Analysis Methods to Control Performance Variability and Costs in Turbine Engine Manufacturing

Karl E. Sheldon

Thesis submitted to the faculty of the Virginia Polytechnic and State University in the partial fulfillment of the requirements for the degree

Master of Science
in
Mechanical Engineering

Graduate Committee Members:

Dr. Walter F. O'Brien, Chair
Dr. Christine M. Anderson-Cook
Dr. Clinton L. Dancey
Dr. Peter S. King
Dr. Bob L. Zimering

May 4, 2001
Blacksburg, Virginia

Keywords: Variation Simulation, Gas Turbine Tolerances, Tip Clearance, Optimization

Analysis Methods to Control Performance Variability and Costs in Turbine Engine Manufacturing

Karl E. Sheldon
Mechanical Engineering Department
Virginia Tech

(ABSTRACT)

Few aircraft engine manufacturers are able to consistently achieve high levels of performance reliability in newly manufactured engines. Much of the variation in performance reliability is due to the combined effect of tolerances of key engine components, including tip clearances of rotating components and flow areas in turbine nozzles. This research presents system analysis methods for determining the maximum possible tolerances of these key components that will allow a turbine engine to pass a number of specified performance constraints at a selected level of reliability.

Through the combined use of a state-of-the-art engine performance code, component clearance loss models, and stochastic simulations, regions of feasible design space can be explored that allow for a pre-determined level of engine reliability. As expected, constraints such as spool speed and fuel consumption that are highly sensitive to certain component tolerances can significantly limit the feasible design space of the component in question. Discussed are methods for determining the bounds of any components feasible design space and for selecting the most economical combinations of component tolerances.

Unique to this research is the method that determines the tolerances of engine components as a system while maintaining the geometric constraints of individual components. The methods presented in this work allow for any number of component tolerances to be varied or held fixed while providing solutions that satisfy all performance criteria. The algorithms presented in this research also allow for an individual specification of reliability on any number of performance parameters and geometric constraints.

This work also serves as a foundation for an even larger algorithm that can include stochastic simulations and reliability prediction of an engine over its entire life cycle. By incorporating information such as time dependent performance data, known mission profiles, and the influence of maintenance into the component models, it would be possible to predict the reliability of an engine over time. Ultimately, a time-variant simulation such as this could help predict the timing and levels of maintenance that will maximize the life of an engine for a minimum cost.

Acknowledgments

First, I would like to thank Dr. O'Brien for giving me the opportunity to return to graduate school and earn an advanced degree after serving in the military and being away from academia for a number of years. In addition, the gentlemen with whom I worked with in the turbomachinery laboratory at Virginia Tech earn special mention in the acknowledgement of this research. Most notably, MAJ Keith Boyer and Dr. Peter King who offered advice and guidance throughout the development of this work played a major role in shaping the final product. Additionally, MAJ Boyer was instrumental in exposing applications of this research to the United States Air Force where the potential for implementation is vast.

Dr. Anderson-Cook of the Virginia Tech Statistics Department was an invaluable addition to the committee. Without her input, this research could not have been completed. Her ability to convey concepts and illustrate the value of statistical applications to engineering problems provided the basis by which this research was possible.

The engineers at Honeywell Engines and Systems in Phoenix, Arizona also played a key role in the development of this research. In addition to providing Virginia Tech with their performance code, a number of the Honeywell engineers repeatedly took time from their daily schedule to assist in the development of this work. Rich Bonis, Tom Cunningham, Bob Zimering, Tom Elliot, Ed Palmreuter, and Kevin Shepherd all provided unique contributions to this work.

Lastly, and most importantly, my family deserves credit. The support and guidance I receive from them helps me chose the best path to take when the road leads in many different directions. Thank you.

Table of Contents

List of Figures

List of Tables

Nomenclature

α	Degree of Blade Taper
A_{throat}	Turbine Nozzle Throat Area (in^2)
C	Ratio of Blade Geometries
C_p	Specific Heat (BTU/lbm °R)
Δl_t	Change in Turbine Blade Length (inches)
γ	Ratio of Specific Heats
ε	Coefficient of Thermal Expansion (1/K)
h	Blade Height (hub to tip)
η	Adiabatic Efficiency
HPC	High Pressure Compressor
HPT	High Pressure Turbine
HPTA	High Pressure Turbine Nozzle Area
k	Blade Loading Constant
LPC	Low Pressure Compressor
LPT	Low Pressure Turbine
LPTA	Low Pressure Turbine Nozzle Area
$\dot{m}$	Choked Mass Flow Rate (lbm/sec)
μ	Mean Value
NH	Actual High Speed Shaft Rotational Speed (RPM)
NL	Actual Low Speed Shaft Rotational Speed (RPM)
P_o	Total Pressure (psi)
PR	Total-to-Total Pressure Ratio
R	Universal Gas Constant (BTU/mol °R)
ρ	Metal Density (lbm/ft^3)
r_h	Hub Radius (inches)
r_t	Tip Radius (inches)
σ	Standard Deviation
SFC	Specific Fuel Consumption (lbm/hr/lbf)
SM	Surge Margin (%)

$\sigma_{\tau o \tau}$	Total Blade Stress (psi)
τ	Tip Clearance (inches)
T41	Total Turbine Inlet Temperature (°F)
T_{o1}	Atmospheric Total Temperature (°F)
ω	Rotational Speed (rad/sec)
Wa	Engine Intake Mass Flow Rate (lbm/sec)
W	Work per Unit Mass (HP/lbm)

1 Introduction

1.1 Motivation

High levels of performance reliability in newly manufactured gas turbine engines are difficult to achieve and maintain. Much of the variation in performance is due to the tolerances of critical engine dimensions including the tip clearances of rotating components and flow areas in turbine nozzles. This research presents optimization and analysis methods for determining the maximum possible tolerances of these key components that will allow a turbine engine to pass a number of specified performance parameters at a user-selected level of reliability. Furthermore, maximizing the tolerances of a specific component infers that the manufacturing cost will be minimized since tolerances are inversely related to manufacturing cost.

Reliability, in this work, is defined as the percent chance that a given engine will pass all the required performance criteria immediately after assembly. Unique to this work is the fact that the tolerances for a single component are determined by the simultaneous evaluation of all other tolerances. It is common for engine component designers to set the tolerances on the component that they are designing based solely on the performance of the isolated component. A major assertion of this work is that the tolerances established for individual components should be set with respect to the tolerances of all other components affecting the performance of the engine. Only when the tolerances for individual components are set with respect to the entire system's reliability can the maximum allowable tolerances of each component be accurately determined.

The concept governing design for variation and reliability-based optimization is that the inputs to a model, a turbine engine performance model in this case, take the form of a known distribution of numbers rather than a single point nominal value. Since for most performance models each discrete input will yield a discrete output, a distribution of inputs will logically yield a distribution of outputs. It is the statistics of the output distributions as a function of the input distributions that is the focus of this work.

Simulations using probabilistic distributions often represent a more realistic model in a number of situations, since producing and assembling parts to exactly the same dimension every time is not possible. An excellent example of statistical variation is found in gas turbine manufacturing where no manufactured part or clearance is ever at the *exact* nominal dimensions specified in the design. There is always a tolerance around the nominal value that results from the variation that is inherent to any manufacturing process. For a number of years, a wide range of industries have been using statistical models to control the quality of their product. More recently, industries have begun to formalize their statistical quality control, and have been putting forth a great deal of time and money to educate employees on the use of statistical methods. "Six Sigma" is an example of this formalization of statistical control, and is now in the forefront of many corporate training agendas.

Also presented with this work are methods to determine which level of reliability is most cost effective for an engine manufacturer. The cost-tolerance models used in this research are simple and intended to illustrate qualitative trends only. Lastly, the analysis methods presented in this work will be shown to be the foundation for future work that involves predicting the performance reliability of engines over time.

1.2 Approach

The approach used to conduct this research was modular in nature, with each successive step building upon the previous step. As such, this document is organized in a manner to walk the reader through each step in detail. This section gives a brief overview of each step that will later be discussed in more detail.

The first stage in being able to predict the performance reliability of an engine is simply being able to predict the performance of the engine with single point nominal inputs. Performance prediction models are common in aero-engine manufacturing industries and often provide very accurate predictions of engine performance. For variability-analysis methods such as those presented in this research, accurate

performance prediction of an engine is critical because the nominal point performance model greatly affects the results. Once a performance prediction model is made available, the ability to model geometric dimensions within the engine must be incorporated (if it is not already included in the performance model). The geometric dimensions for which the user wishes to determine the maximum tolerances for must be included in the engine model. An example of this geometric modeling is the ability to predict engine performance parameters as a function of the tip clearance between the casing of the inlet duct and the fan blade tip. By modeling the effect of this tip gap, the user will ultimately be able to determine the maximum tolerances that can be associated with that tip gap while maintaining an acceptable level of engine performance. There are several ways to do this type of geometric modeling, some of which will be discussed in more detail in Chapter 3. The two most common methods of determining performance as a function of tip clearance are the use of either experimental data or mathematical models that can predict the efficiency, pressure rise, and mass flow rate as a function of the tip clearance. The geometric models presented in this work are derived from both experimental work and analytical predictions.

After incorporating the clearance models with the performance model, the user will be able to predict the performance of an engine as a function of the change in component geometry. Although the example described in the paragraph above is a tip clearance loss model of a rotating component, tip clearances are not the only dimensions that can be included in this work. Any dimension that can be declared as an input to a performance model can be analyzed and the maximum tolerances determined. Other components may include nozzle and stator areas, bypass areas, engine casing diameters, and even shaft diameters. It is common for modern performance models to include tip clearance losses that are experimentally derived since tip clearances have long been known to have a significant effect on gas turbine performance.

The next step in determining the maximum possible tolerances of the engine components is to determine the distribution that governs the probability that a specific dimension will occur. Again, the dimensions in question vary because of manufacturing

or assembly imperfections. The assumed distribution will influence the output and should be supported by evidence that substantiates the assumption. Often, Normal (also called Gaussian) or Weibull distributions are used to describe the probability that a dimension will occur because of a manufacturing process. There is a large amount of data supporting the assumption that the variability of manufacturing processes can be accurately modeled with a Normal distribution [1, 2]. A Normal distribution is a familiar bell shaped curve and can be described by:

$$f(x) = \frac{1}{\sqrt{2\pi\sigma}} e^{\frac{-(x-\mu)^2}{2\sigma^2}} \tag{1-1}$$

where μ is the mean value and σ is the standard deviation. The mean represents that value on which the bell curve is centered and has the highest probability of occurrence given any standard deviation. Standard deviation is a measure of how far the data are spread on either side of the mean. Figure 1-1 shows a Normal distribution centered on a mean of zero and the relative location of the first standard deviation that contains 66.3% of all the data.

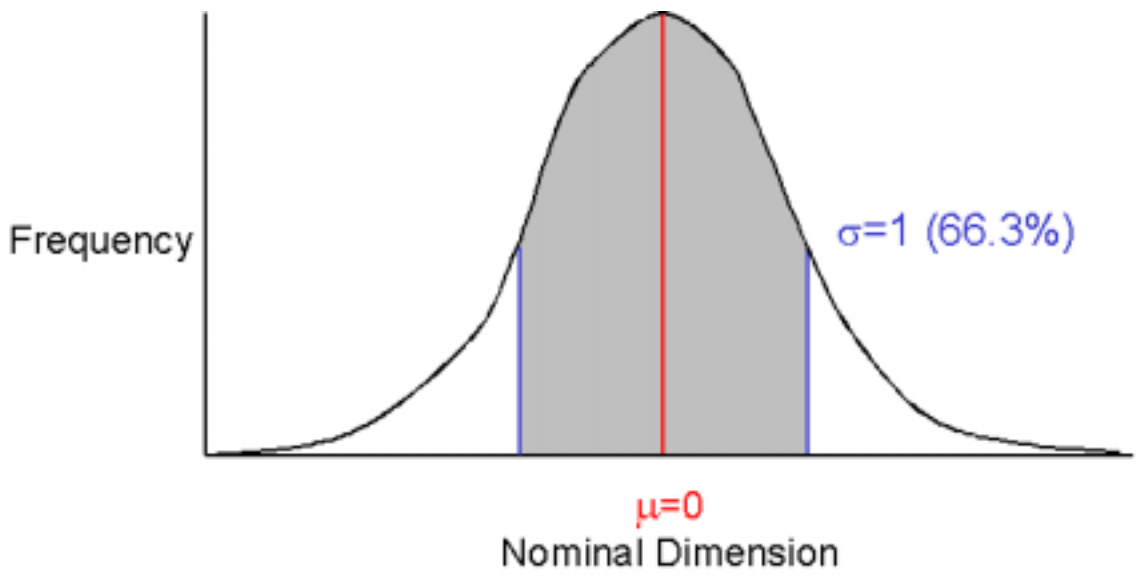

Figure 1-1: Normal Distribution Showing One Standard Deviation

Once the distributions governing the probabilities have been determined, ranges of the variables describing the probabilities must be established. For example, given that

the process governing the clearance of a fan blade is well modeled under a Normal distribution of mean 30 mils (1 mil = 0.001 inches) and a standard deviation of 2 mils, what range of means and standard deviations should be established to search for the most reliable solution? Regardless of how the optimum is defined in this case, there are two design variables (mean and standard deviation) that must be manipulated while searching for an optimum solution. Since many optimization routines can take hours or days to converge, it is in the best interest of time to constrain the ranges of the design variables as tightly as possible. Conversely, if the design ranges are too tightly constrained, there is a risk that the design space that contains the optimum solution is never considered. In other words, the larger the design space, the longer it will take to find a solution but the higher the chance of finding a true optimum. In many cases, assumptions can be made that suggest reasonable ranges to search for optimum solutions. These assumptions will be outlined in detail in Chapter 4.

After adequately constraining the ranges of the design variables and defining an objective function, an optimization algorithm must be chosen. Specific to this work, the design variables are the means and standard deviations of the tip clearances and turbine nozzle areas, while the objective function is to maximize the standard deviations of the design variables while maintaining specified performance criteria. Again, by maximizing the standard deviations of the variables describing a specific dimension, we are in essence determining the maximum tolerances that can be applied to the nominal design value of that specific component. Furthermore, since tolerances are a significant driver of cost, maximizing the standard deviation is also helping to lower the cost.

The choice of optimization algorithms, or combinations of algorithms, is an important factor to consider. Different algorithms offer different benefits, and each carries its own inherent limitations. Exploratory algorithms are well suited for looking over the entirety of a design space and trying to determine where feasible designs may exist. Gradient-based algorithms generally search over a more tightly constrained design space where a known point of initial feasibility exists. This work uses a combination of exploratory and gradient-based algorithms to arrive at the best possible solution. While

no algorithm is ideal for all cases, some algorithms are better suited for maximizing standard deviations than others. Specific details regarding the algorithms used for this work and the reasons that they were selected are outlined in Chapter 4.

Although the engine performance code used in this analysis can run a single case in about 2 seconds, it is not fast enough to do the entire analysis in a time effective manner. The algorithms will be required to run thousands or even millions of iterations to provide enough data for a thorough analysis, so any methods that can reduce the computational time to well below 2 seconds can be well worth exploring. In order to reduce the simulation execution times, this research employs the use of response surface models (RSM) to estimate the performance model output. A response surface model is an algebraic estimation of the performance model and is made from data generated by the performance code itself. The response surface model has execution times on the order of microseconds and can be highly accurate for even non-linear responses. Specifics regarding the creation, use, and limitations of response surface models will be discussed further in Chapter 4.

Once all the algorithms are complete and the simulations executed, a host of post-processing operations can be performed. Some of these operations include sensitivity analyses, response surface model / design space interactions, and multi-dimensional analysis that will be discussed in specific detail in Chapter 5.

