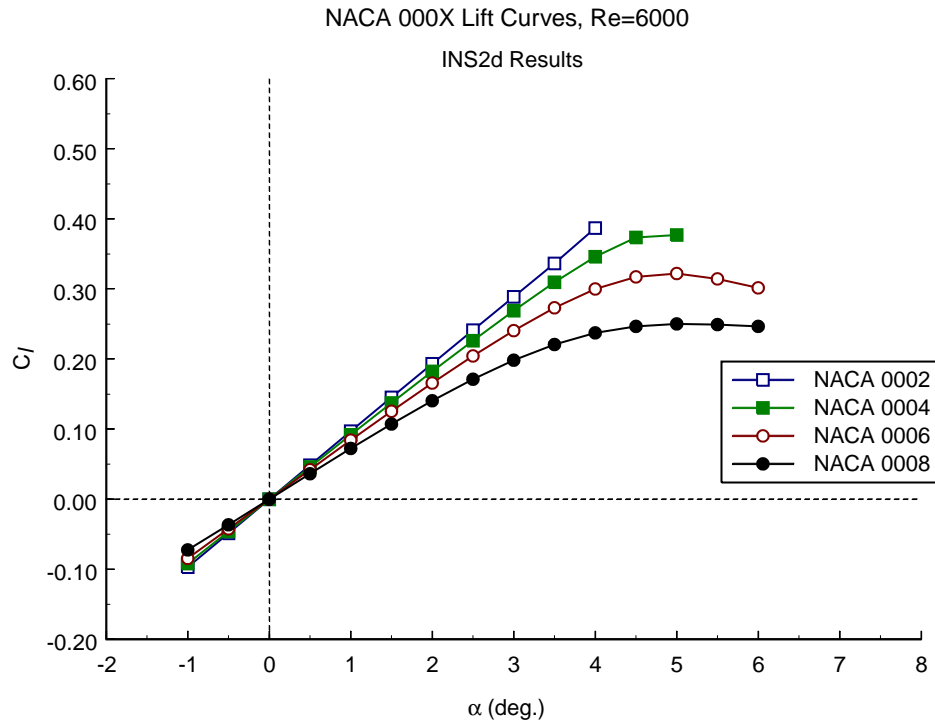


FIGURE 8.




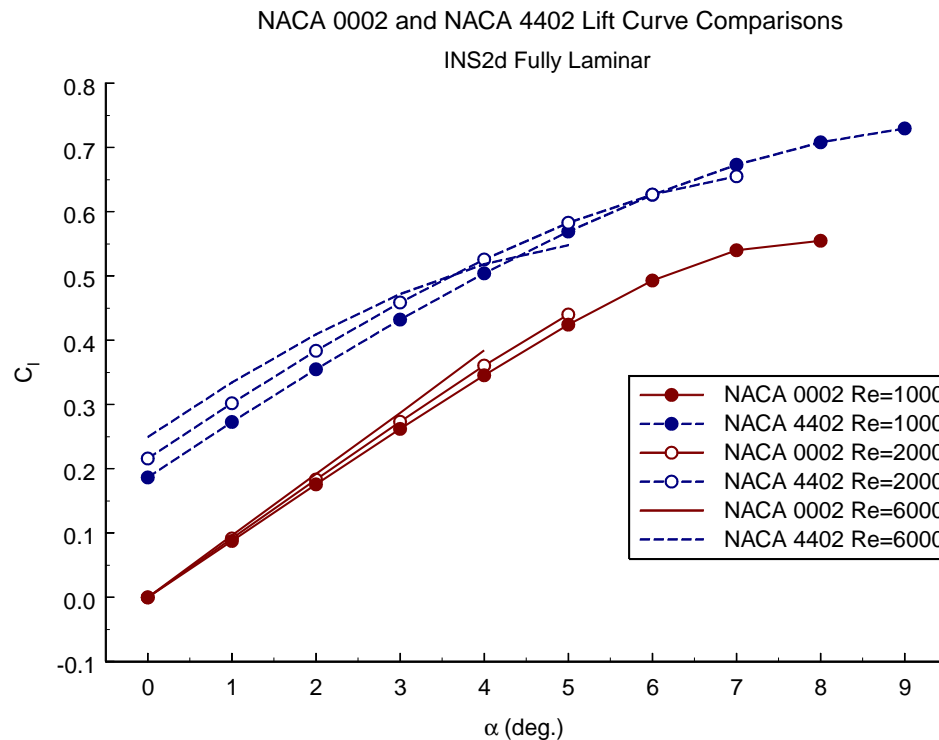
FIGURE 9.

## 2.4 Effects of Camber on Airfoil Performance

The effects of camber do not differ significantly from those at much higher Reynolds numbers, but the discovery that the detailed geometry is still an effective driver of performance at such low Reynolds numbers is itself a useful conclusion. The first order effect on the lift curve is a translation towards lower zero lift angles of attack with

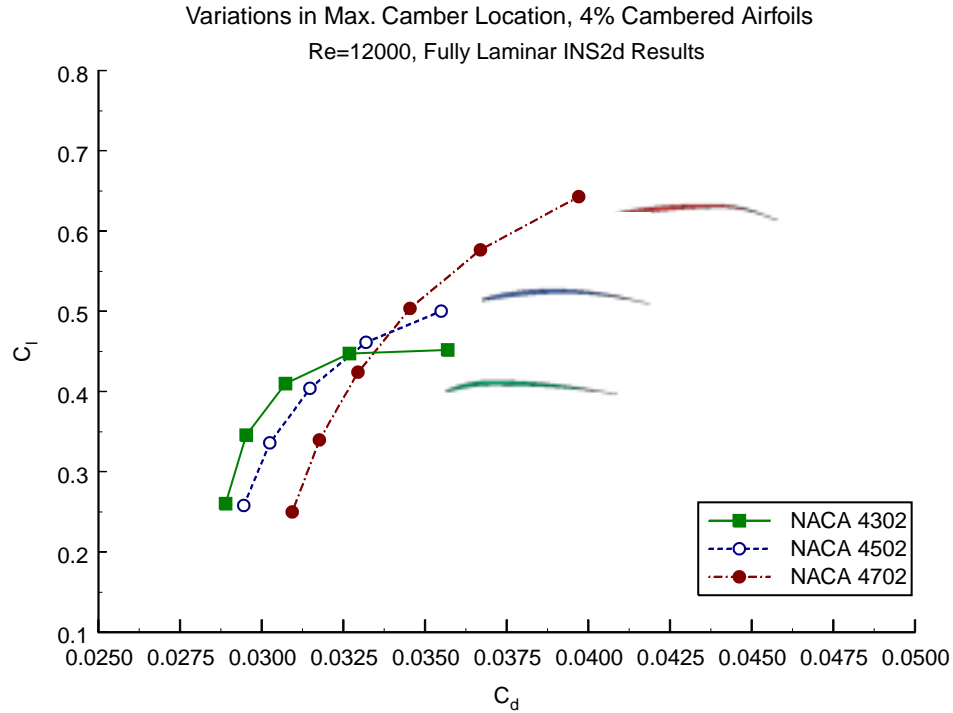


increasing camber. The maximum steady-state lift coefficients also increase. Comparison of the NACA 0002 and NACA 4402 airfoils indicates the gross effects of camber on performance. Referring to the lift curves provided in Figure 10, there is a 30% increase in the maximum steady-state lift coefficient. Although the drag also increases, there is a net gain in lift-to-drag ratio. The maximum lift-to-drag ratio increases from 4.5 to 5.4 at  $Re=1000$  and from 9.3 to 11.0 at  $Re=6000$ . The improvement in maximum lift is aided by the increase in the ideal angle of attack that comes with the introduction of camber, delaying leading edge separation in the thin sections being considered.



**FIGURE 10.**

Further analyses have investigated the possible benefits of varying the distribution of camber. The design space is explored using 2% thick NACA 4-digit profiles. The lift curves in Figure 11 represent 4% camber at three different chord locations. The aft shift of maximum camber results in a less severe reduction of lift past the 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 initial 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 effect on separation is to contain it 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<sup>[7]</sup>.



**FIGURE 11.**

## 2.5 Airfoil Design Optimization

The study of Reynolds number and geometry effects on airfoil performance is not meant to be a detailed indicator for design. It is however, indicative of the large variations in performance that exist within the design space and some of the physical trends responsible. Airfoil optimization tools have been developed to further explore the design space and to develop airfoils specifically suited to the mesicopter application.

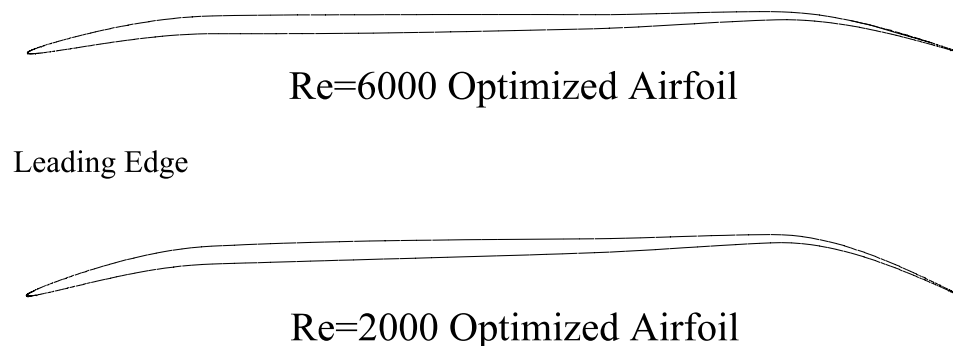
The design method makes use of the previous results to simplify the problem to its most essential elements. The maximum section thickness is fixed at 2% and the NACA 4-digit thickness distribution is used. This should not affect the utility of the results and greatly facilitates automated grid generation. 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 modeled with an Akima spline<sup>[8]</sup> anchored at the leading and trailing edges. Four interior knots are used to define the camber-line. These are evenly distributed and their chord-wise locations are fixed. The knots move perpendicular to the chord-line constrained by upper and lower camber limits.

The optimizer is a constrained simplex optimizer, a modified Nelder-Mead simplex. It is

coupled with the INS2d code and a grid generator. With 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.

Two airfoils have been developed using this approach for  $Re=6000$  (R6) and  $Re=2000$  (R2). Both sections are shown in Figure 12. 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.



**FIGURE 12.**

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 tested 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 13 shows the  $L/D$  versus geometric angle of attack for these three airfoils. The optimized section begins to show gains past  $\alpha=3.0$ , increasing until the maximum  $L/D$  is reached at  $\alpha=5.0$ .

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.

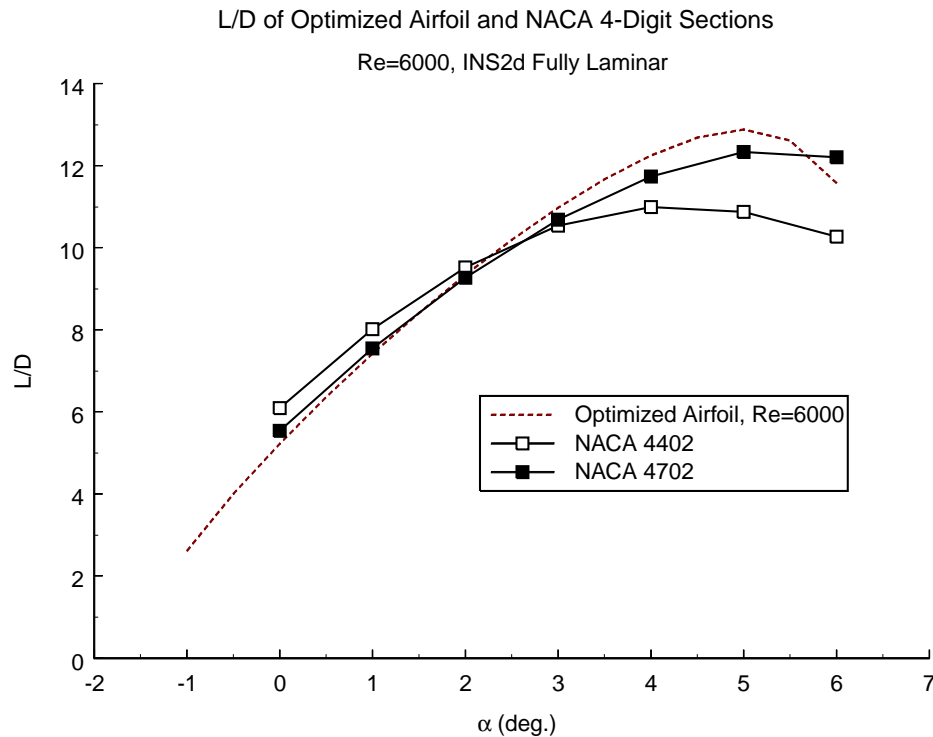


FIGURE 13.

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 effective and important tools, not only for design, but also for enhancing our understanding of this flight regime.

### 3.0 Rotor Design Methodology and Computational Tools

A low order rotor design code and a closely related rotor analysis code have been created for the development of very low Reynolds number rotors. The tools consist of a rotor performance package coupled with a nonlinear optimizer and 2-D section data from 2-D Navier Stokes analyses. The theoretical development and a comparison of analysis results and test data are provided. This code is currently applicable for the static thrust (hover) case only, but the development of relations for steady climbing rotors is also described.

#### 3.1 Theoretical Development

The theory applied in these codes is based on the combination of momentum flux concepts

and the blade element approach. Thrust and torque expressions are developed for a differential blade element. These expressions are then integrated along the length of the blade. Two expressions are developed for thrust and two for torque. One is based on the conservation of momentum across a differential annulus; the other is based on the forces developed by the bound circulation on a differential blade element and known 2-D section properties.

#### *Momentum Theory Equations:*

The momentum theory relations for thrust and torque are based on the conservation of momentum across a differential annulus of the rotor disk. Initially, uniform velocities are assumed across the entire annulus and the flow is assumed inviscid. This is essentially actuator disk theory applied to a differential element. Corrections for these assumptions will be applied later in the formulation. Primes denote quantities per unit length.

**$v$  = Induced horizontal velocity ( + in the direction of blade motion)**

**$U_{\infty}$  = Free Stream Vertical Velocity ( $U_{\infty} = 0$  for hover)**

**Thrust' =  $T' = [ 2 \rho u ( u + U_{\infty} ) (2\pi r)]$**

**Torque' =  $Q' = [ 2 \rho v ( u + U_{\infty} ) (2\pi r) r]$**

Viscous effects are incorporated by adding the two-dimensional section drag to each blade element:

**Thrust' =  $T' = [ 2 \rho u ( u + U_{\infty} ) (2\pi r)] - [B q c C_d \sin(\phi)]$**

**Torque' =  $Q' = [ 2 \rho v ( u + U_{\infty} ) (2\pi r) r] + [B q c C_d r \cos(\phi)]$**

**Where:             $B$  = Number of Blades**

**$c$  = Local Chord**

**$\phi = \alpha_i + A \tan(U_{\infty} / \omega r)$    or    $\tan(\phi) = (u + U_{\infty}) / (\omega r - v)$**

**$\alpha_i$  = Induced angle of attack due to  $u$  and  $v$**

#### *Blade Element Equations:*

The blade element formulation is a strip theory method applied to rotor blades. At any given differential blade element, the generated forces are assumed determined from two-dimensional section properties. 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. In many formulations, the lift force is expressed in terms of the lift curve slope and angle of attack. In this formulation the lift is represented by the Kutta-Joukowski relation: Force per unit length =  $\rho (V \times \Gamma)$ . At higher

Reynolds numbers the lift curve slope is nearly constant over the operating range, but operation at Ultra-low Reynolds numbers below 10,000 results in highly nonlinear lift curves. This invalidates the use of a constant lift curve slope.

$$\text{Thrust}' = T' = [B \rho (\omega r - v) \Gamma] - [B q c C_d \sin(\phi)]$$

$$\text{Torque}' = Q' = [B \rho r (u + U_\infty) \Gamma] + [B q c C_d r \cos(\phi)]$$

The trigonometric terms are eliminated in all four relations by utilizing the following:

$$L' \sin(\phi) = \rho (u + U_\infty) \Gamma$$

$$L' \cos(\phi) = \rho (\omega r - v) \Gamma$$

Substitution yields:

*Momentum:*

$$\text{Thrust}' = T' = [2 \rho u (u + U_\infty) (2\pi r)] - [B (C_d / C_l) (u + U_\infty) r \Gamma]$$

$$\text{Torque}' = Q' = [2 \rho v (u + U_\infty) (2\pi r) r] + [B (C_d / C_l) (\omega r - v) \rho \Gamma r]$$

*Blade Element:*

$$\text{Thrust}' = T' = [B \rho (\omega r - v) \Gamma] - [B (C_d / C_l) (u + U_\infty) \rho \Gamma]$$

$$\text{Torque}' = Q' = [B \rho r (u + U_\infty) \Gamma] + [B (C_d / C_l) (\omega r - v) \rho \Gamma r]$$

The actuator surface approach used in the momentum equations assumes constant induced velocities across any particular annulus. The presence of a finite number of blades results 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<sup>[9]</sup>. This is a correction for finite blade numbers based on a cylindrical vortex helices in the wake. Contraction of the wake is not considered. It has the desired effect of driving the circulation to zero at the tip as required for a finite span wing. The  $\phi_{\text{tip}}$  value represents the tip helix angle.

$$\kappa = (2/\pi) A \cos(e^{-f})$$

$$f = (B/2) (1 - (r/R)) (1 / \sin(\phi_{\text{tip}}))$$

This correction is applied the local bound circulation:

$$B \Gamma = \kappa \Gamma_{\infty \text{ blades}}$$

Applying this to the inviscid portions of the two momentum equations yields:



*Momentum:*

$$\text{Thrust}' = T' = [ 2 \rho u (u + U_\infty) (2\pi r) \kappa ] - [B (C_d / C_l) (u + U_\infty) \rho \Gamma ]$$

$$\text{Torque}' = Q' = [ 2 \rho v (u + U_\infty) (2\pi r) \kappa ] + [B (C_d / C_l) (\omega r - v) \rho \Gamma r ]$$

*Blade Element:*

$$\text{Thrust}' = T' = [ B \rho (\omega r - v) \Gamma ] - [B (C_d / C_l) (u + U_\infty) \rho \Gamma ]$$

$$\text{Torque}' = Q' = [ B \rho r (u + U_\infty) \Gamma ] + [B (C_d / C_l) (\omega r - v) \rho \Gamma r ]$$

The last issue is the treatment of the tangential induced velocity term ( $v$ ). The formulation to this point incorporated 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 becomes an important consideration. The thick wake regions generated by each blade produce a significant ‘viscous swirl’ effect. The model used for this phenomena is described later, but for now this term is incorporated by separating the tangential velocity into  $v_{\text{inviscid}}$  and  $v_{\text{viscous}}$ . No viscous correction is applied to the vertical induced velocity. The viscous swirl is proportional to  $\cos(\phi)$  while any viscous downwash term would be proportional to  $\sin(\phi)$  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 momentum equation for torque. The viscous losses are already accounted for in the viscous drag portion of the torque equation. These substitutions result in the final version of the four basic relations:

*Momentum:*

$$\text{Thrust}' = T' = [ 2 \rho u (u + U_\infty) (2\pi r) \kappa ] - [B (C_d / C_l) (u + U_\infty) \rho \Gamma ]$$

$$\text{Torque}' = Q' = [2\rho v_{\text{inviscid}} (u + U_\infty)(2\pi r) \kappa] + [B(C_d / C_l)(\omega r - v_{\text{inviscid}} - v_{\text{viscous}}) \rho \Gamma r ]$$

*Blade Element:*

$$\text{Thrust}' = T' = [ B \rho (\omega r - v_{\text{inviscid}} - v_{\text{viscous}}) \Gamma ] - [B (C_d / C_l) (u + U_\infty) \rho \Gamma ]$$

$$\text{Torque}' = Q' = [ B \rho r (u + U_\infty) \Gamma ] + [B (C_d / C_l) (\omega r - v_{\text{inviscid}} - v_{\text{viscous}}) \rho \Gamma r ]$$

These four relations yield two equations for two unknowns ( $u$ ,  $v_{\text{inviscid}}$ ) for each differential blade element. The other three ‘unknowns’ ( $\Gamma$ ,  $\kappa$ , and  $v_{\text{viscous}}$ ) are all treated as dependent variables. The circulation can be expressed as a function of the known

geometry, 2-D airfoil characteristics, and the local flow velocity. The classical tip loss factor for a cylindrical wake is based on the tip helix angle. Since there is no coupling between blade elements, the tip section may be solved first, yielding a closed set. This tip velocity may then be used in solving the rest of the blade elements.

### 3.2 Uncoupled equations for Rotor Induced Velocities and Solution of the Design Problem

After a lot of manipulation of the thrust and torque equations, two uncoupled equations are obtained for the  $u$  and  $v_{\text{inviscid}}$  induced velocities. This formulation allows for a direct solution of the rotor performance problem. The required inputs are the rotor speed, chord distribution, section characteristics, and lift distribution. The incidence distribution is an output of the method along with thrust, torque, and power required. This provides a simple method for design of a rotor from a blank sheet, but analysis of an existing rotor, where the incidence is known, but the lift distribution is not, forces an iterative solution.

The equations are still closed for the analysis case, but highly coupled. The determination of the  $\Gamma$  distribution becomes an assortment of inverse trigonometric functions of the induced velocities and the incidence angles. In the design case,  $\Gamma$  may be expressed as function of the input lift distribution and the local velocity vector. After a bit of work the following uncoupled equations for  $u$  and  $v_{\text{inviscid}}$  are obtained:

$$u = (-U_{\infty} + \text{Sqrt} [U_{\infty}^2 + 4 v_{\text{inviscid}} (b_1 - v_{\text{inviscid}})]) / 2$$

$$b_1 = (\omega r - v_{\text{viscous}})$$

The equation for the tangential induced velocity is a quartic equation for arbitrary  $U_{\infty}$ , but for the limiting case of hover ( $U_{\infty} = 0$ ) the equation reduces to a quadratic with the following solution:

$$v_{\text{inviscid}} = ((-b_1 (b_2^2)) + (b_1 b_2 \text{Sqrt} [b_2^2 + 4])) / 2$$

$$b_2 = (B / (8 \pi r \kappa)) c C_l$$

With the induced velocities determined, the incidence angle may be calculated using the lift coefficient and lift curve slope of the airfoil. Either set of thrust and torque equations may be used to determine the rotor thrust, torque, and power required. The main program flow-chart is provided in Figure 14.

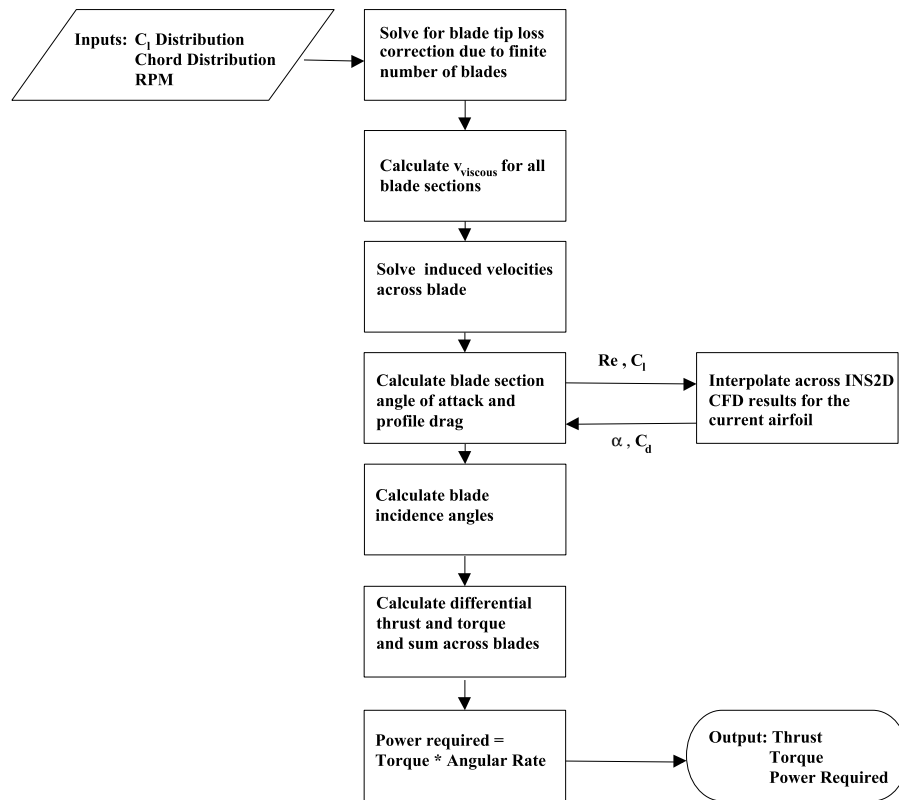


FIGURE 14.

### 3.3 Viscous Swirl Models

The original viscous swirl model described in the NIAC Phase II Interim Report was based on airfoil wake profiles at ultra-low Reynolds numbers computed with INS-2d. A single representative section and angle of attack was analyzed across a range of Reynolds numbers and a model for the mean wake deficit velocity was determined. The resulting model has been found to significantly overestimate the wake deficit velocity, particularly at inboard stations where the Reynolds number due to rotation is lower and the blade spacing is reduced. This model did not account for the velocity distribution across the section wake or the rotor downwash, both of which would have an alleviating effect. This model initially appeared to be adequate and predictions of total rotor thrust that agreed well with the initial experiments for the 15g rotor-craft. Three-dimensional Navier-Stokes analyses have since shown that while the global thrust may be well predicted, the span-wise lift and torque distributions are not, resulting in significant errors in predicting the power required for a given design.

Subsequently this model was modified to account for these two factors. Based on the same INS-2d data, the wake deficit velocity profile was modeled as a Gaussian distribution based on Reynolds number, distance aft from the trailing edge, and distance above the

trailing edge. This distribution was translated downward based upon the local helix angle and blade spacing. The viscous swirl velocity was taken as the value of the Gaussian distribution at the intersection of the translated profile and the next blade's leading edge. This intermediate model also proved unsatisfactory. The application of this model underestimated viscous swirl effects, resulting in negligible corrections for all but the innermost regions of each blade.

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). While empirical corrections could be applied to either model, the loss of generality would diminish their usefulness as design tool components.

The current viscous swirl model<sup>[10]</sup> is consistent with the blade element / actuator disk theory used in the rotor performance model. 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:

$$\begin{aligned} M'_{\text{rotor}} &= B \ D' \ r \ \text{Cos}(\phi) \\ &= B \ q_{\text{local}} \ c \ C_d \ \text{Cos}(\phi) \end{aligned}$$

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

$$\begin{aligned} dH'/dt &= \rho \ v_{\text{viscous}} \ r \ (dV'/dt) \\ &= \rho \ v_{\text{viscous}} \ r \ (2\pi r \ (u+U_{\infty})) \end{aligned}$$

where  $V'$  is the volume of the differential annulus.

