

2007

A technique for shape optimization of ducted fans

David Firel Schaller
Iowa State University

Follow this and additional works at: <https://lib.dr.iastate.edu/rtd>



Part of the [Aerospace Engineering Commons](#), and the [Mechanical Engineering Commons](#)

Recommended Citation

Schaller, David Firel, "A technique for shape optimization of ducted fans" (2007). *Retrospective Theses and Dissertations*. 14799.
<https://lib.dr.iastate.edu/rtd/14799>

This Thesis is brought to you for free and open access by the Iowa State University Capstones, Theses and Dissertations at Iowa State University Digital Repository. It has been accepted for inclusion in Retrospective Theses and Dissertations by an authorized administrator of Iowa State University Digital Repository. For more information, please contact digirep@iastate.edu.

A technique for shape optimization of ducted fans

by

David Firel Schaller

A thesis submitted to the graduate faculty
in partial fulfillment of the requirements for the degree of
MASTER OF SCIENCE

Major: Aerospace Engineering

Program of Study Committee:
R. Ganesh Rajagopalan, Major Professor
Hui Hu
Richard Pletcher

Iowa State University

Ames, Iowa

2007

Copyright © David Firel Schaller, 2007. All rights reserved.

UMI Number: 1443112



UMI Microform 1443112

Copyright 2007 by ProQuest Information and Learning Company.
All rights reserved. This microform edition is protected against
unauthorized copying under Title 17, United States Code.

ProQuest Information and Learning Company
300 North Zeeb Road
P.O. Box 1346
Ann Arbor, MI 48106-1346

TABLE OF CONTENTS

LIST OF TABLES	vi
LIST OF FIGURES	vii
ACKNOWLEDGMENTS	xii
ABSTRACT	xiii
CHAPTER 1. Background	1
1.1 Ducted Fan History	1
1.2 Genetic Algorithm Optimization	3
1.3 Goal of Research	4
CHAPTER 2. Formulation	5
2.1 Genetic Algorithm Primer	5
2.1.1 Basic GA Terminology	6
2.1.2 Schemata	9
2.1.3 Selection	10

2.1.4	Crossover	11
2.1.5	Mutation	12
2.2	Genetic Algorithm Enhancements	13
2.2.1	Selection and Gene Pool	14
2.2.2	Modified Elitist Strategy	16
2.2.3	Population Regeneration	17
2.2.4	Population Scaling	18
2.3	Fitness Evaluation	18
2.4	Duct Parameterization	19
2.5	Non-Uniform Rational B-Splines	21
2.6	Axis-Symmetric Flow Solver	24
2.6.1	Derivation and Discretization of the Flow Governing Equations	25
2.6.2	Discretization of the Generic Governing Equation	26
2.6.3	Numerical Algorithm	27
2.6.4	Rotor Modeling	29
2.6.5	Boundary Conditions	30
2.7	Duct Thrust	30
2.8	Solution Mapping	33
CHAPTER 3. Methodology		36

3.1	Flow Field Similarity	36
3.2	Multi-Modal Test Functions	39
3.2.1	Ackley's Function	40
3.2.2	Schwefel's Function	40
3.2.3	Genetic Algorithm Comparisons	41
3.3	Panel Method Evaluation	43
CHAPTER 4. Results		49
4.1	Computational Domain	50
4.2	Duct Shape Control	50
4.3	Single Rotor Hover	53
4.3.1	Set-Up	53
4.3.2	Optimizer Metrics	54
4.3.3	Optimum and Baseline Flowfield Comparisons	54
4.3.4	Verification of Optimum Conditions	62
4.4	Coaxial Rotor Hover	67
4.4.1	Set-Up	67
4.4.2	Optimizer Metrics	68
4.4.3	Optimum and Baseline Flowfield Comparisons	69
4.4.4	Verification of Optimum Conditions	76

4.5	Single Rotor Axial Flight	80
4.5.1	Set-Up	80
4.5.2	Optimizer Metrics	80
4.5.3	Optimum and Baseline Flowfield Comparisons	81
4.5.4	Verification of Optimum Conditions	88
CHAPTER 5. Conclusions		94
5.1	Optimizer Performance	95
5.2	Ducted Fan Optimization	95
5.3	Future Research Opportunities	98
APPENDIX . Gradient Based Routine		100
BIBLIOGRAPHY		102

LIST OF TABLES

Table 4.1	Effect of tip gap size for hovering duct and rotor thrust.	63
Table 4.2	Effect of tip gap size for hovering system thrust.	63
Table 4.3	Effect of tip gap size for coaxial duct and rotor thrust.	76
Table 4.4	Effect of tip gap size for coaxial system thrust.	77
Table 4.5	Effect of tip gap size for axial flight duct and rotor thrust.	89
Table 4.6	Effect of tip gap size for axial flight system thrust.	89
Table 5.1	Comparison of ducted fan lift.	96
Table 5.2	Total system thrust and contributions.	97
Table 5.3	Performance increase of optimized ducts and rotors.	97
Table 5.4	Performance of optimized ducted fan systems.	98

LIST OF FIGURES

Figure 2.1	Basic genetic algorithm process flow.	7
Figure 2.2	Schematic of roulette wheel for selection.	11
Figure 2.3	Schematic of the crossover scheme.	12
Figure 2.4	GenII-GA process flow chart.	14
Figure 2.5	Bit-matching logic.	15
Figure 2.6	Fitness evaluation process flow.	19
Figure 2.7	Gene locations within the chromosome.	21
Figure 2.8	NURBS curve with control polygon.	24
Figure 2.9	Duct normal vectors.	31
Figure 2.10	Duct panel orientation possibilities.	33
Figure 2.11	Grid mapping formulation.	35
Figure 3.1	Original baseline solution.	38
Figure 3.2	Differenced flow fields.	39
Figure 3.3	Solution space for Ackley's function.	41

Figure 3.4	Solution space for Schwefel's function.	42
Figure 3.5	Comparison of simple GA and GenII-GA for Ackley's function. . .	43
Figure 3.6	Comparison of simple GA and GenII-GA for Schwefel's function.	44
Figure 3.7	Control polygon for the panel method airfoil optimization.	46
Figure 3.8	Fitness history for panel method evaluation.	47
Figure 3.9	Strongest performing airfoils for panel method evaluation.	48
Figure 3.10	Weakest performing airfoils for panel method evaluation.	48
Figure 4.1	Grid generation schematic.	51
Figure 4.2	Single rotor generalized control polygon.	52
Figure 4.3	Coaxial rotor generalized control polygon.	53
Figure 4.4	The maximum and average population fitness as a function of generation.	55
Figure 4.5	Axis-symmetric view of the velocity field for the baseline shape. .	57
Figure 4.6	Axis-symmetric view of the velocity field for the optimized shape.	57
Figure 4.7	Axis-symmetric view of the velocity field near the tip gap on the baseline shape.	58
Figure 4.8	Axis-symmetric view of the velocity field near the tip gap on the optimized shape.	58
Figure 4.9	3-D view of the velocity field for the baseline shape.	59

Figure 4.10	3-D view of the velocity field for the optimized shape.	59
Figure 4.11	Axis-symmetric view of pressure field near the baseline duct. . .	60
Figure 4.12	Axis-symmetric view of pressure field near the optimized duct. .	60
Figure 4.13	3-D view of the body pressure on the baseline duct.	61
Figure 4.14	3-D view of the body pressure on the optimized duct.	61
Figure 4.15	Flowfield of the optimized duct moved 1ε closer to the rotor. . .	65
Figure 4.16	Flowfield of the optimized duct.	65
Figure 4.17	Flowfield of the optimized duct moved 1ε away from the rotor. .	66
Figure 4.18	Flowfield of the optimized duct moved 2ε away from the rotor. .	66
Figure 4.19	The maximum and average population fitness for the coaxial rotor optimization.	68
Figure 4.20	Axis-symmetric view of the velocity field for the ISU baseline coaxial duct.	71
Figure 4.21	Axis-symmetric view of the velocity field for the optimized coaxial shape.	71
Figure 4.22	Axis-symmetric view of the velocity field near the tip gap on the ISU baseline coaxial duct.	72
Figure 4.23	Axis-symmetric view of the velocity field near the tip gap on the optimized coaxial shape.	72
Figure 4.24	3-D view of the velocity field for the ISU baseline coaxial duct. .	73

Figure 4.25	3-D view of the velocity field for the optimized coaxial shape. . .	73
Figure 4.26	Axis-symmetric view of pressure field near the ISU baseline coaxial duct.	74
Figure 4.27	Axis-symmetric view of pressure field near the optimized coaxial duct.	74
Figure 4.28	3-D view of the body pressure on the ISU baseline coaxial duct.	75
Figure 4.29	3-D view of the body pressure on the optimized coaxial duct. . .	75
Figure 4.30	Flowfield of the optimized coaxial duct.	78
Figure 4.31	Flowfield of the optimized coaxial duct moved 1ϵ away from the rotor.	79
Figure 4.32	Flowfield of the optimized coaxial duct moved 2ϵ away from the rotor.	79
Figure 4.33	The maximum and average population fitness for axial flight optimization.	81
Figure 4.34	Axis-symmetric view of the velocity field for the baseline duct. .	83
Figure 4.35	Axis-symmetric view of the velocity field for the optimized duct in climb.	83
Figure 4.36	Axis-symmetric view of the velocity field near the tip gap on the baseline duct in climb.	84
Figure 4.37	Axis-symmetric view of the velocity field near the tip gap on the optimized duct in climb.	84

Figure 4.38	3-D view of the velocity field for the baseline duct in climb. . . .	85
Figure 4.39	3-D view of the velocity field for the optimized duct in climb. . . .	85
Figure 4.40	Axis-symmetric view of pressure field near the baseline duct in climb.	86
Figure 4.41	Axis-symmetric view of pressure field near the optimized duct in climb.	86
Figure 4.42	3-D view of the body pressure on the baseline duct in climb. . . .	87
Figure 4.43	3-D view of the body pressure on the optimized duct in climb. . . .	87
Figure 4.44	Flowfield of the optimized axial flight duct moved 2ε closer to the rotor.	91
Figure 4.45	Flowfield of the optimized axial flight duct moved 1ε closer to the rotor.	91
Figure 4.46	Flowfield of the optimized coaxial duct.	92
Figure 4.47	Flowfield of the optimized axial flight duct moved 1ε away from the rotor.	92
Figure 4.48	Flowfield of the optimized axial flight duct moved 2ε away from the rotor.	93

ACKNOWLEDGMENTS

I would like to thank my wife and parents for all of their patience and support. I would also like to thank Sukra Helitek, Inc. for providing software, and the Sukra staff for providing technical support. I would like to thank Dr. Rajagopalan for his technical guidance on this research. Finally, I would like to thank my committee members Dr. Hu and Dr. Pletcher.

ABSTRACT

Ducted fans have found renewed interest particularly in the field of micro aerial vehicles. The complex flow interactions between the rotor, shroud and nacelle hinder the design of optimal lifting duct shapes. Further, the optimum duct shape changes with freestream velocity and rotor configuration. Analytic equations are insufficient to model and predict the performance of ducted fans, yet general CFD solvers are too slow to be used for optimization. For this research, a genetic algorithm is coupled with a simplified CFD solver based on a momentum source rotor model, creating an efficient shape optimization technique for ducted fan systems. The performance gains of the genetic algorithm are documented and the optimum duct shape for several configurations are presented as proof of concept.

CHAPTER 1. Background

In the past, Computational Fluid Dynamics (CFD) has been a research and analysis tool used by highly trained specialists on state-of-the-art machines. As processor speed and RAM increase in desktop PCs, it is becoming more realistic to expect CFD to be used as a design tool available to engineers. While some geometric configurations push the bounds of even the largest cluster of CPUs and require specialized knowledge to simulate, many geometries of interest to industry can easily be solved on the average desktop workstation. Ducted fans are one such configuration that is too complex for analytical solutions, but solvable with current desktop CFD technology. This research is done in pursuit of developing a GA based shape optimization technique for ducted fan configurations.

1.1 Ducted Fan History

Ducted fans have been modeled and studied extensively and have been used in marine applications [1], rotorcraft / VSTOL applications [2, 3, 4] and HVAC. Much of this work has been done with simplified assumptions [5, 6] or analytical models [7]. A more recent and mainstream application of ducted fans is the anti-torque device of the RAH-66. The FANTAIL of the RAH-66 is a notable configuration that has benefited from extensive computational analysis [8, 9]

In the 1940s and 50s Hiller pioneered a flying platform using a fan situated within annular wings. It was found that the airfoil duct accelerated airflow into the rotors. Additionally, Hiller found that the duct could generate 40% more thrust than a propeller of the same diameter. Hiller's aircraft, the Model 1031 used coaxially mounted rotors within a single fiberglass duct.

The fan in wing design has been around since the 1960s with the Ryan XV-5 and the Lockheed XV-4 were explored as aircraft with short take-off and landing (STOL) capabilities. This application of ducted fans stagnated because the fan system was heavy, took a lot of internal volume and produced less lift than was expected.

Recently, ducted fans have found favor with micro-aerial vehicles, providing amongst other things, protection to delicate fan blades, allowing MAVs to operate in tight spaces. Ducts can provide an increase in total thrust as compared to a free rotor. Composite materials allow ducts to be stiff and strong, yet also lightweight. Given the advances in materials, it is relatively simple to design and construct a basic duct that yields a moderate increase in overall thrust.

It is important when studying ducted fans, or rotorcraft in general, that a model which accurately predicts the flow induced by the rotor is used. Accurate tools are available to designers and engineers in industry, yet often a uniform disc approach is employed. While an actuator disc, or variant thereof, is generally sufficient for the axial component of the rotor wake, rotor performance data must be known *a priori*. However, for a ducted fan configuration, the duct influences the rotor in-flow and out-flow and thus, the rotor performance. The actuator disc model is not sufficient as a propulsive model because the duct and rotor are coupled. CFD is an obvious choice for evaluating ducted fan performance and is used as the function call by the optimization technique developed in this research.

1.2 Genetic Algorithm Optimization

Genetic algorithms have been used to solve a number of Industry problems. Related to this research, GAs have been used to optimize rotor design [10], industrial duct design [11], and used in conjunction with CFD [12]. Customizing a genetic algorithm to fit the problem of concern is essential to the performance and results obtained. A GA scheme which performs well in deceptive multi-modal design spaces will likely not converge quickly in less deceptive spaces. Genetic algorithms require a balance between the capability of global search and the requirement for a solution within a reasonable time frame.

The design space for the ducted fan presents hurdles for more traditional optimization methods such as gradient based routines. The design space may or may not be multi-modal depending on the rotor configuration, and the constraints placed on the optimization control parameters. For multi-modal design spaces, an optimization routine must search globally or risk getting trapped by a local maximum. In addition, using CFD as the method to evaluate the ducted fan performance is computationally intensive. Genetic algorithms generate many potential solutions during the optimization process and therefore, require many calls to the performance evaluation routine. The search through the design space must be efficient in order to obtain results in a reasonable amount of time.

The Darwinian class of optimization routines, of which GAs are a member, are rooted in the notion of survival of the fittest. The optimization parameters are generated for each design, called an individual, and a solution based on the parameters is calculated by a performance evaluation routine. The value of the solution returned by the performance evaluation routine is the fitness of an individual. Individuals with higher fitness values

are encouraged to merge with other fit individuals to form the next generation of a population. Merging the fittest individuals in a population creates strong offspring for the next generation. That is the basis of genetic optimization.

1.3 Goal of Research

The goal of this research is to create a procedure for ducted fan optimization. A key step for success is coupling a commercial CFD solver with an efficient GA routine. The resulting tool will ultimately be used to find a duct shape which maximizes the thrust of a ducted fan configuration. This research is not concerned with the merits or shortcomings of a ducted fan configuration with respect to a free propeller design. Instead it is assumed that a ducted fan configuration is required as a design constraint.