1.3 Scope

This research was conducted using a performance model of a small military turbofan engine capable of delivering approximately 1300 lbf of sea level thrust. Although the numeric values are specific to this engine, the methods and tools used to conduct the research are general and can be applied to any turbine engine performance model. While the work presented here is analyzed in detail, it by no means represents the limits that this type of work can achieve. Admittedly, the dimensions of only seven components were included in this analysis when certainly many more dimensions exist

that influence the performance of an engine and could be included. Figure 1-2 shows the design variables and the constraints that are included in this work. The first five constraints in Figure 1-2 are geometric constraints included in the algorithm to ensure that cold build clearances of the components are not so small that they will scrape the side of the engine casing under extreme operating conditions. The last six constraints are engine performance constraints that are commonly used to assess the quality of an engine.

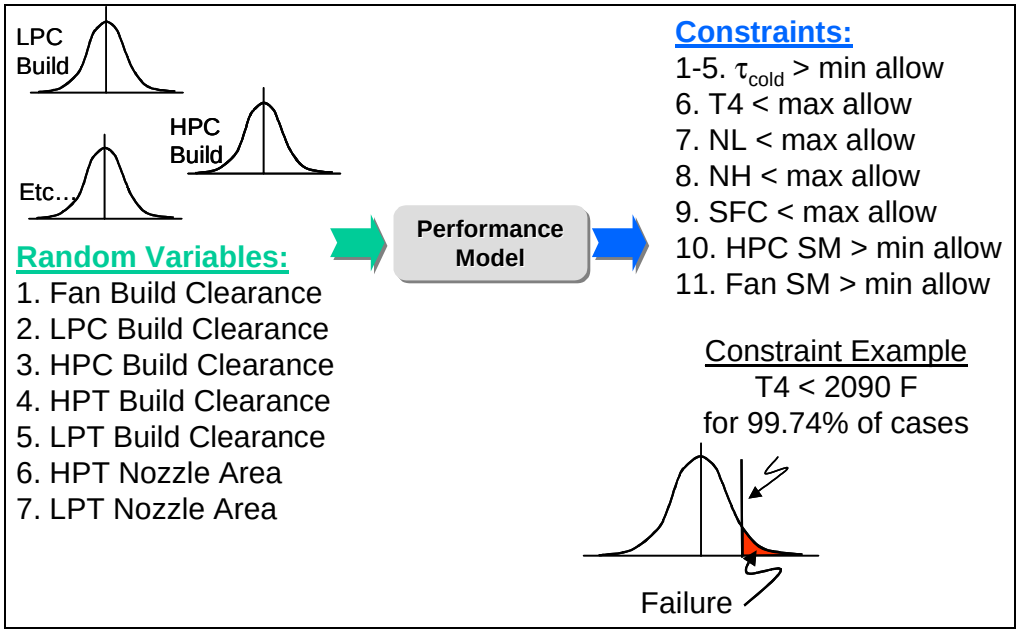

Figure 1-2: Random Variables and Constraints of Variation Simulation Model

Figure 1-2 also shows a schematic of the basic idea behind the variation simulation conducted in this work. A series of distributions representing the probabilities of tip clearances and nozzle areas is fed through a performance model, and the output distributions are examined. If the output distributions are found to pass all of the criteria, then the input distributions are considered feasible, and another attempt of distributions with even higher standard deviations is attempted. This process is repeated until the optimization algorithm is not able to make any further increases to any of the standard deviations without violating a constraint. Also shown in Figure 1-2 is an example of a constraint. The limit of T4 (turbine inlet temperature) is declared as 2090 °F, and the

desired pass level is three-sigma (99.74%). This assumption implies that if 1 - 0.9974 = 0.26% of the engines have a T4 greater than 2090 °F, then the solution is not acceptable.

Given the time and computational resources, the algorithms developed in this work could include any reasonable number of engine component dimensions and constraints. The seven components (labeled Random Variables in Figure 1-2) that were analyzed in this work were included because of the belief that they represent the components that have the most influence on the overall performance of the engine

The results presented in later chapters of this document were developed at a single level of reliability. A three-sigma (99.74%) pass rate of the constraints was chosen for all of the computational algorithms. The algorithms are built to include the level of reliability as an input and are therefore easy to change and to explore the results of other reliability levels as well. Three-sigma was chosen because it represents an initial pass rate at which most engine manufacturers would be very satisfied with, but very few manufacturers achieve. Moreover, Normal distributions were used throughout this analysis for all components. A Normal distribution represents the most ideal case of probabilistic occurrence, when in fact some processes may exhibit a degree of skewness and may be more accurately modeled with a Weibull distribution. The algorithms used in this work will easily accept the modeling of any distribution with only minor changes.

Other assumptions associated with the simulations presented in this work include the fact that the distribution of cold build clearances in a component is a combination of the manufacturing and assembly clearances. In other words, the variance of a component may have many different sources but the solutions from this work will yield the required tolerances of the final product and not necessarily the tolerances required for each individual step of production. Secondly, all the blade clearances represent an average value for the entire component. While each blade will have a unique tip clearance, this work is limited to establishing the tolerances for the average of the tip gaps of all the blades. Lastly, there has been no time-series analysis done in this work. All the

performance conditions were calculated assuming no time dependent wear or degradation.

2 Review of Related Research

Today, turbine engine components and system designs are done on a computer with virtually limitless dimensional accuracy. The manufacturing world, however, does not provide for such ideal conditions. A certain amount of variation is inherent to any manufacturing process and the elimination of all variation is simply not possible. Understanding that variation is a physical and quantifiable factor introduced by manufacturing is key to developing methods to model and analyze variation. Once the variation in a process can be accurately modeled, it can be predicted, controlled, and reduced.

Variation can occur in several forms but is generally classified into two main categories: manufacturing variation and assembly variation. Manufacturing variation is the factor that is responsible for the production of individual parts that do not all have the exact same dimensions. This type of variation is largely due to the fact that there are limitations to the process used to produce the part. Grinding, drilling, milling and all other manufacturing processes have limitations to their accuracy which show up as variation in the parts produced by these processes.

Assembly variation is caused by inconsistencies in the manner in which assemblies and subassemblies are put together. The introduction of assembly variation can occur in a number of different ways. Coupling point alignment, assembly and tooling sequence, gravity and weld distortions, and even measurement error can be sources of assembly variation [3]. Although assembly variation is relatively easy to measure, it is much more difficult to control than manufacturing variation.

The inherent variation in manufacturing processes requires that component design engineers set tolerances on any dimensions critical to the performance or fit of a part in question. It is the tolerances set by the design engineer that the production engineer must adhere to in order to ensure the quality and reliability of the product. In the past, it was common for a design engineer to set tolerances on a part that did not account for the

inherent variability of the process that was used to produce the part. It was this discrepancy between design and production specifications that paved the way for modern methods of identifying and predicting manufacturing variation. With the ability to identify and predict the variation in manufacturing, engineers have developed ways to maximize the reliability of a product while minimizing the cost.

To conduct this research, elements from a number of different disciplines were integrated. The work presented here is not commonly found in literature as a whole, yet the pieces that comprise this work come from efforts of individuals in the field of:

1. Turbine Engine Performance Modeling
2. Modeling Tip Clearance Losses in Turbine Engine Components
3. Variation Simulation Modeling
4. Non-Linear Programming

There is no doubt that, while this type of work may not be commonly published, it is rapidly becoming a tool developed for proprietary use by many engine manufacturers. As a result, the literature review for this work will primarily focus on the development of the individual pieces of the overall algorithm.

2.1 Performance Models

Performance modeling has been employed since the development of turbine engines themselves. Today, engine performance models serve as the backbone to any engine manufacturing company. Not only are the performance models used to design and troubleshoot performance issues, they are also used as the basis for negotiations with potential customers. It is common for engine manufacturers to base the sale of a yet-to-be developed engine on the predictions of their proprietary performance model. The accuracy of these models is critical for a manufacturer as they serve as the basis for bid negotiations insofar as to ensure that the performance specified by the customer can actually be met. Often times an engine manufacturer will pay a significant penalty if an engine fails to meet the performance predicted by their model and agreed to under the development contract.

The performance model used in this research is a proprietary engine code used by a leading producer of gas turbine engines. The engine code used in this research has been continually developed since the early 1990s and is now approximately three million lines of FORTRAN code designed to model any open gas turbine cycle. The engine code was loaned to Virginia Tech for research purposes, but the source code itself is not accessible to Virginia Tech users. Although though the source code is not accessible, a Virginia Tech user does have the authority to edit the engine model and incorporate additional modifications. The engine model can be modified by commands similar to those used in FORTRAN and does allow for a great deal of latitude when exploring the effect of tip clearances and turbine nozzle areas on engine performance [4].

2.2 Clearance Loss Models

Clearance loss models are used to describe the change in performance of an engine component as a function of the clearance between the outer tip of the component and the casing (tip clearance). Figure 2-1 illustrates a simple example of tip clearance geometry specific to an axial compressor.

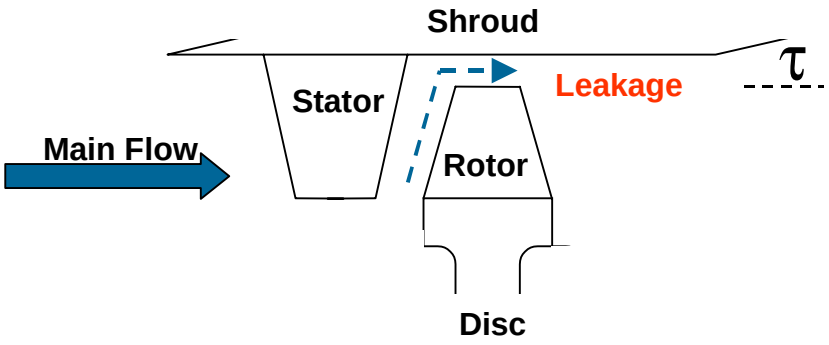

Figure 2-1: Axial Compressor Tip Clearance Geometry

The usefulness of a clearance loss model is to describe the change in efficiency, pressure rise, and flow rate in a component as a function of the tip clearance. As the sensitivity analysis in Chapter 3 will show, tip clearances can have a major effect on the performance of a component. The degradation in performance occurs simply due to the

fact that there is no work done by the compressor blade on the leakage flow, which in effect lowers the amount of useful work produced by the engine. The degrading effects of tip clearances have been known for years and have thus been the focus of intense study for quite some time. Some of the earliest models describing the change in performance of a component as a function of the tip clearance dates back to the early 1940s. These early correlations were linear relations based on simple geometry and empirical constants.

An example of a simple loss correlation is shown in Figure 2-2 where $\Delta\eta$ is the change in efficiency of the turbine, k is a constant, τ is the tip clearance, and h is the mean blade height from hub to tip. The constant 'k' was usually determined empirically from a range of different experiments with similar component types. While this method was suitable for *estimating* the change in efficiency for a variety of components, it did a poor job of accurately describing the efficiency change of a specific component.

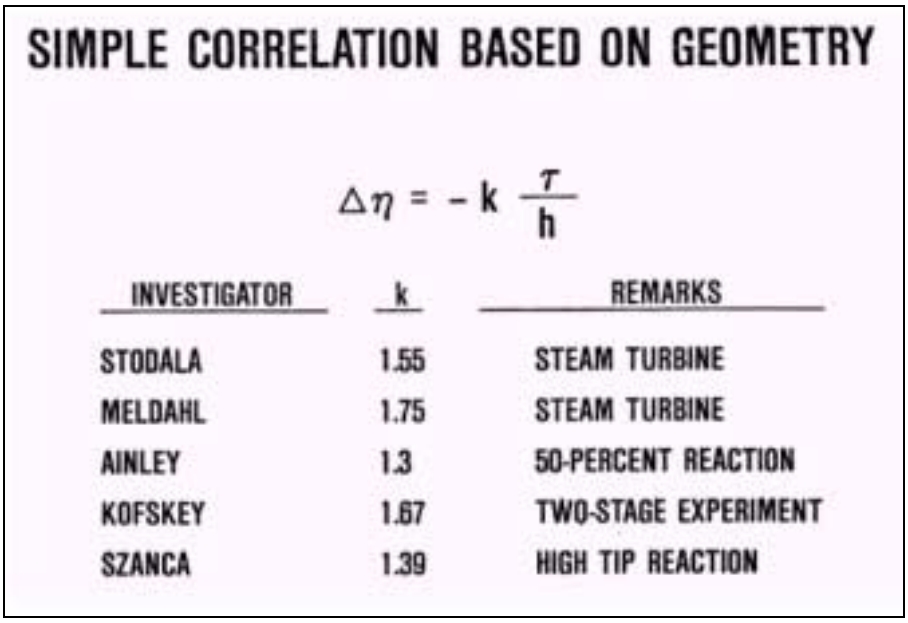

Figure 2-2: Simple Loss Correlations [5]

Much of what the early correlations lacked was consideration for the physics and geometry of unique components. As a result, more advanced correlations based on geometry and blade loading were developed. Figure 2-3 shows a table of advanced

correlations that were developed from the mid 1950s through the mid 1970s. While the advanced correlations look similar to the simple correlations, they offer some significant advantages.

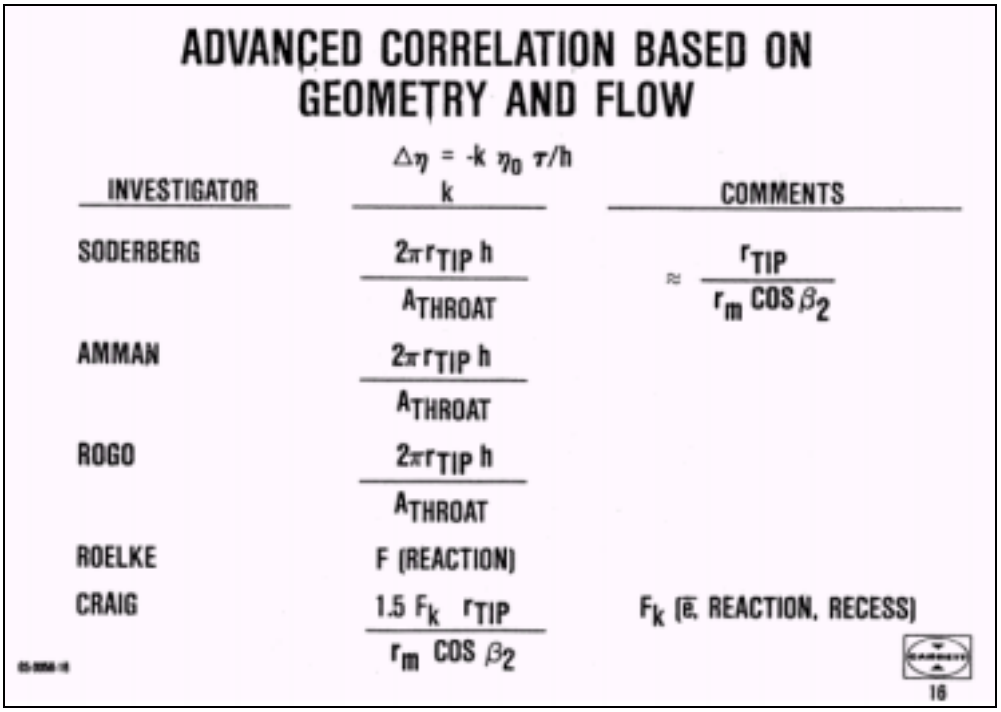

Figure 2-3: Advanced Loss Correlations [5]

The term, "advanced correlation" simply means that more physics is introduced into the model. The "k" term of the advanced loss correlations is determined from the unique geometry of the blade in question rather than an empirically based constant. Furthermore, the η_o (reference efficiency) term accounts for the initial efficiency of the blade without any additional tip clearance losses. Even more advanced correlations have been developed by Lakshminarayana [6] that involve flow coefficients, blade loading coefficients, and turning angle. While this type of correlation is the most detailed of all the correlations mentioned, it also requires a great deal of information regarding the blade and flow conditions.