Solving for  $v_{\text{viscous}}$  yields:

$$v_{\text{viscous}} = [B \ q_{\text{local}} \ c \ C_d \ \text{Cos}(\phi)] / [\rho \ 2\pi r \ (u+U_{\infty})]$$

From momentum theory (neglecting tip loss corrections):

$$T' = [2 \ \rho \ u \ (u + U_{\infty}) \ (2\pi r)] - [B \ q_{\text{local}} \ c \ C_d \ \text{Sin}(\phi)]$$

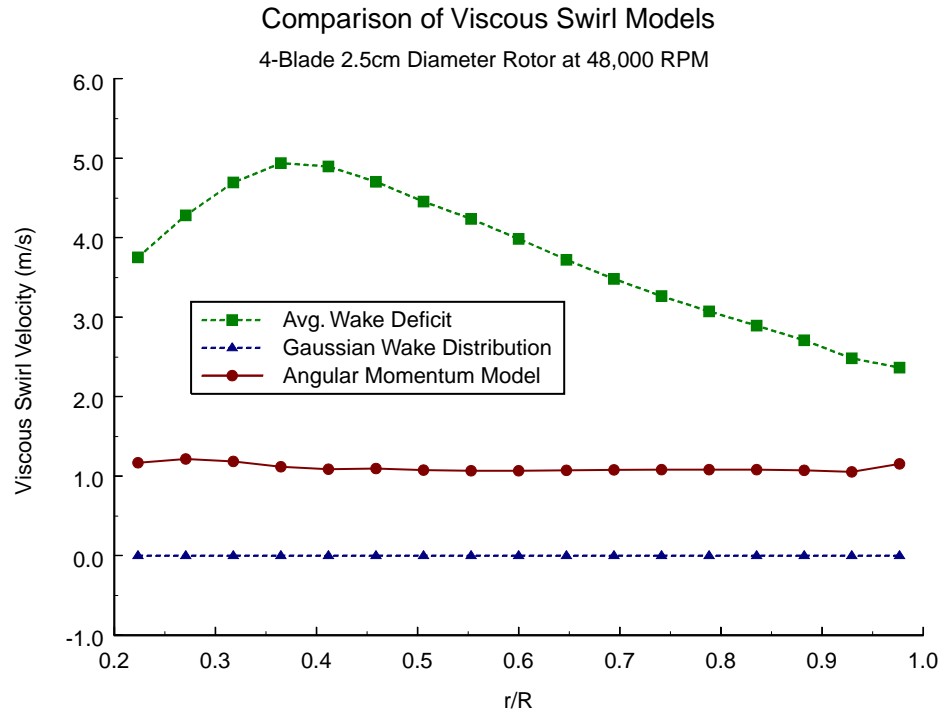
From blade element theory:

$$T' = B q_{local} c [C_l \cos(\phi) - C_d \sin(\phi)]$$

Substituting into the denominator of the viscous equation yields the deceptively simple relation:

$$v_{viscous} = 2u (C_d / C_l)$$

Beyond the assumptions intrinsic in blade element / momentum theory, the only additional assumptions are that lift is inviscid and plays no direct role in the viscous swirl and any tip loss corrections are neglected in the substitution of the thrust equation. Both are reasonable and have a minimal influence on the model. There are no small angle approximations made 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 viscous swirl velocities generated by all three models are compared in Figure 15 for a 4-blade, 2.5cm diameter rotor operating in hover at 48,000 RPM. The swirl velocities generated by the average wake deficit model are too large, causing an excessive reduction in Reynolds number, dynamics pressure, and angle of attack. The Gaussian model has almost no effect and so it was also discarded in favor of the current angular momentum model.



**FIGURE 15.**

### 3.4 Rotor Wake Modeling

Blade-element theory and momentum 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 vortex 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 ( $\kappa$ ) to the infinite blade solution. This model well suited for lightly loaded rotors and rotors with large advance ratios. In these cases, the assumption of a cylindrical wake is accurate. The tendency of the helical wake to contract as it moves downstream is unimportant in the near field either due to lower vorticity or highly pitched helices.

The meso-scale rotor designs typically have a disk loading 2 to 3 times lower than the value for a full-scale helicopter, but high rotor solidity, as much as 30%, coupled with a primary interest in hover performance increases the importance of wake contraction. The first-order effects of wake contraction are captured using a contracting wake model based on an axis-symmetric streamline solution for a vortex ring. Vortex rings are initially stacked as in the Prandtl method, 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 may be expressed in terms of the complete elliptic integrals  $F_1$  and  $E_1$  as:

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

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

**(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 ( $\Gamma$ ) and fixed spacing are determined from the inviscid constant down-wash rotor having equivalent thrust to an inviscid rotor with the specified  $C_1$  distribution:



$$\Gamma = \frac{2\pi T}{\rho\omega B\pi R^2}$$

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

The contracted wake is obtained iteratively by calculating the mass flux through each ring due to the entire wake structure and resizing it 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:

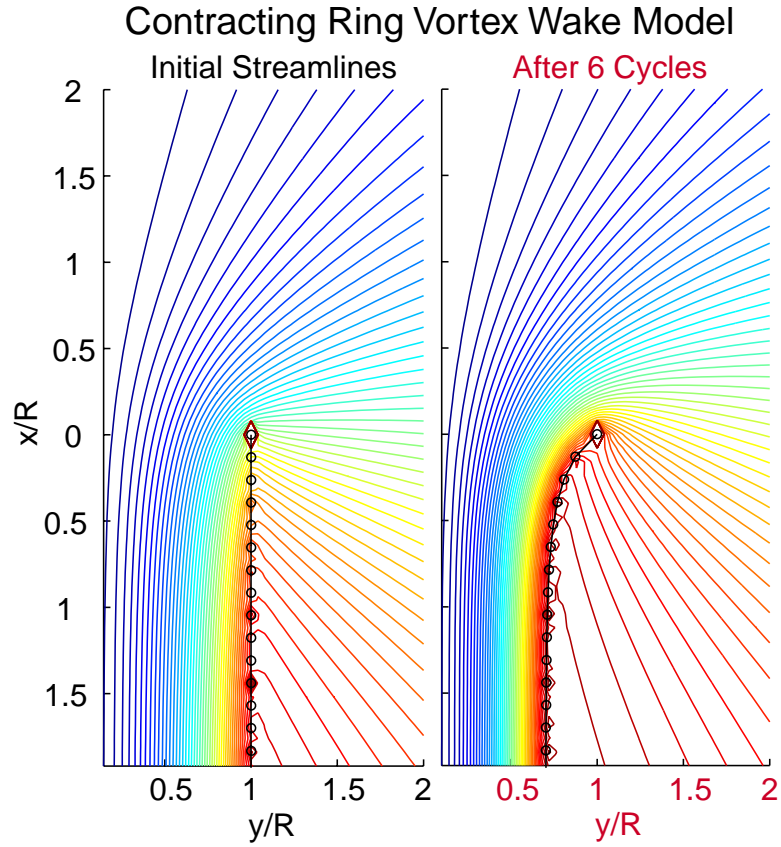
$$S = -2\pi\psi(x, r)$$

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

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

The resized ring is modelling the horizontal component of a helical filament with a modified pitch. The strength of the ring is modified by the ratio of the cosine of the local pitch angle to the cosine of the pitch angle at the rotor.

This procedure is repeated until the wake structure reaches equilibrium. This typically takes only 3 to 5 cycles. The ring model extends downstream 5 rotor radii. The equilibrium wake form is axis-symmetric and symmetric from end to end, with a ‘bell-mouth’ at the stream-tube exit equal to the rotor disk area. The initial and converged ring configuration for a candidate rotor are displayed in Figure 16, but the stream-tube exit is below the lower extent of the figure. The figure shows half of the cutting plane across the rotor diameter with the ring locations signified by black markers and the streamlines in color. The rotor plane is at  $x=0$ . Some flow can still be seen entering the stream tube below the disk. Additional wake length resolves this issue but does not significantly affect the solution and adds to the computational expense.



The rotor inflow velocities are calculated using 2<sup>nd</sup> order central-differencing of the stream-function along the blade. These velocities are utilized to derive and modified  $\kappa$  distribution. Eliminating  $\kappa$  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. Utilizing a corrected  $\kappa$  allows the classically derived equations to be used in an unmodified form. The assumptions made are consistent with the wake model. Beginning with the inviscid momentum and blade element thrust equations:

$$\mathbf{T}' = [ 2 \rho \mathbf{u} ( \mathbf{u} + \mathbf{U}_\infty ) (2\pi\mathbf{r}) ] \quad \text{(Momentum)}$$

$$\mathbf{T}' = [ \mathbf{B} \rho ( \boldsymbol{\omega}\mathbf{r} - \mathbf{v} ) \Gamma ] \quad \text{(Blade Element)}$$

These may be solved for  $B\Gamma$  and equivalently  $\Gamma_\infty$  :

$$B\Gamma = \Gamma_\infty = \frac{4\pi r u (u + U_\infty)}{(\boldsymbol{\omega}\mathbf{r} - \mathbf{v})}$$

Here  $u$  is the constant down-wash velocity predicted by actuator disk theory:

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

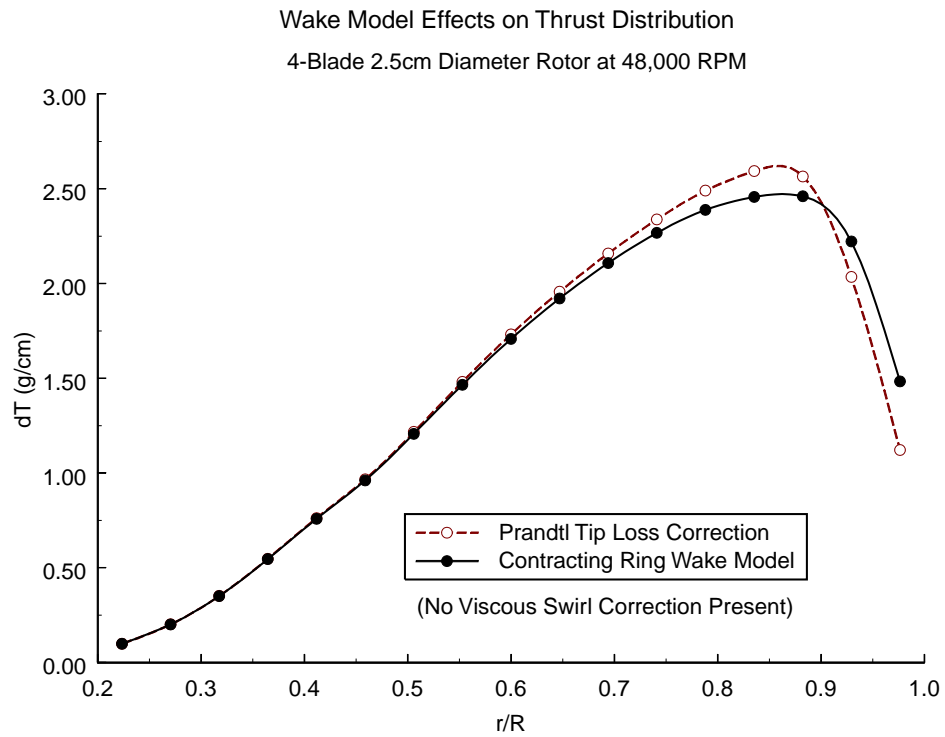
Using the down-wash velocities ( $u'$ ) calculated from the contracted wake model the corrected value of  $B\Gamma$  may be expressed as:

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

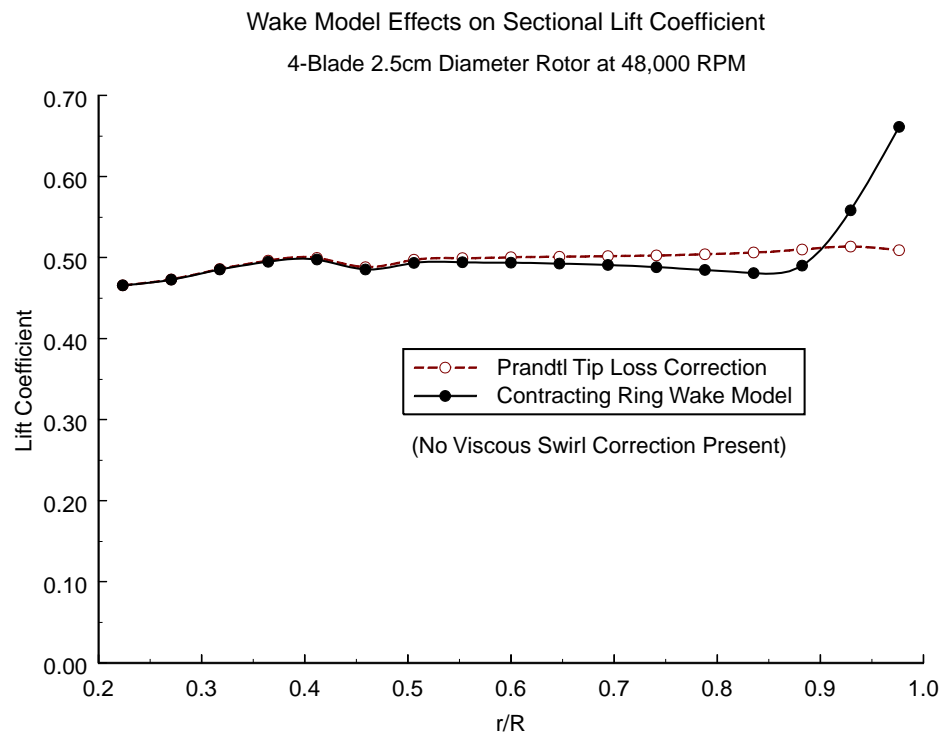
The final expression for  $\kappa$  simplifies to:

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

The primary effect of wake contraction is a shift in the blade loading towards the tip. Predicted thrust distributions comparing the Prandtl and contracted wake models are shown in Figure 17. The contracted wake model generates more thrust on the outboard 10% of the rotor blade, but there is little difference in the global thrust and torque values. The important aspect of incorporating the contracting wake is in predicting tip stall during analysis. The lift coefficient distributions for the same two cases are provided in Figure 18. The contracted wake predicts the tip is operating at  $C_l=0.66$ , equal to the  $C_{l_{\max}}$  value at  $Re=6000$ , while the Prandtl result predicts  $C_l=0.51$ . Three-dimensional RANS analysis of this rotor confirms that the tip is on the verge of stall. In general, the high profile power and low efficiency of stalled sections make it important to predict this phenomena during the design stage.



**FIGURE 17.**



**FIGURE 18.**

## **4.0 Rotor Design and Analysis Utilizing Non-Linear Optimization**

The rotor analysis method described is only capable of estimating the performance of a given geometry and operating condition. Alone, it is incapable of developing and improving a rotor design for a particular application. The analysis 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.

Complete rotor analysis and design tools have 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<sup>[11]</sup>. In the design mode, the goal is to maximize thrust for a given radius and for a particular motor. 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.


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 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, and only, solution that also satisfies the incidence constraints provides the correct lift distribution for the analysis of rotor performance. There is no constraint on the power required since in the analysis mode it is assumed that the case represents a physically realizable system. 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.

## **5.0 Three-Dimensional Rotor Analysis using OVERFLOW-D**

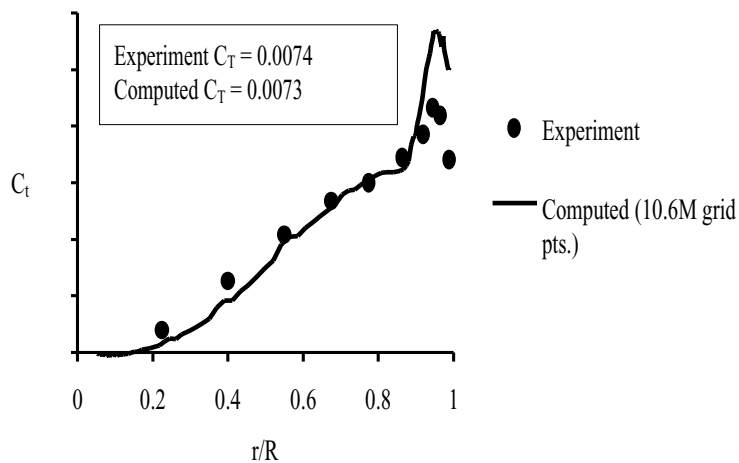
The Reynolds-averaged Navier-Stokes flow solver used to validate the results from the low-order design method is based on a version of the OVERFLOW code developed by Buning et al<sup>[12]</sup>. OVERFLOW is a general-purpose Navier-Stokes code designed for fixed wing overset-grid computations on static grids. Meakin<sup>[13]</sup> has generalized 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. This group-wise structure has been exploited to facilitate efficient parallel computations of multi-body problems on scalable computer platforms<sup>[14]</sup>. 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 and has been extensively

tested on IBM-SP and SGI Origin 2000 parallel computers. The solution of hovering-rotor problems requires a number of modifications to the OVERFLOW-D flow solver. Reference 15 describes these modifications in detail. Post-processing of the sectional and global rotor forces uses the FOMOCO force integration code<sup>[16]</sup>.

This modified version of OVERFLOW-D has been validated with experimental data for a model UH-60A rotor. Figure 19 compares the experimental and computational global thrust and torque coefficients. Figure 20 compares the computed and sectional thrust distribution for the same model.




**FIGURE 19.**



**FIGURE 20.**

The case considered is the 4-blade 2.5cm diameter rotor operating at 48,000 RPM at sea level. The computational intensity 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 for this case, the computational grid still contains over 5 million grid points. The overset near-body grids comprise 40% of the cell count and are divided into four subsets consisting of the hub, blade root / hub collar, blade, and blade tip. The resulting rotor model is pictured in Figure 21 with a superimposed color-map of surface pressures at the design point. The off-body cartesian grids use three levels of spacing refinement for the outside of the domain inwards, with the innermost domain matching the outer cell spacing of the near-body volume grids.



**FIGURE 21.**

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 1mm high 2.5mm diameter disk as a hub. In use, the underside of the hub is effectively faired by the presence of the 5mm Smoovy direct drive motor. The flat upper and lower surface 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. A hub vortex model has been added the analysis method based on the total circulation, but this does not fully address the hub's viscous entrainment. A semi-infinite vortex line is extended from the rotor-plane downward along the rotor axis. The strength of the filament is set equal to the total circulation required for a given thrust as calculated for the vortex ring strength in the contracting wake model. It is chosen to be qualitatively correct and physically consistent with the wake model used for analysis and design.

## **6.0 Experimental Testing**

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 milligram-accurate balance on the other. It is supported and attached in the middle by a shaft and ball-bearings providing one axis of rotation. Upward thrust is transferred to a downward force on the balance that can be read directly. The test equipment configured for thrust measurement is pictured in Figure 22. Rotor torque is 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. The ability to measure rotor torque is limited by the accuracy of the electronic balance and physical constraints that limit the minimum length of the moment arm. This has only allowed torque measurement for the larger rotors.





**FIGURE 22.**

The evaluation of motor performance is accomplished by coupling the shafts of two motors. The one being tested is mounted to the test stand in the same manner as for rotor torque testing. The second motor acts as a torque load and is mounted in a second fixture such that the drive shafts are co-linear. The motor being tested 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.

## **7.0 Results and Discussion**

Over the course of this NIAC research program a variety of rotors have been fabricated and tested in the development of the four different rotor-craft concepts, all operating at low to ultra-low Reynolds numbers. The concepts included the smallest 5g-10g vehicle, a 15g vehicle, 60g vehicle, and the largest, a 150g vehicle that has built and flown under its own power. In addition to these four prototypes, experiments have been completed exploring the feasibility of using rotors for thrust under Martian atmospheric conditions. This work was carried out jointly by Stanford University and NASA-JPL in Pasadena, California. A summary of the rotors fabricated and tested is provided in Table 1.

## Summary of Rotor Fabrication and Experimental Performance

Powerplant	Rotor Details	# of Blades	Dia. (cm)	Voltage (V)	Current (A)	Power Input (W)	Thrust (g)	Design RPM	Notes
3mm Smoovy		3	NA	8	NA	NA	0.75	50,000	
5mm Smoovy	Cast polyurethane SDM fabrication	5	NA	NA	NA	NA	NA	NA	Tests Aborted. Poor performance due to lack of structural stiffness. Limited by 150mA controller rating. RPM indicated is limited by 150mA controller rating.
	Baseline 5-Blade Rotor	5	2.2	8.7	0.15	1.305	1.3	45000	
	Cast epoxy SDM Fabrication								
	Baseline + 4 deg. Incidence	5	2.2	9	0.15	1.35	1.235	45000	
	NACA 4402 camberline								
	Baseline + 6 deg. Incidence	5	2.2	8.8	0.16	1.408	1.84	44800	
	NACA 4402 camberline								
	Baseline + 8 deg. Incidence	5	2.2	8.4	0.16	1.344	1.85	35220	
	NACA 4402 camberline								
	Baseline + 3 deg. Incidence	5	2.2	8	0.15	1.2	1.75	36000	
	NACA 4402 camberline								RPM indicated is limited by 150mA controller rating. Controller overloaded to reach design RPM. RPM indicated is limited by 150mA controller rating. Controller overloaded to reach design RPM. Tests Aborted. Inconsistent performance, lack of structural stiffness. RPM indicated is limited by 150mA controller rating. Controller overloaded to reach design RPM.
	Optimized design using current toolset	4	2.5	10	0.19	1.9	2	41500	
	Re 6000 optimized airfoil								
	Cast epoxy SDM fabrication								
	Optimized design using current toolset	4	2.5	8	0.15	1.2	2.4	34000	
	Increased hub diameter								
	Otherwise same as above								
		4	2.5	12	0.24	2.88	3.7	47850	
Mabuchi B-2 Motor	Cast polyurethane SDM fabrication	2		6	0.76	4.56	21	17,800	Operating condition for 60g Test Vehicle.
Westec DC-5	JPL Low Pressure Test Rotor for Mars Applications	5	20	NA	NA	NA	NA	NA	Fabrication method is unsuitable. Rotor was porous and unbalanced. Rotor suffered structural failure during testing.
	FDM Rapid Prototyping using ABS-400								
	JPL Low Pressure Test Rotor for Mars Applications	5	20	Approx. 6.0	Approx. 1.0	Approx. 6.0	Approx. 4-5 @ 11 Torr	unknown	Improved rotor, well balanced. Low pressure (11 Torr) testing was qualitative. Results indicate that rotor thrust is possible under such conditions.
	Carbon / Kevlar Composite Construction								
Astrolight Freely 16:1 Gearbox	Rotor designed for the 150g prototype aircraft	2	25.4	8.3	0.33	2.74	50	1920	
	Carbon / Kevlar Composite Construction								

TABLE 1.

A large portion of the aerodynamic research and development program has centered on the 2.5cm, 4-blade rotor for the 15g vehicle, chosen as a reasonable target for 5 to 10 years past the current state-of-the-art. This rotor is used in the remaining sections to discuss the implementation and accuracy of the developed design and analysis tools.

## 7.1 Development and Testing of the 4-Blade 2.5cm Diameter Rotor

### 7.1.1 Power plant Selection and Modelling

The rotors of the 15g rotor-craft is powered by the 5mm Smoovy 3-phase stepper motor. 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 designer is free to generate this curve whatever fashion suits the application. The simplest to generate experimentally are curves based on a constant input voltage. Multiple curves may be generated across the range of permissible motor voltages. A single composite curve may also be constructed based on a desired input current or input power. The best method depends on the limiting factor, current or voltage, of the power storage and motor systems.

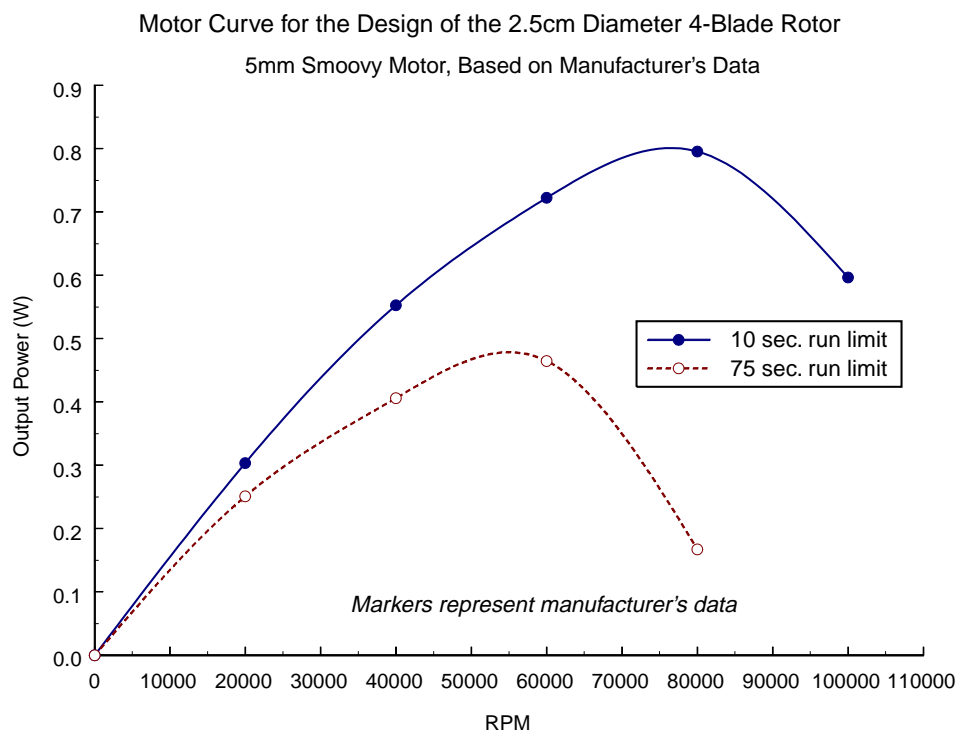


FIGURE 23.

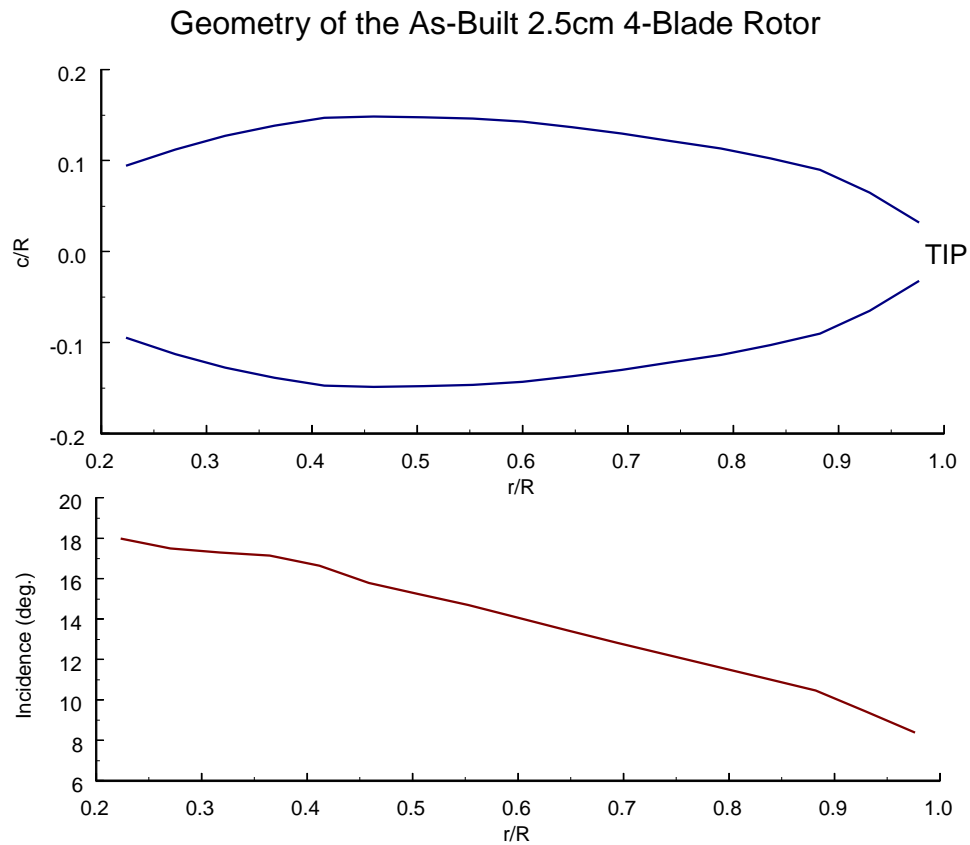
The output power / RPM curves shown in Figure 23 are based on manufacture's data for a maximum 10 second and 75 second run time. These motors provide very good performance both in terms of torque production and high rotational speeds. One limitation of this type of motor is the need for additional circuitry to cycle power through the motor coils. This adds mass to the 15g vehicle and in the final configuration represents a mass fraction of 20% to 30%. Simple circuitry is required to run the motor without feedback, but efficiency is greatly reduced (high efficiency is benefit of this motor) and the

maximum RPM under load drops considerably, in this case from over 50,000 RPM to 25,000 RPM. Feedback is required to account for phase lags in the coil/shaft position due to the rotational mass of the shaft and rotor, and aerodynamic drag. The feedback circuitry does not add considerably to the mass, but the complexity of the circuitry increases. Small traditional or coreless DC motors, such as pager motors, were not a viable alternative. Although they require no additional circuitry, these motors provided insufficient torque and rotational speeds. Their efficiencies are also much lower.

### **7.1.2 Design of the 2.5cm 4-Blade Rotor**

This particular rotor for the 5mm Smoovy motor was designed using a previous version of the design code. It incorporated the wake-deficit viscous swirl model and the Prandtl tip-loss correction. The diameter and blade count are based on a parametric study completed with the same code. The goal was the smallest rotor capable of generating 4g of thrust using the 5mm Smoovy motor. Subsequent 3-D CFD analyses, leading to modification of the rotor code, occurred later and sufficient time was not available to manufacture and test an updated design.

The airfoil is the 2-D optimized design for  $Re=6000$ . The maximum lift coefficient from INS-2d calculations at  $Re=10,000$  is 0.62, increasing to 0.66 at  $Re=6,000$ . The operating 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). Lift coefficients are bounded to be positive and below 0.5. 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 and incidence distributions are shown in Figure 24.