The following list of tools is required to find a duct shape that produces maximum system thrust:

- An optimization routine
- A method to modify duct geometry
- A CFD solver
- A routine for duct thrust calculations

The optimization routine, duct geometry algorithm, and duct thrust calculations are developed as a part of this research. The commercial CFD solver, Rot3DC, is provided courtesy of Sukra Helitek, Inc., and is chosen because of the ease with which ducted fans can be analyzed.

CHAPTER 2. Formulation

The tools required to complete the proposed duct optimization are developed in this chapter. The building blocks of the simple Genetic Algorithm scheme are presented as well as enhancements for coupling with an axis-symmetric CFD solver, part of the commercial tool-kit, Rot3DC [13]. The equations used by the commercial CFD software are detailed and explained. Additional tools required to reduce the CFD data are then presented. An algorithm to create smooth geometric shapes for use as input to the CFD is discussed, followed by a method to determine the minimum allowable spatial and temporal mesh fidelity for acceptable results from the CFD solver.

2.1 Genetic Algorithm Primer

The simple Genetic Algorithm (sGA) scheme was first outlined by Goldberg [14]. This most basic scheme for GA optimization contains some building block steps that all methodologies of the same genre follow. The sGA is presented here to illustrate the basics of GA methodology and also to contrast with the enhancements made as part of this research. The outline of the sGA process is as follows:

- Randomly create individuals to form a population.
- Evaluate and rank the individuals in the population.

- Select individuals to mate.
- Combine the characteristics of the individuals chosen to mate.
- Mutate newly created individuals at some specified rate.

Arranging these basic steps into an iterative loop creates a powerful tool for searching a multi-modal design space.

A flow chart of the genetic algorithm routine used for this research is shown in Figure 2.1. This is the basic structure for genetic algorithms. How each step in the process is handled is what differentiates one GA routine from another. Each step in the process along with some terminology is explained next.

2.1.1 Basic GA Terminology

Genetic algorithms emulate the theory of evolution to find the maximum solution to a set of equations. As a result, much of the terminology used to describe the process is borrowed from biology and genetics.

2.1.1.1 Chromosome

A chromosome is the binary string which represents the input parameters of a particular solution. The input parameters for the ducted fan optimization are the spatial locations of control points which define the control polygon. The duct shape is the spline that is drawn inside of the control polygon. The genetic algorithm is required to convert the binary string values into real value locations of the control points. With respect to the ducted fan optimization, a duct shape encoded in a form that the GA can interpret is a chromosome.




Figure 2.1 Basic genetic algorithm process flow.

2.1.1.2 Allele

An allele is the value of a particular bit of the chromosome. For a simple binary encoding, as is used for this research, the value of each allele can be a 0 or a 1.

2.1.1.3 Fitness

The fitness is the value returned when a chromosome is decoded and evaluated by the set of equations that govern the optimization. For example, the fitness used for the ducted fan optimization is the thrust produced by the duct added to the thrust produced by

the rotor. The fitness evaluation for the ducted fan requires a function call to a CFD ducted fan model.

2.1.1.4 Individual

An individual consists of a chromosome and the resulting fitness. Each individual is potentially the optimized solution to the set of governing equations. For the current research, each individual consists of an encoded duct shape (chromosome) and the combined rotor and duct thrust (fitness) of the resulting ducted fan system.

2.1.1.5 Population

A population is a group of individuals. In the case of the ducted fan optimization, the population is a set of candidate duct shapes and their corresponding thrusts. The fitness of a population is simply the summation of the fitnesses of each individual that exists in the population.

2.1.1.6 Generation

A generation is a subset of the population where all individuals are created at a particular optimization iteration. There may be times during the optimization process where parent individuals and their offspring coexist. Every existing member is part of the population, but the parent individuals are in a different generation than the offspring.

2.1.1.7 Convergence

Convergence occurs when every individual in the population has an identical chromosome. Applied to the ducted fan optimization, this means that every duct shape in the current optimization step is identical. This may occur because each individual is at the global maximum fitness value, or because of a local maximum fitness value. If it is the latter, then the state is referred to as premature convergence. In general, genetic algorithms are expected to maintain genetic diversity and not converge. Exceptions are made when the solution space of the set of equations to optimize is not complex, such as uni-modal space or simple multi-modal space.

2.1.2 Schemata

The first concept to be discussed is schemata. Unlike the following genetic algorithm topics, schemata are not a step in the GA process. They are the binary strings within a chromosome that specify a trait or traits of an individual, much like building blocks fit together to form a structure. They can be thought of as templates where a particular combination of 0's and 1's control specific characteristics of an individual. For the ducted fan optimization, an example of a schema is the the combination of 1's and 0's that cause the leading edge of the duct to be a semi-circle with a radius of 0.1 ft.

Assume, for example, that the chromosome of an individual is three bits long. If there are K values available for each bit then there are K^l possible unique values for chromosomes, where l is the number of bits. For a binary alphabet, $K = 2$ so there are eight unique chromosome combinations for the three bit assumption: {100, 110, 101, 111, 010, 011, 001, 000}.

Now to specify a schemata, add a third character, *, to the possible values of each bit.

The * character means "don't care" as in the value can be either 0 or 1. With the addition of *, $K = 3$ and there are sixteen possible combinations of schemata. The schema 1^*1 represents two chromosome combinations from the example set of eight: {101, 111}. Likewise, the schema 1^{**} represents four chromosome combinations: {100, 110, 101, 111}. Although this increases the number of meaningful character combinations, it also increases the available information about patterns within the chromosome.

Schemata tend to lengthen as the optimization moves forward, making individuals robust and specialized. Protecting the good schemata is an important consideration when choosing the best way to implement the following routines.

2.1.3 Selection

Selection is the process of choosing individuals to combine. In the sGA approach, individuals are selected via a weighted roulette wheel. Essentially, each individual is allotted slots on a computational roulette wheel according to their level of fitness. The most fit individuals get the most slots and are most likely to be selected to mate. However, even the least fit individual has a chance to be selected. The roulette wheel, idealized in Figure 2.2, is implemented in practice as follows:

1. Randomly select a percentage of the population fitness, called the wheel value.
2. Parse the population adding the fitness of each individual to a temporary fitness value.
3. Stop the search when the temporary fitness value exceeds the wheel value.

The individual that raises the temporary fitness value past the wheel value is selected as a mating candidate. With this formulation, an individual whose fitness is a significant

percentage of the total population fitness will likely have multiple opportunities to mate.




Figure 2.2 Schematic of roulette wheel for selection.

2.1.4 Crossover

The process of combining the chromosomes of two individuals that have been selected to mate is called crossover. The most effective form of crossover for the present study is to choose one common location on the chromosome of both individuals and cross the chromosomes there. This method is considered most effective because of the relative protection of schemata compared to other methods where important traits, such as the inlet shape of a duct, can be lost solely due to an aggressive crossover technique. Figure 2.3 shows the crossover event for two individuals, the first has all alleles set to zero, while the other has all alleles set to 1.




Figure 2.3 Schematic of the crossover scheme.

The first step in the crossover process is to determine if crossover should even occur. The crossover probability dictates the likelihood of crossover occurring through a biased coin toss. Next, the loci at which to cut both chromosomes is determined. The location at which to cut the chromosome, i_c , is chosen randomly and is bound by $2 \leq i_c \leq n - 1$, where n is the number of alleles in an individual's chromosome. The two individuals then swap the chunks of chromosomes that have been cut. The result is two new individuals that have characteristics of both parents.

2.1.5 Mutation

For evolutionary optimization, mutation allows alternate schemata to appear. Mutation can be very effective if a particular schemata or chromosome begins the optimization with a significant fitness advantage. In this scenario the dominant schemata will likely choke any alternate gene sequences out. Mutation randomly selects individual alleles

to mutate, for binary chromosomes this is as simple as changing a single 0 to 1 or vice versa.

The rate of mutation can be a variable, determined by population and crossover conditions or as in the case of this research, fixed at the beginning of the optimization cycle. The effect of mutation is to add randomness in the population to achieve a near global search area. However, a mutation rate set to 1 is equivalent to a Monte Carlo scheme where no traits are passed on from generation to generation. Therefore, some care must be taken to leverage global search without degenerating into random search.

2.2 Genetic Algorithm Enhancements

An evolutionary search routine, such as a genetic algorithm, is an efficient way to search an unknown design space. However, for optimizations where a lengthy fitness evaluation is employed, it is prudent to seek ways to limit the total number of evaluations and to maximize the effect of each chromosomal recombination event that occurs.

For the present research, the simple Genetic Algorithm described by Goldberg [14] has been modified to provide usable solutions with fewer fitness evaluations by taking full advantage of each crossover event. The modified algorithm is referred to as the GenII-GA, for second generation Genetic Algorithm. The process flow of the GenII-GA, shown in Figure 2.4, is similar to the simple Genetic Algorithm flow. However, now prior to the selection process, the population is filtered and placed in a gene pool from which individuals are selected. Additionally, if the gene pool is not sufficiently unique, then the population is stagnant and must be regenerated. The new routines, which include the selection process, population regeneration, a new elitist strategy, and population scaling are detailed next.

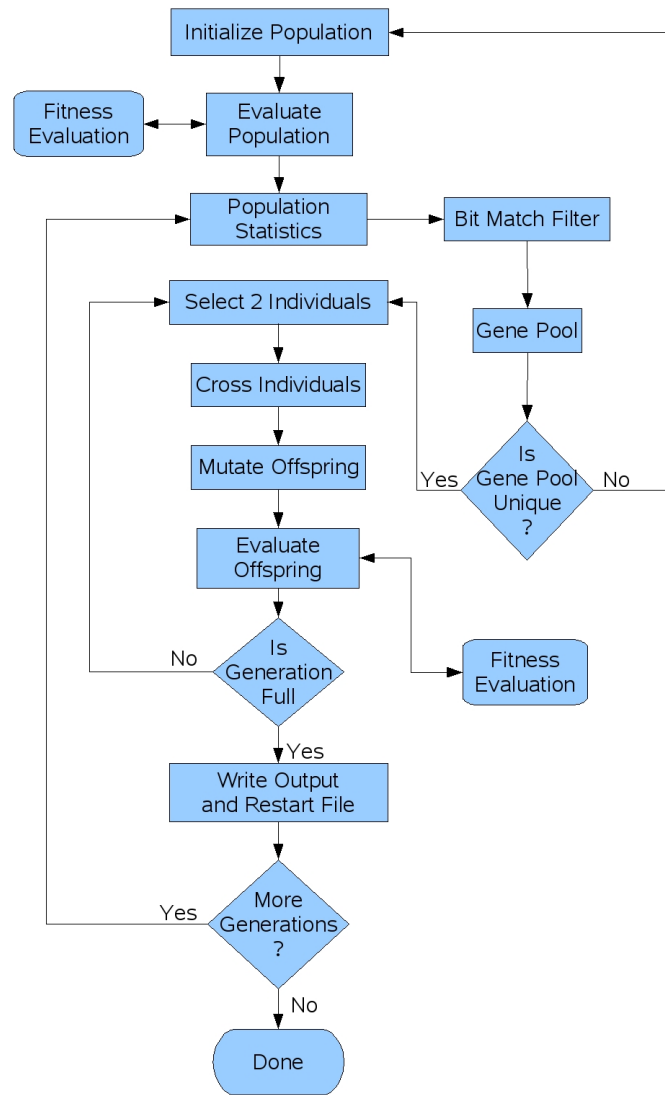


Figure 2.4 GenII-GA process flow chart.

2.2.1 Selection and Gene Pool

The selection process used by the sGA scheme can be improved by adding a subset of the population, dubbed the gene pool. Gene pool rules dictate that all individuals admitted into the pool are candidates that will have the greatest chance of improving the next generation's fitness level. To be admitted to the gene pool, an individual must be sufficiently different from all other individuals currently in the pool. Provided that

requirement is met, individuals are added based on a rank until the gene pool is full.

First, all the individuals of a population are copied into a temporary container. Then the temporary population is ordered from highest to lowest fitness. The algorithm then starts adding individuals to the breeding population from highest to lowest. Each candidate individual must pass a bit-matching filter to be placed in the gene pool.

Figure 2.5 shows the bit-matching requirement for admittance into the gene pool. If the percentage of matching alleles of a chromosome exceeds the threshold set at the beginning of the optimization then the individual is excluded from the gene pool. Although not utilized in this duct optimization research, the GA tool allows for a variable bit-matching threshold.




Figure 2.5 Bit-matching logic.

In essence, the gene pool adds a selection process to the sGA. First, the individual must prove to be genetically diverse enough. Then, the roulette wheel selection takes place as in the sGA. This process postpones premature convergence by favoring population diversity in the selection process.

2.2.2 Modified Elitist Strategy

A common pitfall associated with the sGA scheme is the possibility that the optimum solution will be found and then lost in subsequent generations. It is likely that the current fittest individual will be crossed with many other candidate individuals, diluting or completely destroying the schemata of the current best individual. A common solution to this problem is to employ an elitist strategy. Often, the fittest individual or individuals of a particular generation will be stored and passed on to the next generation unaltered, thereby ensuring that good schema are not destroyed until superior schema exist.

While this notion of preserving the elite individuals remedies the problem as stated, a new issue, namely premature convergence of the solution, may emerge. The role that chance plays in genetic algorithms cannot be overstated. Using a typical elitist strategy gives preference to individuals which may only be very slightly fitter than the rest of the population. As the optimization matures and convergence is being approached, such preference is vital to finding the optimal solution.

During the early stages of the optimization when many different schemata exist, it can be detrimental to allow rather weak schemata, globally speaking, the influence afforded to an elite schema. Doing so encourages weak schemata, which are likely local fitness maximums, to preemptively eradicate other schemata by dominating the selection wheel before other schema can grow and develop. Once the balance has been tipped in favor of a globally weak schema it becomes difficult for the force of chance, via mutation, to correct the situation.

In order to prevent this scenario, common in the basic elitist strategy, an Elite Factor is introduced. The Elite Factor is a real number greater than or equal to 1, and is determined at the beginning of the optimization routine. If, for example, the Elite

Factor is set to 1.3 and the current best individual has a fitness value that is 30% (or more) better than the previous generation's fittest individual, then the Elite Factor requirement is satisfied and the current best individual qualifies for elite status. This way, the elitist strategy can be tuned to find the balance between convergence and global search by ensuring that individuals chosen for elite status are deserving.

If significant improvements in the fitness of the best individual are not realized from one generation to the next, then an elite individual is not selected and no individual's chromosome is guaranteed to exist in the next generation. Note that the Elite Factor set to a value of 1 becomes equivalent to the basic elitist strategy.

2.2.3 Population Regeneration

Despite the attempts of the bit-matching filter to keep the population genetically diverse, it is still likely that the selection pressure will drive the population to genetic convergence. This trend is more evident with smaller populations when an elitist strategy is used. The goal of the optimization is to drive the population to the best fitness which eventually will cause all individuals to exist in a very small region of the solution space. Ideally, this is the global optimum, but for problems with an unknown optimal value, it is nearly impossible to make that judgment.

As a remedy, population regeneration occurs when the genetic diversity of a population dips below a specified threshold. Once a need for regeneration has been detected, the top individuals, with significant genetic diversity, are stored in a container. The chromosomes for the remaining individuals are regenerated using the random initialization process. After regeneration the optimization scheme continues on as before.

2.2.4 Population Scaling

After regeneration, most of the new population has been randomly created. New and unique schemata stand ready for competition. However, the fittest individuals from the old generation are also present in the population and have the advantage of mature schemata. By this point in the schemata evolution, the mature solutions are likely to be significantly superior to all the fledgling new schemata. It is unlikely that this new schemata will have a chance to grow and replace the old solution.