Although no loss correlation is ideal for all cases, some models are consistently better than others. In an effort to determine which parameters are most influential in

determining the accuracy of the loss correlations, Booth [5] provides a compilation of turbine clearance loss models developed by other authors. He then uses these loss models on experimental data from a number of different turbines in an attempt to discover which model parameters appear to be the most influential in predicting the effects of tip clearance losses. He discovers that the tip Zweifel coefficient has little effect on the accuracy of any of the loss models, but that the reaction plays a significant role on the accuracy of all the turbine loss models.

Since the late 1970s, many tip clearance loss investigations have been done with numerical simulations and computational fluid dynamics (CFD). CFD results allow for a very detailed look at the complex flow physics involved around the tip of turbomachinery components. Figure 2-4 illustrates some examples of complex flow physics resulting from axial turbomachine tip clearances.

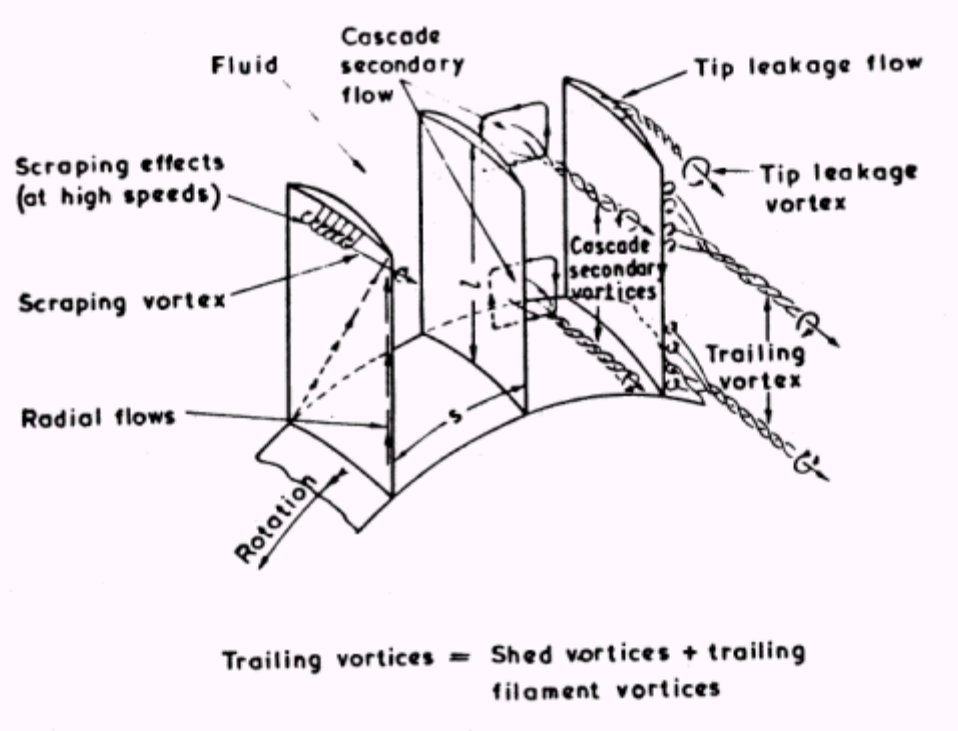

Figure 2-4: Flow Loss Mechanisms in Axial Turbomachines [6]

A number of researchers have done extensive CFD work in an attempt to illustrate the effect of tip clearances in turbomachinery. Using CFD, Tallman and

Lakshminarayana [7] were able to show that the flow does not necessarily need to go through the tip gap to experience secondary flow losses. They showed that a tip gap leakage vortex can exist as a result of the flow through the gap and that this vortex can obstruct the main flow path and cause even further losses. Copenhaver et al. [8] was able to quantify the effect of the tip clearance on a swept transonic compressor rotor with CFD. His work showed that the major losses in efficiency occurred at the outermost 30% of the blade span with very little impact noted at the inner 50% of the blade span.

Modeling tip clearance effects is not the main effort of this research, and as such the tip clearance models used were not generated with CFD. While the loss models used for this work are similar to the advanced correlations shown in Figure 2-3 and do not provide results as detailed as the work performed by Lakshminarayana or Copenhaver, it is entirely possible to generate models from CFD results and use these models in the algorithms developed under this research.

2.3 Variation Simulation Modeling

The development of statistical methods is paramount to understanding and controlling the inherent variability of a manufacturing and assembly process. The use of statistics to analyze and attempt to control product variability is not a new concept. However, only relatively recently have the tools become available that can accurately handle the number of variables and assumptions that contribute to the performance reliability of gas turbine engines.

One of the first, and most rudimentary methods developed to analyze manufacturing variation is called the Worst Case Limit Stack Method. This method involves assuming that each component used to build an assembly is going to be either at its best or worst possible state. Then, by analyzing the assembly at the extremes of each possibility, a range for the variation of the assembly can be determined. While this method will certainly define the extreme limits of an assembly's variation, the chances of simultaneously having all the individual components at their extremes is astronomically

low. This method will often yield a very conservative design with overly stringent tolerances that are designed to accommodate a very rare worst-case condition.

An improvement to the Worst Case Limit Stack Method is called Probability Analysis (Root-Mean-Square). Probability Analysis is a statistical tolerancing method that accounts for the low probability of a worst-case scenario occurring. The probability method uses the square root of the sum of the squares of the variation of individual components to describe the variation of an entire assembly. This method accounts for the statistical nature of component variation and provides a much more realistic analysis of the variation than the Worst Case Limit Stack Method. Although more accurate, the probability method does have limitations. Probability Analysis assumes a Normal distribution, can be difficult to implement in two or three dimensions, and can only be applied if the variation of the individual components is independent [3].

The low cost and abundant availability of high-speed computers has made Monte Carlo Simulations (MCS) an increasingly popular tool to analyze variance. Monte Carlo Analysis is a method by which random numbers are generated under a known distribution to simulate the probability of variation that would be found if the actual process were undertaken. Assuming an accurate model is used to describe the assembly process for two parts, a Monte Carlo Analysis can be used to simulate the manufacturing variance of each part individually and then analyze the assembly of the parts as a whole. Depending on the complexity, Monte Carlo simulations can be required to execute tens of thousands of iterations to get a good statistical average of a process. Monte Carlo Analysis holds many advantages over methods previously employed. First, a MCS does not rely on some of the "classical statistics" assumptions such as a Normal distribution or independence. Monte Carlo Simulations can be used to model any known distribution and are much better suited to handle large numbers of independent variables than classical approaches.

Monte Carlo Simulations have been used to model the variance of manufacturing or assembly processes, but do not traditionally involve modeling the operation or

performance of an assembly as a result of the associated variances. Variation Simulation Modeling (VSM) extends the usefulness of a MCS by including a model of the operational variability that results from a MCS solution. VSM involves modeling how an assembly will perform after being built and manufactured under a variety of individual component variances. Specific to this work, an example of VSM would be to determine the performances of a turbine engine resulting from individual components built and manufactured under distributions with known variances. The utility of a VSM is that it accounts for the performance of the engine as a result of all the independent component variances interacting simultaneously. Modeling the interactions of all the tolerances is a much more realistic prediction of the overall engine performance than simply modeling the performance of the engine by varying only a single component at a time. This is a decided advantage over MCS alone due to the fact that both the engine fitting together correctly as well as the performance of the engine as a whole is of key interested to any engine business. Researchers such as Early and Thompson [3] and Craig [9] have presented VSM results for simple component assemblies and illustrated the usefulness of such a tool.

The present work builds on the VSM method. The next step is the optimization of the independent variables that serve as the inputs to the VSM. The objective function of the optimization routine is to maximize the standard deviations (and thus the tolerances) of the distributions that govern the occurrence of the geometric variability. Although analysis focused on maximizing the engine reliability as well as minimizing the manufacturing cost was performed in this research, most of the emphasis of this work remains on maximizing the tolerances and minimizing the variability.

Although simple to implement as an optimization objective function, accuracy in cost-tolerance models is difficult to achieve. Ostwald and Blake [10] showed that simple cost-tolerance models could result in a maximum error of 1536% over detailed cost-tolerance models that are generated using accurate, real-world data. Some of the data that is key to cost-tolerance modeling and often not considered in simple cost-tolerance models is the time required to set-up the manufacturing process, the lot size, the cycle

hours per unit, and the productive hour rate of the operations. By choosing to maximize the reliability of the engine rather than minimize the cost, a potentially large error is removed from the process and the results are not contingent upon any cost-tolerance assumptions. Furthermore, choosing to maximize the standard deviations of the components is not entirely without regard to the manufacturing cost. It is a well-accepted fact that the minimum cost of manufacturing will be achieved when the tolerances are set to their maximum allowable value. In other words, choosing to maximize the standard deviations (with weighting functions proportional to the cost of the part) will yield the tolerances that allow for the lowest manufacturing cost but will not yield what that cost will be. Some simple cost modeling was performed in this work for illustrative purposes and will be discussed in detail in Chapter 5.

2.4 Optimization Algorithms

As stated earlier, there are a number of different optimization algorithms that all have their own inherent strengths and weaknesses. Optimization algorithms can be classified in many different ways, with one of the most general classifications describing whether the technique is exploratory or gradient based. Due to the fact that each type of technique has a unique strength relative to this research, a combination of techniques was employed for this work. Below are brief explanations of the different types of algorithms.

Exploratory algorithms are useful when the design space in question is not well understood or is discontinuous. The advantage of an exploratory algorithm is that it searches over the entire range of the allowable limits and may reveal solution spaces not previously known to the user. Furthermore, since the technique searches over the entirety of the design space, it is not highly dependent upon the initial starting position of the design variables. This is useful when searching for feasible solutions when no single feasible solution is known.

The exploratory technique used in this research is a genetic algorithm. Genetic algorithms are based on the concept of Darwin's survival of the fittest theory. The idea of natural selection is employed in these algorithms where the design variables are selected and allowed to mutate and crossover based on their fitness relative to the entire population. Weak (highly non-feasible) combinations of variables are not selected for further iterations while the most promising combinations of variables are allowed to continue and mate with other feasible solutions. There is a stochastic element involved in genetic algorithms insomuch as feasible combinations can be randomly mated with other feasible combinations to determine if the resultant combination also yields a feasible result. A genetic algorithm will continue to attempt to strengthen each generation until it determines that further changes do not improve the overall feasibility. Common programming considerations for generic algorithms include the maximum number of allowable iterations and the population size. The maximum number of iterations is used to define the level of detail in which the design space is covered and is a factor that drives the execution time of the algorithm. The population defines the number of times that different combinations of independent variables will be attempted per iteration. A population of 2 to 4 times the number if independent variables is recommended for most applications of a genetic algorithm [11].

Gradient-based algorithms are useful when trying to determine the absolute best fit of an objective function. Generally, gradient-based techniques can find solutions with much smaller tolerance of convergence criteria than exploratory algorithms can. The disadvantage of gradient-based algorithms is that they are highly dependent upon the initial starting position. If a gradient-based technique is initiated in an infeasible design space, it is highly likely that no feasible solutions will be found since it is not capable of exploring the design space very well. Gradient-based methods are analogous to climbing a hill with a blindfold. If you were placed on a hill blindfolded, it would be possible to make it to the top of the hill simply by feeling around your immediate surrounds and determining which way leads up to higher ground. You would know you were on top of the hill when every other path you chose leads down. On the other hand, if you were placed on a flat field adjacent to a hill, no amount of feeling around is going to tell you

where the hill is unless you got lucky and happened to initially walk in the right direction. Below is a description of a gradient-based method known as the sequential quadratic programming algorithm NLPQL.

Sequential quadratic programming – NLPQL is a gradient based numerical optimization technique that was developed by Prof. K. Schittkowski at the University of Bayreuth, Germany in 1985 [12]. NLPQL was originally written in FORTRAN to solve non-linear programming problems that take the form:

$$x \in R^n: \quad \begin{aligned} &\min \ f(x) \\ &g_j(x) = 0 \quad , j = 1,\ldots,m_e \ , \\ &g_j(x) \geq 0 \quad , j = m_e + 1,\ldots,m \ , \\ &x_l \leq x \leq x_u \ . \end{aligned} \tag{2-1}$$

The algorithm assumes that the objective function and constraints are continuously differentiable. The process for a numerical technique would be to determine a search direction and step size and then search along that direction until the minimum (or maximum) is found or a constraint is violated. This process continues with respect to each design variable until all of the objective functions are matched as closely as possible. NLPQL is well suited for smooth non-linear design spaces and requires that user specify only the maximum number of iterations and the final accuracy. NLPQL also has some unique capabilities that make it suitable for optimization of stochastic problems. Stochastic problems are difficult for some numerical techniques to handle since the output of a stochastic model will differ each time even if the inputs are the same. Most numerical techniques expect repeatability of the fitness function given the same set of inputs. In a case where the objective function is to maximize the standard deviation of a Normal distribution, the model is driven stochastically so the outputs may differ each time even though the inputs are identical. Following a recommendation from Dr. Patrick N. Koch of Engineous Software, the NLPQL algorithm was used in this research due to the fact that it appears to handle the issue of non-repeatability better than most direct numerical techniques.

3 Turbine Engine Modeling

3.1 Engine Performance Model

As mentioned in Chapter 2, a state-of-the-art performance code was used for this work. The specific engine model used in this research is similar to that of a moderate bypass subsonic military trainer. Figure 3-1 shows a picture of such an engine.

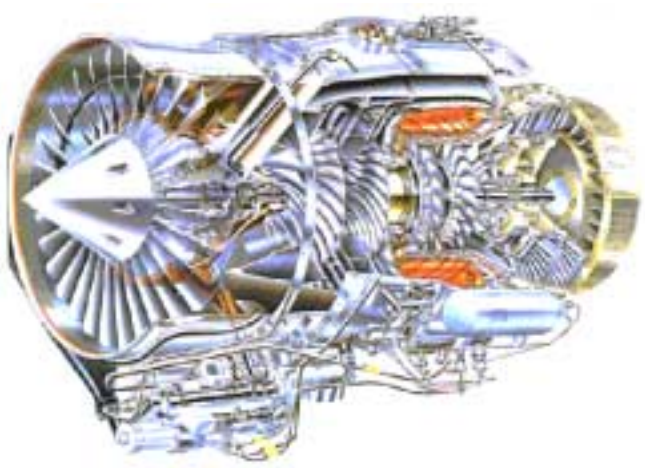

Figure 3-1: Moderate Bypass Military Turbofan Engine

The engine model used for the simulations conducted in this research was tuned to match experimental data. After building and testing the engine in question, the performance model was adjusted to match the experimental data so as to capture the performance trends resulting from the unique build specifications of that engine. The result is an engine model that is very accurate and that can be used to extrapolate realistic performance trends. There was, however, no experimental work done to confirm the results of this research because any confirmation would require an alteration to the tip clearances that would permanently affect the performance of the engine. Again, the results presented here are specific to the engine used in the model, but the methods developed in the research are generic and applicable to any turbine engine.