**FIGURE 24.**

### **7.1.3 Predicted Performance and Experimental Testing**

The predicted thrust of this rotor using the original design code was 3.90g thrust at 48,000 RPM requiring 0.42W of power. Re-analysis of the rotor with the current version of the code, incorporating the angular momentum swirl model and the contracting wake model, results in a predicted thrust of 4.38g at 47,400 RPM requiring 0.45W of power. The change in RPM is due to motor matching. If the motor is ignored, Re-analysis at 48,000 RPM yields 4.49g of thrust requiring 0.47W. Experimental thrust testing of several rotors yielded thrust values 15% to 20% lower than predicted at any given RPM. In Figure 25, these results are compared with the predicted performance both from the current low-order analysis method and a single 3-D CFD solution.

Two factors, manufacturing errors and the design geometry, both come into play to account for this discrepancy. The experimental results for three rotors show a 10% to 15% variation in thrust at a given RPM. This is indicative of incidence variations induced during manufacturing or assembly and testing. Variations as small as a few tenths of a degree in incidence will have a large effect on rotor performance. The design of the rotor using an older version of the design tools resulted in the tip of the blade operating at or

near stall. This was described in a previous section and is supported by the 3-D CFD analysis. This is accounted for in the current design and analysis tools only as a large increase in the tip section drag coefficient. In reality the stall of the tip will have a wider effect on the rotor performance, affecting the lift and drag of nearby sections and influencing the wake formation. The rotor is also designed in isolation, absent any effects from the hub or adjoining motor beyond the hub vortex model described earlier. The aerodynamic issues can be explored as an avenue for further research, but establishment and maintenance of dimensional accuracy in the manufactured part poses the greatest difficulty at such small scales and is a significant obstacle to further development.



**FIGURE 25.**

#### 7.1.4 Aerodynamic Modelling Effects in the Current Toolset

The 2.5cm diameter 4-blade rotor has been redesigned using the latest version of the design and analysis software. This provides insight into the effects of several aspects of the aerodynamic modelling on predicted rotor performance the resulting final rotor geometry. The primary modifications are the addition of the contracting wake model and the replacement of the previously described wake deficit swirl model with the model based on the conservation of angular momentum in the wake.

The re-designed geometry is superimposed over the original in Figure 26. The most obvious changes are in increase in inboard chord-lengths and a reduction in tip incidence.

The chord increase is due to the change in swirl models. The overestimation of viscous swirl velocities resulted in large reductions in the predicted local velocity and Reynolds number. The viscous swirl velocity increases as the spacing between blades is reduced, so the effect is more pronounced on inboard stations. The optimization process was forced to reduce chords inboard to keep the viscous swirl model from bringing the flow nearly to rest in the rotating frame of reference. Replacing this model with the angular momentum model reduces the effect of viscous swirl, resulting in an increase in the inboard chord-lengths, incidence and loadings. In fact, this new model has very little influence on the solution compared to excluding the viscous swirl model entirely. Excluding viscous swirl effects results in a small increase in chord at the widest part of the blade, but the incidence distribution and the rest of the chord distribution see only negligible changes. This model results in relatively constant viscous swirl velocities along the span of the blade and the overall effect on diminishes rapidly from blade root to tip.

The reduction in tip incidence is due to the contracted wake model. The contraction of the wake results in a decrease in the total downwash near the tip, increasing the local angle of attack for a given incidence. In the original design this resulted in the blade tip operating at or near stall. The contracted wake model may be applied directly for the analysis of a given blade, but for design direct application is problematic. The model represents the distributed shed vorticity as discrete vortex rings. If applied directly to a design case, the optimizer exploits this difference in the near-field of the blade and produces chord distributions which are not physically realistic. To work around this, a rotor is initially designed using the Prandtl correction, then the incidence and RPM are optimized with the contracted wake model active and the chord distribution fixed.

Comparing both rotors after re-matching RPM to the 5mm Smoovy, only slight differences in performance are predicted by the latest version of the software. The changes in the aerodynamic models and the large differences in the chord distribution result in a predicted 0.5% increase in thrust and a predicted 1.0% reduction in power required. As discussed earlier, this does not fully account for the tip stall in the original design. Although the viscous swirl model may have negligible impact, the reduction in tip incidence mandated by the contracting wake model is an important design feature, potentially having a strong effect on actual performance.

## Geometry Comparison of the 2.5cm 4-Blade Rotors

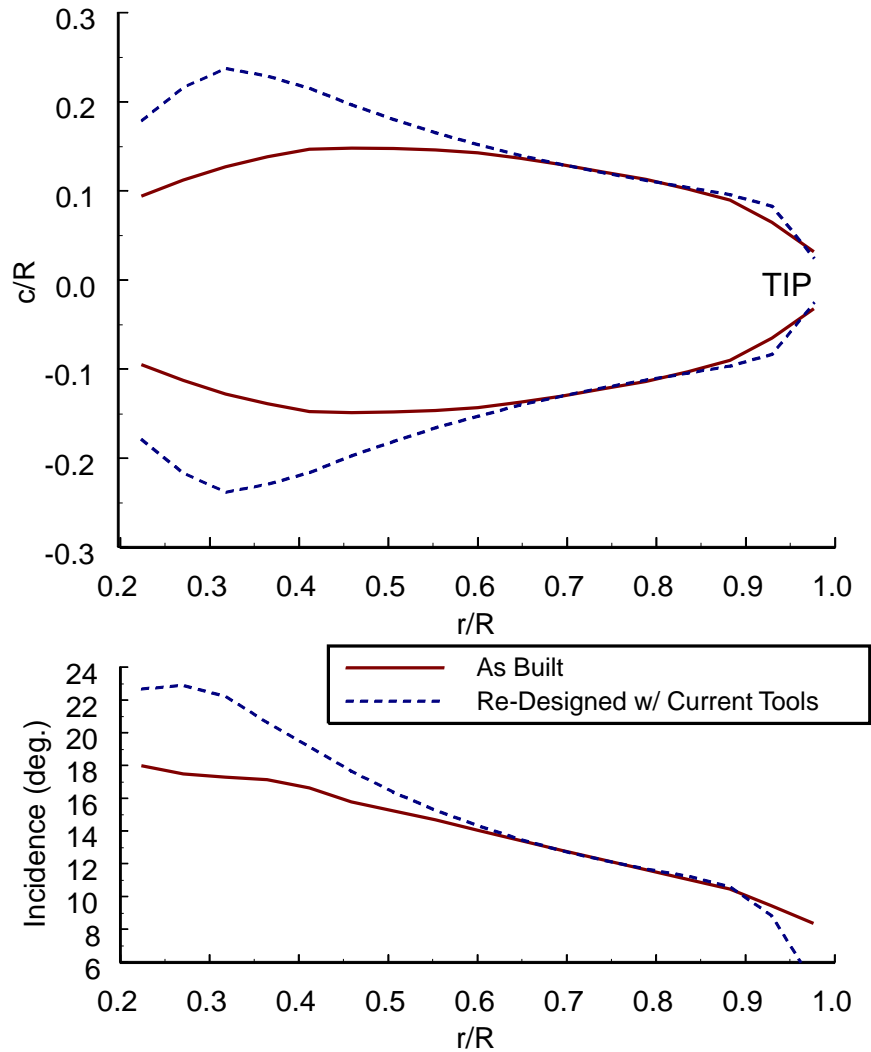


FIGURE 26.

### 7.1.5 Overflow-D Results

The results of the analysis with the Overflow-D code indicate that the developed low-order methods do provide a reasonably accurate means of rapid rotor design and analysis, even at ultra-low Reynolds numbers. The 2.5cm 4-Blade rotor was analyzed at the original operating point of 48,000 RPM. The analysis code compared to the 3-D CFD predicts 9.2% greater thrust at 7.5% less power. The spanwise thrust and torque distributions are shown in Figures 27 and 28. In both figures the low-order distributions are very similar to the CFD results, but the 3-D consistently predicts lower thrust and higher torque across the inner 80% of the span. Larger differences are visible in the tip region due to the complex effects of the tip-stall captured by the CFD but largely unaccounted for in the design and analysis codes.






FIGURE 27.





FIGURE 28.

Blade-vortex interactions and strong wake contraction in the near-field are also evident in the CFD results. This is visible in the vorticity contour plot shown in Figure 28. The cutting plane is directly in front of the blade advancing into the page. The tip vortex from the blade ahead is seen to impact the blade at about 85% of the span. In the design and analysis code, the general effects of contraction are captured by the latest wake model, but the close proximity of the vortex and blade introduce additional complexities such as high local velocities, separation, and vortex breakdown that are not easily accounted for. Eliminating the tip-stall might alleviate this strong interaction, but it is difficult to draw a conclusion from this single calculation.

Additional calculations are warranted, but the expense of a single 3-D CFD calculation is enormously large compared to the low-order tools that have been developed. A converged solution required 24 hours of run time on 16 processors where analysis with the developed toolset takes less than 30 seconds and design less than 5 minutes on a typical desktop computer. This makes iterative design prohibitive using 3-D CFD and is a principal motivation for developing this lower-order method.



**FIGURE 29.**

## **8.0 Conclusions**

The aerodynamics research completed under this NIAC Phase II study has increased the knowledge base both in general applied aerodynamics at ultra-low Reynolds

number and in the particular area of rotor design for this flight regime. Operating at ultra-low Reynolds numbers results in very high drag coefficients. 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. The study of NACA 4-digit sections indicates that at ultra-low Reynolds numbers, geometry variations still have a strong effect on the aerodynamic performance of an airfoil. The impetus for design at these Reynolds numbers is further strengthened by the performance gains achieved with a simple airfoil optimization study.

The rotor design studies indicate that classical momentum / blade element methods, properly modified for ultra-low Reynolds numbers, provide a reasonably accurate and efficient rotor design and analysis tool for this unique regime. In the analysis role inclusion of lower order effects such as viscous swirl have little effect on global performance metrics, but during an optimization based design process these subtle changes may lead to distinctly different solutions. The judicious use of two-dimensional Navier-Stokes analyses in the construction of airfoil performance databases, in conjunction with classical rotor theory, results in a hybrid method that suffers a modest reduction in fidelity, but tremendously reduces computational cost and design cycle times relative to three-dimensional Navier-Stokes analyses.

## 9.0 References

- [1] Kunz, P.J., Kroo, I., "Analysis and Design of Airfoils for Use at Ultra-Low Reynolds Numbers," *Conference on Fixed, Flapping, and Rotary Vehicles at Very Low Reynolds Numbers*, Notre Dame, IN, June 2000.
- [2] Chorin, Alexandre Joel, "A Numerical Method for Solving Incompressible Viscous Flow Problems," *Journal of Computational Physics*, Vol. 2, No. 1, 1967, pp. 12-26.
- [3] Rogers, S. E. and Kwak, D., "An Upwind Differencing Scheme for the Steady-state Incompressible Navier-Stokes Equations," NASA TM 101051, November 1988.
- [4] 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.
- [5] Thom, A., Swart, P., "The Forces on an Aerofoil at Very Low Speeds," *Journal of the Royal Aeronautical Society*, Vol. 44, 1940, pp. 761-770.
- [6] Drela, M., Giles, M. B., "ISES - A Two-Dimensional Viscous Aerodynamic Design and Analysis Code," AIAA Paper 87-0424, January 1987.
- [7] 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.
- [8] 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.
- [9] McCormick, Barnes W. Jr., *Aerodynamics of V/STOL Flight*, Dover Publications Inc., Mineola, New York, 1999, pp. 82-84.
- [10] Johnson, W., *Helicopter Theory*, Dover Publications Inc., New York, 1994, pp.44-45.

- [11] 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.
- [12] Buning, P. G., Jespersen, 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.
- [13] 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.
- [14] 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.
- [15] Strawn, R. C. and Ahmad, J. U., "Computational Modeling of Hovering Rotors and Wakes," AIAA 2000-0110, AIAA 38th Aerospace Sciences Meeting, Reno, NV, January 10-13, 2000.
- [16] Chan, William M., and Buning, Pieter G., "User's Manual for FOMOCO Utilities- Force and Moment Computation for Overset Grids," NASA TM-110408, July 1996.

# Chapter 3

## Rotor Fabrication

### 1.0 Introduction

Three-dimensional smooth surfaces and small blade thickness are the main challenges in mesoscale rotor fabrication. Due to the relatively low weight of rotors with respect to the other components and their better strength-to-weight and stiffness-to-weight ratios on the smaller scales, we were able to consider a variety of materials for their fabrication. This section presents the process of material selection and associated manufacturing processes. We found that among the engineering materials we considered, polymers and aluminum alloys stood out and we further investigated them. Before the fabrication, we needed to construct solid models of rotors based on the design specifications. We determined the appropriate processes for polymer rotor fabrication, experimented on material combinations, applied special strategies to improve manufacturing, and tested the lift of fabricated rotors. We followed the same procedures for aluminum rotors.

### 2.0 Challenges and Opportunities

The rotor is the most challenging part of the mesoscopic flying vehicle. At the mesoscale, the rotor operates at low Reynolds numbers and in a viscous-dominated flow regime that requires unconventional approaches to rotor design. From a manufacturing point of view, the fabrication of rotors that fulfill the requirements of aerodynamic design is also very difficult. This section discusses the general requirements for the rotor's aerodynamic design. Based on these requirements, challenges in manufacturing and opportunities for potential solutions are identified.

When a body is moving in a fluid, the pressure and shear distributions on its surfaces result in lift and drag forces. In the rotor geometry, chords distribution, incidences, and airfoil sections determine these forces. As we integrate all design specifications, the resultant rotor becomes a fully three-dimensional thin structure. In each blade, a surface that transits from one cross section to another should be smooth and continuous, and the thickness varies along the radial direction (see Figure 1). On a large scale these features may not cause problems in manufacturing. However, as parts scale down, the combination of small thickness and three-dimensional shape challenges the thickness limit in traditional molding techniques such as injection molding for plastic parts. For example, in a 22mm diameter rotor where an NACA4402 airfoil section is used, the chords are between 1.1mm and 5.3mm. The thickness of the blade, which is 2% of the length of the chords, will range from 22  $\mu\text{m}$  to 106  $\mu\text{m}$ . This thickness range is smaller than what injection molding can achieve.




**FIGURE 1. The rotor models, which show 3D surfaces and small thickness**

Given the difficulty, an obvious question is whether the blade can be thicker. The results in CFD show that it cannot<sup>[1]</sup>. According to the performance investigation on NACA 4402 and 4404 airfoil sections, an increase in thickness from 2% to 4% results in a 15% drag penalty around maximum L/D (lift-to-drag ratio) and in a similar reduction in the magnitude of the maximum L/D. Therefore, it is not wise to compromise on thickness. Micromachining techniques seem to offer a solution, since MEMS techniques or silicon processes are able to achieve small thickness and small parts. Unfortunately, most of the micromachining methods are only 2.5-dimensional. That is, they introduce steps between layers (see Figure 2), which means they are incapable of producing smooth, three-dimensional surfaces that satisfy aerodynamic performance. In summary, these two requirements for rotor design: fully three-dimensional smooth surfaces and small blade thickness, are major concerns in selecting materials and manufacturing processes.

Besides the geometry requirements, several other material issues need to be addressed in manufacturing. The weight budget is the most important concern in a small flying vehicle; hence, the material should be as light as possible. On the other hand, the thin-blade rotor spins at high RPM and undergoes severe loading conditions, including aerodynamic forces, torques, and centrifugal forces. Therefore, the strength and stiffness of rotor material should be sufficient to retain the geometry under loading and to survive during fabrication, handling, and assembling. Moreover, candidate materials are highly dependent on feasible manufacturing processes. In other words, the material selection and process selection should be integrated.

The challenges of manufacturing rotors at this scale are considerable, but opportunities for solutions exist. The total volume of rotors is relatively low. In the example of a five blade rotor, the total volume, including the hub, is 23.65mm<sup>3</sup>. The weight of different densities for this volume is shown in the following Table 1 for a number of materials. Compared to the total weight of the mesicopter system, 15g, the increases in the weight of these different rotor materials are negligible. Hence, we were not restricted to using traditional light-weight materials. Moreover, the combination of materials and fabrication processes is also flexible.



**FIGURE 2. Steps between layers**

**TABLE 1. Density and weight of various materials**

	Metal (Al alloy)	Polymer (Epoxy)	CFRP	Ceramics ( $\text{Si}_3\text{N}_4$ )
Density ( $\text{g/cm}^3$ )	2.7	1.3	1.5	3.2
Weight (g)	0.064	0.031	0.035	0.076

### 3.0 Material and Process Selection

The first step of fabrication is to find suitable materials and processes that satisfy our criteria. The procedure includes narrowing down the selection criteria, identifying candidate materials that fit the criteria, evaluating the possible manufacturing processes for these candidates, and finalizing decisions.

The materials of mechanical and structural engineering fall into nine broad classes<sup>[2]</sup> — engineering alloys, engineering polymers, engineering ceramics, engineering composites, porous ceramics, glasses, woods, elastomers, and foams. In the mesoscale rotor application, engineering polymers, engineering alloys, and engineering ceramics are the main categories to be considered. The selection criteria, in general, can be separated into bulk mechanical properties, bulk non-mechanical properties, surface properties, production properties, and aesthetic properties (see Table 2). In particular, we are interested in density, stiffness, and strength when the concern is bulk mechanical properties, and ease of manufacturing, fabrication, handling, and finishing when the concern is production properties. Moreover, time and cost in fabrication are also taken into account. Since manufacturing approaches are the most important concern, the production properties are considered first priority, but the mechanical properties must be sufficient for operation and handling.

#### 3.1 Polymers

Engineering polymers can be shaped and formed from solid or liquid raw materials. The solid raw materials, such as pellets or powders, are melted just before the shaping process, while the liquid materials are mainly used in casting or in making reinforced plastic parts. Polymers are also available as sheet, plate, rod, and tubing, which may be formed into various products. Among these shaping and forming techniques, three possible processes

can be used to shape complex parts—parts can be machined from available bulk materials, can be injection molded, or can be cast. Since the rotor is three dimensional, the machining approach will require machining of both surfaces of the part; that is, an accurate indexing system and a special fixture are necessary. Injection molding is the most popular process for complex shapes and high dimensional accuracy. However, the high cost of dies<sup>[3]</sup>, ranging from \$20,00 to \$200,000, is not economically favored in low-volume production. Moreover, essentially, the smallest thickness that dies can achieve is above the range of the blade thickness required for mesoscale rotors. Casting is a feasible approach if suitable molding techniques are applied.

**TABLE 2. Classes of Property<sup>[4]</sup>**

Bulk Mechanical Properties	Price and availability; Density; Modulus and damping; Strength; Hardness; Fracture toughness; Fatigue strength; Creep strength
Bulk Non-mechanical Properties	Thermal properties; Optical properties; Magnetic properties; Electrical properties
Surface Properties	Oxidation and corrosion; Friction, abrasion and wear
Production Properties	Ease of manufacturing; Fabrication; Joining; Finishing
Aesthetic Properties	Appearance; Texture; Feel

Polymers melt or cure at lower temperatures (110-330°C)<sup>[5]</sup> than metals, and hence are easier to handle and require less energy in fabrication. The density of polymers is between 0.9-2 g/cm<sup>3</sup>, which is lower than that of metal alloys or ceramics. The modulus of elasticity (E) and ultimate tensile strength (UTS) of thermoplastics and thermosets are listed in Table 3.

**TABLE 3. Material properties of polymers<sup>[6]</sup>**

Polymer Type	E (GPa)	UTS (MPa)
Thermoplastics	1.4-3.4	7-80
Thermosets	3.5-17	35-170

## 3.2 Metals and Alloys

To use metals and alloys as parts material for mesoscale rotors, the manufacturing approaches can be either material removal by machining or material deposition. In machining, both surfaces of the rotor have to be dealt with, requiring accurate alignment and a special fixture. Material deposition means to add onto substrates materials that can be removed afterward. The deposition approach is widely used in micro-fabrication and rapid prototyping. In mesoscale projects, possible depositions are sputtering by physical vapor deposition (PVD) and electroplating. The thickness that sputtering can attain is very limited (less than 10 μm). Electroplating, on the other hand, can deposit a thick layer. However, most materials that can be electroplated are pure metals, which lack the superior



mechanical properties of alloys. Besides, both methods result in uniform thickness of the deposited material; hence, neither the thickness distribution nor the shapes at the edges can be controlled unless additional electropolishing or EDM is applied. Machining is therefore a better option, although the need for proper alignment and fixture may increase the fabrication time and cost.

Material removal being the preferable approach for metals and alloys, machinability becomes the critical issue in selecting alloys. Light-weight metals in particular, beryllium, magnesium, titanium alloys and aluminum alloys, are promising candidates. The machinability of beryllium is similar to that of cast irons, but beryllium is more abrasive and is toxic. It requires machining in a controlled environment. Magnesium is very easy to machine and has a good surface finish. However, special care must be exercised due to its high rate of oxidation and the danger of fire. Titanium and its alloys have poor thermal conductivity, which can cause significant temperature rise and built-up on edges, and are therefore difficult to machine. Consequently, aluminum alloys are the best candidates from this list. Their density, modulus of elasticity (E), yield stress (Y), and ultimate tensile strength (UTS), are listed in Table 4.

**TABLE 4. Material properties of aluminum and its alloys**

	Density (g/cm <sup>3</sup> )	E (GPa)	Y (MPa)	UTS (MPa)
Aluminum (pure)	2.71	70	20	70
Aluminum Alloys	2.63-2.82	70-79	35-550	90-600

### 3.3 Ceramics

Ceramics are stiff and strong materials but are weak under tension. Ceramics parts are mainly used when light weight, high strength, and a high-temperature operating environment are required. The density of ceramics is between 2.3 and 5.5 g/cm<sup>3</sup>, which is not superior to that of polymers or aluminum alloys. The rotor normally operates at room temperature, and the blades do not need to be especially strong. Therefore, we do not benefit significantly from using ceramics for the rotor. Besides, manufacturing is also an important issue. Molding and casting are the most common processes to shape ceramics. Ceramic green parts that are just taken out of the molds are fragile and require sintering before they turn into dense ceramic parts. After sintering, grinding or lapping may be performed to remove surface flaws and to improve surface finish and tolerances. The sintering process requires careful temperature and pressure control in order to obtain successful parts. Shrinkage and warpage happen during sintering, thus the shrinking effect should be compensated for in the mold design. The amount of shrinkage is usually determined experimentally and varies with materials. Moreover, a mold for thin features may trap air bubbles inside the mold and may also reduce fluidity during casting. Obviously, the fabrication of ceramic parts is time consuming, and every sub-step may cause fatal failure if it is not carefully controlled. Therefore, we considered ceramic rotors


only where we found that other approaches did not work.

### **3.4 Final Decisions**

From this preliminary analysis, polymers are considered the top choice for rotor prototypes. Aluminum alloys are stiffer and stronger and can be the second candidate if polymer rotors do not deliver the required performance. However, extra efforts will be needed to align and fix the parts during machining.

## **4.0 Solid Model Generation**

The designs of rotors are given in terms of the rotors' geometrical specifications—the distribution of chord/radius ( $c/R$ ) and incidence along the radial direction for a single rotor blade, cross section shapes, number of blades, and diameter of the hub. Chords/radius and incidences are given at different percentages of radius ( $r/R$ ). The cross section is proportional to the chord. Therefore, from the radius of the rotor, a simple spreadsheet can be used to calculate radial positions ( $r$ ) and chords ( $c$ ). These data will be manually entered into the CAD modeling system for the generation of a solid model of the rotor. Due to the structural considerations and the reasonable assumption that there is little influence on aerodynamic performance within 25% of the radius, as measured from the center of the hub, the rotor model can be modified as necessary in this regime. A completed rotor model is shown in Figure 3.



**FIGURE 3. A completed 5-blade rotor model**

## **5.0 Polymer Rotor Fabrication**

### **5.1 General Concepts**

Polymer rotors can be either machined from bulk polymer materials or cast. To avoid the complicated alignment necessary for machining, we chose casting. To use this technique, the materials for parts have to be castable and easy to remove from molds. The acceptable

operating and cure temperatures depend on the mold materials. Room-temperature-cured polymers are preferred because they reduce the need for furnace and special temperature control systems. Molds can be permanent or fugitive. Permanent molds are usually made of metal and require long lead times. The temperature that metal molds can endure is much higher than most polymers' handling temperature. However, mold release is difficult, and mechanical release methods tend to destroy thin blades. Moreover, the long lead time and higher cost of permanent molds mean less flexibility in the prototyping stage because any modification of design will take longer to implement. Fugitive molds are more flexible. Mold release can be by chemical or thermal means. The only drawback is the number of molds needed for mass production since in this approach the mold number is equal to the part number. Therefore, lower-cost mold materials should be considered, and the mold making cycle should be short enough to compensate for this disadvantage. Besides, the temperature that fugitive molds can sustain may also limit the selection of candidate polymers.

## 5.2 Manufacturing Processes

The process Shape Deposition Manufacturing (SDM), developed by Carnegie Mellon University and Stanford University, provides a solution for room-temperature polymer casting. SDM (Figure 4) is a layer manufacturing technique with a sequence of additive and subtractive processing steps for fabricating complex 3D parts. In each layer, part material and support material are deposited in approximate volume and machined to the net shapes. After the part is completely built, the support material is removed chemically or thermally, depending on the material characteristics. The substrate and support material serve as one side of the molds, and, unlike in the case of traditional complete molds, no further part material subtraction is required after casting. Various materials, such as metals and polymers, can be fabricated by SDM. Mold SDM (Figure 5) is a variation of this process for the creation of complex-shaped full fugitive molds. A large variety of polymer and ceramic parts have been made using these fugitive molds.






FIGURE 4. SDM




**FIGURE 5. Mold SDM**

For mesoscopic rotor fabrication, the concept of SDM or Mold SDM can be applied and simplified. Since rotors do not have any interlock or embedded features, the model can be easily divided into upper and lower surfaces. According to the definition of SDM, only a single layer of support and part materials is required. This single-layer attribute avoids common problems, such as material shrinkage between layers<sup>[7]</sup> and interlayer weakness<sup>[8]</sup>, common to multi-layer SDM parts. Mold SDM can provide full molds for casting and is a possible approach. However, the blade thinness that molds can achieve is a constraint, and limited fluidity and trapped air bubbles in thin, wide cavities are also potential problems. An additional drawback is that Mold SDM requires the material inside the cavity to be washed out completely before casting. This step is time consuming and no method can assure the material is completely washed out. It leads to longer cycle time and possible defects. Considering its simplicity in manufacturing, the SDM approach was clearly the better option. A detailed sequence of the manufacturing steps we took is illustrated in Figure 6 and outlined below.




**FIGURE 6. Rotor fabrication steps**

*1. Substrate preparation.* For machining, substrates can first be cut into precise individual blocks that fit into the precut area in a pallet (see Figure 7). After machining, the individual blocks are removed from the fixture for casting. An alternative method is to fix a large block of wax on a pallet, after which the substrate material is deposited as blocks on the wax (see Figure 8). The amount of the deposited substrate can be varied, and several rotors can be built out of a single substrate block. The machining of the substrate is done on a 3-axis CNC mill. The purpose of machining is to generate the geometry of the rotor's bottom surface.



**FIGURE 7. Removable single block substrate setup**



**FIGURE 8. Fixed substrates on the pallet**

2. *Polymer casting.* The material for the part is cast to fill the cavity. The substrate and mixed polymer are placed in a vacuum chamber for degassing. This ensures that no air bubbles are trapped in the mixture and cavity. The cure time varies from one material to another. Since room-temperature-cured polymers are selected, no temperature-controlled environment is required.

3. *Surface flattening.* In polymer casting, excess part material is deposited on the substrate. It is more efficient to remove most excess materials by a face mill than by small scale end mills. When small-scale end mills are used, their flute length is usually inadequate to cut anything deep. Therefore, although it may be acceptable to machine the cast part material to net shape directly, flattening the surface by large-size mills first to remove excess materials is more efficient in mesoscale part fabrication.

4. *Part Material shaping to net shape.* Part material that already conformed to the bottom geometry of the support material is now CNC machined to attain the desired geometry of the top rotor surface. The polymers are easy to machine and tool wear is minimal. On the other hand, bulk polymer bonds to the substrate better than a thin polymer layer does. Due to the bonding consideration, rough cutting is usually avoided if the maximum cut-in depth is smaller than the flute length. The elimination of rough cutting also supports the necessity of the previous step.