To give the new schema a chance to grow, the population fitness is scaled. The maximum fitness value of all individuals is capped so that the new, weaker individuals get a spot on the roulette wheel. As the optimization run progresses, the scaling is relaxed and the new schema must compete to survive.

2.3 Fitness Evaluation

Up to this point, the formulation has focused on the processes used by the genetic algorithm. The fitness evaluation was treated as a black box where information about an individual is passed in, and the fitness is returned. The focus of the following sections shifts to the steps required to convert the information given by the optimization routine into a fitness value. The process flow of the fitness evaluation used for the ducted fan optimization is shown in Figure 2.6. The steps in this process are presented next.

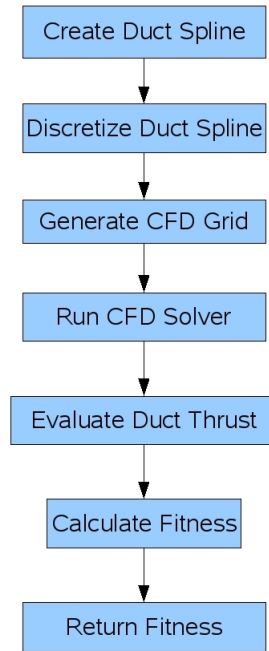


Figure 2.6 Fitness evaluation process flow.

2.4 Duct Parameterization

The GenII-GA is equipped to handle chromosomes that are simple binary strings. The process used to store and convert a string to a point in two dimensional space is explained here. For this research, only uni-directional movement of the control points is considered. The position of the control points in the chordwise direction of the duct is assumed fixed, allowing the points to move only in the direction of the duct thickness. This assumption reduces the size of each individual's chromosome by a factor of two and simplifies the constraints placed on the control points. Increasing the number or density of control points in the chordwise direction compensates for the loss of control fidelity that otherwise results from uni-directional movement.

The airfoil parameterized for the airfoil optimization via panel method in Section 3.3 and the ducts parameterized for the ducted fan optimization in Sections 4.3 through 4.5

share a common strategy to represent a point as a binary string. The 2-D point is already fixed in the chordwise direction, but also must be bounded in the thickness direction to prevent the lower surface from intersecting the upper surface, and to focus the search to a realistic space. Each control point is assigned a minimum value and range. The minimum value of an inner duct surface control point is set by the rotor radius, ensuring there is some clearance between the rotor tip and duct wall. The minimum value of an outer duct surface control point is set to prevent the outer surface intersecting the inner surface when the inner surface is at its maximum.

Each control point is allotted 10 bits on the chromosome. The largest integer that can be represented by a 10 bit binary number is 1024. Therefore, the range of each control point can be divided into 1024 increments. If a control point increment is represented as a percentage of the maximum possible integer value for a 10 bit binary string, 1024, then the deviation from the minimum value is determined by multiplying the range and the incremental percentage. The location of a control point is now determined by adding the deviation to the minimum value. The GenII-GA formulation requires binary strings to be converted into real values for the control points, but does not require a conversion from the real values to the binary string.

When the 10 bit genes representing each control point are placed in the chromosome, the control point storage location is not obvious. Figure 2.7 shows where the control points are stored in the chromosome. Each chromosome segment shown in the figure represents the 10 bit sequence used to determine the location of the control point. The segment G1 is the gene of the first control point. Likewise, Gn is the gene of the n^{th} control point.

With this storage scheme, the real value location of a control point is determined by the following sequence:




Figure 2.7 Gene locations within the chromosome.

1. Divide the chromosome into individual genes
2. Determine the integer value of a binary gene sequence
3. Convert the integer value into a percentage of the maximum integer value (1024 for the 10 bit sequence)
4. Multiply range by percentage and add to minimum value for the gene
5. Repeat steps 2-4 until all genes have been converted

This method of storage is simple and flexible. The size of the number of genes and number of bits per gene can be given as inputs to the GenII-GA at the start of an optimization run.

2.5 Non-Uniform Rational B-Splines

The method of parameterization chosen to define the duct can directly influence the success of an optimization routine. Automating the process of geometry modification is often a difficult task for computational projects. For this research, the duct is parameterized by a set of control points, upon which a spline is fitted. The shape of the duct is controlled by the location of these discrete points, which form a control polygon. Since a spline is a continuous curve, the duct discretization can be non-uniform, allowing for a high fidelity geometric representation, despite having only a few control points. The procedure adopted here is from reference [15].

Non-Uniform Rational B-Splines, NURBS, have become the standard for representing complex computational geometry. The NURBS control points are the parameters of the optimization and thus, as these points change the shape of the duct changes. Fitting a spline to a set of control points allows for a thorough search of the design space with a minimum number of control points. NURBS enable a higher degree of shape control than other spline methodologies and enable a robust method of geometric control.

The formulation of NURBS begins with Bezier curves. An n^{th} -degree Bezier curve is defined by:

$$C(u) = \sum_{i=0}^n B_{i,n}(u) P_i \quad (2.1)$$

where $0 \leq u \leq 1$, the position of the control points are the geometric coefficients P_i , and the basis functions, $B_{i,n}$ are the n^{th} -degree Bernstein polynomials:

$$B_{i,n}(u) = \frac{n!}{i!(n-i)!} u^i (1-u)^{n-i} \quad (2.2)$$

The useful properties of the Bernstein polynomials include: they are positive for $0 \leq u \leq 1$, $\sum_{i=0}^n B_{i,n}(u) = 1$, and they have a single local maximum for all $i \neq 0$. The Bezier curve as written cannot represent conic sections exactly.

To represent conic sections, the rational Bezier curve is defined as:

$$C(u) = \frac{\sum_{i=0}^n B_{i,n}(u) w_i P_i}{\sum_{i=0}^n B_{i,n}(u) w_i} \quad (2.3)$$

where n is the order, $B_{i,n}$ and P_i are as before, and w_i are the weights attracting or repelling the curve to the control points. This curve passes through the end points $\{P_0, P_n\}$, and is tangent to the line segments $P_1 - P_0$ and $P_n - P_{n-1}$. Bezier curves are numerically unstable for a large numbers of control points and changing one control point changes the shape of the curve.

B-splines are Bezier curves in a generalized form with a knot vector. A knot vector, U , is defined as $U = \{u_0, u_1, \dots, u_m\}$, where U is a nondecreasing sequence with $u_i \in [0, 1]$. The degree of the B-spline is defined as $p = m - n - 1$. The basis functions of B-splines are written as:

$$N_{i,0}(u) = \begin{cases} 1 & \text{if } u_i \leq u < u_{i+1} \\ 0 & \text{otherwise} \end{cases} \quad (2.4)$$

$$N_{i,p}(u) = \frac{u - u_i}{u_{i+p} - u_i} N_{i,p-1}(u) + \frac{u_{i+p+1} - u}{u_{i+p+1} - u_{i+1}} N_{i+1,p-1}(u) \quad (2.5)$$

The equation for a B-spline curve is then:

$$C(u) = \sum_{i=0}^n P_i N_{i,p}(u) \quad (2.6)$$

To make the B-spline nonuniform, the knots, u_i , are not uniformly spaced. A B-spline is made rational much the same way as the Bezier curve in Equation 2.3. The equation for a p^{th} -degree NURBS curve is:

$$C(u) = \frac{\sum_{i=0}^n N_{i,p}(u) w_i P_i}{\sum_{i=0}^n N_{i,p}(u) w_i} \quad (2.7)$$

where $0 \leq u \leq 1$, \mathbf{P}_i are the control points, w_i are the weights, and $N_{i,p}(u)$ are the p^{th} -degree B-spline basis functions. $N_{i,p}(u)$ are defined on the knot vector:

$$U = a, \dots, a, u_{p+1}, \dots, u_{m-p-1}, b, \dots, b \quad (2.8)$$

where there are $p + 1$ of both the a and b coefficients.

The weight of each point is the strength which the point pulls the spline towards it. If the weight of a point is set high enough, the spline will pass through that point. Figure 2.8 shows a spline fitted to a control polygon where the weight of all points is set

to 1.0. This weight ensures the spline will pass through the end points of the control polygon, yet retain a level of aerodynamic smoothness. In the case of the airfoil and duct optimizations presented later, the control polygon end points are at the trailing edge.

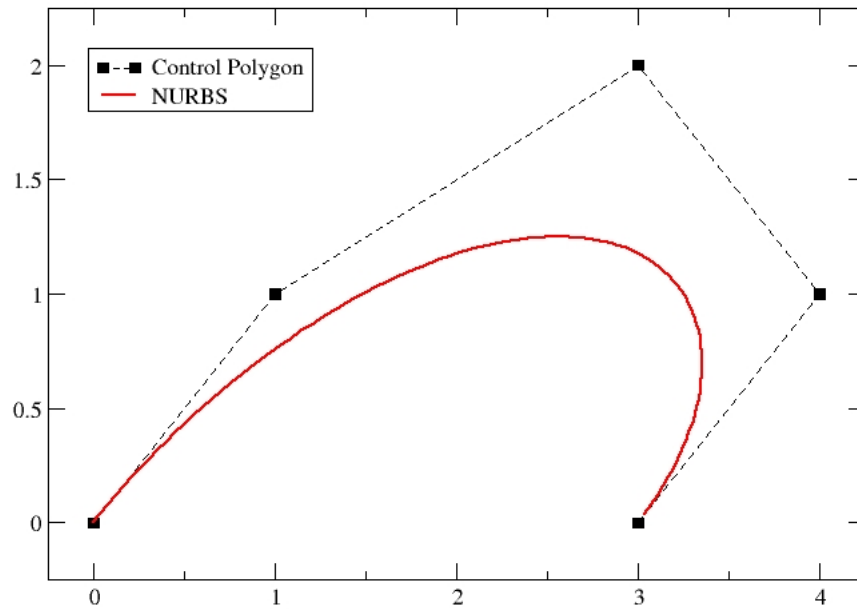


Figure 2.8 NURBS curve with control polygon.

2.6 Axis-Symmetric Flow Solver

The axis-symmetric CFD solver from the Rot3DC tool-kit [13], was chosen to evaluate the performance of the ducted fan systems during the optimization loops. This software was chosen because of the simplicity with which it handles the rotor in computational space. The rotor is the driving force of fluid through this system, and the rotor modeling is fully integrated into the software package. This rotor modeling technique employed by

this software has demonstrated accuracy of the rotor flow with minimal computational cost [16], making it the ideal candidate for determining the fitness of a given ducted fan system.

2.6.1 Derivation and Discretization of the Flow Governing Equations

The fluid flow in the present research is modeled by the unsteady, laminar, incompressible Navier-Stokes equations. For incompressible flow, the conservation of mass and momentum are sufficient conditions for solving the flow field. The conservation of mass applied to a fluid passing through an infinitesimal fixed control volume yields the following equation of continuity:

$$\frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \vec{V}) = 0 \quad (2.9)$$

Newton's second law applied to a fluid passing through an infinitesimal, fixed control volume yields the following momentum equation:

$$\frac{\partial}{\partial t} (\rho \vec{V}) + \nabla \cdot \rho \vec{V} \vec{V} = \rho \vec{f} + \nabla \cdot \Pi_{ij} + S' \quad (2.10)$$

For an axis-symmetric system, the equations in scalar form are as follows:

continuity equation:

$$\frac{1}{r} \left[\frac{\partial}{\partial r} (r \rho V_r) + \frac{\partial}{\partial z} (r \rho V_z) \right] = 0 \quad (2.11)$$

r momentum equation:

$$\begin{aligned} \frac{\partial}{\partial t} (\rho V_r) + \frac{1}{r} \left[\frac{\partial}{\partial r} \left(r \rho V_r^2 - \mu r \frac{\partial V_r}{\partial r} \right) + \frac{\partial}{\partial z} \left(r \rho V_z V_r - \mu r \frac{\partial V_r}{\partial z} \right) \right] = \\ - \frac{\partial p}{\partial r} + \frac{\rho V_\theta^2}{r} - \frac{\mu V_r}{r^2} + S'_r \end{aligned} \quad (2.12)$$

θ momentum equation:

$$\begin{aligned} \frac{\partial}{\partial t}(\rho V_\theta) + \frac{1}{r} \left[\frac{\partial}{\partial r} \left(r \rho V_r V_\theta - \mu r \frac{\partial V_\theta}{\partial r} \right) + \frac{\partial}{\partial z} \left(r \rho V_z V_\theta - \mu r \frac{\partial V_\theta}{\partial z} \right) \right] = \\ -\frac{\rho V_r V_\theta}{r} - \frac{\mu V_\theta}{r^2} + S'_\theta \end{aligned} \quad (2.13)$$

z momentum equation:

$$\begin{aligned} \frac{\partial}{\partial t}(\rho V_z) + \frac{1}{r} \left[\frac{\partial}{\partial r} \left(r \rho V_r V_z - \mu r \frac{\partial V_z}{\partial r} \right) + \frac{\partial}{\partial z} \left(r \rho V_z^2 - \mu r \frac{\partial V_z}{\partial z} \right) \right] = \\ -\frac{\partial p}{\partial z} + S'_z \end{aligned} \quad (2.14)$$

where S'_r , S'_θ and S'_z are the rotor source field terms through which the rotor's influence is introduced into the surrounding flow field.

2.6.2 Discretization of the Generic Governing Equation

Consider the following generic governing equation in an axis-symmetric coordinate system:

$$\frac{\partial}{\partial t}(\rho \Phi) + \frac{1}{r} \left[\frac{\partial}{\partial r} \left(r \rho V_r \Phi - \mu r \frac{\partial \Phi}{\partial r} \right) + \frac{\partial}{\partial z} \left(r \rho V_z \Phi - \mu r \frac{\partial \Phi}{\partial z} \right) \right] = s \quad (2.15)$$

where s represents the terms on the right hand side of the momentum equations, excluding the rotor sources.

The radial and axial fluxes can be written as:

$$\begin{aligned} J_r &= \left(r \rho V_r \Phi - \mu r \frac{\partial \Phi}{\partial r} \right) \\ J_z &= \left(r \rho V_z \Phi - \mu r \frac{\partial \Phi}{\partial z} \right) \end{aligned}$$

Substituting these flux terms into Equation 2.15 yields:

$$\frac{\partial}{\partial t} (\rho\Phi) + \frac{1}{r} \left[\frac{\partial}{\partial r} (J_r) + \frac{\partial}{\partial z} (J_z) \right] = s \quad (2.16)$$

The integration over the time interval can be generalized by the following assumptions:

$$\int_{t_o}^{t_1} \Phi dt = [f\Phi^1 + (1-f)\Phi^o] \Delta t \quad (2.17)$$

where f is a weighting factor between 0 and 1, superscript 1 indicates the new values of the corresponding variables at $t=t_1$ seconds, and superscript o refers to old values of the variables at t_o .

In addition, taking advantage of axis-symmetric system, only one control volume in the θ direction is used. For convenience, the term $\Delta\theta$ is taken to be 1 radian.

2.6.3 Numerical Algorithm

The numerical procedure to solve the fluid equations is based on Patankar's SIMPLER algorithm [17]. This numerical scheme seeks an iterative solution by solving the discretized conservation equations in a sequential line-by-line procedure. In this procedure, the primitive variables (static pressure and the components of velocity) are obtained directly by solving the mass and momentum equations. For a general variable Φ , representing any of u, v, w and p , the discretized equation at a grid point (i, j, k) is found to be:

$$\begin{aligned} a_{i,j,k}\Phi_{i,j,k} &= a_{i+1,j,k}\Phi_{i+1,j,k} + a_{i-1,j,k}\Phi_{i-1,j,k} \\ &+ a_{i,j+1,k}\Phi_{i,j+1,k} + a_{i,j-1,k}\Phi_{i,j-1,k} \\ &+ a_{i,j,k+1}\Phi_{i,j,k+1} + a_{i,j,k-1}\Phi_{i,j,k-1} \\ &+ d_{i,j,k} \end{aligned} \quad (2.18)$$

where i, j, k are the grid indices, a 's are the coefficients that link the neighboring Φ 's to $\Phi_{i,j,k}$ and $d_{i,j,k}$ is the discretized form of the source term that consists of contributions from the specific governing differential equation being discretized.