The engine performance model is built in stations similar to a method commonly employed for cycle analysis. For example, the point immediately in front of the fan face is assigned a station number, *x*, while the point at the exit of the fan is assigned the next logical station number, *x+1*. This type of model allows for the performance of the entire engine to be explored as a function of manipulating specific engine components. To perform this work, the efficiency, pressure rise, and flow rate of each component was altered within the engine model to simulate the effect of changing tip clearances. The tip clearance loss models used in this research were developed in MathCAD™ (Academic Version 8), and the results were then programmed into the engine model.

Similar to the example engine shown in Figure 3-1, the engine used in this research is a turbofan with a high and low pressure centrifugal compressor connected to the same shaft (NH). There are also two high pressure and two low pressure turbines each connected to separate shafts. The design point bypass ratio of the engine is about 6.0 and the maximum rated thrust of the engine is 1330 lbf at sea level.

3.2 Individual Component Models

The engine used in this analysis is a relatively small engine when compared with other military turbofan engines. The size and configuration of this engine requires that unique aspects be considered when modeling the performance. The modeling and special considerations for each component are described in detail under their respective sub-headings below.

3.2.1 Fan Clearance Loss Model

The modeling of the fan tip clearances was done by scaling the fan efficiency while assuming constant work. Fan tip clearances can be accurately simulated over small ranges of tip clearance variation by scaling the efficiency while assuming the work is constant, as described by

$$W = C_p \cdot T_{o1} \cdot \frac{1}{\eta} \cdot \left(PR^{\frac{\gamma-1}{\gamma}} - 1 \right) \qquad (3\text{-}1)$$

The work in Equation 3-1 is calculated directly from the fan performance maps that describe changes in efficiency and pressure rise as a function of mass flow rates and spool speeds. Figure 3-2 shows an example of a 100% speed line (solid line) of the fan tip map along with contour lines of constant work. The dotted line represents a 1.0% efficiency degradation from the solid line. Given the line of scaled efficiency and the constant work assumption, a new flow rate can be determined directly from the performance maps. For example, assume that the fan is operating on the solid line at the coordinates described by point number one. To determine how the fan will operate as a result of a 1.0% efficiency degradation at the tip, a line of constant work is followed from point number one until it meets the 1.0% scaled efficiency line (point number two). The new flow rate is then determined from the coordinates of point number two. With the newly determined flow rate and scaled efficiency, a new pressure rise can be predicted with Equation 3-1.

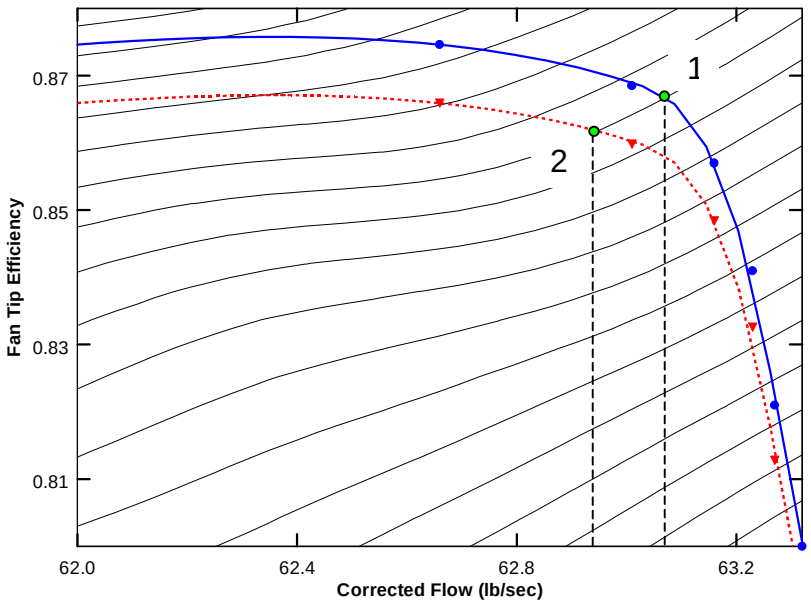

Figure 3-2: Changes in Fan Tip Efficiency Assuming Constant Work

The constant work assumption for modeling fan tip clearances is reasonable because the tip losses in axial machines are proportional to the clearance (τ) over the blade height (h). Figure 3-3 illustrates this concept for a simple stator / rotor row. Because the fan of a turbofan engine is generally very large in comparison to the tip clearance, the ratio of lost work to work is very small. The blade height of a fan is usually at least three orders of magnitude greater than the tip gap making the loss in work due to of the tip gap of the fan nearly negligible.

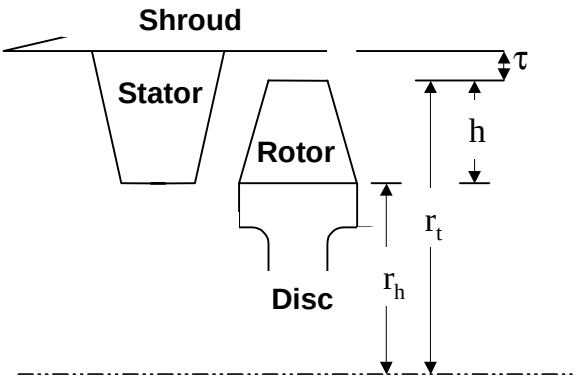

Work:
$\sim U \Delta C_\theta \rho C_z \pi (r_t^2 - r_h^2)$

Lost Work:
$\sim U \Delta C_\theta \rho C_z \pi (2 r_t \tau)$

Lost Work / Work:
$\sim (2 r_t / r_t + r_h) (\tau / h)$
$\sim \tau / h$

Figure 3-3: Tip Clearance Losses Proportional to Blade Height [5]

To describe the change in fan efficiency as a function of the tip gap, non-dimensional experimental data from a similar military turbofan engine were used. After the data were scaled to match the dimensions of the fan used in this research, the model seen in Figure 3-4 was generated.

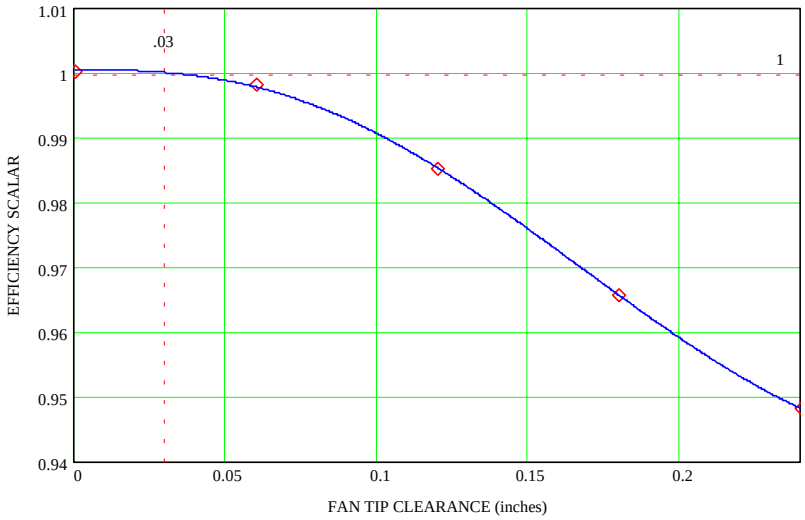

Figure 3-4: Fan Tip Clearance Loss Model

Only five experimental data points were available to model the change in efficiency as a function of the tip clearance of the fan. The five data points were fit with a polynomial, and the model was used in the algorithm describing the change in efficiency as a function of the fan tip clearance. The product of the efficiency scalar and the reference efficiency is the new efficiency resulting from a change in the tip gap from the reference tip gap. The reference tip gap for the fan model is the nominal cold build clearance of the fan, assumed to be 0.03 inches (30 mils). The dashed lines in Figure 3-4 show the reference tip gap corresponding to an efficiency scalar of 1.0. As the tip gap grows, the efficiency scalar becomes smaller which in turn results in a lower fan tip efficiency. Conversely, as the tip gap gets smaller, the efficiency increases. Figure 3-4 also supports the discussion regarding Figure 3-3 in that the tip gap losses are proportional to the ratio of the tip clearance to the blade height. The blade height of the fan is assumed to be approximately 6 inches and the nominal tip gap only 0.03 inches. As a result, a relatively large tip gap must exist before the efficiency loss in the fan becomes a significant factor. Deviations in the tip gap (± 20 mils) from the nominal of 30 mils result in very little efficiency change (approximately 0.02%). This relatively small change in efficiency will not be the case with much smaller blades such as those used for the HPT.

Due to the fact that the length of the fan blades will grow as a result of thermal and centrifugal loading during operation, there is a limit to how small the cold build tip gap can be before the fan scrapes the side of the inlet duct while running. The cold build clearance is named as such because this is the tip clearance that exists after the engine is built and the blade tip clearances are measured at 68°F. The length of the tip gap while the engine is at full power is known as the hot running clearance. The difference between the cold build clearance and the hot running clearance is the amount the blade grows under full power. If the blade growth exceeds the cold build clearance, the fan blades will scrape the side of the casing resulting in performance degradation or potential engine failure.

It is possible to do some simple calculations that determine the order of magnitude of the blade growth under full operation. After performing these calculations and conferring with fan blade design engineers, it was determined that the hot running clearance of the fan being modeled was about 10 mils. This implies that there are 20 mils of fan blade growth while the engine is running under full power. As a result, the practical usefulness of Figure 3-4 does not extend below 20 mils. While Figure 3-4 indicates continually increasing efficiency down to 0 mils cold build clearance, the reality is that any cold build clearance below 20 mils will result in the fan scraping the side of the casing under full power and is not a satisfactory solution.

3.2.2 HPC / LPC Clearance Loss Models

Centrifugal compressors like those of the LPC and HPC of this particular engine require some unique modeling criteria that differ somewhat from the criteria used for axial turbomachinery. In this work, the tip clearance of the centrifugal compressor is defined as the gap between the impeller and the casing that is found at the end of the impeller blade and the entrance to the diffuser as shown in Figure 3-5. Although the tip gaps at the entrance to the impeller and at the mid-span are most likely factors that contribute to losses as well, they are not considered in this work. The experimental data to support any relation between the flow losses and the tip gap at the entrance or mid-

span of the impeller was not available. There was, however, a significant amount of experimental data available for modeling the effect of the tip gap found at the exit of the impeller.

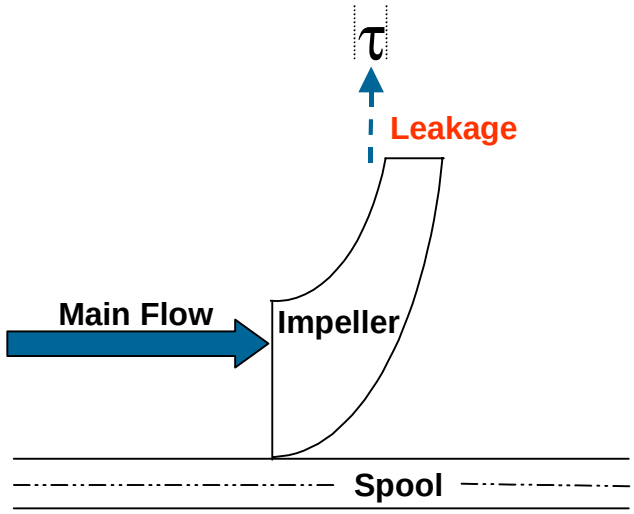

Figure 3-5: Centrifugal Compressor Tip Clearance Geometry

The reason for the abundant experimental data regarding the centrifugal compressors is due to the fact that the engine manufacturer was experiencing some difficulty with the centrifugal compressors during development and decided it was worthwhile to generate models that relate the compressor performance to the tip gap at the exit of the impeller. As a result, there is no constant work or constant pressure rise assumption used in the modeling of the centrifugal compressors. Figure 3-6 shows the linear models generated from experimental data that relates the LPC and HPC efficiency to the cold build tip clearance.

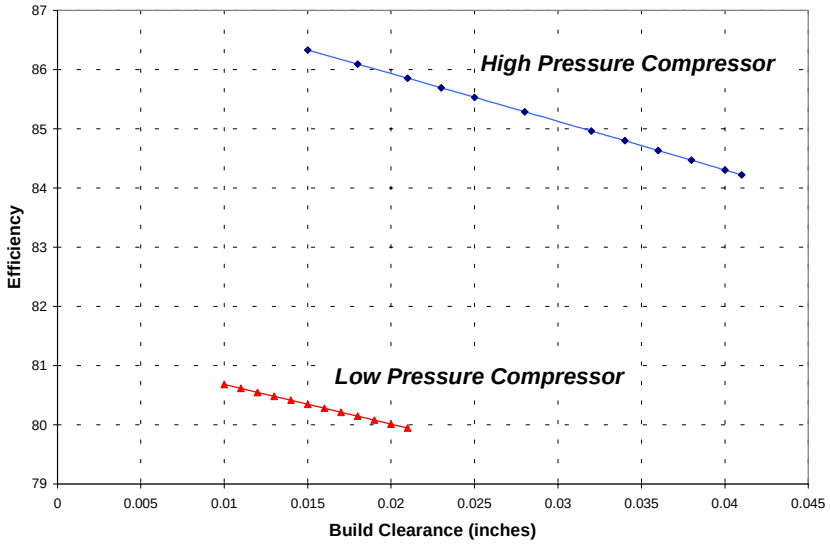

Figure 3-6: Compressor Efficiency Loss Versus Cold Build Tip Clearance

Although the tip gap in a centrifugal compressor is certainly a factor that contributes to a loss in efficiency, the tip gap does not affect a radial compressor as much as an axial compressor. Only about 60% of the losses in a centrifugal compressor actually come from the impeller itself. The other losses come from the vaned diffuser space (~20%) as well as the flow bends and vaneless space (~10%). Of the impeller losses, only about 10% is attributed to clearance losses with nearly 25% attributed to secondary losses, 25% to blade loading, and 25% to skin friction. While the percentage of losses that result from radial compressor tip clearances may not appear to be much, even a 1% loss in compressor efficiency is considered a significant detriment. Furthermore, tip clearances are one of the few aspects that contribute to losses that can be controlled by manufacturing. The other factors that contribute to losses in a radial compressor, such as blade loading and secondary flows, are largely functions of flow physics and not easily controllable.

Similar to axial compressors, radial compressors also experience growth due to thermal and centrifugal loading during operation. When a centrifugal compressor grows under thermal and centrifugal loading, it tends to bloom outward like a flower. The swept blades of the impeller are forced to straighten outward in the radial direction while

at the same time stretching the impeller disc. Figure 3-7 shows the relation of the cold build clearances to the hot running clearances of the HPC and LPC under full power. As Figure 3-7 shows, the hot running clearance of the LPC is nearly zero when the cold build clearance is 10 mils, implying that the tip clearance will shrink approximately 10 mils under full power. After conferring with the radial compressor designers, it was decided that the minimum tip clearance to allow for full power growth of the LPC and HPC is 10 mils and 14 mils respectively.