5. *Substrate removal.* The substrate can be removed mechanically, thermally, or chemically. Some polymers do not bond to the substrate firmly, and in the case of these polymers any force applied at the hub can pull the part out easily. In most cases, the wax substrate must be melted at 100-130°C. Remaining wax traces can be removed chemically with BIOACT 280 <sup>[9]</sup> at 90°C. If the substrate is a water-soluble material, the whole block is placed in agitated water until the substrate is dissolved.

### **5.3 Materials in Polymer SDM**

Once we decided on the basic manufacturing procedures, we needed to determine what part and support materials to use. Some results from previous research in the Rapid Prototyping Laboratory at Stanford University about polymer SDM<sup>[10]</sup> provided us with a

basic knowledge database for selecting part and support materials and their combinations. According to the results of polymer SDM investigation, there are three thermoset part materials that are suitable for the SDM process—Adtech LUC-4180 polyurethane<sup>[11]</sup>, Ciba TDT 205-3 polyurethane<sup>[12]</sup>, and Adtech EE-501/530 epoxy<sup>[13]</sup>. Their physical properties are compared in Table 5. LUC-4180 has highest ultimate tensile strength and largest elongation to failure. TDT 205-3 needs the shortest time to cure and has better interlayer strength than LUC-4180. EE-501/530 has the best interlayer strength of the three candidates. These properties have mostly been determined for bulk materials. As for thin-blade applications, experiments were still required to decide on the most suitable materials since other issues, such as bonding to support materials and stiffness in handling/operation, may also enter and become critical.

Support materials should be completely removable, and we considered three candidates: machinable wax, water-soluble wax, and water-soluble solder masks. Waxes are thermoplastic materials, whereas solder masks are thermosets. Basic properties of support materials are summarized in Table 6. Kindt-Collins Master Protowax<sup>[14]</sup>, known as red wax, has the lowest shrinkage but poor machinability; Kindt-Collins Master File-a-wax<sup>[15]</sup>, known as blue wax, is easy to machine but has high shrinkage. Since we do not need to cast a second or more layers of support material on the first one, shrinkage does not matter much. The melting point is the lowest temperature one has to achieve to remove wax by melting it. The part materials have to endure the melting temperature for support material removal and the casting temperature if the next layer of support material is cast. The mixture of 25% Kindt-Collins Master File-a-wax and 75% Kindt-Collins Master Protowax, known as purple wax, has a good trade-off between machinability and shrinkage. Another preferred wax, yellow wax, is the same mix without dyes and antioxidants.

**TABLE 5. Properties of polymer part materials**

Material	Adtech LUC-4180 Polyurethane	Ciba TDT 205-3 Polyurethane	Adtech EE-501/530 Epoxy
Ultimate Tensile Strength (MPa)	55	23	42
Elongation to Failure	15%	9%	1%
Hardness (Shore D Scale)	78-80	70	86-89
Cure Time (hours)	12	2	24
Density (g/cm <sup>3</sup> )	1.09	1.12	1.6

**TABLE 6. Properties of support materials**

	Solder Mask	Water-soluble Wax	File-a-wax (blue)	Protowax (red)	25% blue and 75% red Wax (purple/yellow)
Removal	water	water	melt + BIOACT	melt + BIOACT	melt + BIOACT
Melting point	UV cure	60	105	60	80
Casting temperature	room temperature	70	150	80	120
Shrinkage (%)	N/A	unknown	2.60	0.46	0.80
Machinability	fair	fair	good	poor	fair

## 5.4 Part and Support Materials Combinations

After we identified possible part and support materials, the next task was to determine the best material combinations. Most of the successful polymer SDM parts use wax as support material; therefore, we evaluated combinations of three part materials (LUC-4180, TDT 205-3, and EE-501/530) and three machinable waxes (red wax, blue wax, and yellow wax). Blue wax is very easy to machine without any gumming up. However, only the epoxy rotor bonds well to blue wax during the machining. Rotors made of LUC-4180 and TDT 205-3 detached from the blue wax substrate at the thin blade tip areas, showing that the bonding between thin blades and blue wax is not strong enough to survive cutting forces. Red wax has the worst machinability among the three, but bonding to part materials is not an issue. Wax chips tend to gum up the cutting tool even if cold air is blown on the cut surface and the tool. Yellow wax has the best trade-off between machining and bonding. If cold air is used during the machining, gumming is significantly reduced. Yellow wax also provides sufficient bonding strength during machining, and finished polyurethane parts can be easily pulled out of the wax substrate. In summary, we found yellow wax to be the best choice for support material for epoxy and polyurethane rotors.

The wax melting process in substrate removal may introduce warping or other distortion if the temperature goes higher than the heat deflection temperature of polymers. We therefore considered water-soluble support materials to replace thermoplastic waxes. In particular, we experimented with solder mask and water-soluble wax in conjunction with epoxy rotors. Solder mask is a UV-cured material, and water-soluble wax is a lower temperature wax. Both have fair machinability and very good bonding to epoxy. We found that the residual stress from the strong bonding between water-soluble materials and the epoxy causes severe distortion when the substrate is dissolved, making water-soluble support materials inadequate in the fabrication of mesoscale rotors.


## 5.5 Issues in Manufacturing and Special Strategies for Improvements

Even once all the choices regarding manufacturing procedures and suitable part and support material combinations have been made, several manufacturing issues remain to be resolved. First, the leading and trailing edges of rotors with airfoil cross sections are hardly ever survive during machining because they are too thin. To overcome this limitation, cambered plate cross sections are used. The aerodynamic performance of such a cross section is within the acceptable range (5-9%) of deviation from an airfoil cross section<sup>[16]</sup>. Also, because of the over-cut strategy described below, we can obtain wedge shapes at leading and trailing edges. Second, the use of ball end mills may introduce flash at the edges when exact part boundaries are specified in the machining codes. This problem is illustrated in Figure 9.

When we specify the boundaries for cutting, the “tangent” technique is usually chosen. This technique ensures that no extra material is removed, especially from molds. However, if the same approach is applied on the top surface, there may be small areas of material left at leading and trailing edges. In such a situation, the cross section is not




smooth at these edges and the roughness will introduce unexpected drag. To solve this problem, we can overcut the edges on the top surface to help clean off the flash. This overcut strategy also produces a wedge shape at the leading and trailing edges (Figure 10), which makes the cross section closer to the airfoil shape.




**FIGURE 9. The flash problem due to tangent cutting**


The resulting wedge shapes do not have sharp corners but are located at the less favorable side. Another strategy, which changes the building direction from +Z to -Z (Figure 11), can be applied instead. Following this strategy, we reversed the wedge shapes to the favorable side (Figure 12), and the cambered plate shape became much closer to an airfoil shape at the leading and trailing edges.



**FIGURE 10. Wedge shapes at leading edge and trailing edge**



**FIGURE 11. Rotor's building direction**



**FIGURE 12. Better wedge shapes in the opposite building direction (-z)**

High melting temperature in the wax removal process, including melting the wax and completely removing it with BIOACT 280, tends to generate warping and distortion. Precisely monitoring the temperature during these processes to keep it at no more than 90°C can reduce warping and distortion. Careful handling is also important. Epoxy is more flexible at higher temperature, and, hence, external forces may cause permanent deformation.

## 5.6 Resultant Parts and Lift Test Results


We fabricated three types of rotors. Their diameters, designed performance, and corresponding motors are listed in Table 7.

**TABLE 7. Summary of three types of rotors**

Rotor	Diameter (mm)	Smoothy Motor	Design RPM	Design Thrust (g)	Cross section
3-Blade	15	3 mm	50,000	1.23	NACA 4402
5-Blade	22	5 mm	45,000	4.06	NACA 4402
4-Blade	25	5 mm	48,000	3.95	Optimized


We experimented on the three part materials discussed earlier, investigating their applicability for the 3-blade rotor (see Figure 13). We found that LUC-4180 polyurethane is too flexible to retain the 3D geometry and hence eliminated it from the part materials list. TDT 205-3 polyurethane has fair stiffness at this size, and can be pulled out of the wax substrate directly. EE-501/530 epoxy is brittle, but retains the 3D shape better than the other two materials. Its drawback is that epoxy rotors cannot be removed from the wax substrate directly. Of the two apparently acceptable part materials for mesoscale rotors, TDT 205-3 and EE-501/530, only the latter ended up being appropriate because in later experiments on 5-blade rotors, TDT 205-3 polyurethane rotors experienced thrust degradation during the test and could not compete with the performance of epoxy rotors with the same geometry. This indicates that polyurethane rotors experience such large

structural deformation during operation that design lift cannot be achieved. Therefore, for newer versions of rotors, only epoxy is used as polymer part material.



**FIGURE 13. Three polymer rotors and a mosquito**

We designed a 5-blade, 22mm diameter rotor for a 5mm Smoovy motor (see Figure 14). In order to verify the effect of incidence and optimize the design experimentally, we built and tested a set of 5-blade epoxy rotors with various incidences. We increased the incidence by 3, 4, 6, and 8 degrees from the original design. The design thrust of a 5-blade rotor should be 4.06 grams at 45,000 RPM. The thrusts generated by the above rotors did not deliver the expected lift. The 3 degrees+ and 8 degrees+ cases have better lift outputs but deliver still only about 75% of the design value. Several factors may account for these results. First, from the point of view of aerodynamic design, there may be some assumptions in the design model that we did not take into account. The other concerns are whether we can actually obtain the required geometry in the manufactured parts and whether structural deformations of the rotors under loading degrade the lift performance.



**FIGURE 14. The 5-blade rotor (real part)**

The follow-on to the five blade rotor is 4-blade rotor also designed 5mm Smoovy motor, but using updated aerodynamic modeling and designing codes. The cross section of this rotor is the optimized design at a Reynolds number of 6000. The same size hub as that used for the 5-blade rotor was first adopted. However, due to the smaller chords in the 4-

blade design, the rotor lacked structural stiffness, resulting in inconsistent performance. We therefore fabricated a modified version with a larger hub (see Figure 15). The hub is still within 25% of the radius, so the lift performance was affected very little. As for the design thrust of 3.95 grams at 48,000 RPM, the fabricated rotors can generate about 90% of the design thrust, which is much better than the thrust achieved in earlier designs.



**FIGURE 15. The bigger hub version of the 4-blade rotor**

## **6.0 Aluminum Rotor Fabrication**

The lift test on epoxy rotors showed that the design lifts are not achieved. The best result is the 4-blade version, which achieves about 90% of the design lift. Although the improvement in aerodynamic modeling and coding did help to shrink the gap between design lift and test results, doubts remain about the accuracy of the geometry of manufactured rotors. Epoxy rotors may warp or deform during the fabrication processes, resulting in more or less incidence than expected. Therefore, we considered metals and alloys, which have higher strength and stiffness than epoxy does. This section discusses selection of the appropriate aluminum alloy, possible manufacturing procedures, and special fixture and alignment approaches.

### **6.1 Material Selection**

Various types of aluminum alloys are available for different applications. Wrought alloys are cast as ingot or billet, and mechanically worked by rolling, extrusion, or forging into final form. Cast alloys are cast to final or near final form without any mechanical working. The elongation and strength in most cast products are lower than those in wrought products due to unpreventable casting defects<sup>[17]</sup>. Therefore, in our application, we are more interested in wrought alloys. Table 8 compares the characteristics and applications for different series of wrought aluminum alloys. According to the strength data, the 7-series has the highest strength, and it is a suitable material. In particular, we selected 7075-T6 alloy for rotor fabrication. This alloy contains 5.5% Zn, 2.5% Mg, 1.5% Cu, and 0.3% Cr. Applications are where very good strength and good resistance to corrosion are required.

**TABLE 8. Wrought Aluminum Alloys<sup>[18]</sup>**

Series	Alloys	Example	UTS (ksi)	Characteristics	Applications
1xxx	Pure Al	1100,1350	10-27	High formability, corrosion resistance, and electrical conductivity	Electrical, chemical
2xxx	Al-Cu	2014,2017, 2024,2219	27-62	High strength at room & elevated temperature	Aircraft, transportation
3xxx	Al-Mn	3003,3004, 3005	16-41	High formability, corrosion resistance, and joinability	Heat transfer, packaging, roofing-siding
4xxx	Al-Si	4032,4043	25-55	Good flow characteristics	Pistons, complex-shaped forgings
5xxx	Al-Mg	5052,5083, 5754	18-51	Excellent corrosion resistance, toughness, weldability	Building & construction, automotive, cryogenic, marine
6xxx	Al-Mg-Si	6061,6063, 6111	18-58	High corrosion resistance, excellent extrudibility	Building & construction, highway, automotive, marine
7xxx	Al-Zn	7005,7075, 7475,7150	32-88	Very high strength; special high toughness versions	Aerospace, Automotive
8xxx	Al+others	8017,8176, 8081,8280	17-35	High conductivity, strength, and hardness	Electrical, Aerospace, Bearing

## 6.2 Manufacturing Approaches

For fabricating rotors out of metals and alloys, it is more advantageous to shape the stock materials. High temperature metal casting and complicated temper processes can therefore be avoided. Moreover, the material properties are more reliable. However, shaping bulk materials to obtain rotor geometry is not at all straightforward. Since both top and bottom surfaces require machining, no single operation by a 3-axis CNC machine can accomplish the task because such a machine can only access the part from one direction. Therefore, the strategy is to machine the two surfaces separately, and a special fixing alignment system is necessary when the second surface is being machined. The basic fabrication can be described in three steps—machining of the first surface, fixture and alignment, and machining of the second surface. The three steps are outlined below.

### *First Surface Machining*

- 1. Cut aluminum plates into small pieces.* Aluminum plates are pre-cut into square pieces whose edges are slightly larger than the rotor diameter.
- 2. Attach the plates to a large flat reference substrate.* Loctite glue is used to bond the small plates to a substrate. The bonding strength is tested to assure that the plates can be held during machining.
- 3. Plane the top surface of the plates to a suitable thickness.* Since the glue layer may not be uniformly distributed over the plates, it is necessary to plane the top surface after the plates are bonded to the flat substrate. This planed surface also defines the z-axis


coordinate offset.

*4. Machine the first surface of the rotor on the plates and cut the contour of the rotor till the substrate is reached.* The contour cutting ensures that the rotor is separated from the original plate and cleans the edges.

*5. Debond the one-sided rotor from the substrate.* Suitable Loctite solvent is used to remove the glue. A sharp razor blade applies force at the corners to help detach the pieces from the substrate. The pieces are further cleaned in acetone solution with ultrasonic agitation.

### *Fixture and Alignment*

A practical alignment setup is a mold built to fit the first machined surface. This mold provides a reference and a solid support for the rotor for machining the second surface. There are two solutions for attaching the rotor to the mold—glue or vacuum holding. The glue approach is easier to implement, and we tested it first. We placed a very thin layer of glue between the mold and the machined surface of the rotor, and evenly applied pressure to assure that the two surfaces bonded to each other. The result of this trial was not satisfying: two blades of the resulting rotor failed during the machining. This may be due to the uneven glue layer having tilted the first machined surface relative to the mold. Some areas became so thin that they did not survive the machining. We therefore used the second approach, vacuum holding, to ensure perfect attachment. The mold (see Figure 16) functions as a vacuum chuck, and the rotor is held tightly and completely by the vacuum. The mold making process was more complicated. A plateau was first machined from the backside of the mold, sealed carefully, and connected to a vacuum pump. On the top surface of the mold, holes and grooves were distributed evenly over the blade areas to ensure sufficient vacuum holding. We successfully manufactured rotors through this approach.



**FIGURE 16. Vacuum chuck mold**

### *Second Surface Machining*

Once the rotor is fixed on the alignment setup, the second surface can be machined

without difficulty. After the rotor was released from the mold, we did some manual deburring. 7075-T6 was strong enough to survive the deburring.

### 6.3 Resultant Parts and Lift Test Results

The successful aluminum rotor is shown in Figure 17. We performed the same lift test we did on the other rotors. Aluminum rotors are stiffer and stronger than polymer rotors, and thus were expected to show improved performance. However, the test results did not support this assumption. Lift, RPM, voltage, and current were similar to those of epoxy rotors. In order to explain this lack of improvement, we needed to understand the geometry of the actual rotor and its effects on aerodynamic performance.



**FIGURE 17. Aluminum rotor**

## 7.0 Geometry Verification

We used two different approaches, non-destructive and destructive, to verify the geometry of fabricated rotors. A non-destructive method is laser scanning, which can provide 3D location information of the blade surfaces. Destructive methods cut blade cross sections and observe the 2D geometry at the cut locations.

### 7.1 Non-destructive Approach

#### 7.1.1 Methodology

There are several non-destructive measuring systems—co-ordinate measuring machines (CMM) <sup>[19]</sup>, laser interferometry <sup>[20]</sup>, confocal scanning <sup>[21]</sup>, and laser triangulation scanning <sup>[22]</sup>. CMM systems mostly operate by making contact with the parts. Optical triggered probes exist, but their resolution is not satisfying. Laser interferometry or a confocal approach is mainly used on microscale applications such as in biology. The resolution of these approaches is good, but rotors are bigger than the capacity of the working platform in the general setups. Laser triangulation scanning has different applications in measurements of different-sized objects. The sensor technique and resolution determine the resolution of the measurement and its range of applications. The Stanford Graphic Lab<sup>[23]</sup> owns a laser triangulation scanning system, but the z-axis



resolution is about 100  $\mu\text{m}$ , which is larger than the thickness of our rotor blades at most locations. After assiduous searching, we located a company in Michigan, GKS Inspection Services, which was able to conduct measurements for us. The manufactured rotors were sent to the company, laser-scanned with more than 40,000 points from the top surface, and compared with the CAD model of the original design.

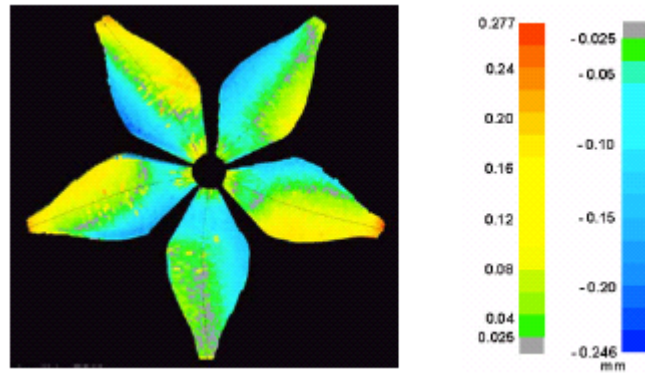
GKS used Laser Design's Surveyor 3D Laser Digitizing Systems<sup>[24]</sup> to provide the service. The system emits a Class II laser light ( $<1\text{mW}$ ) to scan the object. Using triangulation based technologies, 3D coordinate locations lying on the surface of the object were captured by a rapid profile line sensor<sup>[25]</sup>, RPS-450, on a CCD in the probe. The sensor can detect 0.035mm or larger distance in the z direction. After the digitizing was done, the raw scan data were exported to a CAD system for comparison with the original design model.

The company returned two types of output results to us. One is the original point data in the IGES format, which we were able to use for further analysis. The other is the results of the comparison with the CAD model in the image file format. In the image files, colors show the variations between the scanned dimensions of a physical part and the dimensions of its CAD model. Red shading represents excess material or bent areas of parts; blue shading represents missing material. From these image format results, we were able to observe qualitative deviations of the parts from their CAD model. On the other hand, the original point data could be further used to compute the incidences of the cross sections and gave us quantitative information about the actual variations.


### **7.1.2 Results and Discussions**

The 5-blade and 4-blade rotors were inspected via laser scanning. For the 5-blade rotor, the results compared with the CAD file, which had 4 more degrees of incidence than the original design, are shown in Figure 18. We further analyzed the deviations of incidence at 50% and 75% of the radius for each blade based on the color readings from the figure. By measuring the difference between the leading and the trailing edges and taking the chord values from the CAD model, we can calculate the deviation of incidences (see Table 9). From the images and this preliminary analysis, we found the deviations to be inconsistent at different blades and locations. The substantial variations we saw indicated that the blades were twisted clockwise, resulting in higher incidences than expected, which perhaps explained the excessive input powers required for lift that we measured.





**FIGURE 18. Laser scanning results of the 5-blade epoxy rotor#1**




**FIGURE 19. Laser scanning results of the 5-blade epoxy rotor#2**


**TABLE 9. Deviation of the 5-blade epoxy rotor #1**

	Blade 1	Blade 2	Blade 3	Blade 4	Blade 5
deviation at 50% R	1.41°	10.27°	3.05°	4.03°	2.72°
deviation at 75% R	8.62°	2.28°	2.07°	11.38°	0.98°

We assumed that an increase of thickness inboard of 25% of the radius would improve the stiffness of the structures and reduce the twist introduced during manufacturing. Therefore, a modified version with thickened inboard area was manufactured and laser scanned for comparison. As shown in Figure 19 the anomalous twist has been substantially reduced. Two 4-blade epoxy rotors were laser scanned. The images of the comparison with the CAD models are shown in Figure 20 and Figure 21.



**FIGURE 20. Laser scanning results of the 4-blade epoxy rotor#1**



**FIGURE 21. Laser scanning results of the 4-blade epoxy rotor#2**

Since the blades of the 4-blade design are located at  $0^\circ$ ,  $90^\circ$ ,  $180^\circ$ , and  $270^\circ$ , it is easy to identify the radial locations by their x or y coordinates. According to the coordinates of measured points at leading and trailing edges, incidences at different cross sections can be calculated. The following plots (Figure 22 and Figure 23) illustrate the incidence distributions along the radii of two different rotors.




FIGURE 22. 4-blade epoxy rotor#1 incidence results





FIGURE 23. 4-blade epoxy rotor#2 incidence results

The results show the inconsistency of the incidence of different blades on the same rotor. To explain the distortion we observed, we considered the residual internal stress of cured epoxy and the heat deflection temperature. We initially speculated that the internal stress may induce warpage and distortion of blades after the substrate is removed. Since the internal stress data of Adtech EE 501/530 is not available from the manufacturer, we used data of a similar bishenol-A based epoxy found in the literature<sup>[26]</sup> for evaluation. According to the reference, the internal stress measured was about 7 MPa for the high

temperature-cured epoxy, and it was built up during the cooling process from 180°C to room temperature. Therefore, for our room-temperature-cured epoxy with a peak exotherm of 70°C, the internal stress should be about 3 MPa. If we assume this value contributes only to the shear stress in the worst scenario for distortion, then the distortional angle due to this stress can be calculated from  $\tau = G\gamma$ , where  $G$  is from the manufacturer's data. The result is  $\gamma = 0.0014 = 0.08^\circ$ . The distortion due to the internal stress should hence be insignificant. Therefore, the distortion is most likely created by manipulation and handling. Because the melting temperature of the substrate (80°C-90°C) is higher than the heat deflection temperature (48°C) of the epoxy, the epoxy rotors soften during the substrate removal process. Even with very careful handling, the step of transferring the unsupported rotors from the oven to the BIOACT 280 solution is very likely to induce a certain amount of distortion. In conclusion, unless we can eliminate the step of melting the substrate at a temperature higher than the heat deflection temperature, it may be impossible to avoid distortion during the manufacturing.

These results prompted us to fabricate an aluminum rotor. The lift test results of the aluminum rotor, which show a performance similar to that of epoxy rotors, have been previously described. We were particularly interested in the laser scanning results to see whether the deviation of the incidence was lessened in the aluminum rotor. If the deviation in the geometry was lessened significantly but the aerodynamic performance stayed similar, factors other than manufacturing had to be taken into account. The laser scanning image results and the incidence analysis results are shown in Figure 24 and Figure 25. The incidences of the design model are also illustrated in the same plot. The deviation from the design incidence is much smaller in the aluminum rotor than that in the epoxy rotor. Apparently, the aluminum rotor is stiffer and therefore better able to retain the geometry than the epoxy rotor is. Since the aluminum rotor is closer to the design geometry, we expected to obtain better lift performance, provided the improvement in geometry had a significant effect on the performance. Unfortunately, the lift test results of the aluminum rotor demonstrate very little improvements in lift, frequency, and input power.



**FIGURE 24. Laser scanning results of the 4-blade aluminum rotor**




FIGURE 25. 4-blade aluminum rotor incidence results

## 7.2 Destructive Approach (Cross section Geometry Verification)


### 7.2.1 Methodology

The destructive approach is to cut the blades and inspect the cut cross sections under an optical microscope. The rotors were embedded in another material and cut at the desired positions. Pictures can be taken at these positions under a microscope and can provide cross section information about the fabricated rotors. The information obtained from these pictures includes the shape of leading and trailing edges, chord, camber, and thickness. This method gives feedback about the geometry of the fabricated parts and is especially helpful for improvement of the fabrication process. For example, the selection of the best direction in which to build the rotor is based on observations of the measured shapes of leading and trailing edges. We tried several embedding methods. The first is to use wax as substrate material. The substrate serves as a mold. The fabricated rotor is placed into the mold and a room-temperature-cured polymer is cast on top of it. The polymer should be different from the part material. After the polymer is cured, the block is milled down until the tip of blade is revealed. The tip position is then used as reference for other radial positions. In this approach, we do not know where the actual center of the rotor is located since the polymer is not transparent. Since the tip of the rotor blade is determined by gradual milling down, its location may not be precisely determined. Therefore, the cut cross sections may deviate from the desired positions. As a result, the information about the chord may not match that at the desired positions. However, the information obtained regarding thickness and camber is still valid.

The second approach is to manufacture the individual blades only as shown in Figure 26. The blades were aligned at the tip, and at 75% and 50% of the radius. The blades are not

removed from the wax substrate after the top surfaces are machined, and polymer is cast on top to hold the blades. The resulting block is held by a fixture on the pallet; hence it is easy to cut precisely at the desired positions. Only one cut is required in this configuration. This approach does not have the problems of the first approach. However, since the blades are not removed from the substrate, the measurement merely reveals the mold shape and machining results. Any effect that might be introduced by substrate removal is not considered. Nevertheless, the approach is useful if only two-dimensional geometry is being considered. The third embedding approach uses transparent epoxy as substrate. The fabricated rotor sat at the center of a square reservoir where we cast it in epoxy. We carefully aligned one blade perpendicular to the side of the reservoir.

After the epoxy is cured, we can take out the square block and mark approximate desired locations to be cut. Since the block is square and the blades are aligned with the sides, it is quite convenient to mill the four sides down to the marked positions. From the top view of the rotor, we can measure the exact positions of the cut cross sections. Figure 27 illustrates how this is implemented. Information gathered from the above methods were basically two dimensional. For three-dimensional information, we mainly relied on laser scanning results. However, the data points of the incidence distribution fluctuated widely. We suspected the resolution of laser scanning was not sufficient to correctly reveal small incidence variations. Therefore, especially for the aluminum rotor, we investigated cross sections based on the third approach, with special strategies to gather incidence data. Each blade was cut at 18 locations, and each cross section was polished and observed under the optical microscope.



**FIGURE 26. Second embedding method**

Three major issues should be considered and planned before embedding the rotor: (1) how to ensure that the cutting planes are perpendicular to the blades, (2) how to define the horizontal plane for incidence measurement, and (3) how to measure the radial locations correctly. To solve the first problem, a square pocket was machined into the center of the rotor hub, and a square post of the same size was located at the center of the substrate. Fitting the rotor into the post of the substrate ensured that the four blades were placed at 0, 90, 180, and 270 degrees. To solve the second problem, a polyurethane layer was cast and machined flat to provide horizontal reference. Moreover, the cast transparent epoxy block was cut into a square before being detached from the substrate. The center of the block can be determined by measuring the distance between two opposite sides. When cutting the


cross sections, we only cut two adjacent sides and left the other two as reference. At every cross section, the width of the block was measured, and the radial location of the blade could be determined from the measurement. After two adjacent sides were cut down until they were close to the hub, they were used as reference sides for the other two blades.




**FIGURE 27. The embedded rotor (left) was cut at four different locations (right).**

### **7.2.2 Results and Discussion**


We investigated 3-blade rotors designed for Smoovy's 3mm motor and 4-blade rotors designed for Smoovy's 5mm motors by this destructive approach. For the 3-blade design, we used the second method described above, while for the newer 4-blade rotor we used the improved third embedding approach. Microscope images of 3-blade PU rotor blades at the tip and at 75% and 50% of the radius are shown in Figure 28-30. From the images, we measured thickness/chord and camber information at these locations (see Table 10).