The sequence of steps for the SIMPLER algorithm can be summarized as follows:

1. Start with a given initial flow field.
2. Calculate unsteady portion of the center coefficients.
3. Calculate the unsteady portion of the source terms.
4. Calculate the coefficients of the momentum equations and the pseudo-velocities.
5. Using the pseudo-velocities, calculate the source term for the pressure equation.
6. Calculate the coefficients for the pressure equation and solve to obtain the pressure field.
7. Using the calculated pressure field, solve the momentum equations to obtain the velocity field.
8. Calculate the source terms of the pressure correction equation and solve for the pressure corrections.
9. Correct the velocities using the velocity correction equations.
10. Return to step 4 and repeat until convergence.
11. Go to step 2 and start with the next time level.

2.6.4 Rotor Modeling

The rotor modeling is based on the momentum source concept developed by Rajagopalan [18, 19, 20]. In order to obtain the wake of the rotor, not known *a priori*, the action of the rotating blades has to be implicitly introduced into the governing equations. In other words, the effect of the spinning blades is in the form of the force $\mathbf{F}(x,y,z,t)$ imparted by the blade to the fluid particles in the path of the rotor. Realizing that the momentum equation governs the balance of the rate of change of momentum (velocity) and the external forces experienced by the fluid element, the effects of the spinning blade on its path at a given time can be modeled by including the force \mathbf{F} (imparted by the rotor blades at that particular time) to the momentum source at the cells occupied by the rotor at that particular time. The force $\mathbf{F}(x,y,z,t)$ can be described by its components in each of the coordinate directions. These components of $\mathbf{F}(x,y,z,t)$ are introduced in the scalar momentum equations as implicit sources.

The force $-\mathbf{F}$ exerted by the fluid on the blade varies along the span of the blade and in general, may vary as a function of time as well. These variations are largely due to the local flow conditions, variations in the blade's chord, aerodynamic twist, and geometric twist along the span. As a result, the different segments of the blade will also experience different strengths and directions of the relative wind, and the flow field around the rotor itself may be inherently unsteady due to the presence of other bodies in the flow.

For a time-accurate calculation, the source terms in functional notation can be written as:

$$S_i = S_i(C_l, C_d, \alpha, \dot{\alpha}, v_{abs}, \omega, \mathbf{R}, t, c, \rho, \mu_{eff}, Re, M) \quad (2.19)$$

where C_l and C_d are airfoil characteristics of the rotor blade, α is the angle of attack made by the rotor blade to the relative velocity vector, $\dot{\alpha}$ is the time rate of change of

α as the blade moves through a revolution, v_{abs} is the absolute velocity of the fluid at the instantaneous blade location (\mathbf{R}, t) , ω is the angular velocity of the rotor, and c is the chord of the blade. Even though the complete Navier-Stokes equations are solved everywhere in the flow field, the dependence of S_i on μ and Re are considered only implicitly through the airfoil sectional characteristics C_l and C_d in this analysis.

2.6.5 Boundary Conditions

The boundary conditions at the edge of the global computational domain are easy to specify. At the inlet, freestream conditions are imposed. The axis of symmetry is assumed to be an inviscid wall. All solid bodies are treated as blocked cells with no flow through them. The outlet boundary, a plane parallel to the inlet boundary, is left for the numerical scheme to compute with mass conservation strictly enforced.

2.7 Duct Thrust

The thrust produced by the duct is calculated from the pressure field determined by the CFD solver. The duct body is computationally represented by n discrete panels determined from the shape spline. Figure 2.9 shows a discretized two dimensional airfoil with normals extending from each panel. The pressure acting on each panel is the pressure in the first CFD grid cell intersected by the panel normal. This is based on the assumption that the variation of pressure normal to the surface within the boundary layer is negligible, Equation 2.20.

$$\frac{\partial P}{\partial n} = 0 \quad (2.20)$$

Assume that coordinates of the panel end points shown in Figure 2.9 are known. The




Figure 2.9 Duct normal vectors.

i^{th} panel is bound by the end points $\{x_i, y_i\}$ and $\{x_{i+1}, y_{i+1}\}$, where x and y are the components in the Cartesian coordinate system shown in the figure. The length of the panel can be written as:

$$dS_i = \sqrt{(x_{i+1} - x_i)^2 + (y_{i+1} - y_i)^2} \quad (2.21)$$

For convenience, the angle of the panel, θ , is measured clockwise from the X axis and is written for the i^{th} panel as the following.

$$\tan \theta_i = \frac{y_{i+1} - y_i}{x_{i+1} - x_i} \quad (2.22)$$

The 2-D force vector acting on the i^{th} panel can now be written as:

$$\vec{F}_i = P_i \vec{N}_i dS_i \quad (2.23)$$

where \vec{N}_i is the normal vector for the panel.

The sum of all the panels yields the force vector on the 2-D airfoil:

$$\vec{F}_{2D} = \sum_{i=1}^n P_i \vec{N}_i dS_i \quad (2.24)$$

where P_i is the pressure at the i^{th} panel, dS_i is the length of the i^{th} panel, and \vec{N}_i is the

normal vector of the i^{th} panel. Using the coordinate system shown in Figure 2.9, the components of the i^{th} normal vector are:

$$\begin{aligned} N_{x,i} &= \sin \theta_i \\ N_{y,i} &= \cos \theta_i \end{aligned} \quad (2.25)$$

so the components of the force in Equation 2.24 can be written as follows.

$$\begin{aligned} F_x &= \sum_{i=1}^n P_i dS_i \sin \theta_i \\ F_y &= \sum_{i=1}^n P_i dS_i \cos \theta_i \end{aligned} \quad (2.26)$$

The two dimensional duct is extended to 3-D by revolving around the centerline, which in this case is the X-axis of Figure 2.9. Revolving each panel will yield an area which is the panel length, dS_i times the circumference of the revolution, $2\pi r_i$, where r_i is the distance from the axis of symmetry to the i^{th} panel center.

For the three dimensional duct force vector formulation, equation 2.24 is rewritten as:

$$\vec{F} = 2\pi \sum_{i=1}^n P_i \vec{N}_i dS_i r_i \quad (2.27)$$

Since the X-axis in Figure 2.9 is the center of rotation, every panel will be mirrored across it. A positive Y-direction force on a panel will be a negative Y-direction force on the mirror image panel. Therefore, after revolving the duct about the axis of symmetry, the Y forces on the duct cancel out in three dimensions. The X-component of the 3-D forces is:

$$F_x = 2\pi \sum_{i=1}^n P_i dS_i r_i \sin \theta_i \quad (2.28)$$

Equation 2.28 is formulated for pressure and radius values at the center of the panel. These values are known at the panel nodes, but can be linearly interpolated to the center of the panel. Through substitution and interpolation Equation 2.28 can be written as:

$$F_x = 2\pi \sum_{i=1}^n \left(\frac{P_{i+1} + P_i}{2} \right) (y_i - y_{i+1}) \left(\frac{y_{i+1} + y_i}{2} \right) \quad (2.29)$$

where the index i represents the nodes at the ends of the panels. This equation assumes that the airfoil is traversed counter clockwise, as shown in Figure 2.9.

The term $y_i - y_{i+1}$ in Equation 2.29 is written this way to account for the direction θ is measured. Figure 2.10 shows the four possible panel orientations with the end points 1 and 2 representing i and $i + 1$, respectively. As can be seen from this figure, $y_1 - y_2$ will always give the correct sign for the X-force produced by the pressure.

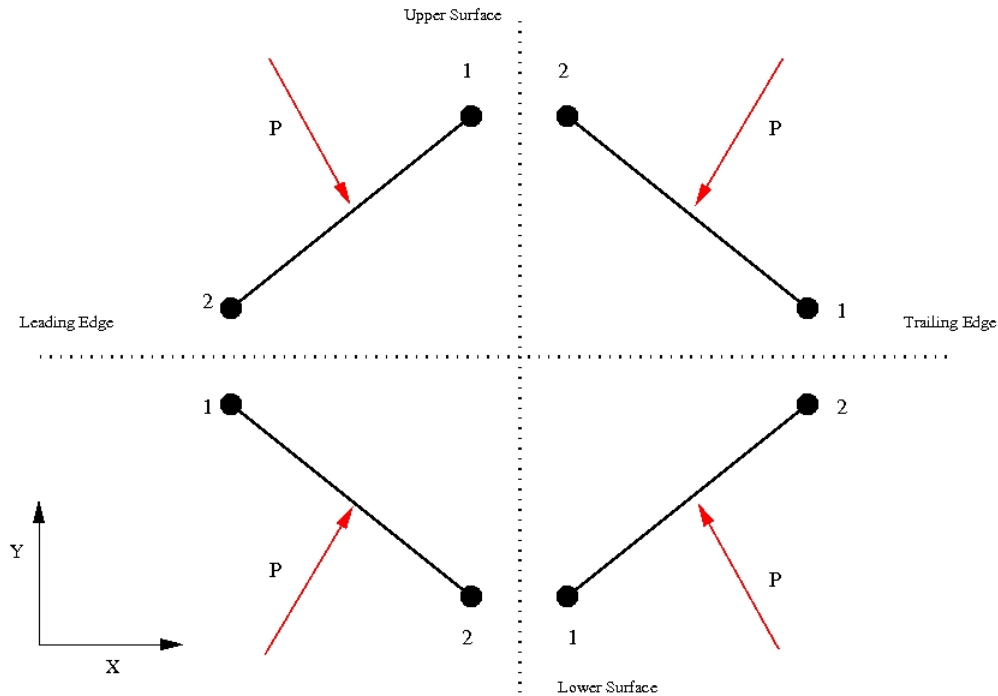


Figure 2.10 Duct panel orientation possibilities.

2.8 Solution Mapping

Flow field differencing, subtracting the velocity field of one simulation from another, is a good way to determine the how similar two solutions are when run with different mesh

fidelities. This is important in determining when grid independence is reached or how much fidelity can be reduced while still producing flow fields of acceptable accuracy. If, for instance, two simulations are run on an identical mesh, then the velocity components in one grid cell can be subtracted from its counterpart in the other simulation. The flow field resulting from equation 2.30 will show the velocity differential between the two solutions.

$$\Delta V (i, j) = V_1 (i, j) - V_2 (i, j) \quad (2.30)$$

This equation is valid so long as:

$$\begin{aligned} x_1 (i, j) &= x_2 (i, j) \\ y_1 (i, j) &= y_2 (i, j) \end{aligned} \quad (2.31)$$

When the size and density of the grid cells are changed for one solution, it is necessary to map the two solutions intended to be differenced to an identical mesh. Assume that the value P , shown as the intersection of the solid lines in Figure 2.11, and the dimensions of the cell, X_w and Y_w are known for all points on a given mesh. The P value can be determined at every point of a different mesh, the intersection of the dashed lines in Figure 2.11, so long as it exists in the domain of the original mesh.

To determine the value of P on the new mesh, $P_{\xi, \eta}$, first look at the influence of the points P_i and P_{i+1} . The distance from P_i to P_{i+1} is noted in Figure 2.11 as X_w and the distance from P_i to P_ξ is X . Notice the η component of P has been dropped for now, so there is only influence from points in the X-direction. The equation for P_ξ is:

$$P_\xi = \left(1 - \frac{X}{X_w}\right) P_i + \left(\frac{X}{X_w}\right) P_{i+1} \quad (2.32)$$

This equation shows a one dimensional average with two points influencing the value of P_ξ . The influence of the j and $j + 1$ points in one dimension can be written similarly:




Figure 2.11 Grid mapping formulation.

$$P_{\eta} = \left(1 - \frac{Y}{Y_w}\right) P_j + \left(\frac{Y}{Y_w}\right) P_{j+1} \quad (2.33)$$

where Y is the distance of P_{η} from P_j and Y_w is the distance from P_j to P_{j+1} . Here the ξ component of P has been dropped indicating only the Y-directional influence is considered.

Combining Equations 2.32 and 2.33 extends the weighted average formulation to two dimensions. An equation for $P_{\xi,\eta}$ can now be written as the following:

$$P_{\xi,\eta} = \left(1 - \frac{X}{X_w}\right) \left(1 - \frac{Y}{Y_w}\right) P_{i,j} + \left(\frac{X}{X_w}\right) \left(1 - \frac{Y}{Y_w}\right) P_{i+1,j} + \left(1 - \frac{X}{X_w}\right) \left(\frac{Y}{Y_w}\right) P_{i,j+1} + \left(\frac{X}{X_w}\right) \left(\frac{Y}{Y_w}\right) P_{i+1,j+1} \quad (2.34)$$

where $i \leq \xi \leq i + 1$ and $j \leq \eta \leq j + 1$. Using this equation as a stencil, it is possible to map areas smaller than the original mesh, however, the original mesh is the limit of the mesh fidelity that can be achieved.

CHAPTER 3. Methodology

This chapter documents the validation of the tools developed for this research. First, the limits of the axis-symmetric solver are explored in an effort to reduce the time required to complete a CFD simulation. Two variable, multi-modal test functions are presented and used to compare the performance of the second generation Genetic Algorithm with the performance of the sGA. A more complex exercise that demonstrates the effectiveness of the duct shape parameterization and the duct thrust calculations is also completed.

3.1 Flow Field Similarity

Optimization techniques that rely on CFD calculations to evaluate the performance of a solution generally require significant computing resources. Small increases in convergence performance can drastically reduce the time required to obtain an optimal value. A 3-D, unsteady CFD solution on a complex geometry, such as a duct, is computationally intensive and may take on the order of weeks to complete. As mentioned previously, the current research constrains the duct optimization to the specific case of axial flight and additionally uses a suitable axis-symmetric solver to reduce the time requirements. Yet, even with this significant reduction in time, more can be done to quickly obtain the optimized results.

The clock time required to perform each CFD evaluation can be further reduced by

decreasing the fidelity of the mesh and increasing the solver time step. Both of these techniques need to be carefully considered and tuned so that the CFD solution remains within acceptable bounds of the high fidelity solution. If it can be shown that a low fidelity mesh and time grid configuration still yields results similar to the high fidelity solution, the configuration will be considered acceptable for use by the optimization routine to evaluate the duct fitness.

A matrix of cases, using an initial geometry, was run to determine the limits of the axis-symmetric solver. The matrix consisted of configurations which changed the mesh and time grid resolution independently, and also in tandem. The solution mapping routine, discussed in Section 2.8, is used to map each CFD solution in the matrix to a uniform grid. The original, high fidelity solution is also mapped to the same uniform grid. Then each matrix solution can be subtracted from the original solution, as shown in Equation 2.30. Flow field differencing shows the changes, if any, that result from changing the grid configuration parameters. The duct is operating in the incompressible flow regime, without heat transfer, so if the velocity field is unchanged, the force on the duct will be unchanged.

Figure 3.1 shows the flow field of the high fidelity solution. The solutions from the test matrix are subtracted from this high fidelity solution. Four solutions from the test matrix are presented. The first matrix solution holds the mesh the same as the baseline case, but increases the time step to 20 times that of the baseline. This is plotted in Figure 3.2a. The second matrix solution, Figure 3.2b uses a non-uniform mesh with 65% of the grid points as the baseline. The time step is 20 times the baseline time step. The third matrix solution uses the baseline time step. The size of the grid cells in the axial direction are 2 times those of the baseline mesh. The third solution is plotted in Figure 3.2c. The final solution is shown in Figure 3.2d. The time step is increased to 50

times that of the baseline and the axial cell widths are 1.25 times the baseline widths.




Figure 3.1 Original baseline solution.