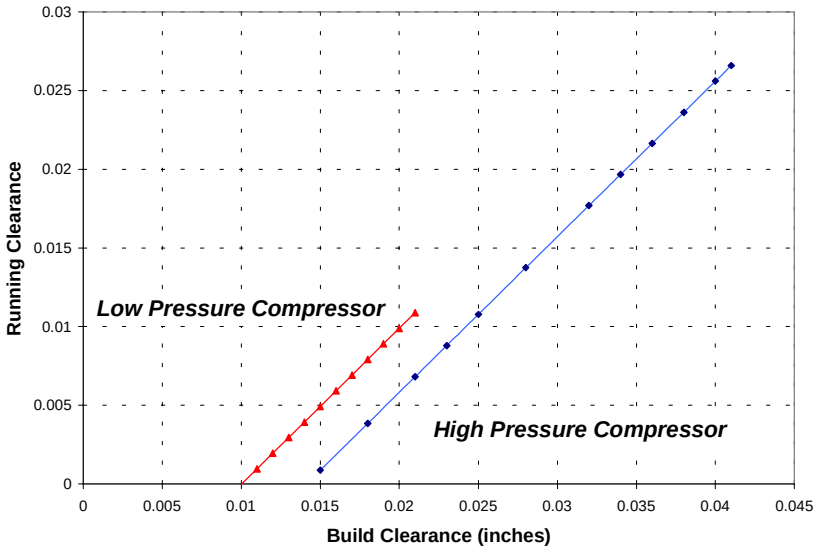

Figure 3-7: Compressor Build Clearances Versus Running Clearances

Another factor that is unique to modeling the tip clearance of a radial compressor is the phenomenon known as spool shift. Referring to Figure 3-5, the pressure on the backside of the impeller is higher than the pressure on the front side of the impeller. Under extreme circumstances such as an engine surge, where the pressure ratio of the compressor can change suddenly, there must be adequate allowance in the tip gap to prevent the impeller from striking the side of the casing. Spool shift can also result from high g maneuver loads such as vertical acceleration / deceleration of the aircraft. Modeling spool shift and the transient flows associated with engine surge was not accounted for when modeling the compressors. Given information on the magnitude of the spool shift resulting from transient flows, it would be a simple task to combine the

maximum spool shift distance with the maximum disc growth and generate a new minimum cold build clearance that would not be safe to exceed.

3.2.3 HPT / LPT Clearance Loss Models

Of all the potential loss mechanism in axial turbines, tip clearances are the most profound. Figure 3-8 shows a number of different turbines manufactured by Garrett Engine Company, and a break down of the loss mechanism seen in each turbine. As the figure shows, tip clearance losses are the single most significant contributor to losses across all the turbines. Tip clearances in axial turbines are also the most detrimental tip clearances found in turbomachinery. The performance of a turbine engine is more dependent upon the tip clearances of the turbines than the tip clearances of the compressors or fan. For this reason, the modeling of the turbine losses must be especially accurate.

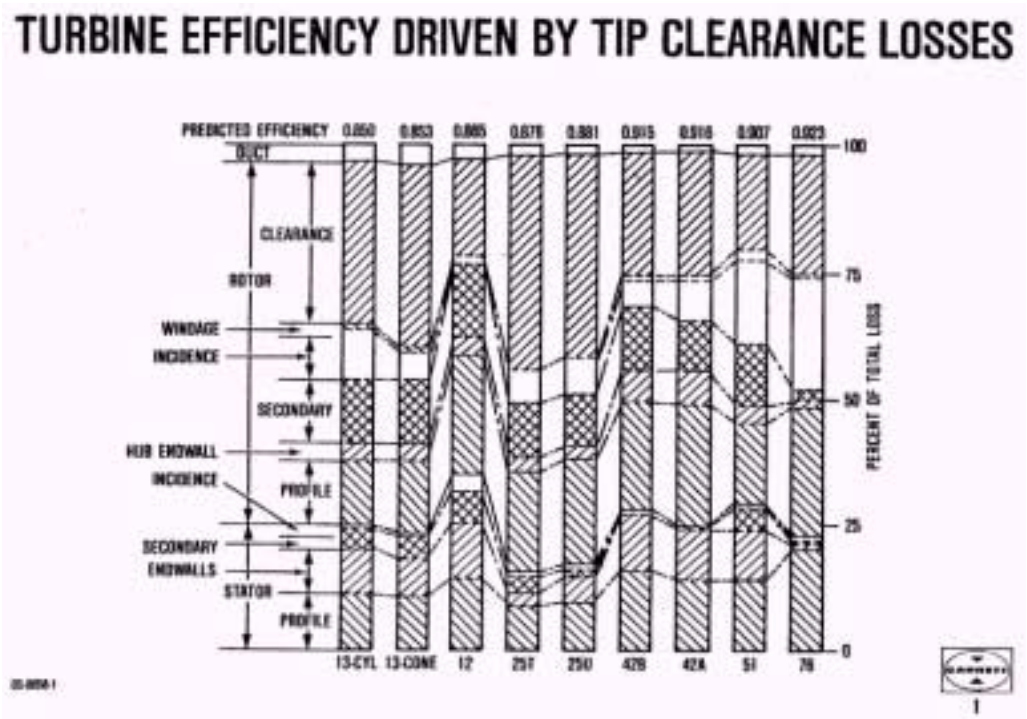

Figure 3-8: Turbine Flow Loss Mechanisms [5]

Unlike the compressors, there was no experimental data available to model the performance of the turbines as a function of the tip clearance. Instead, linear correlations similar to those presented in Figure 2-3 were used. The actual turbine clearance loss correlations that were used in this work were developed by the manufacturer and are proprietary. Equation 3-2 shows the general form of the correlations developed by the manufacturer that model the relationship between the turbine efficiency ratio and the tip clearance.

$$\frac{\eta_{new}}{\eta_{ref}} = 1 - k \cdot C \cdot \tau \tag{3-2}$$

Again, τ is the tip clearance and k is a blade-loading factor that accounts for the aero-loading of the turbine blade. C is a constant that includes the ratios of specific blade geometries. Since the desired result of the turbine clearance loss model is to have an efficiency scalar that is calculated as the difference between some new tip clearance and the reference clearance, Equation 3-2 must be manipulated further.

$$\frac{\eta_{new}}{\eta_{ref}} = \frac{1 - k \cdot C \cdot \tau}{1 - k \cdot C \cdot \tau_{ref}} \tag{3-3}$$

With Equation 3-3, modeling the change in turbine efficiency as a function of the change in the nominal cold build clearance is possible. Equation 3-3 includes τ_{ref} which is the nominal cold build clearance for the turbines and is assumed to be 30 mils and 20 mils for the HPT and LPT respectively. A plot of Equation 3-3 as functions of tip clearance for both the HPT and LPT is shown in Figure 3-9.

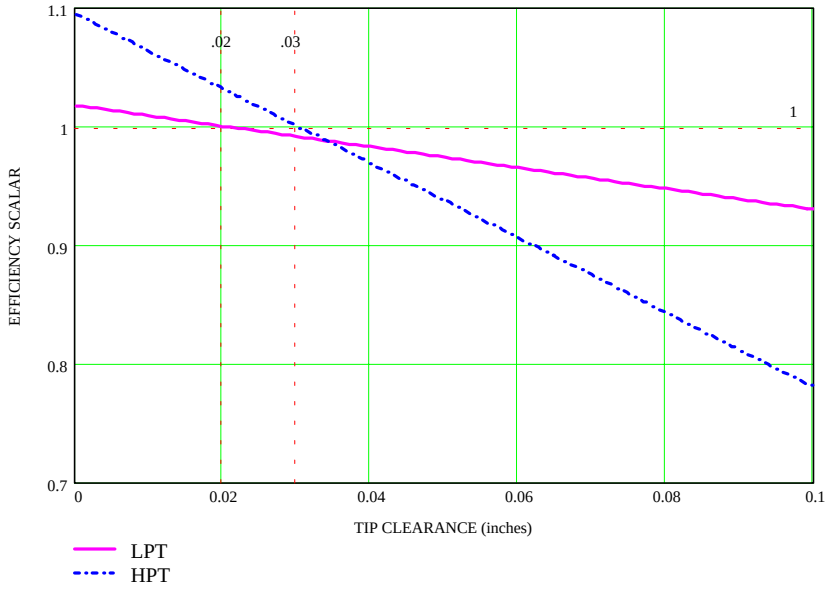

Figure 3-9: Turbine Loss Correlations

As Figure 3-9 shows, the slope of the HPT line is greater than the slope of the LPT line, indicating that the HPT efficiency is more sensitive to tip clearance losses than the LPT. This difference in sensitivity can also be partly attributed to the fact that the LPT blades are shrouded with knife-edge seals that prevents the flow from easily escaping around the tip. Even shrouded blades however, do not block tip leakage entirely. As a result of correcting Equation 3-2 for the nominal cold build clearances, an efficiency scalar of 1.0 is seen for the respective nominal cold build values assumed to be 20 mils and 30 mils for the LPT and HPT. In essence, a scalar of 1.0 tells the engine model to carry out the performance calculation assuming that there is no difference from the nominal cold build clearance. If the clearance is increased above the nominal level, the efficiency scalar will be less than 1.0, resulting in a lower efficiency and a new flow rate and work solved via a constant pressure rise assumption. If the tip gap is less than the nominal cold build clearance, the opposite holds true and the efficiency will be increased.

Similar to the fan model, there is a limit to how low the cold build tip clearance can be before running a risk of scraping the side of the casing. Although turbine blades are much shorter than fan blades, the centrifugal stress is not necessarily any less due to

the fact that HP turbines spin at a much higher speed and LP turbines can house relatively heavy shrouds on their tips. In addition, the thermal loading of turbine blades is much more of a factor than with fan blades because the turbine blades lie in the hot gas path. As with the fan, some simple calculations can reveal the magnitude of growth that can be expected from both the HPT and LPT blades because of thermal and centrifugal loading. Equation 3-4 was used to find the total stress in the HPT and LPT blades under full power conditions.

$$\sigma_{tot} = \int_{r_h}^{r_t} \frac{\rho \cdot \omega^2 \cdot r_t^2}{2 \cdot (1 - \alpha \cdot r)} \cdot \left[1 - \left(\frac{r}{r_t} \right)^2 - \frac{2 \cdot \alpha \cdot r_t}{3} \cdot \left[1 - \left(\frac{r}{r_t} \right)^3 \right] \right] dr \tag{3-4}$$

While not entirely accurate, an assumption was made that the blades were made purely of nickel and had density ρ=8602 kg/m^3. Also, the degree of taper, α, was assumed to be 0.5 divided by the tip radius, r_t. After calculating the total stress and dividing by the modulus of elasticity for nickel, elongation for the turbine blades resulting from the centrifugal forces under full power were found. Next, the growth of the turbine blades was determined as a function of the thermal loading using Equation 3-5.

$$\Delta l_t = \varepsilon \cdot (T_{o4} - T_{o1}) \cdot h \tag{3-5}$$

In Equation 3-5, ε, is the coefficient of thermal expansion for nickel (13.3 x 10^6/K), T_{o1} is assumed to be standard day sea level static ambient temperature, and T_{o4} is the turbine inlet temperature. After assuming a nominal growth in the turbine disc and summing the growth of the turbine blades resulting from the centrifugal and thermal stresses under full power, it was found that the HPT and LPT could be expected to grow approximately 14 and 11 mils respectively.

It is important to note that the tip clearances of all turbomachinery components will vary at different points in the flight envelope and different ambient conditions. Cold days will result in different tip clearances than hot days simply due to the difference in ambient temperature of the intake air. Additionally, takeoff (full power) conditions will result in different tip clearances than cruise conditions. Figure 3-10 shows the magnitude of turbine tip clearances at different points in a flight mission.

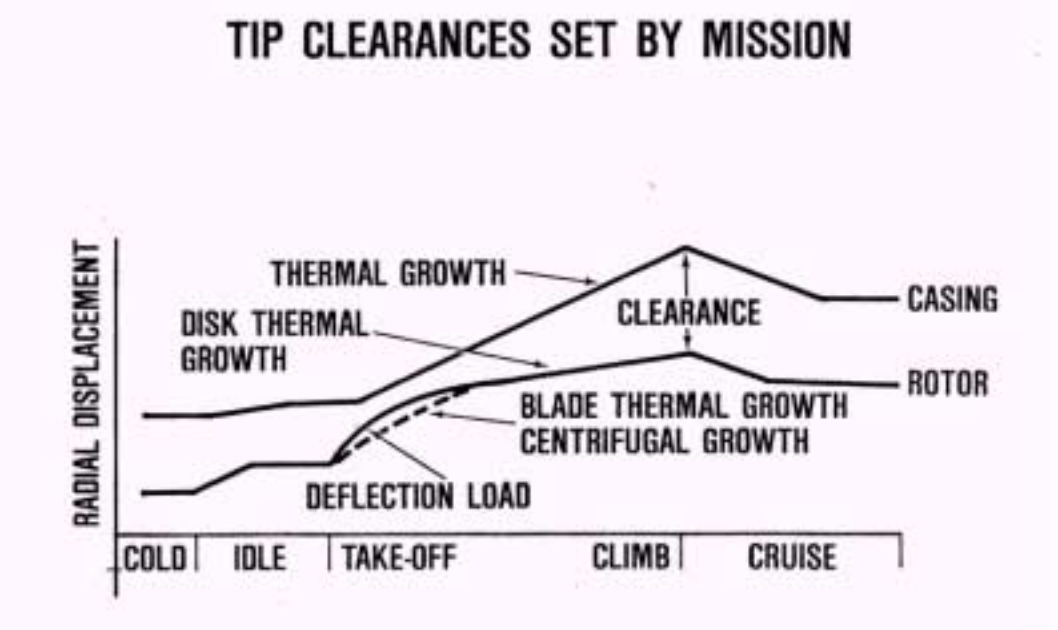

Figure 3-10: Turbine Tip Clearances Set by Mission [5]

The pinch point, or point of minimum tip clearance, is seen immediately after takeoff when the turbine rotor is under full thermal and centrifugal loading while the casing has not begun to expand. Since the casing is more massive than the rotor and has cooler air in contact with some of its surfaces, it does not expand as rapidly as the rotor. Given ample time, the casing will heat up and expand later in the flight envelope. This expansion will result in an increase in the turbine tip gap. Without active cooling schemes to control the thermal expansion of the casing, any effort to minimize the tip gap must be concentrated at the pinch point. Minimizing the tip gap at any other point in the flight envelope will result in the turbine blades scraping the casing at takeoff. Another potential area in the flight envelope where a pinch point may occur is called a re-burst or Bode. A re-burst is most profound when the aircraft is flying at high speeds in cold air

and the engine is taken from full power to idle and then back to full power. When the engine is taken to idle in mid-flight, the cool ambient air will cause the casing to contract more rapidly than the rotor disc. As the engine is taken back to full power, the blades will once again expand quicker than the cooled casing and could result in turbine blade contact with the casing. The reason for a turbine scrape at takeoff and during a re-burst are the same, the rotor expands more quickly than the casing. It should be noted however, that the two circumstances could occur at very different points in the flight mission.

3.2.4 LPT / HPT Nozzle Areas

The only non-rotating components included for analysis in this research were the turbine nozzle areas of the HPT and LPT. The geometric dimensions of the turbine areas are available as direct inputs to the performance code and therefore do not require any additional programming interface as does the clearance loss model. The turbine nozzle areas were included in the analysis due to the fact that their variation can have a significant impact on the performance of the engine. Turbine nozzles are often the first piece of hardware that is changed by an engine manufacturer in an attempt to fix out-of-limits performance of a newly manufactured engine. Turbine nozzle sections have long been known to have a large impact on performance. Another reason for including turbine nozzle areas in the analysis is that they offer a good balancing metric for the stall margin constraints. Stall margin is the only performance parameter included as a constraint that is degraded by decreasing turbine nozzle areas.

A concern when modeling the turbine nozzle areas is the magnitude of the thermal expansion resulting from different turbine inlet temperatures. As the tip clearances of the compressors up-stream of the turbines are increased, the turbine inlet temperature increases. A hotter turbine inlet temperature will result in a greater degree of thermal expansion of the turbine nozzles and thus a larger turbine nozzle area. Although the thickness of the turbine nozzle walls will increase slightly as a function of temperature as well, the change in turbine nozzle area is dominated by the expansion of

the inner and outer radial rings resulting in a net increase in turbine nozzle area as temperature is increased. Experiments conducted by the engine manufacturer indicate that this slight increase in turbine nozzle area as a function of increased temperature is offset by the fact that higher temperatures are associated with lower values of gamma (ratio of specific heats). Assuming a choked turbine nozzle, as the turbine inlet temperature is increased, the throat area will increase, the value of gamma will decrease, and the mass flow rate will remain nearly constant as described by Equation 3-6.

$$\dot{m} = \frac{A_{throat} \cdot P_o}{\sqrt{R \cdot T_o}} \cdot \sqrt{\gamma} \cdot \left(\frac{2}{\gamma + 1} \right)^{\frac{\gamma+1}{2 \cdot (\gamma - 1)}} \tag{3-6}$$

As a result, the effective flow area of the turbine remains nearly constant as turbine inlet temperature increases.