**FIGURE 28. PU rotor blade at tip**



**FIGURE 29. PU rotor blade at 75% r**




**FIGURE 30. PU rotor blade at 50% r**

**TABLE 10. Results of the 3-blade PU rotor**


3-blade rotor cut locations	tip	75%	50%
thickness/chord	2.4%	2.3%	2.8%
camber	4%	4%	4%

The building direction in this version is +Z, and therefore the wedge shape at the leading and trailing edges are on the less favorable side. The feedback we got from the images suggested that we flip the building direction to -Z to get better shapes at these edges. The images of the 4-blade epoxy rotors reveal this improvement (see Figure 31-34). In addition, the measured values of thickness/chord (see Table 11) for the 4-blade rotor are very close to 2%, which agrees with the design.







**FIGURE 31. Epoxy rotor blade at r1**



**FIGURE 32. Epoxy rotor blade at r2**



**FIGURE 33. Epoxy rotor blade at r3**




**FIGURE 34. Epoxy rotor blade at r4**

As for the aluminum rotor with special embedded reference planes, we were particularly interested in the incidence distributions. All photos taken from the optical microscope were analyzed with CorelDRAW<sup>[27]</sup>, a commercial drawing software. Lines were drawn from the leading edge to the trailing edge and on the reference plane. By reading out the coordinates of the start and end points of lines, the angle between the chord and the reference plane can be calculated. The incidence distributions of four blades analyzed through this approach are shown in Figure 35. The results were less noisy and deviated less than those from the laser scanning data. The maximum incidence deviation (except at the tips) was less than one degree. These incidence distribution data were then used as input data in the design codes to recalculate the performance of the rotor. The performance difference due to the deviation of the incidence was found to be less than 5%. That is, the rotor we built was functional and within the allowed tolerance range specified to deliver design performance.

**TABLE 11. Results of the 4-blade epoxy rotor**

	r1 = 11.0 mm	r2 = 9.5 mm	r3 = 7.0 mm	r4 = 5.5 mm
% of R	88%	76%	56%	44%
thickness/chord	2.2%	2.7%	2.1%	2.1%



**FIGURE 35. 4-blade aluminum rotor incidences from cross section investigation**

As for the higher measured input power, this discrepancy may be due to the combination of two other factors. In the CFD calculation, the results showed that the rotor consumed 15% more power than the design value obtained from the design codes. If this is true, then in order to deliver this higher power for the rotor, the motor would operate at a different performance curve, one with lower efficiency. The combination of both problems would result in the higher input power we measured.

## 8.0 Conclusion

This section has identified the challenges and opportunities in rotor fabrication, evaluated the possible materials and processes, discussed the fabrication of rotors based on a number of designs, and presented the results of lift tests and geometry verification on fabricated rotors. The thin blade structure and smooth three-dimensional surfaces necessary for a mesoscale rotor pose challenges for its fabrication. Fortunately, the stiffness-to-weight and strength-to-weight ratios become more favorable as parts are scaled down. Besides, the volume and weight of the rotors are very small compared to those of the whole mesicopter system. As a result, a variety of materials and processes become possible candidates.

We particularly looked at three types of engineering materials—polymers, metals and alloys, and ceramics. After evaluating the possible processes for these materials at the mesoscale, we determined that polymers were the best choice and aluminum alloys the second best. We experimented on the combinations of various part and support materials, and decided that epoxy is the best polymer part material and yellow wax is the most suitable support material. Some special strategies were developed to improve manufacturing. These included selecting the suitable direction in which to build the structure, overcutting the top surface, and monitoring the temperature precisely during the wax substrate removal.

## 9.0 References

- [1] Kunz, P.J., Kroo, I., "Analysis and Design of Airfoils for Use at Ultra-Low Reynolds Numbers," Conference on Fixed, Flapping, and Rotary Vehicles at Very Low Reynolds Numbers, Notre Dame, IN, June 2000.
- [2] M. F. Ashby, *Materials Selection in Mechanical Design*, Pergamon Press, pp. 26.
- [3] S. Kalpakjian, *Manufacturing Engineering and Technology, 3rd Edition*, Addison-Wesley Publishing Company, pp. 538.
- [4] M. F. Ashby and D. R. H. Jones, *Engineering Materials: An Introduction to their Properties and Applications*, Pergamon Press., pp. 2.
- [5] S. Kalpakjian, *Manufacturing Engineering and Technology, 3rd Edition*, Addison-Wesley Publishing Company, pp.104.
- [6] S. Kalpakjian, *Manufacturing Engineering and Technology, 3rd Edition*, Addison-Wesley Publishing Company, pp.64.
- [7] J. Kietzman, "Rapid Prototyping Polymer Parts via Shape Deposition Manufacturing," PhD Thesis, Stanford University, Stanford, California, 1999, pp. 54-57.
- [8] J. Kietzman, "Rapid Prototyping Polymer Parts via Shape Deposition Manufacturing," PhD Thesis, Stanford University, Stanford, California, 1999, pp. 69-71.
- [9] "BIOACT 280 Precision Cleaner Technical Data Sheet," Petroferm, Inc., Fernandin Beach, Florida.
- [10] J. Kietzman, "Rapid Prototyping Polymer Parts via Shape Deposition Manufacturing," PhD Thesis, Stanford University, Stanford, California, 1999.
- [11] "LUC-4180 Lowex Casting System," Adtech Plastic Systems Corp., Charlotte, Michigan.
- [12] "TDT 205-3 Resin and Hardener / DT 081 / DT 082 Three Component Rapid-Curing Polyurethane Casting System," Ciba Specialty Chemicals Corp., East Lansing, Michigan, October 1996.
- [13] "System 5 Epoxy Encapsulation Systems," Adtech Plastic Systems Corp., Charlotte, Michigan.
- [14] "Investment Casting Waxes: Protowax," The Kindt-Collins Company ([http:// www.kindt-collins.com/waxes/casting.html](http://www.kindt-collins.com/waxes/casting.html)).
- [15] "Tooling Waxes: MASTER File-A-Wax," The Kindt-Collins Company ([http:// www.kindt-collins.com/waxes/tooling.html](http://www.kindt-collins.com/waxes/tooling.html)).
- [16] "The Mesicopter: A Miniature Rotorcraft Concept, Phase II Interim Report," Stanford University, pp. 2-3 (available at <http://adg.stanford.edu/mesicopter/ProgressReports/MesicopterInterimReport.pdf>).
- [17] "Aluminum Alloy Selection and Applications," The Aluminum Association, Inc., Washington, D.C., 1998, pp. 14-15.
- [18] "Aluminum Alloy Selection and Applications," The Aluminum Association, Inc., Washington, D.C., 1998, pp. 4-14.
- [19] S. Kalpakjian, *Manufacturing Engineering and Technology, 3rd Edition*, Addison-Wesley Publishing Company, pp. 1049-1051.

- [20] A. Tebo, "Laser interferometry: A tool for many industrial applications," *SPIE OE Report*, November, 1996 (<http://www.spie.org/web/oer/november/nov96/laseri.html>).
- [21] L. Ladic, "3D Laser Scanning Confocal Microscopy: Introduction," Department of Physiology, University of B.C., Vancouver, B.C., Canada (<http://www.cs.ubc.ca/spider/ladic/intro.html>).
- [22] W. Kennedy, "Noncontact Measurement: Can Laser Triangulation Help You?" *Quality Magazine*, September 1998 (<http://www.qualitymag.com/articles/1998/sep98/0998f2.html>).
- [23] Stanford Computer Graphics Laboratory (<http://www-graphics.stanford.edu/>).
- [24] "Surveyor 3D Laser Digitizing System," Laser Design Inc. (<http://www.laserdesign.com/technol.htm>).
- [25] "Rapid Profile Sensor," Laser Design Inc. (<http://www.laserdesign.com/specs/specsheets/laser-sensors/pdf/rps.pdf>).
- [26] M. Ochi, K. Yamashita, and M. Shimbo, "Mechanism for Occurrence of Internal Stress During Curing Epoxide Resins," *Journal of Applied Polymer Science*, December, 1991, vol. 43, no. 11, pp. 2013-2019.
- [27] Corel Corporation (<http://www3.corel.com/>).

# Chapter 4


## Airframe Fabrication and Assembly for the 15g Prototype

### 1.0 Design Considerations


The airframe is a structural component of the mesicopter integrating all subsystems. In addition to the structural connections, it also includes rotor shrouds that not only protect the fragile rotors but conduct airflow to improve the rotors' aerodynamic performance. Airframe design should provide these functions under a limited size and weight budget.

Besides size and weight, other issues need to be considered as well. The first is the arrangement of all subsystems. The layout should utilize the space most efficiently, and the assembly of subsystems should be easy. Little or no assembly of the airframe itself is preferred. The second issue concerns the stability of the vehicle. In order to maximize damping of the system, each rotor should be tilted relative to the center line of the system. Third, the cross sectional geometry of the shrouds should be airfoil-like and any sharp corners and edges should be avoided except at the trailing edges. The second and third issues introduce another important concern—manufacturability. Any complex design that is impossible to fabricate with selected materials and processes must be discarded.

The evolution of the airframe design is shown in Figure 1. The first three models did not implement shrouds in the design, and hence printed circuit board (PCB) was used as structural material. Since shrouds are included in the airframe, PCB, which is limited to planar structures, is inappropriate. The last design condenses the airframe into one single part. The structure in this case is stiffer and the alignment between shrouds and motor holders is more precise. The material selection and manufacturing approach discussed later are based on this final design. As shown in Figure 2, this airframe design consists of several parts—shrouds, motor holders, arms, base, and filler between shrouds and base. Their volumes and contributions to the total volume are listed in Table 1. The overall size of the whole system is 65.6mm. Rotors are attached to motors held by motor holders, while batteries are located at the center of the base region. The cavities in the base depend on the type and number of batteries/capacitors that are used. Moreover, the rotor/motor is tilted by 15° toward the center line of the airframe for stability.



**FIGURE 1. Evolution of the airframe design**



**FIGURE 2. Components of the airframe**

**TABLE 1. Volume of the airframe components**

Part	volume (mm <sup>3</sup> )	% of total volume
shrouds	1841.8	58.79 %
motor holders	125.7	4.01 %
arms	252.0	8.04 %
base	507.6	16.20 %
filler	406.0	12.96 %
total	3133.1	100 %

## 2.0 Material and Process Selection

Printed circuit board was considered as airframe material in the early design. The advantage was to integrate motor controllers and sensor components with the frame structure. However, since 3D, tilted shrouds were included in the final design to improve rotor performance and to provide protection of the rotors, we abandoned PCB because it is impossible to shape PCB into precise 3D features on this small scale and then assemble them. Polymers are better options for light-weight structures. The SDM process is capable of building complex polymer parts, hence is the top choice for airframe manufacturing.

Three polymers were identified as suitable materials for the polymer SDM process in earlier research—Ciba TDT 205-3 polyurethane (PU), Adtech LUC 4180 polyurethane (white PU), and Adtech EE-501/530 epoxy (Table 2). Due to the weight constraint, epoxy, which has the highest density (1.6 g/cm<sup>3</sup>), is less suitable, even though it has better strength. Therefore, we focused on PU and white PU, and also investigated possibilities to reduce the density further. In addition to these two polymers, we also evaluated a rigid polyurethane foam material<sup>[1]</sup> available in hobby shops whose density is not provided and whose curing time is about 30 minutes.

A feasible solution to reduce density is to add glass bubbles, Micro-Balloons (MB)<sup>[2]</sup>, into PU or white PU. The mixture is cast into a wax substrate with a 3.5cm x 2.5cm x 0.4cm cavity. After the cured block is removed from the substrate, the volume and weight of the sample can be measured to calculate the density. The results of tests of these three materials are summarized in Table 3. The weight ratio of PU to MB is 7.7:1, while the weight ratio of white PU to MB is 8.4:1. With these achievable densities, the weight of PU could be reduced to 57% of pure PU and the weight of white PU to 46% of pure white PU.

**TABLE 2. Properties of three polymers**

Material	PU	white PU	Epoxy
Density (g/cm <sup>3</sup> )	1.12	1.09	1.6
Curing time	2-3 hours	12 hours	24 hours

**TABLE 3. Comparison of three candidate airframe materials**

Material	Density (g/cm <sup>3</sup> )	Comment
PU w/ MB	0.64	Short mixing and handling time before being cured
PU foam	<0.15	Foam expansion is not uniform and not controllable
white PU w/ MB	0.43-0.51	Advantage of long handling time before being cured

White PU is more attractive than PU because it has a long handling time and has lower density before and after mixing with MB. However, the longer curing hours required for




white PU increase the total fabrication cycle time. The machinability of both materials remains fair after mixing with MB. Commercially available PU foam is not as good as the other PUs because the foam expanded from the mixture is hard to control and is not uniformly distributed. In density measurements of PU foam samples, we found that the bottom regions are much denser than the top regions and that the parts are so weak that they can be broken by hand.

### **3.0 Fabrication**

We fabricated the airframe by the SDM process, using wax as support material. Through SDM, multiple layers and different materials can be manufactured sequentially on top of previous layers without additional assembly. In this section, manufacturing strategies, including the required layers and building direction, are given. We first conducted a simplified experiment using 3-axis CNC machining to evaluate the feasibility. Finally, we fabricated the final design using 5-axis CNC machining and tested it in the 15-gram Mesicopter.

The airframe can be fabricated in two sequential layers. The first layer includes base, arms, and motor holders, whereas the second layer contains the shrouds and filled-in area. Since the cross-sectional geometry of shrouds is an irregular airfoil shape that is tricky to machine, the building direction is dominated by the strategy for machining the shrouds. According to our experiences, the surface finish of machined wax is better than that of machined polymers mixed with MB. Besides, the inner surface of shrouds, which conducts airflow, should be as smooth as possible. Therefore, of the two possible building directions (see Figure 3), the -z direction was the better choice because the inner surface in the -z building direction is a machined wax surface, whereas in the +z building direction it is a machined polymer surface. The manufacturing strategies are summarized in Figure 4, which shows the building direction opposite to the design direction and the two principal layers. In addition, the tilted rotor/motor and shroud require a 5-axis CNC milling machine to manufacture.



**FIGURE 3. Cross section of shroud and building directions**



**FIGURE 4. Two principal layers in airframe fabrication**

A simplified airframe design, which is similar to the final design without 15° tilt, was first built by using a 3-axis CNC machine to evaluate its feasibility. This helped us to identify material and manufacturing issues. Two tests were run in this investigation. The first one implemented the concept of using white PU with MB for the main structure and PU foam for shrouds and fill-ins. The finished part weighs 1.1g and is shown in Figure 5. The two materials bonded nicely without any compatibility problems, but we found that PU foam is not rigid enough to retain the required geometry when the thickness decreases. The endmills were able to access the regions to be machined as long as appropriate machining steps were applied. In the second test the PU foam was discarded and the whole structure was made of white PU with MB (see Figure 6). In this version, the airframe weighs 1.7g, which is 54% heavier than the previous one. The shrouds are stiffer to retain the shape but still flexible at the bottom, where they are very thin. Besides, we observed some cavities in the airframe due to air bubbles trapped inside during the casting. These cavities may help to reduce weight, but they would pose a problem if they appeared at places where they would endanger the structure. Applying vacuum during casting does not significantly decrease the release of air bubbles in this case because the mixture is too viscous for efficient degassing. Packing the part material firmly to the substrate and manually popping visible air bubbles, however, proved to be practical approaches to minimize this problem, and we observed no air bubbles at the structurally critical places.



**FIGURE 5. 3-axis CNC machined airframe with PU foam shrouds**



**FIGURE 6. 3-axis CNC machined airframe made of white PU w/ MB**


We found that white PU with MB has the necessary reduced weight and sufficient stiffness to be the suitable airframe material. In order to improve the stiffness at the bottom of the shrouds, a thin layer of PU was added during the fabrication. The general fabrication sequences and machining strategies we used in the 3-axis CNC machining trials can mostly be adapted to the fabrication of the final design with the 5-axis CNC machining. The finished airframe, shown in Figure 7, weighs 1.5 grams and is one piece without assembly. The total machining time for the whole airframe is about 6-8 hours, and the material curing time for each layer is 12 hours. Thus, the complete cycle time for the airframe is about 30-32 hours.




**FIGURE 7. Airframe fabricated by 5-axis CNC machine**

## 4.0 Final Assembly and Thrust Testing

The assembled mesicopter is shown in Figure 8. The rotors used in this prototype are the 4-blade epoxy rotors. Two clockwise and two counterclockwise rotors are used to balance the torques, and the motors are Smoovy 5mm motors controlled by the closed-loop controllers. Motor controller circuits and power sources were integrated on-board. Smoovy's closed-loop controller is too large to use directly on our vehicle, and we instead used a simplified version with the same IC (Philips TDA 5145)<sup>[3]</sup>. Figure 9 shows the Philips TDA 5145 IC in a surface mount device packaged in a minimal configuration with three external capacitors and interconnections. The weight of this controller is about 0.7g, which is an acceptable solution for the 15g prototype. To design and build a self-powered prototype for take-off and short flight was the initial goal for vehicle. We considered three types of power source systems—batteries, batteries with a DC/DC converter, and super-capacitors. Comments regarding each solution are summarized in Table 4.



**FIGURE 8. Assembled 15g version of mesicopter**



**FIGURE 9. Closed-loop controllers**

**TABLE 4. Comments on power sources**

Power System	Comments
Batteries	- high energy density, flat discharge - low cell voltage, high single cell weight
Batteries w/ DC/DC converter	- fewer single cells required - minimum converter weight > 1g, limited efficiency
Supercapacitors	- very low internal impedance - low nominal voltage, low energy density, fast discharge

From the lift tests of a single rotor, we estimated that at least 15 V voltage input is required for a 15g mesicopter to be able to take off. The weight budget for the power system is only 6.4g, and the current will be around 1.0A. Under these high voltage and high current specifications, currently available batteries or batteries with a DC/DC converter cannot satisfy the power requirement within the weight constraint. Therefore, we chose super-capacitors. A super-capacitor is a double-layer capacitor, and it has become a viable alternative solution to rechargeable batteries. It weighs less than rechargeable batteries and takes a shorter time to recharge. Besides, a super-capacitor can be charged and discharged more than a million times without decreases in its energy storage capacity. Most importantly, it has high current charge and discharge capacity. To achieve our aim, we stacked four 1.0 Farad super-capacitors in series. The individual capacitors are rated for 2.5 V. Therefore, the stack can theoretically be considered as a 0.25 F capacitor with a maximum rated voltage of 10 V. We overcharged the super-capacitors to deliver the required power. The weight budget and the actual values of our mesicopter components are listed in Table 5.

**TABLE 5. Weights of mesicopter components**


Weight	Airframe	Motors and controllers	Power system	Rotors	Total
Budget	1.5 g	8.0 g	5.0 g	0.1 g	14.6 g
Measured	1.5 g	8.8 g	6.4 g	0.3 g	17.0 g

The assembly sequence was as follows:

- (1) Four super capacitors were first inserted into the airframe and wired in series.
- (2) Four sets of motor controller circuits were individually soldered to the IC's (see Figure 10). Four motors were wired with controllers so that two of them rotated clockwise and two counterclockwise.
- (3) The motor/controller blocks were attached to the airframe (see Figure 11).

(4) The controllers were wired to the super-capacitors.

(5) The rotors were glued to the motor shafts according to their rotation orientations (see Figure 12) and the assembly was finished (see Figure 13).



**FIGURE 10. Preamble motor/controller circuits**



**FIGURE 11. Attach motor/controller blocks**




**FIGURE 12. Attach rotors**



**FIGURE 13. Finished assembly**

Lift tests of the assembled mesicopter with the external power supply are shown in Figure 14. The estimated lift from a single rotor is also plotted for comparison. The mesicopter takes off at 13 V and is able to hover with 16 V draining almost 1 A. The shrouds produce negligible effects on the lift performance due to the large gap between them and the rotor tips. The measurement from the external power supply provided us the lift and controller voltage information, and we were able to use this relation as a reference to evaluate the performance when we used on-board super-capacitors.



**FIGURE 14. Lift vs. controlled voltage**

We performed two experiments on the super-capacitor stack. First, the stack was disconnected from the motor controller, charged to a specific voltage, and then reconnected to the controller. The super-capacitors were charged to 12 V, but once the electrical contact to the controller was established the voltage dropped immediately to about 7-6 V and continued declining at a rate of 0.5 V per second. The result was a startup delay in the first few milliseconds after the contact. To find out how to avoid these initial losses we conducted a second measurement. This time we kept the super-capacitor stack

connected to the controller. An external voltage was applied to the stack and controllers long enough to assure that the rotors spun at a steady speed. The external voltage was then disconnected. As we had seen previously, the voltage dropped immediately to 8-7 V and then declined at a constant rate. In both cases, the maximum mesicopter lift delivered by the super-capacitors was about 6 g, corresponding to 7-8 V at the controller. The results indicate that the prototype at 17 g currently cannot fly with the power available from the on-board super-capacitors, even if overcharged.

## 5.0 References

- [1] NHP Expandable Liquid Foam, NHP Co., Inc., Lowell, Massachusetts (notes are also available at <http://depts.washington.edu/mti/1999/labs/polymers/foam.html>)
- [2] SIG Micro-Balloons Filler Material, SIG Manufacturing Co., Inc., Montezuma, Iowa.
- [3] "TDA5145; Brushless DC motor drive circuit," Philips Semiconductors (<http://www.semiconductors.philips.com/pip/tda5145>)



# Chapter 5

## Stability and Control Issues

### 1.0 Introduction


The Stanford Micro Air Vehicle is a project funded by the Nasa Institute for Advanced Concepts to test the feasibility of operating micro-sized rotorcraft. This class generally refers to vehicles that have a maximum dimension of six inches or less. The design and construction of these small flying vehicles presents new challenges when compared with their larger counterparts. The aerodynamic properties of the vehicle differs greatly because of low Reynolds number effects, while structural considerations are dominated by different goals. This report focuses mainly on the control issues of the Stanford micro air vehicle as well as the development of a computer simulation model to test these control implementations.

### 2.0 Computer Simulation of the Micro Air Vehicle

A dynamic simulation of the micro air vehicle was needed to provide a platform to quickly test stability and control law implementation without jeopardizing actual hardware. The derivation of the equations of motion used in this computer simulation are presented in this section. The derivation is divided into two main parts, the aerodynamic forces acting on the vehicle, and the kinematic equations of motion. In order to develop these equations, the geometry of the vehicle is first described. Finally, the simulation algorithm and some test data are reviewed.

#### 2.1 Geometry

The geometry of the vehicle was decided upon early in the project based on considerations for simplicity, low weight, and stability and control issues. The side and top view of the vehicle are shown in Figures 1(a) and 1(b). For simplicity, a four rotor, rather than a single rotor, design was chosen because trying to implement an articulated rotor in such a small space was not feasible. Pairs of counter-rotating blades are used to cancel out the motor torques in flight. By independently varying the speed of the four individual rotors, the same basic flight controls of an articulated rotor can be maintained. The four independent controls translate into four degrees of freedom for the vehicle. The inputs can be mixed in such a way that these four degrees of freedom correspond to the familiar roll, pitch, yaw, and collective commands present on conventional rotorcraft. Figure 2 describes how this mixing occurs. By differentially varying the front and rear rotor speeds, pitch commands are given. Likewise, differential varying the left and right rotor speeds results in roll commands. Differentially varying the speed of the front/rear pair with the left/right pair maintains constant overall lift while inducing positive or negative yaw torques. Finally,



**Figure 1 Geometry of the Stanford Micro Air Vehicle.**


increasing and decreasing the speeds of all the motors in unison changes the vertical velocity of the vehicle similar to collective on standard rotorcraft.

Based on earlier two-dimensional computer simulations, that can be obtained in last years report, it was determined that an inward cant of the motors resulted in a more inherently stable vehicle. In addition, lowering the center of mass (C.M.) below the rotor plane also added to stability. Therefore, the arms of the vehicle were tilted up at an angle  $\beta$  as seen in Figure 1(a). This not only cants the rotor shafts inwards by an angle  $\beta$ , but it also provides clearance for the tilted rotors and lowers the location of the heavier battery and electronics below the plane of the rotors. Also shown on the diagram are the forces and torques acting on the vehicle, and some important dimensions which locate the point of application of those forces and torques.

## 2.2 Derivation of Aerodynamic Forces and Torques

The derivation of the aerodynamic forces is resolved using a combination of momentum and blade element theory. Momentum theory is used to derive the inflow or induced velocity. The rotor is modeled as an actuator disk across which there is no discontinuity in velocity, but there is a discontinuity in pressure. The inflow velocity of the air when it crosses the disk, and is assumed to be uniform at all points on the disk. Using conservation of mass through the disk and conservation of energy in the wake leads to an equation for thrust of the form  $T = 2\rho A v_1^2$ . In hover, the thrust is equal to the weight so  $T$  is replaced by  $W$ . Solving for the inflow velocity in hover results in the following expression:

$$v_{1_{hov}} = \sqrt{\frac{W}{2\rho A}} \quad [\text{m/sec}].$$



**Figure 2 Diagram of Control Mixing**

Now consider the case of the rotor moving sideways through the air. Momentum theory will produce a different expression for thrust,  $T = 2\rho A\nu_1\sqrt{V^2 + \nu_1^2}$ . Note that the rotor disk is assumed to have zero angle of attack with respect to the incoming flow, i.e. the rotor shaft is perpendicular to the direction of travel. Solving for the inflow velocity in this case results in the final expression that will be used in all future calculations involving inflow:

$$\nu_1 = \sqrt{-\frac{V^2}{2} + \sqrt{\left(\frac{V^2}{2}\right)^2 + \left(\frac{W}{2\rho A}\right)^2}} \quad [\text{m/sec}]. \quad (1)$$

Note that for the case of arbitrary motion in three dimensions, the rotor velocity is most easily described in the body fixed shaft axes; z-axis down parallel to the shaft, x-axis forward, and y-axis out the right side of the vehicle. In this case, the sideways velocity of the rotor is  $V = \sqrt{\dot{x}^2 + \dot{y}^2}$ . Also, the total inflow velocity includes vertical motion of the rotor as well. The inflow ratio,  $\lambda$ , is a dimensionless quantity often used in helicopter literature to relate the inflow velocity to the rotor tip velocity. The inflow ratio is defined as:


$$\lambda = \frac{\nu_1 - \dot{z}}{\Omega R}.$$

Another dimensionless quantity used in helicopter literature is the rotor advance ratio, which relates the horizontal velocity with the tip velocity of a rotor. The definition of the rotor advance ratio is:

$$\mu = \frac{V}{\Omega R},$$

where  $V$  is the horizontal velocity defined earlier.

Blade element theory is used to determine the total aerodynamic forces and torques acting on a rotor. As the name implies, these forces and torques are calculated by integrating the individual forces acting on smaller blade elements over the entire rotor. Figure 3 depicts a blade element as well as the velocities it sees and the forces acting on it. The “horizon” line is the line perpendicular to the rotor shaft, and will be parallel with the true horizon in hover. The geometric pitch or incidence, denoted by  $\theta$ , is the angle between



**Figure 3 Blade Element including Local Velocities and Forces.**

the horizon line and the blade chord line. The angle between the blade chord line and the local velocity vector is the local angle of attack,  $\alpha$ , and the angle between the horizon and the local velocity vector is the local inflow angle,  $\phi$ . The local velocity vector seen by each blade element is further decomposed into vertical and horizontal components,  $U_P$  and  $U_T$ . The vertical component of the local velocity vector is due to inflow and vertical motion of the rotor, so it is uniform for every section. The horizontal component is due to the angular velocity of the blade element and horizontal motion of the rotor; therefore, it is depends on the radial position of the element as well as the current orientation of the blade with respect to the incoming horizontal flow. These quantities are defined as:

$$U_P = v_1 - \dot{z} = \Omega R \lambda, \quad \text{and} \quad (2)$$

$$U_T = \Omega r + V \sin \Psi = \Omega R \left( \frac{r}{R} + \mu \sin \Psi \right), \quad (3)$$

where  $\Psi$  is the azimuth angle. This is the angle the blade makes as it rotates about the shaft such that  $\Psi = 0$  is defined to be in the direction of the horizontal velocity vector  $V$ , i.e. opposite the direction of motion of the vehicle.  $V$  was defined earlier as  $V = \sqrt{\dot{x}^2 + \dot{y}^2}$ .

The goal of blade element theory is to calculate the total forces acting parallel and perpendicular to the rotor shaft as well as the torques acting about the rotor shaft. The force parallel to the rotor shaft is defined as the thrust of the rotor, and the forces perpendicular to the rotor shaft are defined as the hub forces. In the following subsections, the derivation of the thrust, hub forces, and rotor torques will be described.