The solutions presented in Figure 3.2 show the solution is more sensitive to changes in spatial resolution than to changes in temporal resolution. The region of velocity discrepancy near the rotor tip found in the two cases with a scalar change in the axial cell width size is evidence of this conclusion. When the time step is changed and the mesh remains constant, the flow field is nearly identical to the baseline solution. A non-uniform mesh is ideal for reducing the computational time with little change in the flow field.

The test matrix of cases has established the limits of the mesh and time steps to be used for the fitness calculation required by the optimizer. The time step limit is approximately 0.05 seconds. The limit on the mesh cell size near the body is slightly over 0.01 radii. These values balance the speed and accuracy of the CFD solution. While a non-uniform

grid could reduce the time required to obtain a solution, the method of evaluating body pressure is better suited to a uniform mesh. Thus, a uniform grid is used around the duct body.

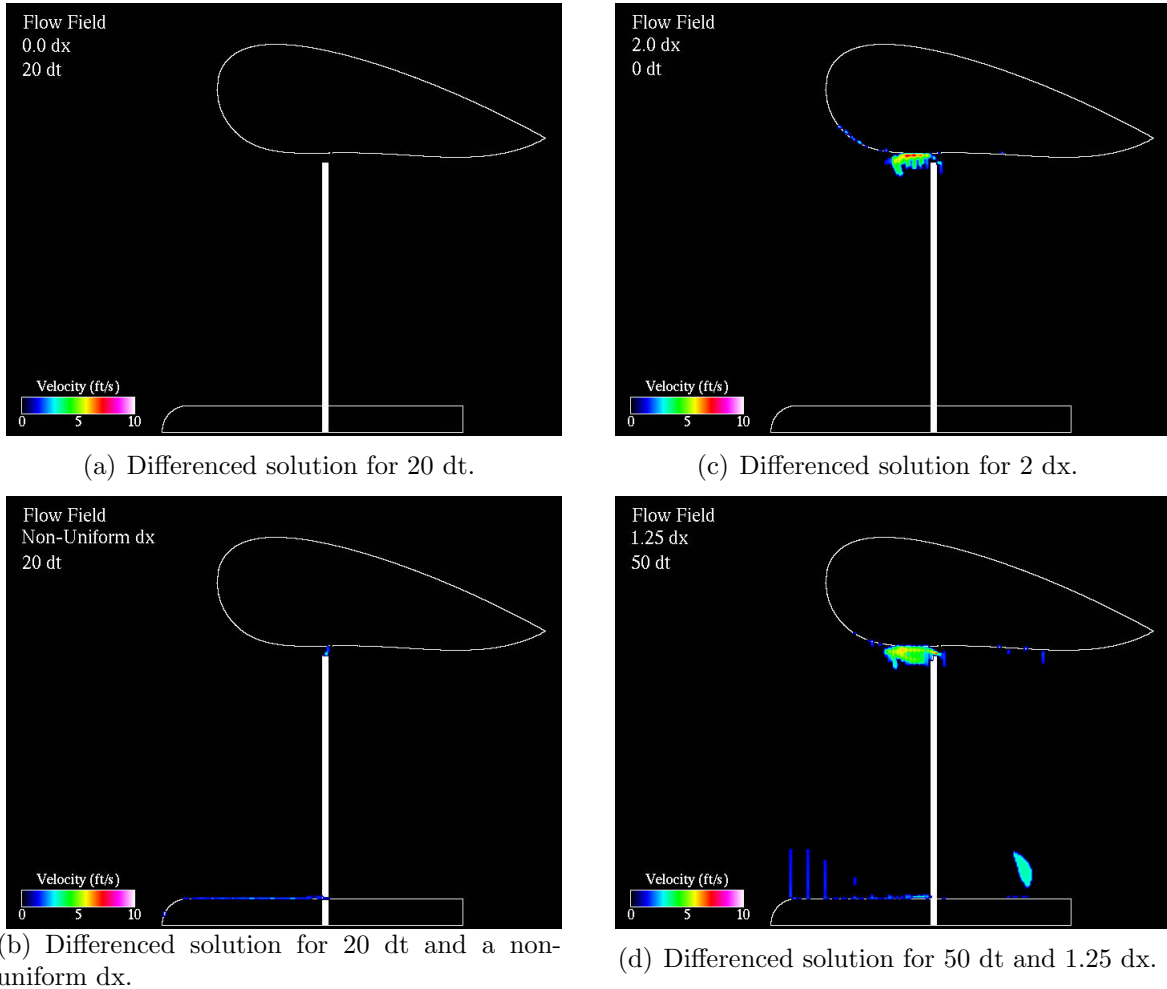


Figure 3.2 Differenced flow fields.

3.2 Multi-Modal Test Functions

Two well known multi-modal test functions were chosen to validate the GenII-GA. Ackley's Function [21] and Schwefel's Function [22] are commonly used to validate evolu-

tionary search algorithms because they force the search algorithm to exhibit both local and global search characteristics. Since the function maximums and the control inputs required to obtain these maximums are known, it is possible to evaluate the performance of the search algorithm and to set termination criteria.

3.2.1 Ackley's Function

Ackley's test function, defined in Equation 3.1, is a very common test for global search methods [23].

$$f(x_n) = 20 + e - 20e^{-0.2\sqrt{\frac{1}{n}\sum_{i=1}^n x_i^2} - e^{\frac{1}{n}\sum_{i=1}^n (\cos(2\pi x_i))}} \quad (3.1)$$

The function minimum, $f_{min} = 0$, occurs when $x_1 = x_2 = 0$. The function is continuous in the range, $x_i \in (-32.7, 32.7)$. The fitness of the solution is defined as:

$$Fit = \frac{1}{f(x_n) + 1} \quad (3.2)$$

which inverts the problem to one of maximization and confines the solution space to between 0 and 1. Figure 3.3 shows the solution to Ackley's function with two parameters.

3.2.2 Schwefel's Function

Schwefel's Function, Equation 3.3 tests the same qualities of search robustness as Ackley's function.

$$f(x_n) = 418.9829n - \sum_{i=1}^n \left(x_i \sin \left(\sqrt{|x_i|} \right) \right) \quad (3.3)$$

The function minimum is again, $f_{min} = 0$ and occurs when $x_1 = x_2 = 420.9687$. The solution is continuous if $x \in (-500, 500)$. As before the fitness is obtained from:




Figure 3.3 Solution space for Ackley's function.

$$Fit = \frac{1}{f(x_n) + 1} \quad (3.4)$$

which turns the problem into one of maximization. The solution space, plotted in Figure 3.4 is also scaled to exist between 0 and 1.

3.2.3 Genetic Algorithm Comparisons

The performance of the sGA is compared with the performance of the GenII-GA to show the benefit of the enhancements documented in Section 2.2. Since chance plays a significant role in the success of genetic algorithms, a single run is not a fair evaluation of the performance of a particular GA scheme. Five optimization runs were completed by both the sGA and GenII-GA routines for Ackley's function and Schwefel's function.

The GenII-GA results for Ackley's function are shown in the top plot of Figure 3.5,




Figure 3.4 Solution space for Schwefel's function.

while the sGA results are shown on the bottom. Both graphs show the maximum fitness plotted as a function of the generation. It is clear from Figure 3.5 that the enhancements in the GenII-GA reduce the number of generations required to find the optimum solution. Fewer generations required to find the optimal solution means that fewer fitness evaluations need to be completed and the optimization efficiency is increased.

Figure 3.6 shows the results for the two GA routines finding the optimal solution to Schwefel's function. This proved to be more difficult for both GA routines, but the performance of the GenII-GA is clearly superior to the sGA. The sGA was not able to find the optimum value for this function. The second generation GA also had difficulty finding the optimum, but faired far better than the sGA.




Figure 3.5 Comparison of simple GA and GenII-GA for Ackley's function.

3.3 Panel Method Evaluation

To test integration of the optimization algorithm with the subsystems required for the ducted fan optimization, a simpler problem was chosen to optimize. The goal was to find the shape of an airfoil which produces the maximum lift. The panel method, detailed by Anderson [24], was chosen to evaluate the lift of the airfoil. Since the panel method is derived from potential flow, a drag force due to friction is not calculated. Instead, a penalty function based on the relative angle of two adjacent panels is introduced to maintain an aerodynamic shape. The equation for the penalty function, PF , can be expressed as:

$$PF = 1 + (2)^m \sum_{j=1}^{n-1} d\theta_j \quad (3.5)$$




Figure 3.6 Comparison of simple GA and GenII-GA for Schwefel's function.

where $d\theta_j$ is written as:

$$d\theta_j = \begin{cases} |\theta_{i+1} - \theta_i| & \text{if } \theta_{i+1} - \theta_i > 8^\circ \\ 0 & \text{otherwise} \end{cases} \quad (3.6)$$

and m is the number of times $d\theta_j$ is not 0. With this formulation the penalty function is bound by $PF \geq 1$.

Additional constraints force the airfoil to be a closed body, meaning at the trailing edge $P_1 = P_n$. The angle between the first and last airfoil panel is not included in the penalty function equation. A direct result of specifying $P_1 = P_n$ is that the penalty function now limits the thickness that can be achieved by the airfoil in addition to the previously mentioned aerodynamic smoothness.

This exercise tests the optimization algorithm, the shape routine, the pressure calculation, and the integration of these components into a single technique. Although lift can be determined directly from the panel method, the lift was determined from a pressure field mapped onto a uniform mesh. The pressure field was found by assuming some far-field pressure, P_∞ , for every grid cell. Then the dynamic pressure at each grid point, found from the velocity induced by the unit sources and unit vortices of the airfoil panels, was subtracted from P_∞ . The lift obtained by mapping the pressure back to the airfoil panels, as outlined in Section 2.7, is within floating point error of the lift obtained using the direct panel method equation.

Five optimization runs were completed with 200 generations set as termination criteria. For the optimization runs, the fitness of an individual is $FIT = lift/PF$, where PF is the penalty function given in Equation 3.5. The maximum and minimum bounds for the NURBS control polygon is plotted in Figure 3.7. The total binary string size was 100 elements representing the 10 control points. Thus, each control point is modeled by 10 elements.

Figure 3.8 is the fitness history for the panel method evaluations. All runs seem to converge very near a maximum fitness of about 0.25. There is a definite and repeatable upward trend in the maximum fitness with respect to generation.

Figure 3.9 shows the shapes of the airfoils that were found to be the fittest at the end of the five optimization runs. The five airfoil shapes are very similar to one another. By comparing the shapes plotted in this figure with the maximum fitness data plotted in Figure 3.8 it is clear that the thicker airfoil shapes of runs 1, 4 and 5 perform better than the thinner shapes found in runs 2 and 3.

For incompressible flow, a higher velocity translates to larger dynamic pressure. As

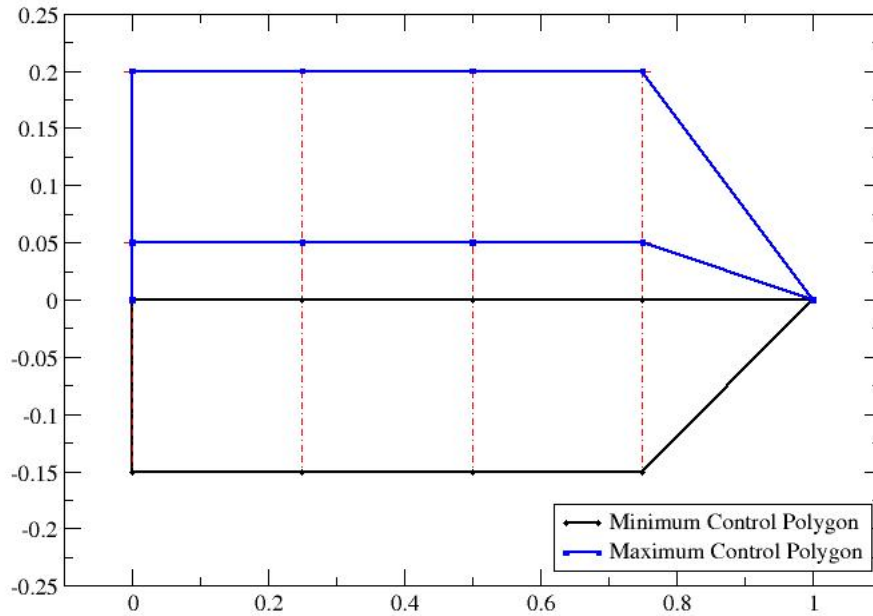


Figure 3.7 Control polygon for the panel method airfoil optimization.

stated previously, the dynamic pressure is subtracted from P_∞ when the pressure field is created. The optimization then, should seek an airfoil with an upper surface that is as far as allowable from the zero line. At the same time, the lower surface should be as close to the zero line as allowable. This is indeed what is happening, as shown in Figure 3.9.

In an effort to show the penalty function is not so confining that only airfoils with a high fitness are produced, the least fit shapes are plotted in Figure 3.10. Contrary to the maximum fitness airfoil plot, the least fit shapes vary widely. This is an indication of the genetic diversity present in the population.

Given the complexity, even of this simplified solution, the results of the optimization rou-




Figure 3.8 Fitness history for panel method evaluation.

tine are taken as a success. The maximum fitness shapes found during the optimization runs are similar in shape. The slight differences that do exist amongst the fittest airfoils affect the fitness in a way that makes sense. Of the fittest shapes, the ones with an upper surface that is farther from the zero line yield a higher fitness. Further, the integration of the optimization code with the systems created to support is demonstrated. With the validation and testing of the optimization tool complete, the ducted fan optimization can begin.




Figure 3.9 Strongest performing airfoils for panel method evaluation.




Figure 3.10 Weakest performing airfoils for panel method evaluation.

CHAPTER 4. Results

The optimization method has been proven effective for searching the global space of multi-modal problems in Sections 3.2.2 and 3.2.1. The auxiliary routines for creating duct geometry and evaluating duct thrust were demonstrated in Section 3.3. The run time required by the CFD solver has been reduced without sacrificing accuracy. The stage is now set to find the optimum duct shapes for a fixed rotor configuration and boundary condition.

Optimum duct shapes for three combinations of rotor configurations and boundary conditions are presented in this chapter as a proof of concept. The optimum duct shapes for each boundary condition/rotor configuration is compared with a baseline duct. The baseline ducts are good initial shapes from a design perspective.

The baseline duct for the first rotor/freestream configuration follows some basic guidelines of duct design. The duct inlet is kept smooth to prevent the flow from separating. The tip gap is small to limit the recirculation of flow from the high pressure rotor outlet back up to the low pressure rotor inlet. The rotor is placed in the minimum radius region of the duct, called the throat.

The baseline duct shape for the second rotor/freestream combination was designed by the 2006 senior class of Aerospace Engineering at Iowa State University. The duct shape was designed for one of the Aero 462 class projects [25]. This duct shape would be found

by a design engineer doing an initial search for available existing shapes.

The third and final baseline duct shape is the optimum duct recycled from the first rotor/freestream combination. Despite changing the boundary conditions, reusing a proven design is common practice for designers. Without information to the contrary, it is reasonable to assume that the optimum shape found for one condition will have acceptable performance in other conditions.

4.1 Computational Domain

The simulations of the following work all use the same computational mesh strategy, shown by schematic in Figure 4.1. A high fidelity grid is created in the region marked Body Mesh. The grid cells in this region are uniform in the axial direction, but are slightly biased in the radially direction from the axis of symmetry to the tip of the rotor.

From the rotor gap to the outer radius of the duct the mesh size is radially uniform. The far field mesh is stretched rather aggressively to keep the boundaries far from the zone of interest without paying a hefty computational price. The performance of the optimizer is also presented for each simulation run to demonstrate the effectiveness of this method.