3.3 Component Sensitivities

After generating the methods to model the geometric variation of the components, it was possible to perform a sensitivity analysis. A sensitivity analysis is useful in determining which component variability has the greatest effect on the engine performance parameters. To generate the sensitivity analysis, commercial software called iSight™ was coupled with the engine performance code. When iSight™ is properly coupled with the engine model it can write to the inputs required by the model, read the outputs, and even drive the execution of the performance code itself [13]. This type of interface and automation allows for a high degree of flexibility and a useful means to quantitatively explore the performance model. The sensitivity analysis performed by iSight™ is a simple parametric algorithm known as percent total effect. To perform the sensitivity analysis, each input parameter (tip clearances and turbine nozzle areas) was independently varied from its nominal value by 5.0% while holding all other values constant. iSight™ then takes the difference (absolute value) of the output parameter value and its nominal value. The differences from each input parameter's +/- variation are summed, and then divided by the overall sum of the differences to get each

parameter's percent total effect. This method will effectively hide any nonlinear behavior within the region of interest, but does provide a reasonable quantitative measure of sensitivity. Blue bars in Figures 3-11 through 3-16 indicate a positive correlation between the input and output parameters, while red indicates a negative correlation. The Pareto Plots (which are essentially sorted histograms) used to illustrate the percent total effect can also offer some insight as to the physics of what the model is describing. For example, in Figure 3-11, the first blue bar indicates that an increase HPT nozzle area will result in an increase in high pressure compressor surge margin while an increase in the HPT tip clearance will cause a decrease in the HPC surge margin.

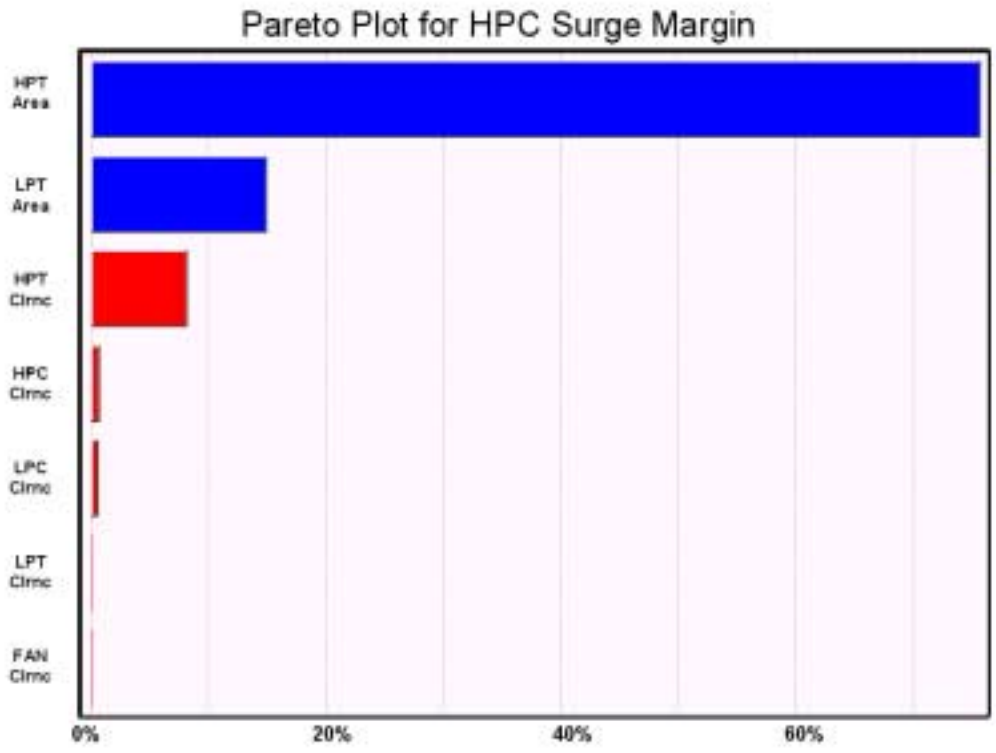

Figure 3-11: Component Percent Total Effect on HPC Surge Margin

Figure 3-11 also supports the claim that the inclusion of turbine nozzle areas offers a good balancing metric to the surge margin since the turbine areas are the only geometries that are positively correlated with surge margin. Furthermore, it is evident that the turbine nozzle areas have the greatest percent total effect on the surge margin.

Figure 3-12 shows the percent total effect on the high speed spool (NH) rotational speed as a function of parametrically changing each of the model geometries by 5.0%. The plot shows that increasing the LPT nozzle area will increase NH while increasing the HPT nozzle area decrease NH. The author theorizes this is the case due to the following: as the LPT nozzle area is increased, the backpressure on the exit of the HPT is decreased, causing an increased pressure gradient across the HPT. This increase in pressure gradient is forcing the NH spool to spin faster. Conversely, as the HPT nozzle area is increased and the flow held constant, the velocity through the nozzle is decreased causing a subsequent reduction in the NH spool speed.

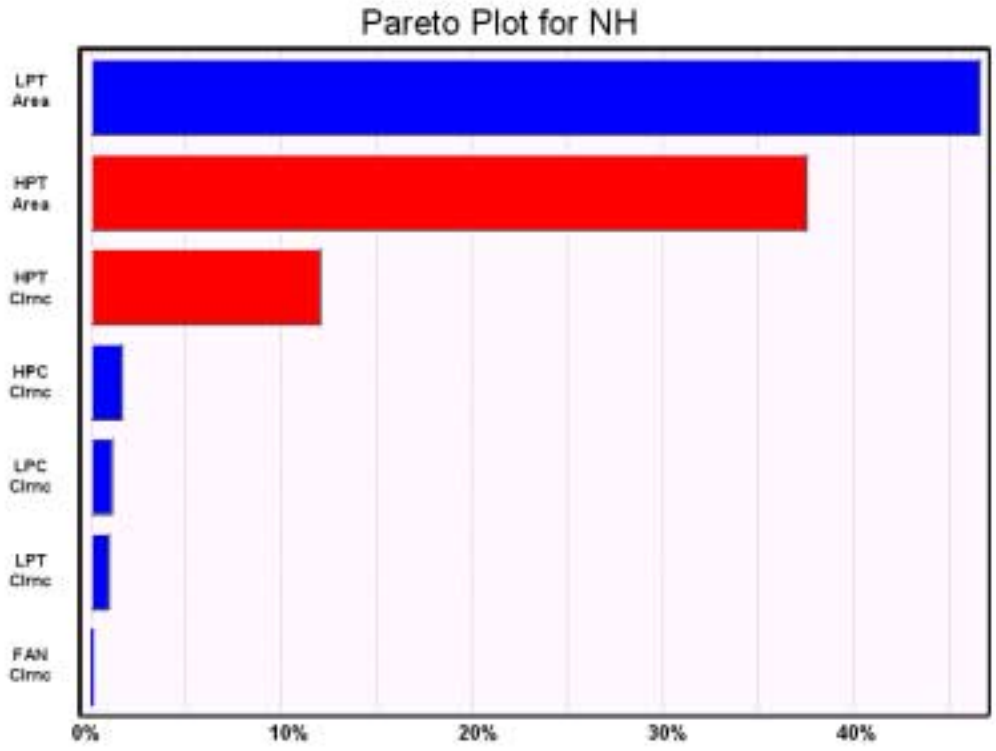

Figure 3-12: Component Percent Total Effect on NH

Another sensitivity analysis of interest is the plot quantifying the percent total effect on the engine's SFC. Figure 3-13 shows that an increase in any of the geometries will have a detrimental increase on the SFC. Another trend that can be seen across plots 3-11 through 3-16 is that the turbine nozzle areas are consistently the most influential geometries affecting the performance of the engine.

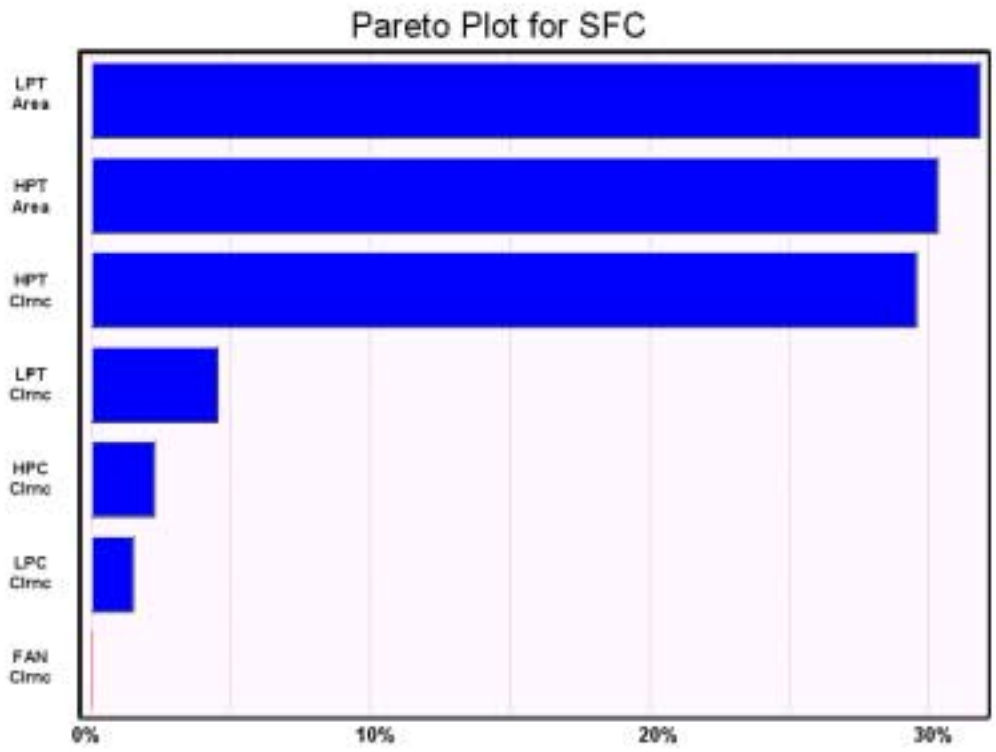

Figure 3-13: Component Percent Total Effect on Specific Fuel Consumption (SFC)

Sensitivities for the remaining performance constraints included in this research are included below.

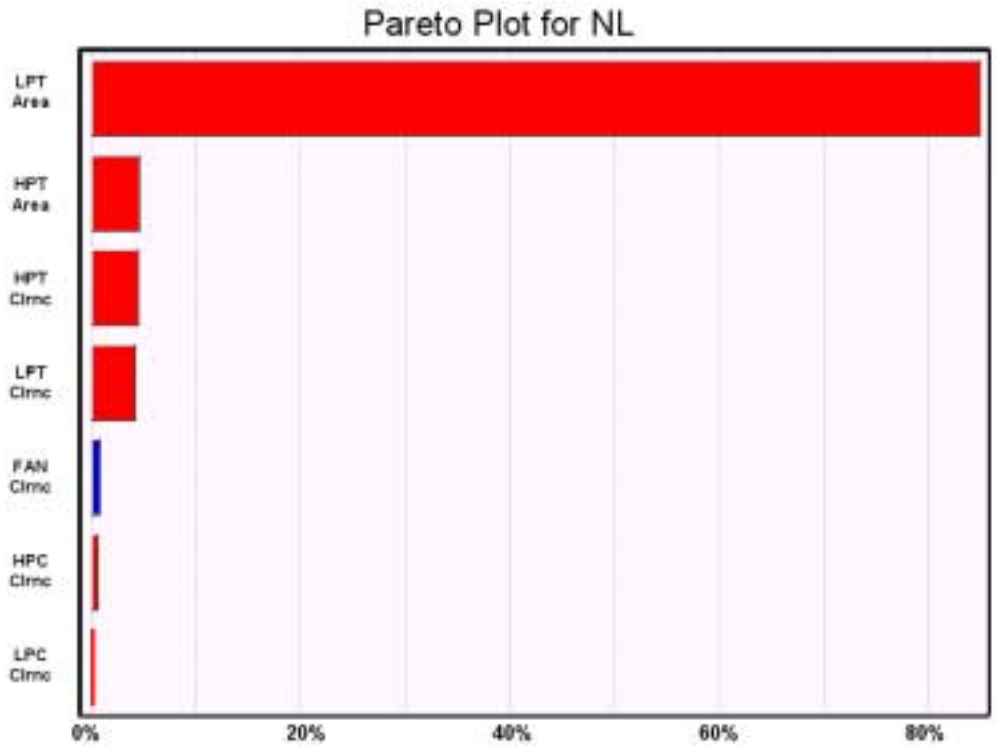

Figure 3-14: Component Percent Total Effect on NL

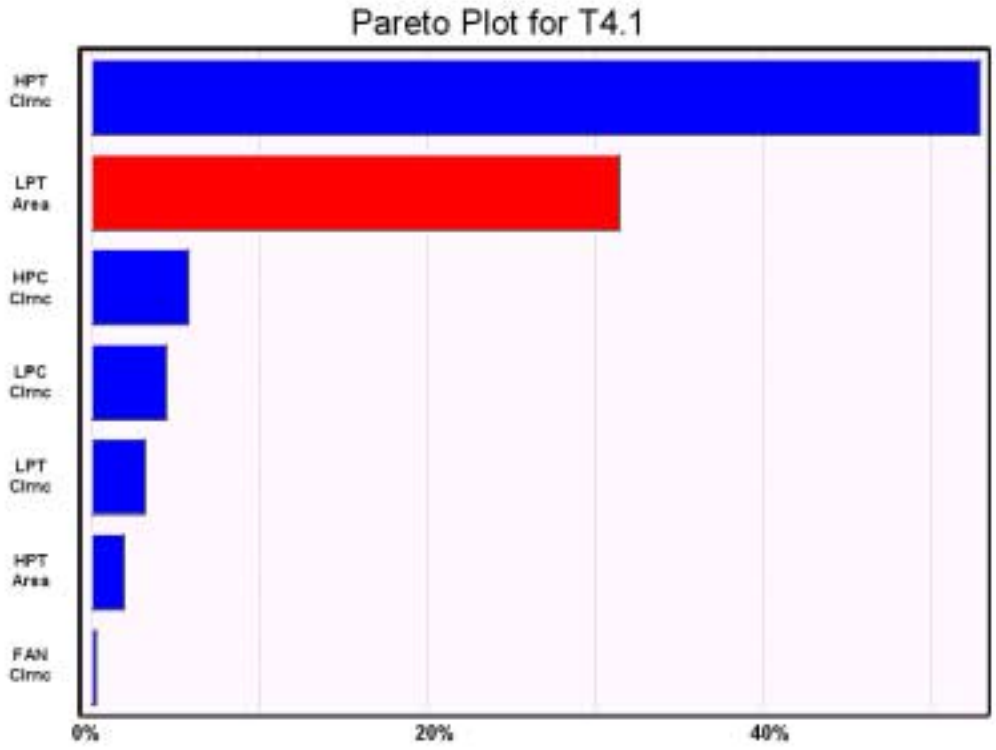

Figure 3-15: Component Percent Total Effect on Turbine Inlet Temperature (T41)