### 2.2.1 Thrust Derivation


Thrust is calculated by integrating the vertical forces acting on all the blade elements. In this case, vertical is defined as perpendicular to the horizon line. The lift force is perpendicular to the local velocity vector, while the drag force is parallel to the local

velocity. Thus the vertical force acting on a blade section is  $\Delta F_V = \Delta L \cos(\phi) - \Delta D \sin(\phi)$ . Here,  $\Delta L$ ,  $\Delta D$ , and  $\Delta F_V$  correspond to the forces acting on a blade element section,  $\Delta r$ . The equations for lift and drag are obtained from any introductory aerodynamic text as:

$$\Delta L = q C_l c \Delta r, \quad \text{and}$$

$$\Delta D = q C_d c \Delta r.$$

At this point many assumptions need to be made to obtain a closed form solution for the thrust. First, only aerodynamic lift and drag forces are assumed to act on a section. Forces out of plane with the section are dominated by the centrifugal force, and cancel if the rotor is balanced. Aerodynamic moments are assumed to be negligible by assuming the shear center and aerodynamic center of the blade lie close together, and that the blades are reasonably stiff. The blades are light enough so that gravity torques may be safely ignored. The blade is assumed to be rigid, so that there is no flapping or coning. Centrifugal stiffening reinforces this assumption, so the entire rotor is assumed to lie in a plane. To define the reference area of the element, an assumption that the rotor blade has constant chord is made. For cases where the chord is not constant, average chord is used. The next assumption, which can be found in any introductory aerodynamics text, is that the coefficient of lift varies linearly with angle of attack by  $C_l = a\alpha = a(\theta - \phi)$ . Here a linear twist distribution is assumed so that the twist varies linearly with radial position according to the relation  $\theta = \theta_0 - \theta_{tw}(r/R)$ . Figure 4 shows a linear approximation of



**Figure 4 Comparison of Actual Blade Twist with Linear Twist Approximation.**

blade twist plotted on top of the actual twist distribution of the designed blade. This

approximation works well from one quarter of the radius out to the tip. The inboard twist is not as important since the lift produced in this region is small because the angular velocity,  $\Omega r$ , approaches zero. The inflow angle is also needed to define the lift coefficient, and may be approximated by assuming that the angular velocity of the blade section is much larger than the total inflow through the blade. Then small angle approximations can be used to define the inflow angle as  $\phi = U_P/U_T$ . One last assumption used later is that the lift acting on the blade is about an order of magnitude higher than the drag.

The terms in the lift and drag equations can now be substituted with quantities that are well known. The dynamic pressure of the section is  $q = \frac{1}{2}\rho U_T^2$ . The reference area is defined using the average chord,  $S = \bar{c}\Delta r$ . And the coefficient of lift is replaced with the expression defined above resulting in the lift for a single blade element:

$$\Delta L = \frac{1}{2}\rho U_T^2 a \left( \theta_0 - \theta_{tw} \frac{r}{R} - \frac{U_P}{U_T} \right) \bar{c} \Delta r. \quad (4)$$

Similarly, an expression for the drag of the section can be developed:

$$\Delta D = \frac{1}{2}\rho U_T^2 \overline{C_d} \bar{c} \Delta r, \quad (5)$$

where  $\overline{C_d}$  is the drag coefficient of the section corresponding to the 70% radial station. This may be assumed to be typical of the lift producing portion of the blade which is also the portion of the blade causing the majority of the drag.

The thrust is finally found by integrating the vertical forces acting on all the blade element sections. The vertical force was defined at the beginning of this section, but may be approximated by  $\Delta F_V = \Delta L$  by applying small angle approximations for  $\phi$  and using the fact that the thrust is an order of magnitude larger than the drag. Making use of the lift defined in equation 4, and the velocity components given by equations 2 and 3, the total thrust is computed:

$$\begin{aligned} T &= \frac{N}{2\pi} \int_0^{2\pi} \int_0^R \frac{\Delta L}{\Delta r} dr d\Psi \\ &= \frac{N\rho a \bar{c}}{4\pi} \int_0^{2\pi} \int_0^R \left[ U_T^2 \left( \theta_0 - \theta_{tw} \frac{r}{R} \right) - U_T U_P \right] dr d\Psi \\ &= \frac{N\rho a \bar{c} (\Omega R)^2}{4\pi} \int_0^{2\pi} \int_0^R \left[ \left( \frac{r}{R} + \mu \sin \Psi \right)^2 \left( \theta_0 - \theta_{tw} \frac{r}{R} \right) \right. \\ &\quad \left. - \lambda \left( \frac{r}{R} + \mu \sin \Psi \right) \right] dr d\Psi \\ &= N\rho a \bar{c} (\Omega R)^2 R \left[ \left( \frac{1}{6} + \frac{1}{4}\mu^2 \right) \theta_0 - (1 + \mu^2) \frac{\theta_{tw}}{8} - \frac{1}{4}\lambda \right]. \end{aligned}$$

It is common to use non-dimensional quantities when describing the aerodynamic forces. Aerodynamicists use the coefficient of thrust defined by:

$$C_T = \frac{T}{\rho A (\Omega R)^2}.$$

Another important non-dimensional quantity used to simplify the thrust equation is the rotor solidity defined as the ratio of blade area to disk area. Using the average blade chord, then the solidity ratio is  $\sigma = N\bar{c}/\pi R$ . The coefficient of thrust can be computed by inserting the formula for thrust calculated above and simplifying:

$$\frac{C_T}{\sigma a} = \left( \frac{1}{6} + \frac{1}{4}\mu^2 \right) \theta_0 - (1 + \mu^2) \frac{\theta_{tw}}{8} - \frac{1}{4}\lambda. \quad (6)$$

This is the relationship that will be used to derive the dynamics of the vehicle.

### 2.2.2 Hub Force Derivation

The derivation of the hub force is nearly identical to the derivation of the thrust above, except that now the forces in the horizontal plane are integrated over all the blade elements. The hub force is actually decomposed into two orthogonal forces both perpendicular to the rotor shaft. The  $H$ -force lies along the  $\Psi = 0$  azimuth, and the  $Y$ -force lies along the azimuth  $\Psi = \frac{\pi}{2}$ . The horizontal force acting on a blade element is  $\Delta F_H = \Delta D \cos(\phi) + \Delta L \sin(\phi)$ . Small angle approximations can still be made, but now both terms are significant, so this expression only simplifies to  $\Delta F_H = \Delta D + \Delta L (U_P/U_T)$ . The lift and drag forces were defined in the previous section along with all the other terms that are needed to compute the hub force. First the  $H$ -force will be calculated by integrating only the component of the hub forces lying in the  $\Psi = 0$  direction:

$$\begin{aligned} H &= \frac{N}{2\pi} \int_0^{2\pi} \int_0^R \left[ \frac{\Delta D}{\Delta r} + \frac{\Delta L}{\Delta r} \frac{U_P}{U_T} \right] \sin \Psi dr d\Psi \\ &= \frac{N\rho\bar{c}}{4\pi} \int_0^{2\pi} \int_0^R \left[ U_T^2 \bar{C}_d + a \left( U_T U_P \left( \theta_0 - \theta_{tw} \frac{r}{R} \right) - U_P^2 \right) \right] \sin \Psi dr d\Psi \\ &= \frac{N\rho\bar{c}(\Omega R)^2}{4\pi} \int_0^{2\pi} \int_0^R \left[ \left( \frac{r}{R} + \mu \sin \Psi \right)^2 \bar{C}_d \right. \\ &\quad \left. + a \left( \left( \frac{r}{R} + \mu \sin \Psi \right) \lambda \left( \theta_0 - \theta_{tw} \frac{r}{R} \right) - \lambda^2 \right) \right] \sin \Psi dr d\Psi \\ &= N\rho a \bar{c} (\Omega R)^2 R \left[ \frac{1}{4a} \mu \bar{C}_d + \frac{1}{4} \lambda \mu \left( \theta_0 - \frac{\theta_{tw}}{2} \right) \right]. \end{aligned}$$

Once again, a non-dimensional coefficient for the hub force can be defined as:

$$C_H = \frac{H}{\rho A (\Omega R)^2}.$$

The final result that will be used when deriving the vehicle dynamics is:

$$\frac{C_H}{\sigma a} = \frac{1}{4a} \mu \bar{C}_d + \frac{1}{4} \lambda \mu \left( \theta_0 - \frac{\theta_{tw}}{2} \right). \quad (7)$$

The side force,  $Y$ , is similarly calculated by integrating the component of the hub forces lying in the  $\Psi = \frac{\pi}{2}$  direction:

$$Y = -\frac{N}{2\pi} \int_0^{2\pi} \int_0^R \left[ \frac{\Delta D}{\Delta r} + \frac{\Delta L}{\Delta r} \frac{U_P}{U_T} \right] \cos \Psi dr d\Psi.$$

Following the same steps as above results in the forces canceling in this direction, so the total side force is zero. Therefore the non-dimensional coefficient is also zero:

$$C_Y = 0. \quad (8)$$

### 2.2.3 Torque Derivation

The aerodynamic forces acting on the blade elements also cause moments about the rotor shaft and hub. This section presents the derivations of the aerodynamic torques that act about the center of the rotor. In particular, the torque acting about the shaft is important because it determines the power required for the motor to keep the rotor spinning. The remaining torques cause the rotor to pitch or roll about the hub.

The rotor shaft and the vertical axis are collinear. Therefore, the torque about the rotor shaft is found by integrating the horizontal aerodynamic forces acting on each rotor section with a moment arm equal to the radius of that section. The horizontal forces acting on a rotor section have already been defined above as  $\Delta F_H$ . In this case, the horizontal force is multiplied by the moment arm,  $r$ , and integrated over the entire rotor:

$$\begin{aligned} Q &= \frac{N}{2\pi} \int_0^{2\pi} \int_0^R \left[ \frac{\Delta D}{\Delta r} + \frac{\Delta L}{\Delta r} \frac{U_P}{U_T} \right] r dr d\Psi \\ &= \frac{N \rho \bar{c} (\Omega R)^2}{4\pi} \int_0^{2\pi} \int_0^R \left[ \left( \frac{r}{R} + \mu \sin \Psi \right)^2 \bar{C}_d \right. \\ &\quad \left. + a \left( \left( \frac{r}{R} + \mu \sin \Psi \right) \lambda \left( \theta_0 - \theta_{tw} \frac{r}{R} \right) - \lambda^2 \right) \right] r dr d\Psi \\ &= N \rho a \bar{c} (\Omega R)^2 R^2 \left[ \frac{1}{8a} (1 + \mu^2) \bar{C}_d + \lambda \left( \frac{1}{6} \theta_0 - \frac{1}{8} \theta_{tw} - \frac{1}{4} \lambda \right) \right]. \end{aligned}$$

Using dimensionless quantities, the rotor torque coefficient is:

$$C_Q = \frac{Q}{\rho A (\Omega R)^2 R}.$$

The final result is:

$$\frac{C_Q}{\sigma a} = \frac{1}{8a} (1 + \mu^2) \bar{C}_d + \lambda \left( \frac{1}{6} \theta_0 - \frac{1}{8} \theta_{tw} - \frac{1}{4} \lambda \right). \quad (9)$$

Finally, the rolling and pitching moments will be considered. Because the blade will be moving horizontally through the air, the advancing blade sections will see larger horizontal velocity components than the retreating blade elements. The affect of this is that the advancing blade elements will produce more lift than the retreating blade elements resulting in an overall rolling moment being produced. The total rolling moment is calculated by integrating over the entire rotor the moments caused by the lift of each



section acting at a moment arm equal to its radial position. The derivation follows:

$$\begin{aligned}
R &= -\frac{N}{2\pi} \int_0^{2\pi} \int_0^R \frac{\Delta L}{\Delta r} r \sin \Psi dr d\Psi \\
&= -\frac{N \rho a \bar{c} (\Omega R)^2}{4\pi} \int_0^{2\pi} \int_0^R \left[ \left( \frac{r}{R} + \mu \sin \Psi \right)^2 \left( \theta_0 - \theta_{tw} \frac{r}{R} \right) \right. \\
&\quad \left. - \lambda \left( \frac{r}{R} + \mu \sin \Psi \right) \right] r \sin \Psi dr d\Psi \\
&= -N \rho a \bar{c} (\Omega R)^2 R^2 \mu \left( \frac{1}{6} \theta_0 - \frac{1}{8} \theta_{tw} - \frac{1}{8} \lambda \right).
\end{aligned}$$

The non-dimensional rolling moment is found using the same method described for the rotor torque coefficient,  $C_Q$ . Take care not to confuse the rolling moment with the total blade radius. The context in which the symbol is used should make the definition obvious. The rolling torque coefficient is then defined as:

$$\frac{C_R}{\sigma a} = -\mu \left( \frac{1}{6} \theta_0 - \frac{1}{8} \theta_{tw} - \frac{1}{8} \lambda \right). \quad (10)$$

The pitching moment calculation is developed in a similar manner except that the cosine of the azimuth angle is taken:

$$P = -\frac{N}{2\pi} \int_0^{2\pi} \int_0^R \frac{\Delta L}{\Delta r} r \cos \Psi dr d\Psi.$$

The result of this integration ends up being zero. Therefore the pitching torque coefficient is also zero:

$$C_P = 0. \quad (11)$$

## 2.3 Equations of Motion

The equations of motion of the micro air vehicle are derived in this section using kinematics. Given the geometric properties of the vehicle, the location of the application of forces and torques can be determined. Assuming the vehicle is a rigid body, and given the mass properties of the vehicle, the dynamic equations of motion can then be written. First, some reference frames must be defined for describing the vehicle motion. For the purposes of the simulation, the ground is assumed to be the inertial reference frame, **I**, denoted by the set of axis  $(X, Y, Z)$ . The second reference frame is a body fixed frame, **B**, which has its origin at the vehicle's center of mass and rotates with the vehicle. The axes of frame **B** were depicted in Figures 1(a) and 1(b) above and are denoted by  $(x, y, z)$ . Each rotor also has its own body fixed frame, **R<sub>i</sub>**, centered at the rotor hub. The rotor frames are used for calculating the aerodynamic forces and torques as described in the previous section.

The time history of the vehicle motion is recorded by a set of state variables. For this simulation, position and velocity of the C.M. as well as vehicle orientation are the variables of interest. The position and velocity are recorded in Cartesian coordinates while the orientation is specified as Eulerian angles. The inertial frame is used as a basis since

the motion of the vehicle in inertial space is desired. In addition, the rotational dynamics will be solved with Euler's equations which use the angular velocity of the vehicle referenced in the body fixed frame. A transformation is applied to get the rate of change of the Eulerian angles directly from the angular velocity. From this set of requirements, twelve state variables are needed to completely track the three-dimensional motion. The state variables are:

$$\vec{\chi} = \begin{bmatrix} \dot{X} & \dot{Y} & \dot{Z} & X & Y & Z & \omega_x & \omega_y & \omega_z & \phi & \theta & \psi \end{bmatrix}^T.$$

The first step is to determine the external forces and torques acting on the vehicle body. The equations for calculating the force and torque coefficients of a single rotor were given in equations 6, 7, 9, and 10 above. Given the geometric and mass properties of the vehicle, the only unknown in the equations is the inertial velocity of the rotor given in the rotor's frame of reference,  $\mathbf{R}_i$ . The following equation is used to calculate the velocity of the rotor based on the motion of the vehicle:

$$\mathbf{I}\vec{V}^{r_i} = \mathbf{I}\vec{V}^{cm} + \mathbf{I}\vec{\omega}^B \times {}^{cm}\vec{r}^{r_i},$$

where  $\mathbf{I}\vec{V}^{r_i}$  is the inertial velocity of the rotor in question,  $\mathbf{I}\vec{V}^{cm}$  is the inertial velocity of the vehicle C.M.,  $\mathbf{I}\vec{\omega}^B$  is the angular velocity of the body fixed frame with respect to the inertial frame, and  ${}^{cm}\vec{r}^{r_i}$  is the vector from the C.M. to the rotor hub. As an aside, this formula holds true as long as the angular velocity vector is taken about the C.G. This is why the physical rate gyro sensors are located as close to the C.G. as possible. If the angular velocity measurements are taken at another location than the C.G., then the equations will become much more complicated. The velocity of the C.M. is given in the state vector  $\vec{\chi}$  with the inertial frame as its basis, so two direction cosine matrices are needed to transform from the inertial frame to the body frame to the rotor frame as follows:

$$\mathbf{I}\vec{V}_{\mathbf{R}_i}^{r_i} = \begin{bmatrix} \cos \beta & 0 & \sin \beta \\ 0 & 1 & 0 \\ -\sin \beta & 0 & \cos \beta \end{bmatrix} \cdot \mathbf{I}\vec{V}_{\mathbf{B}}^{r_i}, \quad \text{and} \quad (12)$$

$$\mathbf{I}\vec{V}_{\mathbf{B}}^{r_i} = \begin{bmatrix} \cos \theta \cos \psi & \cos \theta \sin \psi & -\sin \theta \\ (-\cos \phi \sin \psi & (\cos \phi \cos \psi & \sin \phi \cos \theta \\ + \sin \phi \sin \theta \cos \psi) & + \sin \phi \sin \theta \sin \psi) & \\ (\sin \phi \sin \psi & (-\sin \phi \cos \psi & \cos \phi \cos \theta \\ + \cos \phi \sin \theta \cos \psi) & + \cos \phi \sin \theta \sin \psi) & \end{bmatrix} \cdot \mathbf{I}\vec{V}_{\mathbf{I}}^{r_i}, \quad (13)$$

where  $\mathbf{I}\vec{V}_{\mathbf{R}_i}^{r_i}$  is the inertial velocity of the rotor hub with respect to the rotor frame,  $\mathbf{I}\vec{V}_{\mathbf{B}}^{r_i}$  is the same velocity with respect to the body frame, and  $\mathbf{I}\vec{V}_{\mathbf{I}}^{r_i}$  is the same velocity with respect to the inertial frame. The aerodynamic force and torque coefficients are calculated using the velocities above, then converted to dimensional units using the mass and geometry of the particular vehicle under consideration. In addition to the aerodynamic torques produced by the individual rotors, moments are also produced on the vehicle due to forces acting a distance away from the C.M. These moments are calculated using

the moment equation  ${}^{cm}\vec{M} = {}^{cm}\vec{r}_i \times \vec{F}$ . At this point, all the forces and torques acting on the body can be calculated, but they are given with respect to the rotor frames,  $\mathbf{R}_i$ , and should be converted back to the inertial frame  $\mathbf{I}$ . This is accomplished by using the inverse of the direction cosine matrices given in equations 12 and 13. Finally, the forces and torques are summed, including any gravitational force, resulting in a single equivalent force and torque acting at the C.M. of the vehicle.

The second step is to use dynamic principles to determine the motion of the vehicle when subject to the equivalent force and torque above. These dynamics can be split into translational and rotational parts. The translational dynamics are solved by simply applying Newton's law,

$$\vec{F} = m \cdot {}^I\vec{a}^{cm} = m \cdot [\ddot{X} \quad \ddot{Y} \quad \ddot{Z}]^T, \quad (14)$$

where  $\vec{F}$  is the force acting on the vehicle, and  ${}^I\vec{a}^{cm}$  is the inertial acceleration of the C.M. given as the derivative of state variables. Next, the rotational dynamics are solved using Euler's equations. The rate of change of the angular momentum of the vehicle is equal to the torque acting on the vehicle,  ${}^{cm}\vec{\tau} = {}^{cm}\dot{\vec{H}}$ . The angular momentum is defined by  ${}^{cm}\vec{H} = {}^{cm}\vec{I} \cdot {}^I\vec{\omega}^B$ , where  ${}^{cm}\vec{I}$  is moment of inertia tensor about the C.M. Substituting and solving for the torque results in:

$${}^{cm}\vec{\tau} = {}^{cm}\vec{I} \cdot \dot{{}^I\vec{\omega}^B} + {}^I\vec{\omega}^B \times ({}^{cm}\vec{I} \cdot {}^I\vec{\omega}^B). \quad (15)$$

The last information needed to complete the equations of motion are the rate of change of the Eulerian angles. Another transformation matrix is used to calculate the change in the Eulerian angles directly from the angular velocity:

$$\begin{bmatrix} \dot{\phi} \\ \dot{\theta} \\ \dot{\psi} \end{bmatrix} = \begin{bmatrix} 1 & \sin \phi \tan \theta & \cos \phi \tan \theta \\ 0 & \cos \phi & -\sin \phi \\ 0 & \sin \phi / \cos \theta & \cos \phi / \cos \theta \end{bmatrix} \cdot \begin{bmatrix} \omega_x \\ \omega_y \\ \omega_z \end{bmatrix} \quad (16)$$

This completes the derivation of the various parts of the equations of motion. All that is left is to reorder equations 14, 15, and 16 so that they may be written in familiar form

making use of the state variable matrix:

$$\dot{\vec{\chi}} = f(\vec{\chi})$$

$$\begin{bmatrix} \ddot{X} \\ \ddot{Y} \\ \ddot{Z} \\ \dot{X} \\ \dot{Y} \\ \dot{Z} \\ \dot{\omega}_x \\ \dot{\omega}_y \\ \dot{\omega}_z \\ \dot{\phi} \\ \dot{\theta} \\ \dot{\psi} \end{bmatrix} = \begin{bmatrix} \frac{F_X}{m} \\ \frac{F_Y}{m} \\ \frac{F_Z}{m} + g \\ \dot{X} \\ \dot{Y} \\ \dot{Z} \\ \frac{1}{I_{xx}} ((I_{yy} - I_{zz})\omega_y\omega_z + \tau_x) \\ \frac{1}{I_{yy}} ((I_{zz} - I_{xx})\omega_x\omega_z + \tau_y) \\ \frac{1}{I_{zz}} ((I_{xx} - I_{yy})\omega_x\omega_y + \tau_z) \\ \omega_x + \omega_y \sin \phi \tan \theta + \omega_z \cos \phi \tan \theta \\ \omega_y \cos \phi - \omega_z \sin \phi \\ \omega_y \frac{\sin \phi}{\cos \theta} + \omega_z \frac{\cos \phi}{\cos \theta} \end{bmatrix} \quad (17)$$

Note that the forces above are taken with respect to the inertial frame (as required by Newton's law), and as such, the individual forces are denoted with subscripts corresponding to the axes of the inertial frame. On the other hand, the torques and angular velocities used in Euler's equations are taken with respect to the body frame, and likewise, they are denoted with subscripts corresponding to the axes of the body frame. It can also be seen above that the equations of motion are nonlinear, especially since the driving aerodynamic forces and torques are dependent on the velocity of the vehicle. These nonlinearities are not a problem for the dynamic simulation, however difficulties arise when designing control laws for the vehicle and will be discussed later.

## 2.4 Implementation of the Dynamic Simulation and Control Algorithms

The section above sets the foundations for building a computer simulation of the micro air vehicle. The goal of the simulation is to visualize a time history of the vehicle motion given a set of initial conditions. The implementation should also be flexible enough to allow for future integration of any control law implementations. In order to run a computer simulation, the equations of motion must be discretized and solved at regular time intervals. A fourth order Runge-Kutta integration scheme was chosen to propagate the equations of motion forward at each time step. Given the vehicle state at the previous time step, the equations of motion are numerically integrated to find the state at the next time step. The size of the time step is an important consideration. If the time step is too large, errors in the numerical integration will grow possibly leading to a solution that doesn't converge, but if the time step is too small, then the simulation will take too long to

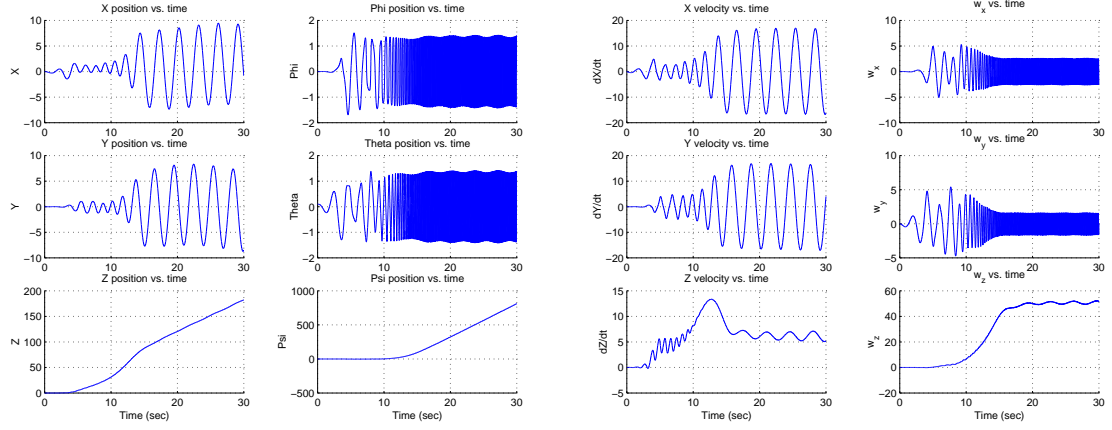
run and is no longer an efficient analysis tool. The Runge-Kutta algorithm follows:

$$\begin{aligned}
K_1 &= h \cdot f(\vec{\chi}_n, t_n) \\
K_2 &= h \cdot f(\vec{\chi}_n + \frac{K_1}{2}, t_n + \frac{h}{2}) \\
K_3 &= h \cdot f(\vec{\chi}_n + \frac{K_2}{2}, t_n + \frac{h}{2}) \\
K_4 &= h \cdot f(\vec{\chi}_n + K_3, t_n + h) \\
\vec{\chi}_{n+1} &= \vec{\chi}_n + \frac{1}{6} \cdot (K_1 + 2K_2 + 2K_3 + K_4)
\end{aligned}$$

where  $f(\vec{\chi}, t)$  is the function that returns  $\dot{\vec{\chi}}$  given the state vector at time  $t_n$ , and  $h$  is the size of the time step. The equations of motion for the micro air vehicle are time independent, and are given in equation 17 above. Thus, the algorithm above is used to compute the state of the vehicle at each time step using the state at the previous time step.

A computer simulation was written using the equations of motion and the Runge-Kutta algorithm. Given the physical properties of the vehicle, the initial state vector, and the size of the time step, the simulation will output the time history of the vehicle by returning the state vectors calculated at each time step. Currently, the only method for analysis of these nonlinear equations is to plot the various individual states versus time and visually examine the results. The simulation was used to examine the affects of a couple of design parameters. These parameters are the cant angle,  $\beta$ , and the vertical distance of the C.G. below the rotor hub. Initially, these parameters were studied using a root locus analysis of the two-dimensional linearized equations of motion. This led to choices of the vertical distance and cant angle that helped to stabilize the motion of the vehicle. (This analysis can be found in last year's report.) By assuming that the dynamic coupling between the roll and pitch axes ( $x$  and  $y$ ) is small in the three-dimensional model, the two-dimensional parameters can be applied directly to the three-dimensional model. These parameters proved to be satisfactory after running several cases on the simulator. A few final details of the simulation code must now be explained. In many of the aerodynamic force and torque equations above, the angular velocity of the rotors shows up. This parameter is neither a mass nor a geometric property of the vehicle, however, a value must be determined to solve the equations of motion. The problem was solved by calculating the angular velocity needed to produce enough lift to counteract gravity in the static hover case. The angular velocities in the simulation were then set constant to this static hover value. A final comment about the simulation is that performance depends on keeping the motion of the vehicle in a reasonable operating regime. For example, if the vertical motion of an individual rotor causes the inflow to drop below zero, then the aerodynamic equations derived above no longer hold and the simulation breaks down. This leads to non-realistic solutions.