4.2 Duct Shape Control

The duct parameterization can greatly affect the outcome of the optimization routine. The bounds of the parameterization must be wide enough to include the optimum shape, but also must be strict enough to keep the duct shape realistic and solvable. One thing




Figure 4.1 Grid generation schematic.

to avoid is self intersection of the duct spline resulting in unrealistic shapes. The danger lies in the fact that values are obtained from the CFD for these shapes. If the bounds are set wide enough, special error checking needs to be in place to identify when the duct shape is self intersecting.

The general control polygon for the single rotor ducts is shown in Figure 4.2. The closed loop shown in the figure is a plot of the minimum values for the control points. The shape for the control polygon shown in the figure is achieved by setting all allele values of a chromosome to zero. The duct spline is a convex hull contained within the control polygon. The error bars show the range available to each point from the minimum value.

There are 13 control points used for the single rotor duct cases, but only twelve are independent inputs. The thirteenth point is set equal to the first point to force shape continuity at the trailing edge of the duct. Notice in Figure 4.2, the clustering of control points on the inner surface of the duct to achieve better shape resolution on this critical area. Also, note the minimum allowable rotor tip gap, necessary to prevent rotor

interference.




Figure 4.2 Single rotor generalized control polygon.

Figure 4.3 shows the general control polygon for the coaxial rotor duct optimization case. There are sixteen control points shown and two locations where the minimum tip gap is controlled. Similar to the previous control polygon, fifteen control points are independent with the sixteenth set equal to the first.

As stated in Section 2.4, each control point is allotted 10 bits in the chromosome. The chromosomes for the first and third duct optimizations are 120 bits long to allow control of the twelve control points. The chromosomes for the second duct optimization are 150 bits long for fifteen control points.




Figure 4.3 Coaxial rotor generalized control polygon.

4.3 Single Rotor Hover

4.3.1 Set-Up

The duct optimization of a single rotor ducted fan system is presented in this section. For this optimization the ducted fan system is set to a hover condition. The freestream pressure and temperature are set to 2116.2 psf and 518.69 R, respectively.

The computational grid for this optimization run has 192 grid cells in the axial direction and 237 in the radial direction. The Body Mesh grid cells are kept uniform in both the axial and radial directions. The size of the grid cells within the uniform region are about 0.01 radii in the axial direction and 0.006 radii in the radial direction. Moving away

from the duct the grid is smooth and never exceeds cell ratios of 1.5. The solver used sixty 0.05 second time steps with 5 sub-iterations per time step.

The rotor used for this optimization runs has two blades and a radius of 0.49 ft. The tip speed of the rotor is 384.85 ft/s, which is equivalent to 7500 RPM. The airfoil sections of the blades are NACA 0012 and the collective pitch of the rotor is 10 degrees. The rotor is assumed to have no cone angle and no out of plane deflection.

The termination criteria for the optimizer was set as 200 generation with 50 individuals per generation, for a total of 10,000 fitness evaluations. On a 2.0 GHz, Intel machine a fitness evaluation can be completed in about 1 minute. The total clock time for this optimization run was about 167 hours or nearly one week.

4.3.2 Optimizer Metrics

Some metrics of this optimization run are presented in Figure 4.4. This figure shows the maximum fitness for a given generation, the solid line, and the average fitness for a given generation, the dashed line. The large drops in the average fitness curve indicate where population regeneration takes place. It can be seen from Figure 4.4 that three regeneration events take place and are generally preceded by an increase in the population maximum fitness. From this plot it is also clear that the maximum fitness, the best performer, occurs around the 175th generation. The pressure and flow fields of the best performer are discussed next, with the baseline duct presented for comparison.

4.3.3 Optimum and Baseline Flowfield Comparisons

Velocity contours of the baseline duct flow field are shown in Figure 4.5. Near the tip of the rotor the velocity is high. The velocity contours for the optimized duct shape are




Figure 4.4 The maximum and average population fitness as a function of generation.

shown in Figure 4.6. The optimized duct creates an inlet flow profile that is smoother than the inlet flow of the baseline shape.

Air is recirculating through the tip gap of the baseline duct, shown in the vector plot of velocity in Figure 4.7. Flow near the surface of the baseline duct inlet is separated. The velocity vectors near the rotor tip of the optimized duct are plotted in Figure 4.8. The vectors in this figure show no sign of recirculation and the flow is attached along the entire inlet surface of the duct.

Another difference between the flow through the original baseline duct shape and the optimized shape is the flow behavior near the nacelle. Comparing Figures 4.5 and 4.6 it is clear that duct shape has an impact on flow behavior throughout the duct. For the

baseline duct shape, there is a large recirculation bubble aft of the nacelle. However, for the optimized shape this reversed flow is significantly reduced. The rotor wash through the optimized duct is not only smoother than the wash of the baseline duct shape, but is also faster from about 0.25 radius to the inner duct wall.

Figures 4.9 and 4.10 show three dimensional views of the baseline and optimized duct flow field, respectively. Flow structure upstream of the duct inlet is similar for both shapes. On the upper surface of both ducts, low velocity is expected given the hover condition. This lack of flow indicates that the outside surface of the duct has little to do with the overall thrust for this condition.

The flow field pressure is plotted for the baseline duct shape in Figure 4.11 and for the optimized duct shape in Figure 4.12. The simulation assumes incompressible flow, so the pressure field should be analogous to the velocity field. For the baseline duct shape there is a large and relatively low pressure region near the tip of the rotor in the same location of the reversed tip gap flow. The optimized duct also has a low pressure region near the tip of the rotor, but not as extreme as the baseline duct pressure. Yet, the optimized duct manages to produce more thrust in the axial direction because nearly the entire inside surface of the duct is oriented such that a component of a normal to the surface points in the axial thrust direction.

Despite having relatively higher pressure on the inner surface of the duct (still less than freestream), the optimized duct is able to convert the low pressure to usable thrust. The baseline duct, by contrast, is flat along much of the inside surface of the duct. Therefore, even though it can produce lower pressure on this surface, the result is a force toward the centerline of the duct, which is canceled out when integrated around a revolved duct. This is illustrated in the three dimensional body pressure plots shown in Figures 4.13 and 4.14 for the baseline and optimized duct, respectively.




Figure 4.5 Axis-symmetric view of the velocity field for the baseline shape.




Figure 4.6 Axis-symmetric view of the velocity field for the optimized shape.




Figure 4.7 Axis-symmetric view of the velocity field near the tip gap on the baseline shape.




Figure 4.8 Axis-symmetric view of the velocity field near the tip gap on the optimized shape.




Figure 4.9 3-D view of the velocity field for the baseline shape.




Figure 4.10 3-D view of the velocity field for the optimized shape.




Figure 4.11 Axis-symmetric view of pressure field near the baseline duct.




Figure 4.12 Axis-symmetric view of pressure field near the optimized duct.




Figure 4.13 3-D view of the body pressure on the baseline duct.




Figure 4.14 3-D view of the body pressure on the optimized duct.

4.3.4 Verification of Optimum Conditions

It is impossible to be certain that a particular design is optimal in an unknown design space. To prove that the duct design, presented here as the optimum, is at least a local best solution for a single design parameter, a set of solutions is explored. The tip gap width is increased and decreased from the optimum value by ε . ε is set to 2.5% of the rotor radius, or 0.01225 feet for this ducted fan configuration. The optimized duct shape moves either toward (-) or away (+) from the rotor by the value ε . For this ducted fan configuration, the duct is limited to a movement of away from the rotor (+) to prevent duct-rotor interference. The performance of each perturbed duct configuration is presented in Tables 4.1 and 4.2.

The case where the tip gap is reduced, -1ε , shows an increase in rotor and duct thrust. Combined, the thrust is increased 1.8% over the optimum thrust. It should be noted that the optimum duct is very near the minimum tip gap allowed by the duct parameterization algorithm. This new location of -1ε is outside the bounds of the design space for the duct, and was not available to the optimization routine. It is included in this variational study to help identify the flowfield characteristics associated with duct performance.

Moving the duct away from the rotor ($1\varepsilon, 2\varepsilon$) decreases the thrust of the system and the components. At 1ε , the duct thrust is very nearly the same as the optimum duct thrust. The rotor thrust is slightly lower than the optimum. The result is a 0.6% loss from the optimum system thrust. At 2ε , 1.7% of the system thrust is lost.

These results show the duct is at least a local optimum shape. Moving away from the rotor decreases system thrust and the bounds of the optimization prevent moving towards the rotor. Further, the design is robust with respect to the tip gap parameter. Small manufacturing defects on the duct would minimally affect the performance of this

system.

Table 4.1 Effect of tip gap size for hovering duct and rotor thrust.

Case	Duct Thrust (lbs)	Rotor Thrust (lbs)
-2ε	N/A	N/A
-1ε	0.62	1.03
Optimum	0.61	1.01
1ε	0.61	1.0
2ε	0.563	1.03

Table 4.2 Effect of tip gap size for hovering system thrust.

Case	System Thrust (lbs)	Percent Change
-2ε	N/A	N/A
-1ε	1.65	1.8%
Optimum	1.62	0.0%
1ε	1.61	-0.6%
2ε	1.593	-1.7%

The flowfields of the four cases presented in the tables are shown in Figures 4.15 through 4.18. Figure 4.15 shows the duct at the -1ε position. Notice the smooth, attached flow along the entire inlet surface. The rotor outflow is uniform in velocity. The flow into the duct is noticeably accelerated near the duct inlet lip.

Figure 4.16 shows the flowfield for the optimum duct. The flow structure at this position is similar to the -1ε duct position, but some key characteristics are beginning to manifest. First the flow is still attached for nearly the entire length of the inner duct surface. The exception lies aft of the rotor where a small region of recirculation has formed. Although the flow is still attached, the magnitude of velocity very near the duct surface has decreased. Next, the rotor wash is not uniform. A high velocity region exists near the tip of the rotor blades. For this optimum position it is still small, extending

the distance of the recirculation region at the surface of the duct. Finally, the velocity on the inlet region has decreased.

The flowfield of the duct at position 1ε is shown in Figure 4.17. The key characteristics identified in the optimum duct position have degraded for this duct position. Most noticeably, the recirculation region aft of the rotor has grown considerably. The flow does not appear to be crossing the tip gap threshold, instead the vortex center has moved downstream from its location in the optimum duct flowfield. The high velocity region from the rotor tip has grown in strength and correlates well with the length of the recirculation region. The magnitude of the inlet flow has decreased further, but still maintains a velocity higher than the rest of the duct inflow.

Lastly, the ducted fan flowfield at 2ε is shown in Figure 4.18. The key characteristics continue to change from the optimum condition. The center of the recirculation vortex has moved further downstream. The vortex has increased in size along with the high velocity rotor tip flow. The recirculation is now large enough to entrain air from outside the duct exit. The high velocity flow along the duct inlet surface has been severely reduced from the optimum condition. It is now difficult to identify this region.




Figure 4.15 Flowfield of the optimized duct moved 1ϵ closer to the rotor.

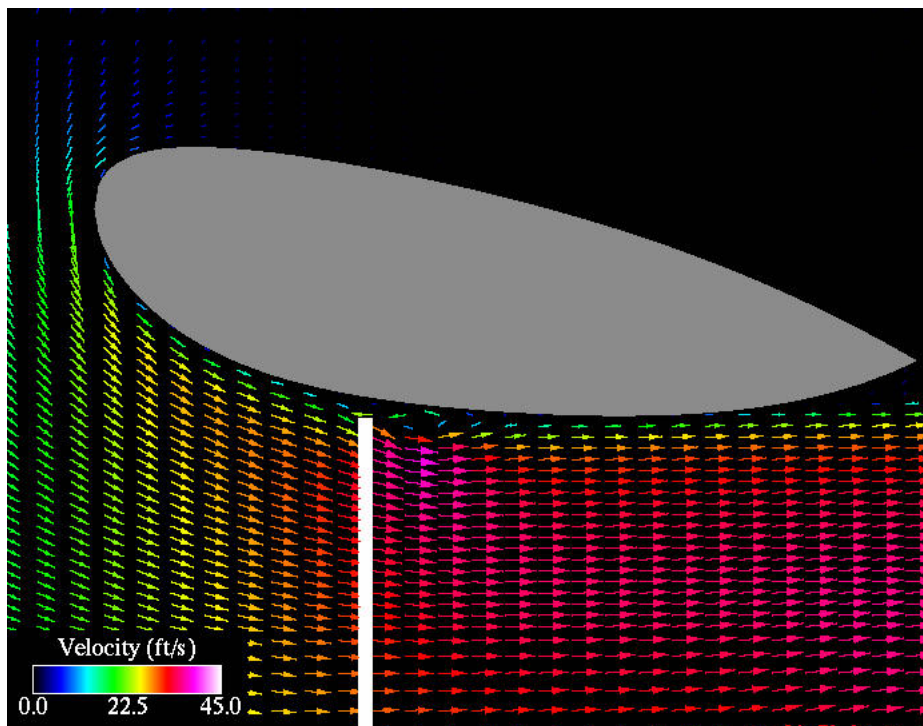


Figure 4.16 Flowfield of the optimized duct.




Figure 4.17 Flowfield of the optimized duct moved 1ϵ away from the rotor.




Figure 4.18 Flowfield of the optimized duct moved 2ϵ away from the rotor.

4.4 Coaxial Rotor Hover

4.4.1 Set-Up

The duct optimization of a coaxial rotor ducted fan system is presented next. The baseline ISU duct was designed for good hover performance. As a result, only the inner duct wall shape was assumed important for duct thrust. The outer surface of the duct is an extrusion of the inner thrust producing surface. The engine nacelle is divided into two halves with the coaxial rotors in the middle. The radius of the rotors is larger than the rotor radius used previously to take advantage of efficiency inherent in lightly loaded rotor blades. Again, a freestream pressure and temperature of 2116.2 psf and 518.69 R, respectively are used.

The computational grid for this optimization run has 200 grid cells in the axial direction and 235 in the radial direction. The size of the grid cells within the uniform region are about 0.0077 radii in the axial direction and 0.0034 radii in the radial direction. Moving away from the duct the grid is smooth and never exceeds cell ratios of 1.5. The solver used sixty 0.05 second time steps with 5 sub-iterations per time step.

The rotors are assumed to have two blades and both have a radius of 0.991667 ft. The tip speed of the rotor is 155.77 ft/s, equivalent to 1500 RPM. The airfoil sections of the blades are NACA 0012 and the collective pitch of both rotors is 10 degrees. The rotors are assumed to have no cone angle and no out-of-plane deflection.

The termination criteria for the optimizer was again set at 200 generations with 50 individuals per generation, for a total of 10,000 fitness evaluations. The small increase in number of mesh cells leads to a small increase in the time required for a fitness evaluation. The total clock time required to complete this optimization run was about

7.5 days on a 2.0 GHz Intel processor.

4.4.2 Optimizer Metrics

The optimizer metrics are shown in Figure 4.19. The solid line in the figure is the maximum fitness for a given generation. The average fitness is shown as a dashed line. It can be seen from Figure 4.19 that three regeneration events take place at nearly the same intervals as found in the previous optimization run, Figure 4.4. The best performer here occurs just after the 50th generation, after which the Elitist Factor formulation of the GenII-GA permits the population to degrade without recovery.




Figure 4.19 The maximum and average population fitness for the coaxial rotor optimization.

4.4.3 Optimum and Baseline Flowfield Comparisons

The velocity contours on the ISU coaxial duct are shown in Figure 4.20 while the velocity vectors are shown in Figure 4.22. There is some flow separation along the duct inlet, but in general, the flow looks smooth. The tip gap for both rotors is small, so despite an adverse pressure gradient only a small amount of air is allowed to recirculate through this region. Flow through the first rotor is relatively low speed which results in efficient rotors, but does not take advantage of the potential to create thrust on the duct.