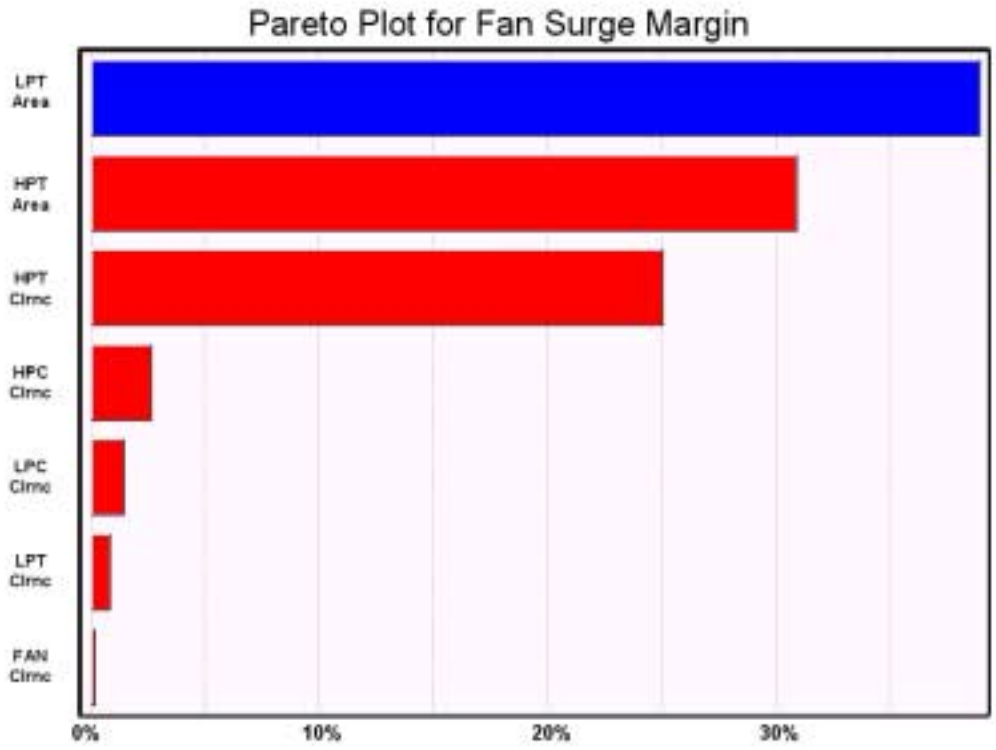

Figure 3-16: Component Percent Total Effect on Fan Surge Margin

Another method of quantifying the sensitivities of the components is to examine the relationship between the build clearance and the efficiency. This method is only defined for those components that have tip clearance models (the rotating components) but does offer a unique comparison. Figure 3-17 shows the relationship between the cold build tip clearance and the efficiency for each of the rotating components described in sections 3.2.1 through 3.2.3. Figure 3-17 was generated by means of varying the nominal build clearance by a constant percentage so the components with smaller cold build clearances have proportionally shorter lines. It is however, the slope of the lines in Figure 3-17 that are of note. The slope is in effect quantifying the sensitivity of the clearance to the efficiency.

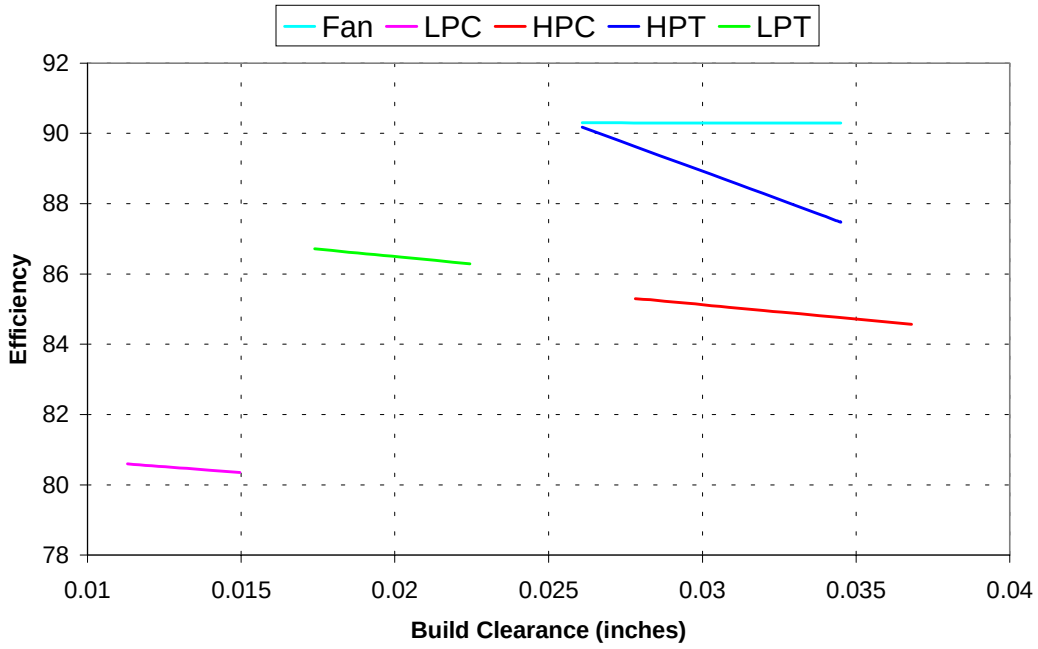

Figure 3-17: Build Clearance Versus Efficiency of Rotating Components

It is clear from Figure 3-17 that the HPT efficiency has the greatest slope of all the components listed. Moreover, the slope of the fan line has the least slope and is the least sensitive to tip clearance for reason explained in section 3.2.1. Although apparently zero, the fan line does in fact have some curvature but is difficult to see over the range that is shown. The fan line also happens to be the only non-linear tip clearance loss model in the group. Table 3-1 quantifies the slope of each of the components as well as the R^2 value.

Table 3-1: R^2 and Slope for Each Clearance Loss Model

	R^2	Slope
Fan	0.97	-0.5
LPC	1.0	-66.5
HPC	1.0	-81.8
HPT	1.0	-321.1
LPT	1.0	-84.9

An R^2 (coefficient of variation) of 1.0 in Table 3-1 indicates that the model is perfect fit to all the data points while anything other than 1.0 indicates that there is some variance between the data and the model. Table 3-1 reveals that the HPT efficiency is 3.8 times

more sensitive to tip clearance than the next most sensitive component (LPT). Also of note is the fact that the sign of each of the slopes shown in Table 3-1 is negative, indicating a loss in efficiency as tip clearance is increased.

4 Methods of Simulating Performance Variation and Costs

4.1 Response Surface Models

Response surface models (RSMs) can be used to approximate the response of an engine performance code with simple algebraic functions. In this work, response surface models were used to approximate the engine performance parameters (T41, SFC, NH, NL, Fan Surge Margin and HPC Surge Margin) as functions of the tip clearances and the turbine nozzle areas. The advantage of using an RSM rather than the engine performance code itself is that the RSM can be executed with only a fraction of the computational power and speed required to run the actual code. The reason for the decrease in computational requirements with an RSM is that the evaluation involves only the calculation of a polynomial given a set of input variables. A second order RSM takes the form:

$$F(x) = a_o + \sum_{i=1}^{N} b_i \cdot x_i + \sum_{i=1}^{N} c_{ii} \cdot x_i^{\,2} + \sum_{ij(i<j)} c_{ij} \cdot x_i \cdot x_j \qquad (4\text{-}1)$$

where F(x) is the predicted value of the response given a particular combination of x, the explanatory variables. The coefficients (a, b, c) of Equation 4-1 are determined by solving a linear system of equations consisting of one equation for each analysis point with the goal of minimizing the difference between the observed values and those predicted by the model. The accuracy of the RSM is highly dependent upon the amount of data available and the size of the design space. The more data that can be used to generate the RSM, the more accurate the model will be at any given point in the design space.

The RSMs used in this work are second order and were generated in iSight™. Each performance parameter approximation requires a separate RSM. For example, one

RSM was used to predict T41 as a function of the tip clearances and nozzle areas; another RSM was used to predict SFC as functions of the same variables. In all, six RSMs were created to predict the six different performance parameters included for analysis in this work.

Since the RSM used to approximate the performance parameters is created directly from the engine model, iSight™ was coupled with the performance code and programmed to randomly generate input values for the tip clearances and turbine nozzle areas. Using a uniform distribution, 400 points were randomly selected between specific ranges of component geometry. A uniform distribution allows for the equal probability of any number between a specified range of values. Table 4-1 shows the ranges of tip clearance or turbine nozzle area for each component from which each RSM was created. For example, the nominal value of the fan tip clearance is 0.0300" and a design space of ±12% was chosen as the limits of the model validity. Thus, 400 points were randomly selected between 0.0264 and 0.0336 and input to the engine model in order to generate the response as a function of the changes in the fan tip clearance. This process was done simultaneously for all seven component geometries to capture the interaction between combinations of different geometries as well.

Table 4-1: Response Surface Model Limits

	% Model Limits	Min Mean	Nom Mean	Max Mean
Fan	12.00%	0.0264	0.0300	0.0336
LPC	12.00%	0.0114	0.0130	0.0146
HPC	12.00%	0.0282	0.0320	0.0358
HPT	12.00%	0.0264	0.0300	0.0336
LPT	12.00%	0.0176	0.0200	0.0224
LPT Area	5.00%	14.5627	15.3292	16.0957
HPT Area	5.00%	3.2140	3.3832	3.5524

400 points were determined to be sufficient based on the fact that the number of points used for the analysis should be considerably greater (three to ten times more) than the number of terms. The number of points must be noticeably greater than the number of polynomial terms to ensure an accurate model and that the R^2 analysis is not deceiving. R^2 (coefficient of determination) is a measure of the goodness of fit between the RSM

and the actual data points output by the performance code. An R^2 of 1.0 indicates that the polynomial is a perfect fit to all of the data points. However, it is always possible to obtain a perfect fit to *N* points with a polynomial that has *N*+1 coefficients. Therefore, it does not necessarily mean that the RSM is accurate in any space between the data points unless the number of points is considerably greater than the number of polynomial coefficients. In this case, there are seven inputs to the model, which means there are thirty-six possible polynomial terms (one intercept, seven linear terms, seven quadratic terms, and twenty-one interaction terms). 400 points were used in the creation of the RSM since 400 exceeds ten times the number of polynomial coefficients (10 x 36 = 360). There is also a method of computing an R^2 that accounts for the number of data points used to generate the model. This value is known as an adjusted R^2 and will yield a low number when compared to the R^2 value if the number of data points used to generate the model is not sufficient. Table 4-2 shows the R^2 and adjusted R^2 values for the RSMs used in this work. The close agreement between the R^2 and the adjusted R^2 values supports the fact that 400 points was a sufficient number of samples for the creation of the RSMs. Furthermore, all the models indicate values of adjusted R^2 very close to 1.0 and are therefore considered excellent approximations.

Table 4-2: R^2 Values of the Response Surface Models

Model Name	R2	Adjusted R2
SFC	0.99968	0.99965
HPCSM	0.99997	0.99997
FANSM	0.99992	0.99991
T41	0.99977	0.99975
NL	0.99978	0.99975
NH	0.99980	0.99978

Although the adjusted R^2 values are very high for the specific ranges of the engine model used in this research, lower adjusted R^2 values would not necessarily invalidate the use of an RSM. A lower adjusted R^2 (on the order of 0.90) simply means that the model does not fit the observed values quite as well as an adjusted R^2 of 0.999. The model will overestimate in some areas, be close in some areas, and underestimate in other areas. A way to compensate for the use of an RSM with a relatively low adjusted R^2 is to carefully

check any solutions with the actual engine model itself. Using this method, any errors in the RSM will be caught by the actual engine model. The disadvantage of having a low adjusted R^2 of an RSM is that it may take longer to find a solution that agrees with the actual engine model since the RSM may repeatedly find solutions in areas where it is a poor predictor of the actual engine code.

The percent model limits column of Table 4-1 indicates that the rotating component geometries have a design space limit of ±12% around the nominal value while the turbine nozzle areas have only a ±5% margin around the nominal values. These limits were set based on the convergence criteria of the performance code itself. The engine model was not able to converge on solutions that exceeded a 5% variation in either turbine nozzle area or more than 12% variation in any of the tip clearances. This limitation is expected given the nature of some of the iterative solution algorithms programmed into the performance code. The engine model attempts to solve for the performance parameters based on a balance of energy, mass flow rate, and other thermodynamic properties. By taking some of the geometries to extreme values, discontinuities in the balance of some of the thermodynamics properties are introduced which the iterative algorithm is not able to resolve. The fact that the code cannot converge on some extreme combinations of geometry is not detrimental to this work. In fact, if the engine model is unable to converge on a solution, it is a strong indication that the geometry in question is not an acceptable solution (e.g. the engine would not run properly given the geometry in question).

Although an RSM is a powerful tool that can be executed with little computational resources, it does have limitations. An RSM cannot be accurately used to predict a response outside of the design space from which it was constructed. Referring to Table 4-1, any value for the fan tip clearance that is less than 0.0264” or more than 0.0336” will not necessarily be accurately modeled by any of the six RSMs describing the engine performance parameters. This limitation of the RSM becomes an issue when defining a design space that consists of both mean values and standard deviations.

4.2 Design Space Definition

Since standard deviations have been included as design variables, the design space that can be explored and accurately modeled by an RSM is limited. This is the case due to the fact that a distribution, or range of values, must now be accurately modeled instead of just a single point in the design space. For example, when exploring the output distribution of T41 (turbine inlet temperature) as a function of the LPC clearance distribution, there must exist a model that can accurately predict T41 over all ranges of the LPC clearance distribution.

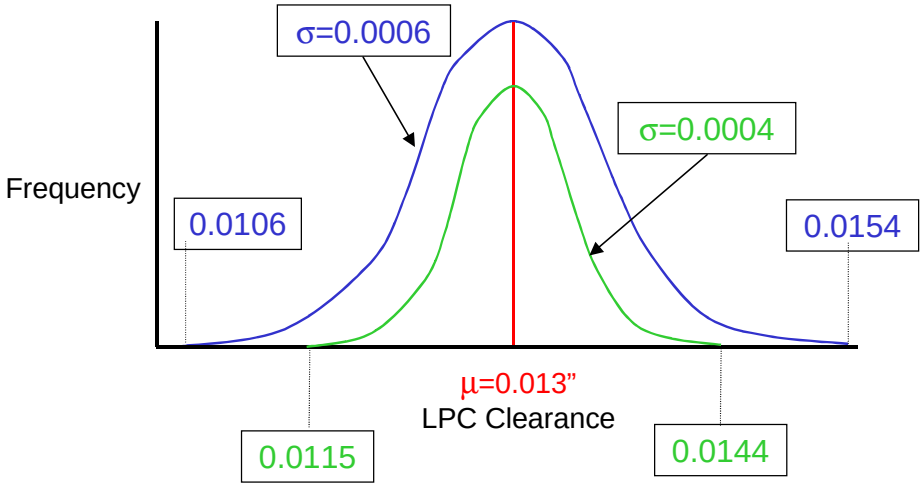

Figure 4-1: Example of Normal Distribution of LPC Tip Clearances

Figure 4-1 illustrates the example given above. If the user wishes to explore the variability of T41 as a function of the variability of the LPC clearance as described by a Normal distribution with mean 0.013” and a standard deviation of 0.0004”, the model describing changes to T41 must be valid over the LPC clearance ranges of 0.0115” to 0.144”. If an even greater variance is explored (0.0006”) then the model describing the changes to T41 as a function of the LPC clearance must be valid from 0.0106” to 0.0154”.