The results of the simulation are plotted in the following figures. Each set of twelve plots represents the time history of the twelve states variables for one run. The state variables are the three components of the inertial position, the Euler angles representing the orientation of the vehicle in the inertial frame, the three components of the inertial velocity, and the three components of the angular velocity in the body frame. Four separate runs are plotted to show the affect of changing the vehicle parameters. Figure 5



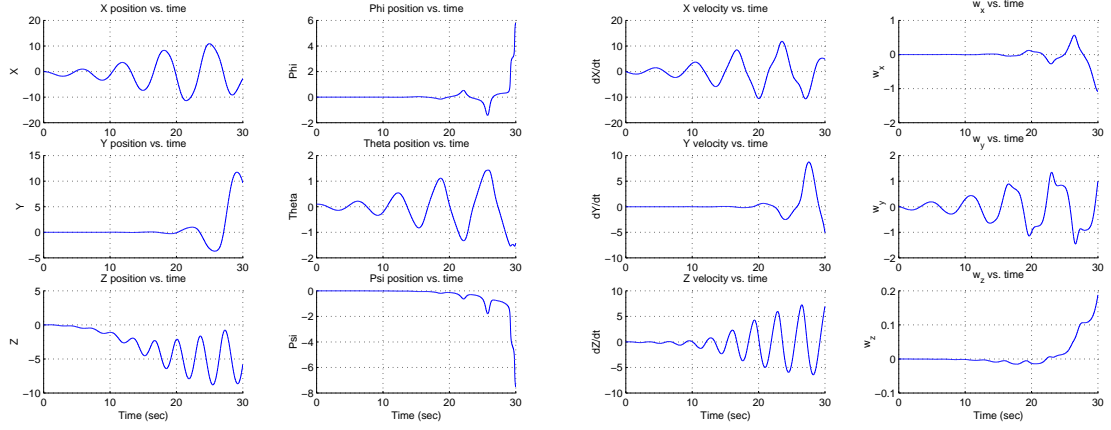
(a) Position and Orientation of the Vehicle.

(b) Linear and Angular Velocities the of Vehicle.

**Figure 5 Open-loop behavior of vehicle as built.**

shows the open-loop simulated behavior of the vehicle after an initial pitch forward. This run uses the parameters determined in the two-dimensional study. In this case, the motion of the vehicle does not damp out, but enters into a steady limit cycle. The assumption that coupling is small between roll and pitch is validated since the rolling motion is not excited in this example for several seconds. In this case, it can also be seen that the simulation breaks down due after about ten seconds. This is due to the problem with the reversed inflow described above. Further analysis showed that at this point the vehicle was indeed in that regime. Finally, note that the  $Z$  velocity oscillates about a steady positive value because the periodic roll and pitch motion causes the lift vector to be directed side to side rather than straight up. This causes a loss in lift and a subsequent decrease in the altitude of the vehicle as seen in the plot of the  $Z$  position (positive down). Figures 6 and 7 show the effect of changing the cant angle,  $\beta$ . The angle  $\beta$  is set to 0 and -15 degrees respectively. With  $\beta = 0$ , the motion no longer settles into a limit cycle, but will continue to grow unstable. With  $\beta = -15^\circ$ , the vehicle seems to have a steady state equilibrium which it converges to. Unfortunately this involves the vehicle rolling upside down  $180^\circ$  and flying straight into the ground. Figure 8 shows more instability when the C.G. is moved above the plane of the rotors. In this case the vehicle flips end over end a few times as seen in the plot of the roll axis ( $\phi$ ).

The next step is to develop a method to incorporate control law algorithms into the simulator for testing before they are actually programmed into hardware. Intuitively, the angular velocity of the rotors seems like an obvious choice for the control input. It would be desirable to come up with some control algorithm that follows the formula  $\dot{\vec{\chi}} = f(\vec{\chi}, \vec{u})$ , where  $\vec{u} = [\Omega_1 \ \Omega_2 \ \Omega_3 \ \Omega_4]^T$ . Unfortunately, without working through the entire nonlinear control analysis, only guesses can be made as to what the function  $f$  should look like. In order to try to quickly have a working platform to test control laws, it was decided to once again fall back on the decoupled linearized two-dimensional problem. In this problem, the only state variables are the  $X$  and  $Y$  position, the tilt angle



(a) Position and Orientation of the Vehicle.

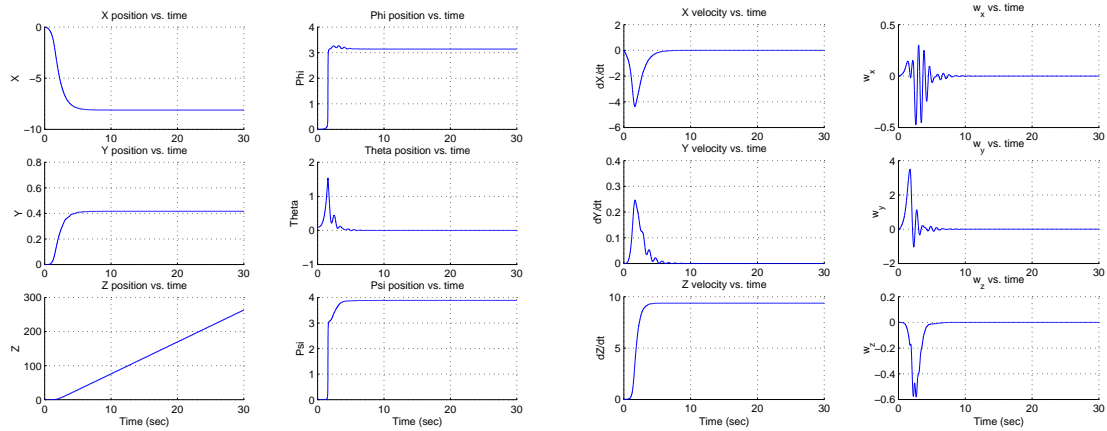
(b) Linear and Angular Velocities of the Vehicle.

**Figure 6 Open-loop behavior of vehicle with zero cant angle.**

$\theta$ , and their derivatives. In this linearized problem, the rotational equations decouple from the translational equations, allowing the tilt angle to be controlled independently. The control law can now be written in the standard form:

$$\begin{bmatrix} \ddot{\theta} \\ \dot{\theta} \end{bmatrix} = \mathbf{A} \cdot \begin{bmatrix} \dot{\theta} \\ \theta \end{bmatrix} + \mathbf{B}\vec{u},$$

with  $\vec{u} = -\mathbf{K} \cdot [\dot{\theta} \ \theta]^T$ . After deriving these two dimensional equations of motion, it turns out that the input vector corresponds to the change in thrust of the individual motors,  $\vec{u} = [\Delta T_1 \ \Delta T_2]^T$ . The gain matrix,  $\mathbf{K}$ , is found using LQR methods. To apply this solution to the three-dimensional case, some assumptions must be made. First, the assumption is again made that the coupling between roll and pitch is small so that the three dimensional problem can be split into a pair of two-dimensional problems. The second assumption made is that the linearized two-dimensional solution approximates the nonlinear two-dimensional problem reasonably for small deviations. This was also shown to be true in the previous two-dimensional analysis. Finally, the assumption that the change in thrust of the individual motors can be controlled directly is not entirely true. Instead, the desired change in thrust is determined using feedback of the angular rates. Then the aerodynamic equations are used to back out the required angular velocity that would produce that thrust. But the angular velocity is once again calculated using the equations for steady state hover. This does not result exactly in the desired feedback when the vehicle is in motion, but it is close enough. One final observation about must be made about using the angular velocities as the control inputs for the simulation. In reality, the hardware on board the physical vehicle will only control the voltage applied to the motors. Although a simple relation between the voltage and the angular velocity of the rotor is easy to find, some delay will occur between the time the voltage command is sent and the time the rotors actually achieve the desired angular velocity. Fortunately, the blade inertia decreases with the square of the radius. Therefore, the time delay decreases significantly



(a) Position and Orientation of the Vehicle.

(b) Linear and Angular Velocities of the Vehicle.


**Figure 7 Open-loop behavior of vehicle with negative cant angle.**

as the vehicle size gets smaller. For the smallest vehicles, this time delay is negligible. For the larger test bed vehicles, the delay starts to become more noticeable. With these assumptions made, it is now possible to insert an implementation of the roll and pitch control algorithms into the dynamic simulation. The feedback gain matrix from the two-dimensional simulation was used directly. Figure 9 shows the closed-loop response of the vehicle given the same initial conditions as the cases above. Compare this plot with with the open-loop motion shown in figure 5 to see the affect of the feedback. The roll and pitch angles are well-damped and quickly converge to zero. The actual test bed vehicle will incorporate rate gyros to measure the angular rates in the body frame. Currently, no method of measuring the actual roll and pitch angles will be incorporated. Another test simulation was performed to determine if it is possible to stabilize the motion using rate feedback alone. The results are plotted in figure 10. The feedback gain on the rates had to be increased by an order of magnitude, and the motion is still not fully damped after 30 seconds. This may cause problems on the actual hardware since any noisy data in the system will be amplified by a large gain before being fed back into the system. The following conclusions are made based on results from the dynamic simulation. Feedback control using the angular rates alone does not appear feasible as a solution to stabilize the vehicle motion. However, a more robust algorithm might be obtained by incorporating an estimate of the roll and pitch angles obtained by integrating the angular rates.

### 3.0 Flight Test Platform

In this last section, the vehicle test platforms are described. In order to test the control algorithms developed using the tools above, a physical flight test vehicle was designed and built. This vehicle was larger than the intended size of the micro air vehicle because the cost of components rises exponentially the smaller they get, and the lead time to order





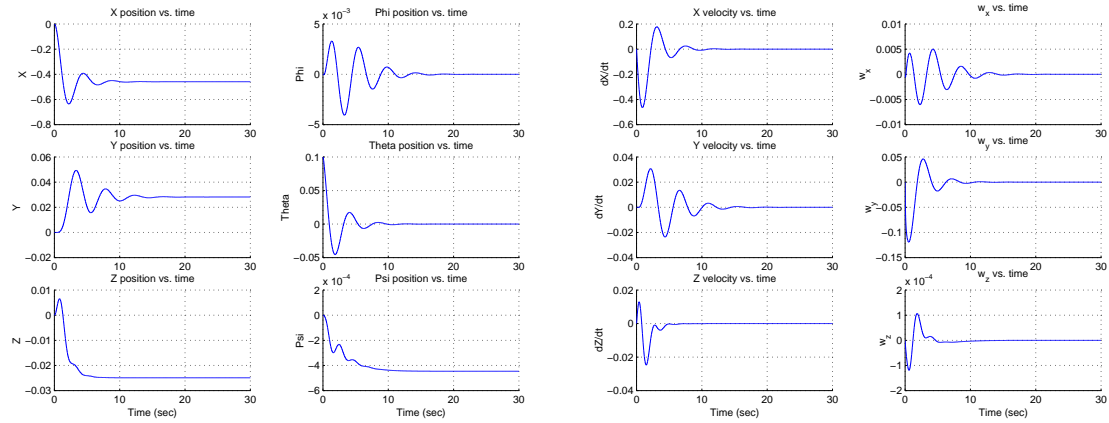
(a) Position and Orientation of the Vehicle.

(b) Linear and Angular Velocities of the Vehicle.

**Figure 8 Open-loop behavior of vehicle with the C.G. about the rotor plane.**

these specialized parts is long. For a larger vehicle, the parts are available over the counter at a reasonable cost. The flight test vehicle was designed to be a proof of concept stage, so that a stable platform could be developed, and the results transferred to a smaller vehicle later. The inherent dynamics would remain the same. Two important factors drove the design of the test vehicle: duration of flight and payload. These factors, and the details of the design implementation will be discussed below.

Flight duration is important because it is directly proportional to the amount of testing that can be done with the vehicle. More time could be spent refining the control algorithms rather than waiting for batteries to charge. Payload is important because it allows the vehicle to lift additional sensors and hardware. Research into currently available battery technologies showed promising candidates in the area of lithium based batteries. These batteries can have up to four times the energy density of NiCd batteries. One problem with lithium batteries is that they cannot output large quantities of current. The power required for the vehicle to hover is on the order of 10-11 watts. Each motor draws 8.5 volts which means the batteries need to continuously supply 1.3 amps. The only rechargeable lithium battery capable of meeting this amperage requirements was the Tadiran<sup>©</sup> INCHARGE Lithium metal rechargeable batteries. These batteries came in both AA and two-thirds AA size. Flight duration is also improved by choosing high efficiency motors with high power to weight ratios. The best off-the-shelf motor that provided the lift required was the Astro Firefly Coreless motor. The motor incorporates a 16:1 planetary gearbox which enables it to spin larger rotors, producing more lift, and increasing payload capacity. The manufacturer suggested mating the motors with Wes-Technik ten inch propellers which were tested using a thrust stand. The thrust produced was adequate, however the manufacturer only made counter-clockwise rotating rotors. In order to negate the motor torques in flight, the pairs of rotors must spin in opposite directions. To solve this problem, the rotor was digitally transferred into a CAD drawing, then be mirrored to produce a clockwise spinning rotor. In addition, the twist of the blade was reduced because the



(a) Position and Orientation of the Vehicle.

(b) Linear and Angular Velocities of the Vehicle.

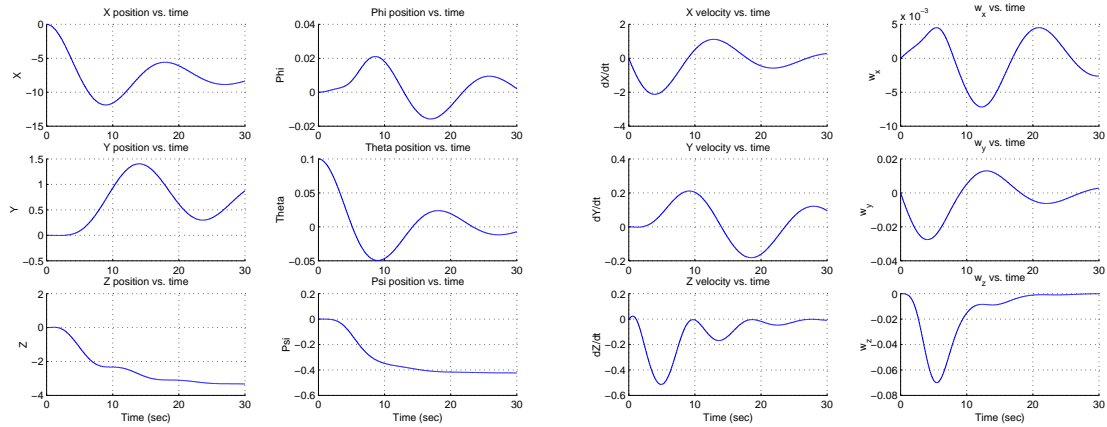
**Figure 9 Closed-loop behavior of vehicle.**

rotors were designed as aircraft propellers which would have a large inflow due to forward motion of the plane. Hovering vehicles, on the other hand, have rotors designed to operate at a stationary position. This refinement alone resulted in a 16% reduction in power which adds to the endurance. Two-sided aluminum molds were made of both the clockwise and counter-clockwise spinning rotors. The rotors are manufactured using laminating resin with multiple layers of carbon-fiber, fiber-glass, and Kevlar. The main arms and braces of the vehicle frame were also manufactured out of carbon-fiber. These were created by molding the carbon-fiber and resin into a beam with a U-channel cross-section, then closing off the open end with a thin piece of balsa. This produced a structure that was both rigid in bending and torsional loading. Using batteries with high energy density, efficient motors with a high power to weight ratio, and light weight materials for the rotors and structure left sufficient payload for mounting sensors and electronics while still maintaining long flight durations. The final design was capable of flight in excess of 15 minutes. Table 1 lists a breakdown of the weight of each component in the final design.

Component	Weight
Batteries	49 g
Motors	50 g
Rotors	19 g
Frame	10 g
Electronics	25 g
<b>TOTAL</b>	<b>153 g</b>

**Table 1 Weight Breakdown of Micro Air Vehicle**

The electronics on board the vehicle control the speed of the motors, and therefore the overall motion of the vehicle. The first iteration of the electronics was designed to test the

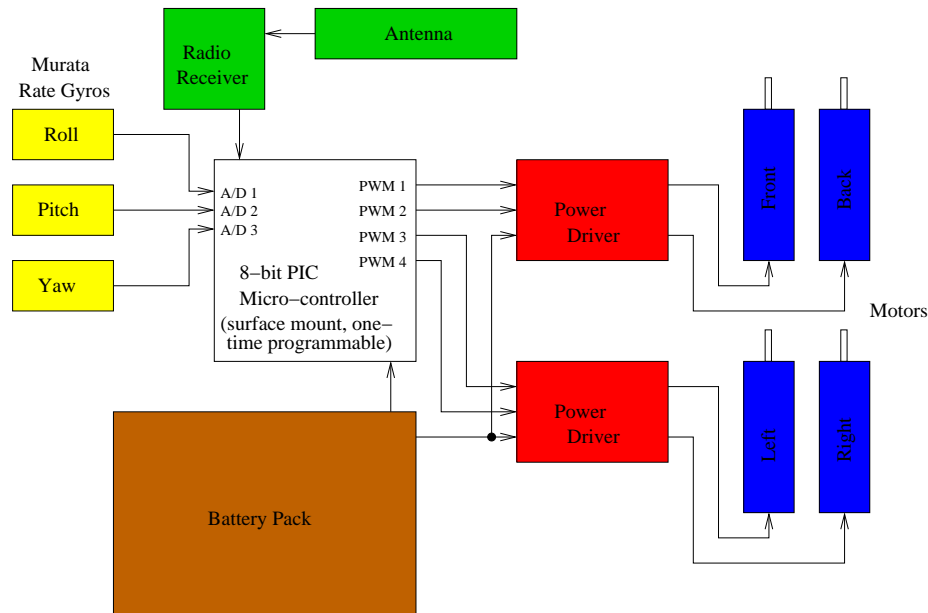


(a) Position and Orientation of the Vehicle.

(b) Linear and Angular Velocities of the Vehicle.

**Figure 10 Closed-loop behavior of vehicle as with rate feedback only.**


components and allow for a simple control law to be implemented. This consisted of only a microcontroller, power regulators to drive the motors at a desired speeds, a radio receiver to allow commands to be sent to the vehicle, and rate gyros to feedback the angular velocities. The rate gyros are Murata MEMS devices which weigh only one gram each, and are extremely sensitive to rate changes. The microcontroller chosen is an 8-bit PIC series chip with built-in 10-bit analog to digital (A/D) converters and pulse width modulation (PWM) capable of operating at 32 MHz. To maximize payload capacity, a one-time programmable version of the chip was used and soldered directly to the board. This chip weighs well under a gram, however in hind-sight, the reprogrammable version of the chip should have been used in conjunction with a socket even though this setup weighs 20 times that of the smaller chip. This would allow changes to be made to the on board software easily. A block diagram of the the interaction of the hardware components is shown in figure 11. The control law that was hard-coded into the microcontroller in this first version of the vehicle was a PID controller. The gains used in the PID algorithm on board the vehicle are sent from the transmitter at start up allowing for some experimentation of different settings. It is possible to convert the LQR gains from the simulation into PID gains enabling the results from the simulation to be tested on the vehicle. Figure 12 shows a photograph of the completed test platform vehicle with its transmitter and controller. Testing of this vehicle produced unsatisfactory results because of poor construction of the circuit boards which were sent to an outside manufacturer. In order to save weight, the thinnest board material was chosen, and evaluation boards were ordered which does not use the most robust manufacturing process. Also, the majority of the components in the circuit were surface mount because that happens to also be the lightest packaging option. These three factors caused the connections in the circuit to become intermittent on start up due to flexing and vibration of the boards caused by the spinning rotors. The circuit was also built on two separate boards connected in the middle by a bus. The bus also provides the only connection holding the boards together. The



**Figure 11 Block Diagram of Component Interaction**

stress at this center joint might also have caused intermittent connections.

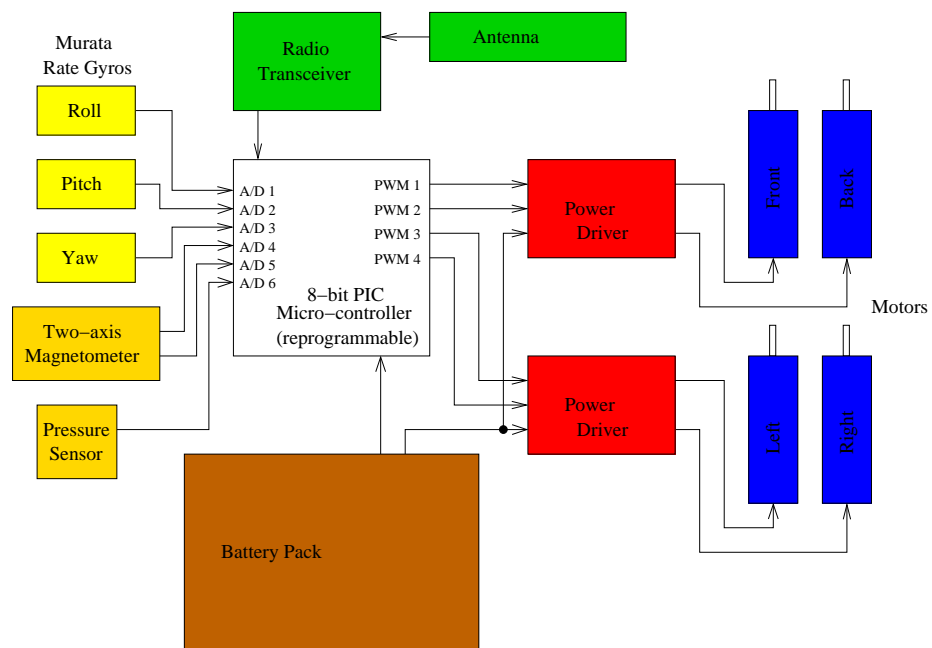
After testing the preliminary vehicle above, it was clear that a redesign was needed. This second revision of the vehicle was designed but never built. The two main goals of the redesign were to make the vehicle more robust and to make software revisions simpler. In addition, the opportunity was used to incorporate additional sensors. Thicker PCB material and more refined manufacturing processes were decided upon. In addition, the two board design would be replaced by a single board layout. Next, if the savings in weight was not significant, then components with leads were to replace surface mount parts. This included replacing the one-time programmable surface mount PIC microcontroller with its reprogrammable counterpart along with a socket. The unused A/D inputs on the PIC would be put to use reading the output of two magnetometers and a pressure sensor. The magnetometers could be used to calculate the magnetic heading providing feedback for stabilizing yaw motion. The pressure sensor would provide a direct measurement of the height, or *Z* position in the inertial frame. Using feedback from would minimize the loss in altitude due to maneuvering the vehicle. The final hardware change planned was to replace the receiver with a transceiver. With the current system, data could be sent to the vehicle, but no information was returned. The addition of a transceiver has two purposes. First it allows debugging of the on board software if something is not working. Currently, it is extremely hard to debug code in the on board software. Second, transmitting the current state of the vehicle back to a ground station allows this data to be stored and later analyzed. The results of a particular control algorithm could then be easily studied. A block diagram of the planned implementation of the second test platform is shown in figure 13. Future revisions of the vehicle beyond this might include a miniaturized GPS sensor if it becomes available. Absolute position feedback used in conjunction with the other sensors would allow the entire low level control of the vehicle to be controlled by the a computer leaving the operator free to



**Figure 12 Photograph of Test Platform Vehicle and Transmitter**

perform higher level task management such as mission planning.

The last paragraph of this report is dedicated to the final vehicle that was built. Once testing of the first vehicle ended, the components were removed from the board leaving the frame intact. The goal was to show that the vehicle itself was airworthy. Therefore, a completely analog solution was built instead of the redesign above. This vehicle was completely built out of off-the-shelf components because these would already be tested for use on board a moving vehicle. These components included a micro-receiver, miniature helicopter gyros, and commercial speed controller's. The gyros directly augment the PWM signals from the receiver, which are then sent to the speed controllers and converted to voltage levels for the motors. The gyros only provide rate feedback (internally using the same Murata MEMS gyros described above) with a single gain adjustment. The existing motors and batteries were kept since they were commercial components. An 8-channel programmable transmitter was used to control the vehicle by programming the mixing scheme described in figure 2. The old vehicle frame was reused to mount the new components. This vehicle was successfully flown, and is capable of hovering nearly 10 feet above the ground. The learning curve for operating the vehicle is very steep, and even then the workload on the pilot is extremely high. In addition, there is no means of upgrading the control systems on board or incorporating digital feedback through a microcontroller. Nevertheless, the successful flight test is promising. Given a robust design with simple methods for making revisions, it is possible to implement a solution using a microcontroller as described above.



**Figure 13 Block Diagram of Redesigned Layout**

# Chapter 6

## Closed Loop Vehicle Control via Optical Tracking and Ground-Based Processing

### 1.0 Introduction

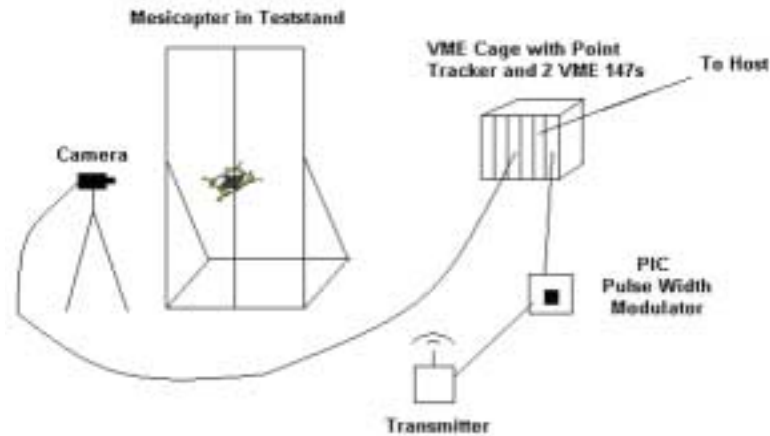
From the onset of this project, one of our goals has been to create a fully autonomous miniature hover vehicle. We focused this year on developing a software infrastructure capable of integrating feedback from multiple sensors, applying advanced control strategies, and commanding the vehicle using various techniques. We were able to complete an infrastructure, which when supplemented with enough sensing capability, can be extended to control all of the vehicle's degrees of freedom (6). A new modular vision sensor was designed to fit within this infrastructure. Utilizing our latest infrastructure and sensing capabilities, we enhanced our closed loop control demonstration, where the hover vehicle's yaw angle and height are now both regulated on a shaft.

In this section, we highlight the inadequacies of the control platform developed last year, which regulated the height of our 30g hover vehicle on a shaft. That infrastructure could not have been scaled to support 6 degree of freedom (DOF) control or possibly even 2 DOF control. We illustrate the changes we made to the system to improve control and sensing and show how the new system was used to provide height and orientation regulation on a shaft. We give a background on the logical control procedure we followed which involved identification, simulation, and experimentation. We give an analysis of how our techniques were successful, but show the difficulties involved in using these techniques to enable 6 DOF control. External vision alone probably could not provide enough sensor information for that application. Nevertheless, without external vision as a supplement, a practical 6 DOF control solution may not be readily available due to the limitations of micro-sensor technology.

### 2.0 Previous Work (Through June 2000)

In May 2000 we had designed a system for closed loop feedback control of a 30g miniature hover vehicle restricted to motion along a wire. In June 2000, we had closed the loop and enabled vertical control of the vehicle, disregarding the orientation for the moment. We called this our 1 DOF closed loop control testbed. The experimental apparatus for this testbed is shown below in Figure 1. For a more detailed description of the 1DOF testbed, see the June 2000 midterm report. We provide here a brief description of the testbed.





**FIGURE 1. The Antiquated 1-DOF Closed Loop Control Testbed**

Infrared LED's were fixed to the vehicle. The camera, monochrome and sensitive to infrared light, was placed in a position suitable to track the LED's. A point tracker board, a hardware vision solution, was dedicated to processing the incoming frames from the camera. The board tracked 1 LED on board the vehicle (although we had two shining at the same height, only 1 was tracked). The LED's position within the frame coordinate system was sent to the first of 2 Motorola 147 boards, acting as a vision server. The vision server had a number of capabilities including geometrical configuration tracking of multiple points (body tracking). Originally, we did not use the advanced features of the board (as we were interested in tracking a single LED). Control software, provided by Real Time Innovations, ran on the other Motorola 147 board. This software polled the location of the LED from the vision server. The control software calculated the state of the global state of the vehicle assisted by various transformations. Using a PD control law and adding input sufficient to eliminate the gravity bias, a desired control input was obtained. This input was output from the VME control board via a serial port and sent to a dedicated PIC processor, which created the associated pulse width modulation (PWM) signal required to command the on-board motors accordingly.

We could have used this testbed to control the robot in 2 DOF, but as will be shown in the next subsection, there were various problems, and we decided that many portions of the architecture needed to be redesigned. Areas that would be problematic included the processing power, sensor expandability, the vision system, and the transmitter/receiver pair.