The velocity contours for the optimized coaxial duct are shown in Figure 4.21. For this duct shape the velocity of the flow is significantly increased near the inlet of the duct as well as at the outlet. The flow structure around the nacelle is generally similar to the ISU duct design, but slightly less severe. The optimal design presented here resembles a traditional airfoil cross section because of the NURBS algorithm used to create the geometry. The outer duct surface provides little benefit to the overall thrust of a duct in hover.

Figure 4.23 shows the velocity vectors on the optimized duct shape for the coaxial rotor configuration. There is little tip gap recirculation through the forward rotor. The vortex seen upstream of the forward rotor entrains fluid from the inlet side of the rotor. Only a small amount of the rotor wake flow passes through the tip gap. The duct flow is attached to the inlet wall. Figures 4.24 and 4.25 show three dimensional views of the velocity contours for the ISU duct shape and optimized duct shape, respectively.

The pressure field for the ISU duct shape is shown in Figure 4.26, while the optimized shape pressure field is shown in Figure 4.27. Both shapes develop lower than freestream pressure on the inlet of the duct up to the first rotor. However, the optimized duct obtains a pressure lower than what is found for the ISU duct shape. The low pressure

region of the optimized duct shape also extends further down the radius of the rotor. The shape used for the ISU duct stretches the low pressure region along more of the duct chord length, but only takes advantage of this low pressure for a portion of this length. Much of the force that would be created by the pressure differential for the ISU duct is canceled when the body is revolved to three dimensions. The pressure discontinuity across the forward rotor, seen in Figures 4.26 and 4.27, is particularly strong because of the small tip gap of these two duct configurations. The high pressure fluid in the forward rotor wake is prevented from flowing around the rotor blade tips back into the rotor inflow. This is verified by the vector plots shown in Figures 4.22 and 4.23.

The pressure mapped onto three dimensional bodies is shown for the ISU duct in Figure 4.28 and for the optimized duct in Figure 4.29. These two figures confirm what was apparent in the pressure field plots. The optimized duct shape achieves its superior results by creating lower pressure where the duct normals have the largest axial component, on the duct inlet.




Figure 4.20 Axis-symmetric view of the velocity field for the ISU baseline coaxial duct.




Figure 4.21 Axis-symmetric view of the velocity field for the optimized coaxial shape.




Figure 4.22 Axis-symmetric view of the velocity field near the tip gap on the ISU baseline coaxial duct.




Figure 4.23 Axis-symmetric view of the velocity field near the tip gap on the optimized coaxial shape.




Figure 4.24 3-D view of the velocity field for the ISU baseline coaxial duct.




Figure 4.25 3-D view of the velocity field for the optimized coaxial shape.




Figure 4.26 Axis-symmetric view of pressure field near the ISU baseline coaxial duct.




Figure 4.27 Axis-symmetric view of pressure field near the optimized coaxial duct.




Figure 4.28 3-D view of the body pressure on the ISU baseline coaxial duct.




Figure 4.29 3-D view of the body pressure on the optimized coaxial duct.

4.4.4 Verification of Optimum Conditions

For the coaxial rotor configuration the ε value is again set to 2.5% of the rotor radius, equating to 0.025 ft for these rotors. For this configuration, the duct cannot move in toward the rotors because the optimum tip gap value is less than ε . Moving the duct away from the rotors decreases the total thrust produced by the system. Tables 4.3 and 4.4 show the component thrusts, the system thrust, and the percent change in thrust.

At 1ε from the optimum placement, the combined rotor and duct thrust have decreased. The result is a loss of 8.9% of the optimum thrust value. At 2ε , the duct thrust has decreased more, but the combined thrust of the two coaxial rotors has increased from the optimum value. The system still loses a net of 8.3% of the optimum system thrust.

The rather large thrust degradation indicates that this design is sensitive to the tip gap. Small defects in the manufacturing may cause measurable decreases in performance. However, this duct does appear to at least be a local maximum with respect to the tip gap parameter. Velocity vector plots of the three duct position are shown next.

Table 4.3 Effect of tip gap size for coaxial duct and rotor thrust.

Case	Duct Thrust (lbs)	Rotor Thrust (lbs)
-2ε	N/A	N/A
-1ε	N/A	N/A
Optimum	0.42	1.15
1ε	0.33	1.1
2ε	0.276	1.16

Figure 4.30 shows the optimum placement for the coaxial duct design. The forward rotor is very near the inner duct wall, preventing the duct from moving toward the rotors.

Table 4.4 Effect of tip gap size for coaxial system thrust.

Case	System Thrust (lbs)	Percent Change
-2ε	N/A	N/A
-1ε	N/A	N/A
Optimum	1.57	0.0%
1ε	1.43	-8.9%
2ε	1.44	-8.3%

A high velocity region extends from the leading edge of the duct to the forward rotor. The flow is attached to the inner duct surface from the forward rotor to the aft rotor. The tip gap of the aft rotor is larger than for the forward rotor and as a result, some separation occurs downstream of the aft rotor.

Figure 4.31 shows the flowfield of the duct moved 1ε away from the rotors. The high velocity region that was on the duct inlet for the optimum placement duct has disappeared. Slightly upstream of the forward rotor, a vortex has manifested. The recirculation zone extends downstream of the forward rotor and reversed flow is allowed to pass through the tip gap. Below this vortex is the familiar increased velocity region passing through the rotor. Flow velocity along the duct wall from the forward rotor to the aft rotor has slowed considerably. The flow separation bubble near the aft rotor has grown.

The flowfield for the duct at 2ε is shown in Figure 4.32. The flow structure has changed considerably from the previous two plots. There is no increase in flow velocity along the inlet of the duct. The vortex upstream of the forward rotor has moved between the forward and aft rotors and dominates the flow near the duct surface in this region. The signature increase in velocity magnitude is seen below the large vortex and extends the length of the visible recirculation zone. The flow on the duct surface is separated from the forward rotor to the trailing edge. The aft rotor recirculation bubble has increased

in width, but is considerably weaker than its upstream counterpart.




Figure 4.30 Flowfield of the optimized coaxial duct.




Figure 4.31 Flowfield of the optimized coaxial duct moved 1ϵ away from the rotor.




Figure 4.32 Flowfield of the optimized coaxial duct moved 2ϵ away from the rotor.

4.5 Single Rotor Axial Flight

4.5.1 Set-Up

The final optimization run attempts to improve upon the single rotor duct found for the hover condition ($V_\infty = 0$ ft/s), Section 4.3, when there is an axial freestream velocity of 25 ft/s. The rotor geometry and parameters are identical to the previous optimization exercise for hover. The optimal duct design for this condition is remarkably different than the design for optimum hover thrust. The new duct design is much thinner and the gap between the rotor tip and inside duct wall is considerable.

Although the computational grid and rotor parameters are identical to the hover optimization run in Section 4.3, the length of simulation time needs to be extended because of the freestream conditions. The time step for this run is still 0.05 seconds, but there are now 80 time steps. The termination criteria for the optimizer is again 200 generations with 50 individuals per generation, but the fitness evaluation now takes 1.25 minutes. The clock time required to complete this run was just over 8.5 days on a 2.0 GHz Intel processor.

4.5.2 Optimizer Metrics

The optimizer metrics are shown in Figure 4.33. The solid line in the figure is the maximum fitness for a given generation. The average fitness is shown as a dashed line. During this optimization run there are four regeneration events. The maximum population fitness increases with generation, finding the fittest individual near the last iteration. The range of variation for this case is much larger than was seen with the previous optimization runs.




Figure 4.33 The maximum and average population fitness for axial flight optimization.

4.5.3 Optimum and Baseline Flowfield Comparisons

Figure 4.34 shows the velocity contours for the baseline duct. The flow appears to be attached to the inner surface of the duct except for a small region aft of the rotor. The flow on the outer surface is definitely separated and shows signs of recirculation near the surface. The velocity contours for the duct optimized for this axial flight condition are shown in Figure 4.35. The flow near the inner duct surface appears to be low speed relative to freestream. The velocity of the rotor outflow is higher for this duct shape than for the baseline shape, but the outflow is far from the inner duct surface. Flow is also detached from the outer surface of this duct.

The velocity vectors near the rotor tip are plotted in Figures 4.36 and 4.37 for the baseline and the optimum ducts, respectively. Despite a large gap between the rotor tip and duct wall for the optimum shape, flow is not reversed. Also, since the optimum shape is much thinner than the baseline shape, the stagnation region on the leading edge is much smaller.

The velocity contour field is shown on the three dimensional ducts in Figures 4.38 and 4.39. Upstream and downstream of the rotor, both ducted fan configurations maintain similar flow velocity. It is worth noting that a strong rotor inflow is present from the nacelle to the inner duct surface for the baseline shape. The optimum shape, by contrast, has a much smaller area of strong inflow.

The pressure field plot for the baseline duct is shown in Figure 4.40. The stagnation point of the flow is on the inlet of this duct and the high pressure that results from bringing flow to rest is present over the most effective portion. Continuing down the inner surface of this duct towards the inlet, low pressure is achieved on portions of the duct that have a normal component in the axial direction, but not like the optimum duct shown in Figure 4.41. The axial flight optimum duct also has a stagnation point on the most effective area of the inlet. However, unlike the optimum case for hover, low pressure exists down the entire inner surface of the duct. This optimum shape is designed to take advantage of the pressure because there is a thrust component to the surface normal vector along the entire inner surface.

The body pressure of the baseline duct is shown in Figure 4.42. This figure shows the large area of high pressure on the inlet. The smaller area of high pressure on the inlet of the optimum duct can be seen in Figure 4.43. Like the baseline duct, the inner and outer surfaces near the inlet produce below freestream pressure. However, the low pressure on the inlet surface of the optimum continues past the rotor to the duct exit.




Figure 4.34 Axis-symmetric view of the velocity field for the baseline duct.




Figure 4.35 Axis-symmetric view of the velocity field for the optimized duct in climb.




Figure 4.36 Axis-symmetric view of the velocity field near the tip gap on the baseline duct in climb.




Figure 4.37 Axis-symmetric view of the velocity field near the tip gap on the optimized duct in climb.




Figure 4.38 3-D view of the velocity field for the baseline duct in climb.




Figure 4.39 3-D view of the velocity field for the optimized duct in climb.




Figure 4.40 Axis-symmetric view of pressure field near the baseline duct in climb.




Figure 4.41 Axis-symmetric view of pressure field near the optimized duct in climb.




Figure 4.42 3-D view of the body pressure on the baseline duct in climb.




Figure 4.43 3-D view of the body pressure on the optimized duct in climb.

4.5.4 Verification of Optimum Conditions

For the axial flight condition, the optimum duct placement is unusually far from the rotor. For this case, ε is 0.01225 ft, or 2.5% of the rotor radius. The large tip gap allows the duct to move to -2ε without duct-rotor interference. The system thrust, component thrusts, and percent change from optimum for each of the duct placements are shown in Tables 4.5 and 4.6.

When the duct is moved toward the rotor to -1ε , the duct thrust is reduced from the optimum value, but still produces a positive force. The rotor thrust increases for this placement slightly, but overall the ducted fan system nets a loss of 23% of the optimum system thrust. Moving the duct to -2ε degrades the system performance further. At this position, the duct is creating drag on the system. The loss of thrust amounts to 71.3% of the optimum thrust. At this location the rotor thrust increases, but not enough to compensate for the duct drag.

Moving the duct further from the rotor than the optimum placement is worse for performance than moving the duct closer. At 1ε , the duct is a drag on the system. The rotor thrust is virtually unchanged from the thrust produced for the optimum position. The loss of thrust at this position totals 44% of the optimum thrust value. At 2ε from the optimum position, the drag on the duct nearly equals the thrust produced by the rotor. The ducted fan system nets a loss of 90% from the optimum thrust.

The axial flight configuration is by far the most sensitive to tip gap. Small changes in geometry produce huge losses in performance. This duct shape is however, at least a local optimum. The flowfields of the duct placement investigation are shown next.

The flowfield for the -2ε duct placement is shown in Figure 4.44. The flow along the

Table 4.5 Effect of tip gap size for axial flight duct and rotor thrust.

Case	Duct Thrust (lbs)	Rotor Thrust (lbs)
-2ε	-0.71	1.12
-1ε	0.07	1.03
Optimum	0.41	1.02
1ε	-0.21	1.01
2ε	-0.87	1.02

Table 4.6 Effect of tip gap size for axial flight system thrust.

Case	System Thrust (lbs)	Percent Change
-2ε	0.41	-71.3%
-1ε	1.1	-23.1%
Optimum	1.43	0.0%
1ε	0.8	-44.1%
2ε	0.147	-90.0%

outer duct surface is low speed, with two vortices inside a separation bubble. Along the inner surface of the duct, the flow is generally smooth. A small vortex exists aft of the rotor tip, but the flow quickly reattaches and continues to the trailing edge. The rotor tip extends out beyond the trailing edge of the duct, forcing some of the rotor flow to be turned by the duct.

Figure 4.45 shows the flowfield for the duct at -1ε . The flow over the outside duct surface appears unchanged. A large separation bubble has developed on the inner duct surface between the rotor and the trailing edge. Multiple vortices are present within this bubble. The high velocity region identified in previous plots can be seen below the separation bubble. The flow reattaches to the inner duct surface prior to reaching the trailing edge. The rotor tip is above the trailing edge, so the duct is still forced to turn some portion of the rotor flow. However, flow from the rotor dips down and around the

separation bubble and then impinges on the duct. Freestream flow is not allowed to pass through the duct without mixing with the rotor wake.

The optimum duct placement flowfield is plotted in Figure 4.46. The flow on the upper surface is unchanged from the previous figures. The structure of the separation bubble on the inner duct surface has fundamentally changed. Low speed flow still occupies a region with roughly the same size and shape as the bubble in the 1ε placement case, but now the vortices have disappeared. Flow on the inner surface is low speed but not reversed. The rotor wash can travel downstream and out the duct exit without turning because the rotor tip is nearly level with the trailing edge. The high speed rotor wake caps the duct exit, preventing flow from being sucked into the duct through the outlet.

The duct is moved to 1ε in Figure 4.47. The outer surface duct flow remains unchanged. The rotor tip is now below the duct trailing edge. The rotor wash does not form a seal at the duct exit. As a result, freestream flow is allowed to pass through the duct without mixing with the rotor wake. The remnants of the low speed region on the inner duct surface still exist but are small. The rotor wake is exhausted from the duct well below the trailing edge.

At a position of 2ε above the optimum duct placement, the flowfield is similar to the 1ε case. Figure 4.48 shows that the flow on the outer surface of the duct has continued to remain unchanged. The boundary of the rotor wash is further below the trailing edge of the duct, allowing freestream flow to pass through the duct unaltered.




Figure 4.44 Flowfield of the optimized axial flight duct moved 2ϵ closer to the rotor.




Figure 4.45 Flowfield of the optimized axial flight duct moved 1ϵ closer to the rotor.

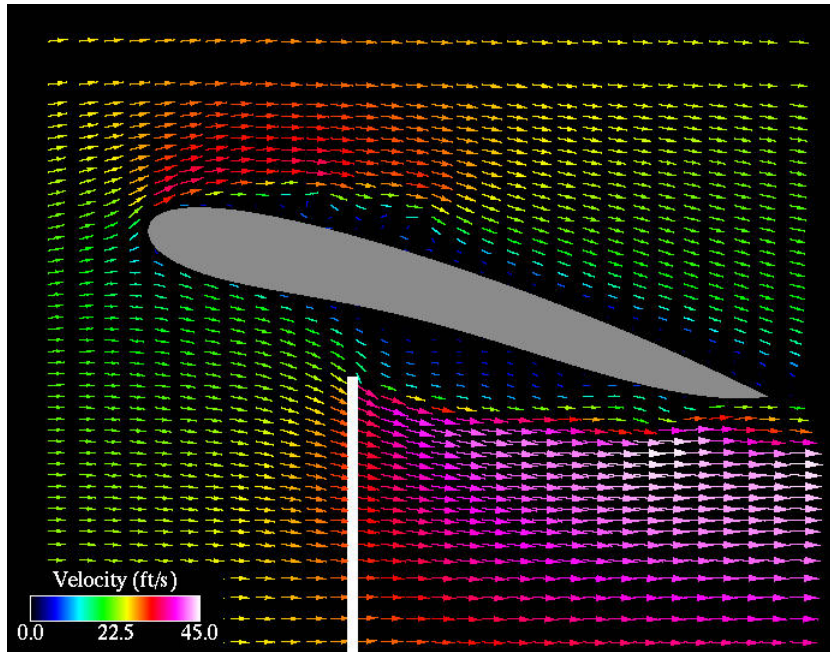


Figure 4.46 Flowfield of the optimized coaxial duct.




Figure 4.47 Flowfield of the optimized axial flight duct moved 1ϵ away from the rotor.




Figure 4.48 Flowfield of the optimized axial flight duct moved 2ε away from the rotor.

CHAPTER 5. Conclusions

The goal of this project was to develop a procedure for shape optimization of a ducted fan system. The technique combines a powerful global search routine, GenII-GA, with a commercial CFD code to optimize duct shapes. Many routines were developed to facilitate the needs of the CFD tool and to automate the tasks for optimization. Included in the auxiliary routines are the NURBS curve automation for parameterizing ducts, the method to evaluate duct thrust, and a flow differencing tool to quantify the flow field differences of non-similar computational domains. Although the work presented in this research employs a particular solver for the fitness evaluation, the technique is robust enough to work with any CFD solver.

The routines and algorithms presented in this work add to the value and enhance the performance of the simple Genetic Algorithm described by Goldberg. The number of fitness evaluations required to achieve an optimum was shown to be less than the sGA methodology. The results of the duct optimization exercises show that the method that has been developed for this research is successful in finding duct shapes that are superior to a baseline shape.

5.1 Optimizer Performance

The performance of the optimizer was generally good for the three duct optimization runs. For all three cases, the average population fitness increased as the simulation progressed. In two of the cases, the duct performance continued to improve as the simulation progressed. The case where the maximum fitness was not constantly increasing still showed the tendency towards improvement during the first 50 generations. The population regeneration routine proved successful in preempting premature convergence of the solution, yielding the greatest chance for the optimum solution to be found.

The two single rotor optimization runs, plotted in Figures 4.4 and 4.33, exemplify how the results from an optimization run should look. The fitness of the solution starts low and increases as the run continues. Further, the tendency of the average fitness to follow the same trend, but on a much shorter frequency due to population regeneration, is encouraging. The results from the coaxial optimization run, Figure 4.19, are not as exemplary. Here the run begins low and reaches its peak early in the optimization. From there, the maximum fitness is allowed to fall because of a lenient elitist strategy, set to search more design space at the expense of propagating strong schemata. Unfortunately, as in most engineering applications, the maximum possible fitness is not known *a priori*, and therefore, the entire duration of the simulation must be completed to be reasonably confident that the best solution is found.

5.2 Ducted Fan Optimization

It is virtually impossible to prove that an optimization strategy has found the global optimum of an unknown design space. Success can be measured by the increase in the performance of a design thought to be optimum versus an existing design. The duct

Table 5.1 Comparison of ducted fan lift.

Case	Duct Thrust (lbs)	Rotor Thrust (lbs)	Rotor Torque (ft-lbs)
Hover Baseline	0.49	0.98	0.0467
Hover Optimum	0.61	1.01	0.0486
Axial Flight Baseline	-0.05	0.69	0.0409
Axial Flight Optimum	0.41	1.02	0.049
Coaxial Baseline	0.29	0.94	0.10385
Coaxial Optimum	0.42	1.15	0.132

thrust, rotor thrust, and rotor torque for the three optimum duct designs and their corresponding baseline ducts are shown in Table 5.1. The three optimum shapes have superior thrust values for the duct and rotor. Also apparent in this table is the increase in torque required for the higher thrusting optimum designs compared to the baselines.

The total system thrust for the six ducts analyzed, along with the duct and rotor contributions to the system thrust, are documented in Table 5.2. From this table, the duct and rotor contributions to system thrust are shown to be similar for the baseline and optimum shapes. For the hovering duct case, the increase in system thrust of 0.15 lbs can be obviously tied to the duct thrust which increased by 0.12 lbs. The rotor thrust for this configuration, seen in Table 5.1, shows little increase from baseline to optimum.

For the axial flight configuration, the gain in system thrust due to duct optimization has more balanced component contributions, from the standpoint of thrust values. The optimum design for this condition produces a thrust increase of 0.46 lbs for the duct and 0.33 lbs for the rotor. The baseline duct was counter-productive to the system thrust because it was creating a drag on the system.

The gains found by optimizing the coaxial rotor duct were more modest than the axial flight optimization. Similar to the hover condition optimum, the duct and rotor

Table 5.2 Total system thrust and contributions.

Case	System Thrust (lbs)	Duct Contribution	Rotor Contribution
Hover Baseline	1.47	33.3%	66.7%
Hover Optimum	1.62	37.7%	62.3%
Axial Flight Baseline	0.64	-7.8%	107.8%
Axial Flight Optimum	1.43	28.7%	71.3%
Coaxial Baseline	1.23	23.6%	76.4%
Coaxial Optimum	1.57	26.8%	73.2%

Table 5.3 Performance increase of optimized ducts and rotors.

Case	Duct Thrust Increase	Rotor Thrust Increase
Hover Optimum	24.5%	3.1%
Axial Flight Optimum	920%	47.8%
Coaxial Optimum	44.8%	22.3%

thrust increases were balanced. The duct thrust increased by 0.13 lbs while both rotors combined had a 0.21 lbs thrust increase.

The increases in rotor and duct thrust for the optimum duct shape are shown as percentages in Table 5.3. The system thrust and required rotor torque increases are given in Table 5.4. For the three configurations, the duct thrust percentage increases were larger than the rotor thrust increases. This is expected since the rotor thrust increases are a secondary effect of the optimization control inputs used in this research. The increases in rotor thrust shown in Table 5.3 are noteworthy considering they result only from changes in the duct geometry. This phenomenon establishes credibility for the argument that ducted fans are complex, interconnected systems that are difficult to model with analytical tools.

An interesting point can be inferred from the modest increase in both the system thrust

Table 5.4 Performance of optimized ducted fan systems.

Case	System Thrust Increase	Rotor Torque Increase
Hover Optimum	10.2%	4.1%
Axial Flight Optimum	123%	19.8%
Coaxial Optimum	27.6%	27.1%

and required rotor torque for the hover condition case shown in Table 5.4. Recall that this duct was created by following some basic guidelines for duct design. The comparatively small increase in these values verifies that previous findings [26] about the inlet shape and rotor placement are applicable to a single rotor hovering duct.

The increase in torque required by the optimum ducts is disappointing, but can be reasoned. All the optimum ducted fan designs benefited from an increase in rotor thrust. In essence, this means the rotors had to process more mass in the form of air. It would be interesting to see the rotor torque included in the fitness evaluation of the optimizer as part of future work on ducted fan design.

5.3 Future Research Opportunities

Future opportunities to expand upon this research are numerous. First, much more can be done with the optimization routine to further improve the speed and efficiency at which an optimum is found. One potential technique to increase the efficacy of the GA is to add a local search algorithm. A good candidate is a hill climbing algorithm such as a gradient based routine. Adding local search would allow the GA to quickly climb local peaks which are discovered. The drawback to such a combination is the increased pressure for premature convergence, explained next.

With the current routines, the GenII-GA will see the peak found by local search algorithm as a promising opportunity for enhancing the population fitness. Globally, the local maximum may be weak, but relative to the current solution will appear strong. After a few generations, the number of solutions with schemata nearly identical to the local maximum will exponentially increase, eliminating population diversity. The fitness history plots, Figures 4.4, 4.19 and 4.33, show that the population already converges and is regenerated at least three times in 200 generations without a local search capability implemented. With local search, the number of required regenerations would increase significantly. Some research is needed to effectively exploit this type of algorithm.

Another area of research is on the solver that performs the fitness evaluations. The solver used for this research employs a structured Cartesian grid and is limited to axial flight. A body fitted grid would undoubtedly increase the accuracy of the body force calculations. A 3-D solver would allow for non-axial flight simulations which are more likely to be of interest to industry. Using currently available technology, both enhancements significantly increase the time required to get a solution.

Finally, a multi-conditional fitness evaluation could be employed. The difference in the optimal duct shape, particularly for the hover and axial flight conditions, is considerable. The optimum hover condition adheres to established duct guidelines with respect to what is a high thrust producing shape. The thrust is produced in a localized area on the duct inlet and maintains a small tip gap to limit reversed flow. The axial flight condition is quite different, leveraging the freestream velocity and a large tip gap to produce thrust down the entire inner surface of the duct. Real world ducts will be required to hover, climb, descend and fly non-axially. Using a multi-conditional fitness evaluation may take on the order of months to get optimized results, but would provide a duct of broader appeal to Industry.

APPENDIX. Gradient Based Routine

Hill climbing optimization techniques, of which gradient based routines are a member, are very efficient when dealing with uni-modal design space. They find the peak of the hill quickly and generally with little wasted computational effort. The computational aggressiveness that enables gradient based routines to be so effective are the very thing that prevent them from being useful on multi-modal problems. gradient based routines suffer from localized blinders. They cannot see the big picture of the design space and because of that climb the first peak they see and finish the optimization loop.

A common way to apply gradient based algorithms is to randomize the starting position. This allows the optimization to eventually find the global optimum, but with significant computational cost. This method of globalizing gradient based algorithms also drastically reduces the computational efficiency because the same peak will likely be climbed repeatedly. At the limit, to ensure the entire design space is searched, gradient based methods become a brute force, randomized guess method.

The last hurdle for a gradient based search algorithm is the formulation and assumptions of the algorithm. By definition, a gradient based search must calculate the gradient of the solution with respect to the optimization parameters. The gradient calculation is simple when the solution is smooth and continuous, but the design space a ducted fan configuration is likely not. For instance, as the duct surface becomes steeper, the flow

will suddenly separate. Finding the gradient when the parameters are in this region of the design space becomes difficult.

The general gradient method optimization routine is now presented. Start with a control vector \vec{X} containing the variables to be optimized. For the ducted fan:

$$\begin{aligned} X_1 &= \text{location of control point 1} \\ X_2 &= \text{location of control point 2} \\ &\dots \\ X_n &= \text{location of control point n} \end{aligned} \tag{A.1}$$

The initial control vector must produce a valid performance index (PI). The PI is analogous to a fitness value from GA notation and is written as $\vec{f}(X)$. Now perturb each component of the control vector and reevaluate the PI. The result is the gradient vector written as:

$$\nabla \vec{f}(X) = \left(\frac{\partial f}{\partial X_1}, \frac{\partial f}{\partial X_2}, \dots, \frac{\partial f}{\partial X_n} \right) \tag{A.2}$$

The direction to move along the virtual hill the algorithm intends to climb is written as $\hat{u} = \frac{\nabla f}{|\nabla f|}$. This is the maximum directional derivative. Now define the size of step, y , to be taken along the directional derivative. The movement in computational space can be written as:

$$\vec{X}_c = y \vec{X}_n \hat{u} \tag{A.3}$$

where \vec{X}_n is the current location, \vec{X}_c is the candidate location, and y is the step size vector. Prior to making the move to \vec{X}_c , check that $\vec{X}_c > \vec{X}_n$. If this is not true reduce the step size. Termination criteria is achieved when y reaches a predetermined value.

BIBLIOGRAPHY

- [1] S. Thurston and R. C. Amsler. Review of marine propellers and ducted propeller propulsive devices. *Journal of Aircraft*, 3, May-June 1966.
- [2] K. P. Spreemann. Wind tunnel investigation of longitudinal aerodynamic characteristics of a powered four-duct-propeller vtol model in transition. *NASA TN D-3192*, April 1966.
- [3] E. E. Davenport and K. P. Spreemann. Transition characteristics of a vtol aircraft powered by four ducted tandem propellers. *NASA TN D-2254*, April 1964.
- [4] W. A. Newsom Jr. Aerodynamic characteristics of four-duct tandem vtol-aircraft configurations. *NASA TN D-1481*, January 1963.
- [5] M. R. Mendenhall and S. B. Spangler. Theoretical study of ducted fan performance. *NASA CR-1494*, January 1970.
- [6] M. R. Mendenhall and S. B. Spangler. Theoretical study of ducted fan performance. *NASA CR-1495*, January 1970.
- [7] R. B. Gray and T. Wright. A vortex wake model for optimum heavily loaded ducted fans. *Journal of Aircraft*, 7, March-April 1970.
- [8] G. C. Ruzicka, R. C. Strawn, and E. T. Meadowcroft. Discrete-blade, navier-stokes computational fluid dynamics analysis of ducted-fan flow. AIAA. Aerospace Sciences Meeting and Exhibit, Reno, Nevada, 2005.

- [9] R. G. Rajagopalan and C. N. Keys. Detailed aerodynamics analysis of the RAH-66 FANTAIL using CFD. *Journal of the American Helicopter Society*, October 1997.
- [10] K. Katakura. Study of deceptive-tolerant genetic algorithms for helicopter rotor design optimization, 2000.
- [11] N. Al-Khalidy. Design optimization of industrial ducts using computational fluid dynamics. 3rd International Conference on CFD in the Minerals and Process Industries, Melbourne, Australia, December 2003.
- [12] A. M. Malkawi, R. S. Srinivasan, Y. K. Yi, and R. Choudary. Performance-based design evolution: The use of genetic algorithms and CFD. 8th International IBPSA Conference, Eindhoven, Netherlands, August 2003.
- [13] Sukra Helitek Inc. A rotor simulation package with graphical user interface. *User Manual*, 1997.
- [14] D. E. Goldberg. *Genetic Algorithms in Search, Optimization and Machine Learning*. Addison Wesley Longman, Inc., New York, 1989.
- [15] L. Piegl and W. Tiller. *The NURBS Book*. Springer-Verlag, Germany, 1995.
- [16] R. G. Rajagopalan. “A Procedure for Rotor Flowfield and Interference: A Perspective”. AIAA-2000-0116, 38th AIAA Aerospace Sciences Meeting, Reno, NV, January 2000.
- [17] S. V. Patankar. *Numerical Heat Transfer and Fluid Flow*. Hemisphere Publishing Corporation, New York, 1980.
- [18] R. G. Rajagopalan and C. K. Lim. Laminar flow analysis of a rotor in hover. *Journal of the American Helicopter Society*, 36:12–23, 1991.

- [19] R. G. Rajagopalan and S. Mathur. Three dimensional analysis of a rotor in forward flight. *Journal of the American Helicopter Society*, 38:14–25, 1993.
- [20] R. G. Rajagopalan and L. Zori. Navier-Stokes calculations of rotor-airframe interaction in forward flight. *Journal of the American Helicopter Society*, 40:57–67, 1995.
- [21] D. H. Ackley. *A connectionist machine for genetic hillclimbing*. Kluwer Academic Publishers, Boston, 1987.
- [22] H. P. Schwefel. *Numerical Optimization of Computer Models*. John Wiley & Sons, 1981.
- [23] A. Simoes and E. Costa. Improving the genetic algorithm's performance when using transformation. 6th International Conference on Neural Networks and Genetic Algorithms, Roanne, France, April 2003.
- [24] J. D. Anderson. *Fundamentals of Aerodynamics*. McGraw-Hill, Inc., New York, 1991.
- [25] Aero 462 Senior Class. CHUAV. Iowa State University, Ames, IA, May 2006.
- [26] D. F. Schaller. Design and analysis of a ducted fan system. Iowa State University, Ames, IA, May 2002.