### 3.0 Latest Control Testbed

The vision system for our original testbed was based on outdated technology. In the late 80s, the Aerospace Robotics Lab at Stanford had developed a dedicated VME hardware board capable of tracking high intensity dots within a camera frame. This hardware board was developed because processors were not capable of performing vision processing in



software at 30Hz. The board developed had special processing chips that parallelized the process of detecting dots. A VME processing board ran vision server software in parallel with the dedicated board. The vision servers task was to request the state of the dots in the vision frame, sometimes process the geometric configuration of the dot patterns, and send the state of the tracked objects to vision clients. The server could provide information about individual dots and/or bodies made up of 3 or more points. For body tracking, the server was given a geometric description of the body, and the server attempted to fit the dot patterns to these bodies. The server returned the 2-D coordinates of a body reference point along with the orientation of that body. Another VME processing board ran the program that used the tracking information. Often this program was controls based, like the program we originally used to regulate the vehicle's height on the shaft.

The point tracker/vision server/dedicated control strategy used in the late 80's was often at the peak of it's capacity when tracking more than a few dots. Already, modern vision tracking and control could be performed on a single processor PC with a frame-grabber. Not only would this system be faster, but one could do much more interesting vision algorithms like blob tracking which one cannot do merely by tracking points. A new system could also track more features. Since we thought we might need more intense algorithms than point tracking, we felt that the old vision server would not be adequate.

Additionally, using Motorola 147 boards would be out of the question when one wanted to do extreme matrix processing for the 6 DOF or even 2 DOF control at 30Hz. A Pentium chip would provide us with all the processing power we needed.

We felt we needed a new solution and decided to implement one. From July 2000 to November 2000, we performed work to replace our infrastructure. Since our control software, ControlShell supported multiple architecture, we could transfer our single DOF control software to Linux, and revamp the vision segment of the software to support our new infrastructure.

Our latest infrastructure would use a Linux based PC instead of the VME cage to perform all processing. The vision tracking hardware was replaced with a PC/frame-grabber pair. All vision tracking algorithms were performed in software making the system much more flexible.

We had to make another temporary change to our system necessary to control the motors on the hover vehicle individually. For the original demo, we had output the control command through the serial port of the VME processing board. This command was received by a PIC processor which output an 8 channel pulse width modulated signal. This signal was sent through the back of a Futaba RC radio, which sent the signal to a receiver on the hover vehicle. The receiver broke apart the channel commands and sent them to their individual motor controllers.

We sent the pulse width signal into a trainer port on the Futaba radio that would nominally be connected to another Futaba radio. To model the trainer port signal, we connected another control into the trainer port and analyzed the pulse width signal sent. We attempted to get a specification of the signal, but Futaba refused to send us one. After

careful analysis, we felt we could recreate the signal without Futaba's help. We implemented a PIC chip, with the purpose of cloning the signal.


In the 1 DOF case, we gave each channel the same command although we only needed the first one. This method worked as we were able to regulate height exactly. When we tried to use separate command on all the channels, we came across severe problems.

When using only the first channel, our imprecise model of the timing injected a small amount of noise into the command, which, in this case, was insignificant. When controlling all eight channels independently, we found that the noise was no longer insignificant. Since the signals were timing based, the error of the first channel not only affected the first channel but affected the second channel and all channels thereafter. The system was subject to an incredible amount of noise.

Without the specification from Futaba, we could not proceed with that transmitter solution any further. Instead, we decided to drive the motors by thin wires. The voltage on these motors would be driven off amplifiers, which would be commanded by a digital to analog board coming off the PC.

Besides being noise resistant, this solution was better for additional reasons. Once we moved to Linux, we could not get the serial port to communicate 8 channels faster than 10 Hz. By using a digital to analog solution, we were able to up our sampling rate to 30 Hz (the camera's rate).

Our latest system is illustrated in Figure 2 below. Comparing it to our original system, you can see that we replaced the VME cage with a Linux box and removed the PIC processor/transmitter from our diagram. We instead used digital to analog converters to drive the motors directly. That solution is adequate until we can integrate our own transmitter/receiver of which a complete hardware specification is available. We kept in mind that we would probably be moving to a larger vehicle for the 6 DOF control and adding additional sensors later. At that time, we would add a new transmitter/receiver solution.



**FIGURE 2. The Latest Feedback Control Testbed**

In the following subsections, we provide analysis on the software and control strategies used within our latest testbed. We start with a description of the modular vision software we recently developed.

### **3.1 Vision System**

Our latest control system heavily depends on the accuracy of our external vision sensor. Not only does our sensor have to give accurate state estimates for the vehicle to be regulated in 2 DOF, it must have the flexibility to incorporate multiple tracking algorithms and cameras if we attempt to use vision to help solve the 6 DOF tracking problem.

Our previous VME hardware vision solution would not be suitable for 6 DOF and probably 2 DOF sensing. We needed to overhaul our vision system, but in doing so still take advantage of the many vision algorithms and drivers already available. We desired to create a vision system that would be modular, extendable into the future, and suitable for open source. We also intended our solution to be relatively architecture independent so it could be used on multiple platforms.

We decided to develop our own modular vision system within the ControlShell software infrastructure. ControlShell supports all major platforms and has graphical reusable components that are actually C++ objects. Some of the code we would write would have to be architecture dependent like components utilizing frame-grabber device drivers, but those objects could be of a special component type and could still connect to the generic processing components.

Decisions were made as to the types of components a vision system would need. We decided on five base types of components, which are inherited from their respective base classes. The component types would be acquisition, processing, state, network, and display. All components would be connected via a common vision interface.


Acquisition components are OS/device specific (the only architecture specific component) and are wrappers for specific frame-grabber device drivers. Acquisition components provide a memory reference to the current image and the sampling rate (based on the frame rate). The interface hides the architecture specific nature of these components so most of the other components can be device independent.

Processing components manipulate an image and output a modified image. For example, we would provide a component to implement a common thresholding technique. The output image of a thresholding component would consist of binary pixels. If an individual pixel value of the entering image was above a certain threshold, the output pixel of the corresponding image would be 1. For pixel below the threshold, the output pixel would be 0.

State components provide information specific to an input image. For the hover vehicle project, state components calculate specific data about position, orientation, etc. of the vehicle from an image. State components are often application specific.

Network components are used to pass an image from one application to another (local or remote). One can functionally pass images. These components generally utilize some high bandwidth connection between two applications. In many applications this component is not useful because of the expense.

Display components are used to display the image on a monitor so one can get visual feedback of what they are tracking. These are often used for calibration and debugging. A common interface connects the different components. Looking at Figure 3 below, we can see an example of how these components may be connected.



**FIGURE 3. A Possible Component Level Layout with the ControlShell Vision Repository**

As hinted at above, the interface between these components is a perfect means for abstraction. The interface provides the pixel dimensions of the image associated with the interface, an API to perform functions necessary to get pixel values or image parameters, and does not change between two components. A processing component, for instance, may shrink an image, or change the individual pixels, but the image that comes out of it, although different, can be described using a new image interface. The processing component therefore, would provide all the abstract functions that an acquisition component might as an output although the component itself is not acquiring a frame from a camera. Many useful data paths can be created by stringing components together with various interfaces.

In ControlShell, we created the base classes that all components would inherit from and defined the interface. Next we created the first components of the ControlShell vision repository that could be used with our application.

Since we were using a Linux workstation as our software engine and had a Matrox Meteor I frame-grabber, we developed a acquisition component that utilized the Linux Matrox Meteor I device driver. In developing this component, we utilized some code previously written for the Yale XVision project under Greg Hager. Since XVision was a popular vision tracking package, we found a large quantity of code that was useful to our cause.

To be able to detect the position of LED's within the frame, we developed single point and multi-point tracking components that used thresholding as a means of detection. The single point tracker could be used to re-implement the single degree of freedom

demonstration on our latest testbed. The multi-point tracker was required to detect two or more LED's in the same image and thus move to 2 DOF tracking.

The 1 DOF tracker uses a single point thresholding algorithm. This algorithm is roughly the simplest piece of code in machine vision. The algorithm is trivial and does not merit much discussion.

To track multiple LED's and indirectly height and yaw, we developed a multi-purpose, multiple point tracking vision state component. The component has threshold, minimum pixel radius, points desired, and max points tracked inputs. We use an algorithm that processes each pixel with a value greater than the threshold. If a candidate pixel is a distance greater than the minimum pixel radius from any tracked point, the candidate point is added to the list of tracked points. The component tracks only one of the pixels within the pixel radius. If a point is close to another tracked point, we increment the intensity value for the tracked point. In the end, the component outputs the number of points desired even though it may have identified many unique tracked points (maximum - max tracked points). The selected output points have the highest intensity values (highest number of valid, adjacent pixels). We use the intensity value as a filter to eliminate noise.


We created a X Windows display component to view images real time. The code for this component was partially gotten from the XVision package. This component tended to be slow (we could not achieve 30Hz control while using it), but the component is useful for calibration and is a good proof of concept.

The components just listed are the ones we created for our use, but, recently members of the Robotics Lab, the Flight Research Group at Stanford, and the Aerospace Robotics Lab at Stanford are taking advantage of the reusable components and are adding to the repository with new Processing, Acquisition, and Display components.

### **3.2 Computing Vehicle State from Pixel State**

In the 2 DOF application the vision system outputs a pair of LED coordinates. Calculating height, yaw angle, and velocities (translational and rotational) from the LED coordinates is relatively straightforward. Before running the software, we had calibrated the system and recorded the pixel/world coordinate transformation matrix in software. During software execution, each tracked pixel's coordinates are converted to the global coordinate equivalent using the transformation matrix. We applied a perspective transformation with constant depth (the pixel/meter transformation is constant) to LED coordinates. Since the hover vehicle rotates about the shaft, it's LED's do not move in a plane. Nevertheless, the depth variations in the LED locations are negligible compared to the depth from camera to vehicle.

To calculate height and yaw, the scaled LED values are passed through geometric calculations. To calculate height, the math is trivial. We find the average height ( $z$ ) of the two pixels. Since they should both be at the same height, the calculation is unnecessary but rather a safe guard. To understand the yaw ( $\psi$ ) calculation, analyze the following diagram (Figure 4):



**FIGURE 4. Geometric Diagram for Yaw Calculation**

We calibrate the  $x$  position of the shaft and set it to zero. We compute the  $x$  position of the midpoint between the two LED's. Having measured the radius from the shaft to the midpoint ( $r$ ), we can calculate  $\psi$ .

$$\psi = \sin^{-1}\left(\frac{x}{r}\right)$$

Since the pixels are discrete, the height and yaw values are indirectly discrete. If we had used a cosine to calculate yaw instead of sine (by measuring the length between the two pixels rather than the midpoint), the resolution would have been poor near 0 yaw. A sine calculation makes smaller increments near 0 where a cosine jumps immensely. Using a sine improved the performance of the system.


To calculate translational and rotational velocities of the vehicle, we applied a first difference to successive yaw and height measurements. For 2 DOF case, the first differences were acceptable, but barely. Any less resolution within frames, and the velocity estimate wouldn't be precise enough. For the 6 DOF case, these filters would not be acceptable.

One last thing to note is that increasing the visible area in which the hover vehicle can be tracked decreases the available resolution. Yaw is extremely sensitive to resolution as it has only  $N$  possible values between  $[\psi/2 - \psi/2]$  where  $N$  is the total pixel length between two LED's. One must be careful how the camera is placed.

### 3.3 Dynamics

To regulate height and yaw we had to clearly define the dynamics of the vehicle on the shaft. Since the vehicle is restricted to motion on the shaft, gyroscopic effects and most rotational effects are eliminated. A simple set of equations govern the motion of the vehicle. The collective thrust provided by the four rotors controls the vertical acceleration. Since adjacent motors counter rotate, the torque differential between the four rotors

angular acceleration around the shaft. Shaft friction was ignored. The following diagram (Figure 5) provides a shaft coordinate reference.



**FIGURE 5. Hover Vehicle in Shaft Coordinate System**

Since the rotors have a significant time constant and are limited by drag, we've considered these effects in our dynamic analysis. Alternately, we assume that the time constant associated with voltage to torque is zero.

Drag and lift are considerable forces in our model, but we approximate them linearly although the effects are nonlinear. Since the curves associated with these effects are generally linear through our operating modes, we deem these approximations to be sufficient.

We've modeled the system with state space equations and an additional gravitational bias. Note that each set of two motors with the same rotational direction receive the same torque command. This constraint simplifies calculations. We use the following plant equations:

$$\dot{x} = Ax + Bu + g$$

$$y = Cx$$

where,

$$\begin{aligned}
 x &= \begin{bmatrix} z \\ \dot{z} \\ \psi \\ \dot{\psi} \\ \omega_A \\ \omega_B \\ \omega_C \\ \omega_D \end{bmatrix} \quad u = \begin{bmatrix} \tau_A \\ \tau_B \end{bmatrix} \quad g = \begin{bmatrix} 0 \\ 9.8 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix} \\
 A &= \begin{bmatrix} 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & \frac{-a}{m} & \frac{-a}{m} & \frac{-a}{m} & \frac{-a}{m} \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & \frac{-b}{I_r} & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & \frac{-b}{I_r} & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & \frac{-b}{I_r} & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & \frac{-b}{I_r} \end{bmatrix} \quad B = \begin{bmatrix} 0 & 0 \\ 0 & 0 \\ 0 & 0 \\ \frac{-2}{I_z} & \frac{2}{I_z} \\ \frac{1}{I_r} & 0 \\ \frac{1}{I_r} & 0 \\ 0 & \frac{1}{I_r} \\ 0 & \frac{1}{I_r} \end{bmatrix}
 \end{aligned}$$

The following parameters are defined:

$z$  – the vertical position of the hover vehicle

$\psi$  - the orientation of the hover vehicle on the shaft

$\omega$  - the rotational velocity of a rotor. There are four rotors labeled A, B, C, and D

$a$  – Thrust per rotational velocity of each individual rotor

$b$  – Drag per rotational velocity of each individual rotor

$\tau$  - Torque input of each motor (Note that  $\tau_A = \tau_C$  and  $\tau_B = \tau_D$ )

$I_z$  – Moment of inertia of the vehicle (excluding the rotors)

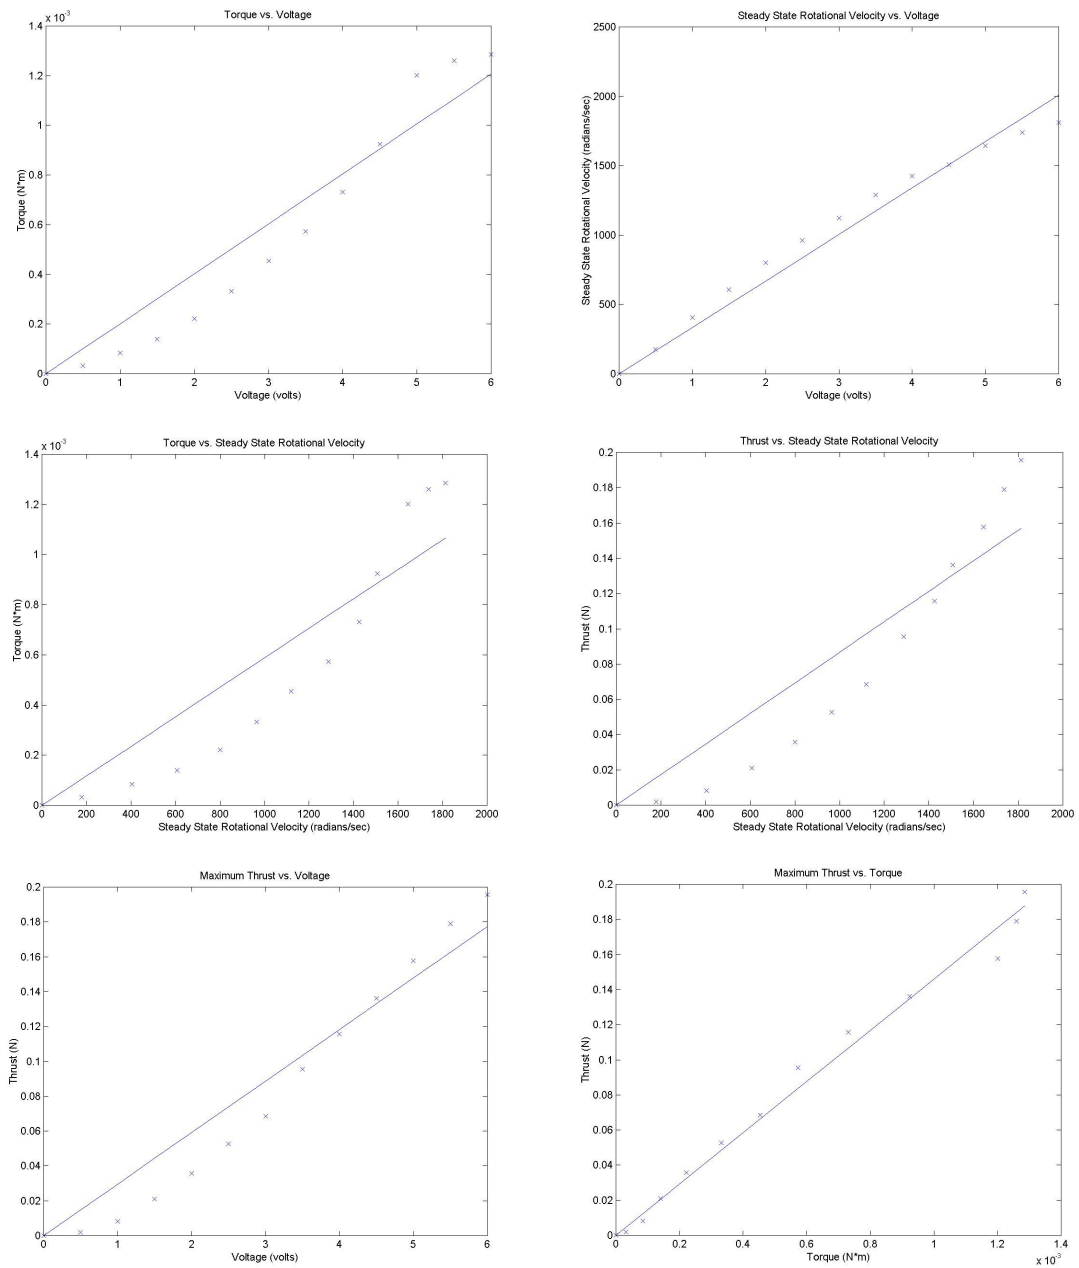


$I_r$  – Moment of inertia of each rotor

$m$  – Mass of the vehicle

### 3.4 Identification

To provide a quantitative model of the vehicle dynamics, we identified  $a$ ,  $b$ ,  $I_r$ ,  $I_z$ , and  $m$ . To extend our control model we also identified the saturation point of the motors. We did not want our gains to contribute to command levels which could saturate near our operating point. We weighed the vehicle and found the  $m$  to be 0.04 kg. We found the  $I_r$  to be approximately  $1.0 \cdot 10^{-7} \text{ kg} \cdot \text{m}^2$  and  $I_z$  to be approximately  $7.0 \cdot 10^{-5} \text{ kg} \cdot \text{m}^2$ . For the motor related parameters, we performed some experiments. We ran the motors (with rotors attached) at various voltage levels and measured the following quantities at various voltage levels: torque, steady state rotational velocity, and maximum thrust. Voltage levels tested were from [0, 6] at increments of 0.5 volts. The following plots in Figure 6 show the key relationships. Although all the information could be displayed in a spreadsheet, we show the graphs to emphasize the best fit lines we've placed to estimate linear relationships within the plots.



**FIGURE 6.**

With least squares we approximate the following slopes for the curves:

Maximum Thrust/Steady State Rotational Velocity of the Rotor( $a$ ):

$$8.6623 \times 10^{-5} \text{ N/(Radians/sec)}$$

Torque/Steady State Rotational Velocity of the Rotor ( $b$ ):

$$5.8817 \cdot 10^{-7} \text{ N}\cdot\text{m}/(\text{Radians}/\text{sec})$$

Torque vs. Voltage:

$$2.0112 \cdot 10^{-4} \text{ N}\cdot\text{m}/\text{Volt}$$

Steady State Rotational Velocity of the Rotor vs. Voltage:

$$334.8841 \text{ Radians}/(\text{sec}\cdot\text{Volt})$$

Maximum Thrust vs. Voltage:

$$0.0296 \text{ N}/\text{Volt}$$

Maximum Thrust vs. Torque:

$$145.9852 \text{ N}/(\text{N}\cdot\text{m})$$

With 6 volts being the maximum voltage we would put into the motors, our maximum thrust is roughly 20 grams per motor. Since the vehicle is 40 grams we can roughly carry the weight of two vehicles. We have enough extra control effort to provide sufficient control.

### 3.5 Feedback Control

To use a linear control law, the bias term had to be eliminated from the equations above. We eliminated the bias by linearizing around an equilibrium point. We define  $x_o$  as the desired equilibrium state and  $u_o$  as the equilibrium command (calculated). Additionally,  $\delta x$  and  $\delta u$  are deviations from those states. We redefine  $x$  and  $u$  as

$$x = x_o + \delta x$$

$$u = u_o + \delta u$$

The dynamic equations can be rewritten as:

$$(\dot{x}_o + \delta \dot{x}) = A(x_o + \delta x) + B(u_o + \delta u) + g$$

$$\delta \dot{x} = A\delta x + B\delta u + Ax_o + Bu_o + g$$

If we can find an  $u_o$  such that:

$$Ax_o + Bu_o = g$$

which we can in this case, we can eliminate  $g$  from the equations which become

$$\partial \dot{x} = A\partial x + B\partial u$$

We focus our efforts on producing a linear control law around these reduced dynamics (calculate  $\delta u$ ) and then add the force  $u_o$  to account for gravity before applying the control. As long as the equilibrium linear control law is stable, the control should be adequate.

From the identified dynamics we find our equilibrium state and control to be:

$$x_o = \begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \\ 1131 \\ 1131 \\ 1131 \\ 1131 \end{bmatrix} \quad u_o = \begin{bmatrix} 6.5 * 10^{-4} \\ 6.5 * 10^{-4} \end{bmatrix}$$

To identify the localized control, we use the familiar state space control law:


$$\partial u = -K\partial x$$

Because we can only sense  $z$  and  $\psi$ , and their first order differences, without an estimator, we cannot provide full state feedback. Regardless, the dynamics are such that feedback of these two states and associated velocities one can produce a stable, robust controller for the 2 DOF hover vehicle.


The first step to finding a controller is to use a zero order hold on the analog plant model. At 30 Hz, we get a digital model in which we design our control around. Since the height and yaw control are coupled only at the inputs, we can create separate controllers for each knowing that, when combined, the system will be made up completely of the eigenvalues from the separate controllers.

We first look at the thrust dynamics associated with the height. Ignoring the dynamics

associated with the time constant of the rotors, we can simulate the effects of command vs. vertical control. With the input directly into  $\ddot{z}$ , we have a second order system. We desire the poles at 0.8 and 0.81. With these pole locations, the Matlab place command returns a set of gain values, [-0.0023 -0.0008], which will be multiplied by  $\partial z$  and  $\partial \dot{z}$ . Figures 7 and 8 display the height response and the input control of this system with zero rotor time constant.




**FIGURE 7. Simulated Vertical Response of Uncoupled System with Zero Rotor Time Constants**




**FIGURE 8. Individual Motor Torques for Simulated Vertical Response of the Uncoupled System with Zero Rotor Time Constants**

When we use the same gain values and add in the real time constants associated with drag of the rotors, we get a different response, as shown in Figures 9 and 10. With the time constants and these gain values, some dampened oscillatory behavior has been induced, but since we don't have feedback information about the rotational velocity of the rotor, we can't eliminate that mode. We stick directly with the gain values found from the zero rotor time constant case.




**FIGURE 9. Simulated Vertical Response of Uncoupled Linear System**




**FIGURE 10. Individual Motor Torques for Simulated Vertical Response of the Uncoupled Linear System**

We next try to set the gains attributed to the yaw portion of the system. Since the yaw is driven directly from torques on the motors, we don't have any time constants. The yaw system is a  $1/s^2$  plant with inputs being the sum of the torques from the 4 motors. Since two of the motors counter rotate, we can deliver a large range of collective torque centered at 0.


With the input directly into  $\ddot{\psi}$ , we have a second order system. We desire the poles at 0.8 and 0.81 as above. With these pole locations, the Matlab place command returns a set of gain values,  $[-0.00060 \ -0.00019]$  for motors A and C and  $[0.00060 \ 0.00019]$  for motors B and D. Each of these gain matrices are multiplied by  $\partial\psi$  and  $\partial\dot{\psi}$  and the output torque is to be applied to each motor. Figures 11, 12, and 13 are plots of the yaw response and the input control of this system.



**FIGURE 11. Simulated Yaw Response of Uncoupled Linear System**



**FIGURE 12. Motor Torques for Motors A and C in Simulated Yaw Response of the Uncoupled Linear System**




**FIGURE 13. Motor Torques for Motors B and D in Simulated Yaw Response of the Uncoupled Linear System**

Combining the gain results for the uncoupled system, we produce a set of gains that can be used for the local control of the real system.

$$K = \begin{bmatrix} -0.0023 & -0.0008 & -0.0006 & -0.00019 \\ -0.0023 & -0.0008 & 0.0006 & 0.00019 \end{bmatrix}$$


## 4.0 Expected Results

With all control inputs available, we simulated the system. Into this simulation not only did we use the dynamic model, but we added in the nonlinear affects of saturation and discretization resulting from the resolution of the sensors. Additionally, we added sensor noise (pixel variance = 2). With an initial offset of 0.5 meters for the height and 0 initial conditions for the rest of the state variables, the results achieved are displayed in Figures 14 and 15.



**FIGURE 14. Torque Inputs for Motors A/C and B/D**





**FIGURE 15. Full Simulation Results**

The simulation shows a stable response by the hover vehicle with a small jitter resulting from the sensor noise. This jitter is relatively insignificant though, and we expected this control law to work on the actual vehicle.

## 5.0 Experimental Results

When the control was applied to the physical testbed, the experiment was a success. The vehicle regulates its orientation and height in a stable fashion. The wires connected to the vehicle restricted some motion in yaw and also provide natural dampening, which adds some error to the system. Regardless, when the wires are lifted away, one can see the vehicle staying stable.

A vertical step response was added into the software and the vehicle moved up and down a 0.3 meters every 5 seconds. This was a proof of concept of the vehicles stability. A video was shown at a recent presentation showing the demonstration I discuss. It clearly shows the stability of the hover vehicle, and one is urged to request the video until more detailed

numerical results from physical experiments becomes available. Contact Chad Partridge at [cpartrid@stanford.edu](mailto:cpartrid@stanford.edu) to request a MPEG copy of the video.

## 6.0 Analysis

Although the development of a 2 DOF testbed was successful, there are a number of key issues to consider before moving to 6 DOF. Although an external vision server was sufficient to track all the necessary states in two DOF, this vision server can only be a supplement to a much larger sensor strategy for the new vehicle. The vision sensor does not have the resolution at large depths to provide solid state information for all moments. The vehicle will have to have more on-board sensors like gyros and accelerometers to supplement the vision information. Two degrees of freedom was difficult to enable using the external vision sensor. The sensor strategy for six must be analyzed carefully before proceeding.

Although more sensors would be needed, vision may have to be depended on for a positioning solution. If calibrating the vision sensor using multiple camera became problematic, a micro GPS solution could be a worthy replacement, but we are unclear as to how soon devices, which meet payload requirements, will be available for these miniature vehicles. For years people could not get helicopters to fly autonomously. It was only when a GPS solution was uncovered could people sense the helicopter's state reliably. Nevertheless, with centimeter level accuracy, a GPS solution might not work for small hover vehicles. The resolution of the states may not be fine enough to accurately predict where the vehicle is. For that reason, vision should be analyzed further as a solution. On board vision could certainly be the best solution, except our vehicles lack the processing power to perform the vision algorithms on board. Sending the image across to an off-board workstation is far too expensive figuring we'd need bandwidths of 10MB/sec. to analyze common image formats (RS-170). Even though off-board vision solutions may be problematic, the current state of technology dictates that off-board vision might be the only solution.

Sensing may be the only problem limiting our ability to fly these vehicles autonomously. Even these 40g vehicles have the capacity for a few more lightweight sensors. We've made vehicles slightly larger to assure we'd have more capacity. For the larger vehicles, power is not as important of an issue. For these smaller vehicles, power has generally been our limiting factor so with the power issue solved, sensing comes to the forefront.

Regardless, the software and computer infrastructure we have created for the two degree of freedom case can easily be extended to six with little difficulty assuming sensor information can be communicated from the vehicle and somehow accessed through the PCI bus (wireless Ethernet). If the sensors do become available, we probably have all the tools necessary to put a 6 DOF hover vehicle autonomously in the air.