The seemingly simple implication is that a larger design space (ranges of means and standard deviations) will require that the RSM be valid over wider ranges of component variability. Referring back to section 4.1, in order for the RSM to be valid

over a specific design space, it must have been created from values in that design space. As table 4-1 illustrates, there are limits to the range that the RSM can predict due to the fact that the engine code will not converge outside of those limits. In essence, the mean values of the design space are set by the convergence criteria of performance model. When considering mean values only, this limitation is not a concern since, as mentioned earlier, any single point nominal value outside of the convergence criteria is most likely an unacceptable solution. However, as Figure 4-2 illustrates, when standard deviation is introduced as a design variable, the region of design space that can be completely and entirely predicted by the RSM is reduced.

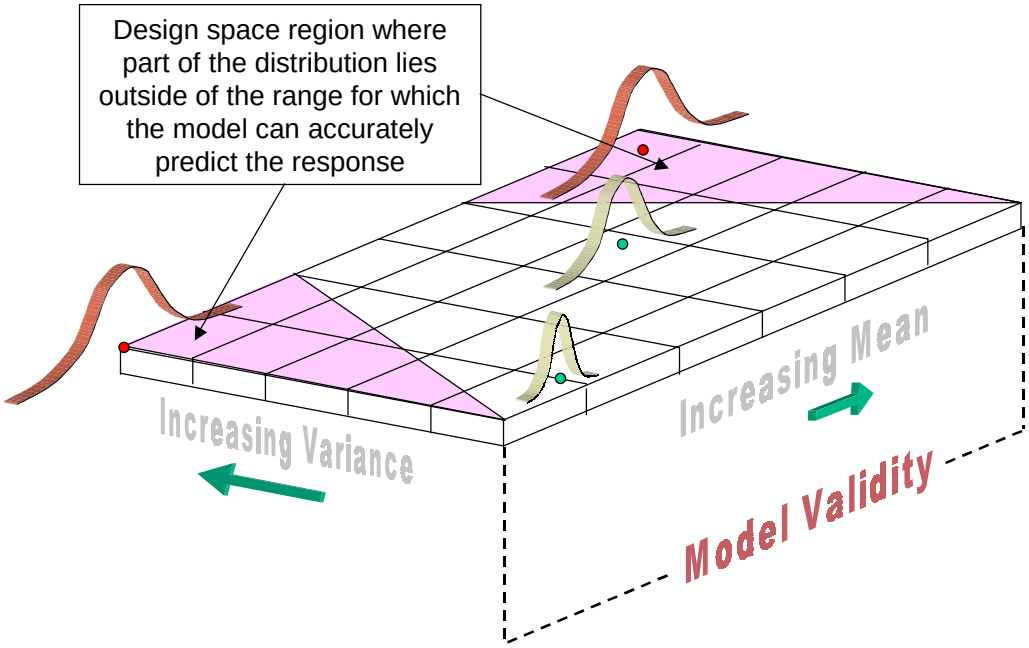

Figure 4-2: Design Space Limitations

Figure 4-2 represents the design space of a single component. The x-axis (labeled Increasing Mean) describes the range of mean values and the y-axis (labeled Increasing Variance) describes the range of standard deviations that will be considered in the design space. At the same time, the range of means shown on the x-axis represents the values from which the RSM was created and for which it is valid. Returning to the previous example, imagine that the RSM used to describe changes in T41 was generated from and valid for the range of LPC clearance values marked as Model Validity in Figure 4-2.

Any LPC clearance value outside of the model validity range was not considered when generating the T41 model and can therefore not be used to predict T41. By associating a standard deviation with a mean value in certain regions of the design space, the resulting distribution will have some of its values that lie outside of the valid RSM space. This space (shaded region) will result in T41 predictions being made by LPC clearances that were not considered when creating the T41 RSM. As much as 50% of the LPC clearance values can be outside of the valid T41 model range if the maximum or minimum mean is selected with the maximum variance.

There are some options available that still allow for a complete exploration of the design space without compromising accuracy. The first, and seemingly simplest method is to generate the response surface model over a much wider range than the design space that will be explored. The range of the RSM validity can be set such that any combination of means and standard deviations in the design space can be entirely contained within the range of the model. However, as mentioned earlier, the convergence criteria of the engine performance model is the limiting factor so the option of expanding the space over which the RSM is valid is not available. Another option, and the one employed in this work, is to assume that the RSM remains valid even outside of the range from which it was created but to confirm all solutions with the performance code itself. Understanding that the RSM does not necessarily go from being completely valid in the design space to being completely invalid immediately outside of the design space is key to this approach. It is true that the further outside of the valid design space a point is evaluated, the worse the prediction will likely be. There is however, a small range outside of the design space where the prediction may still be relatively accurate. As mentioned before, combinations of the lowest (or highest) mean value and the highest standard deviation will result in as much as 50% of the RSM predictions taking place outside of the valid space.

In order to employ the method that assumes the RSM is valid outside of its design space, a careful check of any potential solutions that result from points outside the RSM design space must be made. Consider an optimum solution that is found using an RSM

with some of the responses being predicted outside of the RSM range. By taking the solution of means and standard deviations found with the RSM and plugging those solutions into the engine model, it is possible to determine if the solution is in fact valid. If the RSM solution and the engine model solution are identical, the implication is that the RSM prediction was still accurate even though it was outside of the RSM range and that the solution is truly valid. In summary, the RSM is used to quickly, and as accurately as possible, find an optimum solution, and that solution is then run through the actual performance code to ensure that the RSM did an accurate job of modeling the engine as predicted by the performance model.

Table 4-3 shows the mean and standard deviation design space limits for each of the component geometries covered in this research. Again, the minimum and maximum means shown in Table 4-3 are the same as described in Table 4-1. The range of standard deviations was determined using a coefficient of variation. The minimum standard deviation considered in the design space was determined by taking the product of the minimum coefficient of variation column (Min Coef Vartn*)* and the minimum mean column (Min Mean). Likewise, the maximum standard deviation considered in the design space was determined by taking the product of the maximum coefficient of variation column (Max Coef Vartn) and the maximum mean column (Max Mean). The minimum and maximum coefficients of variation were selected as such to ensure that the design space was sufficiently large yet not so large that infeasible design space significantly dominated the search range. The results, presented in Chapter 5, support the fact that the design space is in fact large enough to contain an optimum solution.

Table 4-3: Component Design Space Limits

	Min Mean	Max Mean	Min Coef Vartn	Max Coef Vartn	Min Std	Max Std
Fan	0.02640	0.03360	0.025	0.05	0.000660	0.001680
LPC	0.01144	0.01456	0.025	0.05	0.000286	0.000728
HPC	0.02816	0.03584	0.025	0.05	0.000704	0.001792
HPT	0.02640	0.03360	0.025	0.05	0.000660	0.001680
LPT	0.01760	0.02240	0.025	0.05	0.000440	0.001120
LPT Area	14.56274	16.09566	0.01	0.02	0.145627	0.321913
HPT Area	3.21404	3.55236	0.01	0.02	0.032140	0.071047

Table 4-4 summarizes the constraints for each of the geometric and performance parameters that must be met by a combination of points from within the design space defined in Table 4-3. The lower limit of the tip clearances represents the minimum cold build clearance required to prevent scraping of the casing under thermal, centrifugal, and maneuvering loads. Of the performance parameters, the only two with lower limits are the surge margin of the Fan and the HPC while the remaining performance parameters have upper limits that cannot be exceeded. There are no lower limits set on SFC, T41, NL, or NH due to the fact that a lower value here represents a better engine, and including a lower limit could potentially cause a very well built engine to be declared infeasible. For the same reason, there is no upper limit on any of the other constraints.

Table 4-4: Design Space Constraints

Constraint	*Lower Limit*	*Nominal*	*Upper Limit*
Fan Tip Clearance (in)	**0.020**	0.030	
LPC Tip Clearance (in)	**0.010**	0.013	
HPC Tip Clearance (in)	**0.014**	0.032	
HPT Tip Clearance (in)	**0.016**	0.030	
LPT Tip Clearance (in)	**0.009**	0.020	
LPT Nozzle Area (in^2)	**NA**	15.3292	
HPT Nozzle Area (in^2)	**NA**	3.3832	
HPC Surge Margin (%)	**20.000**	20.840	
Fan Surge Margin (%)	**12.500**	13.0047	
SFC (lbm/hr/lbf)		0.3915	**0.3950**
T41 (Deg F)		2052.60	**2090.0**
NL (RPM)		14630.94	**14665.0**
NH (RPM)		44519.58	**44595.0**

It is important to note here that the selection of the constraints is equally as important as the definition of the design space limits. For constraints that compete, such as turbine nozzle area and surge margin, it is possible to specify limits that are impossible to satisfy. As turbine nozzle area is decreased, the performance parameters are improved with the exception of surge margin. It is possible to specify a surge margin so high that to meet it means that it will be impossible to meet some of the other performance parameters. There must be, at the very least, a small union of design space where competing constraints can coexist. Without this union, the optimization algorithm will not find feasible solutions, since there are not any.

4.3 Simulation Algorithms

Before proceeding into the details of the simulation algorithms, it is desirable to review the process we have completed thus far that will allow us to perform the simulations. The first step was simply being able to accurately model the engine performance at a single nominal point value with an engine cycle deck. The next step was to build models that describe the changes in performance as functions of the changes in geometry (tip clearances, turbine nozzle areas). The third step was to create an RSM over a suitable range that models the output of the performance code when it is coupled with the clearance loss models. The last step up to this point was to establish the limits of a design space that covers the range of means and standard deviations that will be considered in the simulation. The point of the simulation algorithm is to run different combinations of the means and standard deviations in the design space in order to determine the set of combinations that will yield the highest standard deviations while maintaining a three-sigma pass rate of all the specified performance and geometric constraints. The highest standard deviations will yield the maximum tolerances that can be set on a specific geometry and will represent the least expensive build configuration.

There were two different variation simulation algorithms used to explore the design space. The first method was purely stochastic (randomly generated) and the second method was an optimization routine consisting of a combination of exploratory and gradient-based algorithms.

4.3.1 Stochastic Variation Simulation Model

The stochastic variation simulation model was built on the idea that given enough random combinations of means and standard deviations, there will exist a combination that yields the best possible result. The stochastic variation simulation model was built and executed in Splus™ 4.5. Figure 4-3 illustrates the concept of the stochastic variation simulation model. The outer loop is used to randomly select (under a uniform distribution) combinations of means and standard deviations from the previously defined

design space of each component. Means and standard deviations for each of the seven components are selected simultaneously and then fed to the inner loop. Each component will have an independently generated distribution from which the clearances for that particular component will be sampled. The inner loop simulates building *N* number of engines under the Normal distributions established by the outer loop. For example, the outer loop will randomly select a mean and standard deviation that defines the cold build tip clearance of the fan (τ_{fan}). The mean and standard deviation of τ_{fan} define the distribution from which the probability of any specific geometry will occur. The inner loop then simulates building *N* engines with a τ_{fan} defined by the distribution that governs the geometry of τ_{fan}. This process is done simultaneously for all seven components to ensure that the interactions between the variances of all the components are also modeled.

The method of selecting combinations of means and standard deviations to define distributions by which the clearances for each component will be determined is not limited to Normal distributions. Because the algorithm does not rely on the assumption of normality, other distributions can also be explored by the same means. Parameters describing other distributions can easily be included in the simulation, which provides the capability to define virtually any type of distribution.

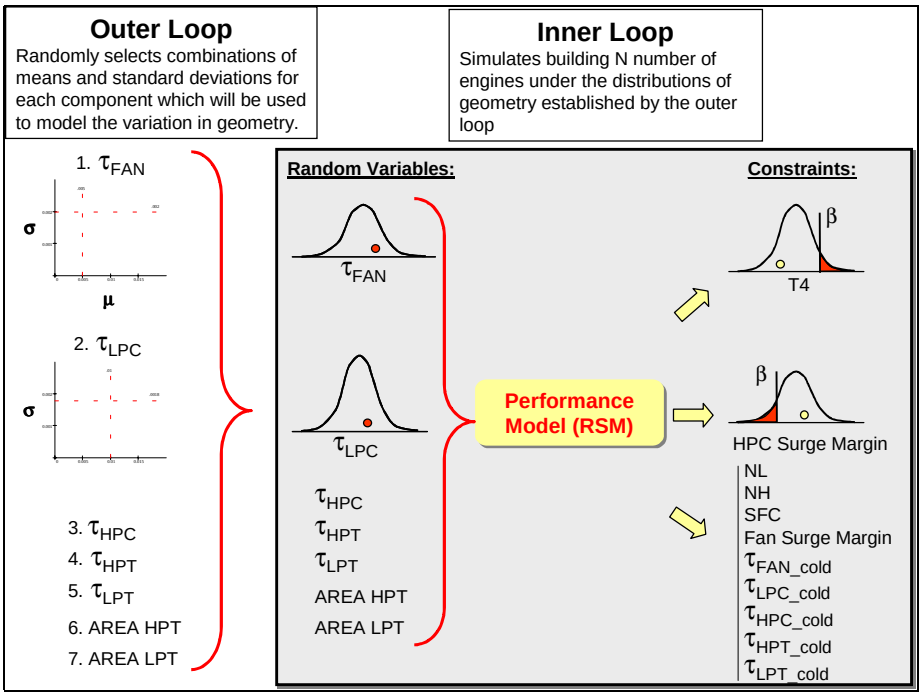

Figure 4-3: Stochastic Simulation Algorithm

The inner loop takes each of the *N* individual tip clearances and nozzle areas from their respective distributions and feeds them into the performance model. The performance model is replaced with the RSM in the inner loop to enable very fast calculation of the performance parameters. The speed of the RSM is essential at this point due to the sheer number of calculations that will need to be made. Because the simulation involves fourteen-dimensional space, a large sample size will be required to account for all the possible interactions and ensure good coverage of the entire design space. Imagine that a respectable sample of 1 million different combinations of means and standard deviations is chosen for the outer loop. Further, assume building $N = 500$ engines under each combination of means and standard deviations. Knowing that the average convergence time of the actual performance code is 2 seconds, the simulation would not be complete for well over 11 years on a single processor PC. By replacing the actual performance code with the RSM, the simulation will finish in less than a day.

When the inner loop has completed building *N* engines, the distributions of the outputs (constraints) are examined. Quantiles of each of the output distributions are taken at three-sigma and compared against the minimum or maximum allowable value

for that parameter. For example, the three-sigma quantile of the output parameter T41 is taken and compared against the value that it is not allowed to exceed (2090 °F). If more than 0.9974 – 1 = 0.26% of the T41 values exceed 2090 °F then the combinations of means and standard deviations that were generated by the outer loop for that particular iteration are dismissed as invalid. All of the constraints for each of the output parameters are systematically compared to their respective limits and only if all the constraints pass, is the solution returned as a valid possibility. At the end of 1 million iterations of the outer loop, only a fraction of the attempted combinations will have satisfied all eleven constraints. By examining the combinations that did satisfy all the criteria, it is possible to determine which satisfactory combinations had the highest standard deviations. Using this method, the actual optimization is done more as a post-processing step given a series of satisfactory results.

4.3.2 Variation Simulation Model Using an Optimization Plan

Performing a variation simulation with an optimization plan is similar to the stochastic method with one important difference. The selection of the next combinations of means and standard deviations from the outer loop is not purely random. The selection of the next combination is based on how well the previous iteration satisfied all of the constraints. The optimization plan variation simulation was built and executed with Splus™ 4.5 and iSight™ 5.5 coupled. Figure 4-4 illustrates the concept of the variation simulation that includes an optimization plan.

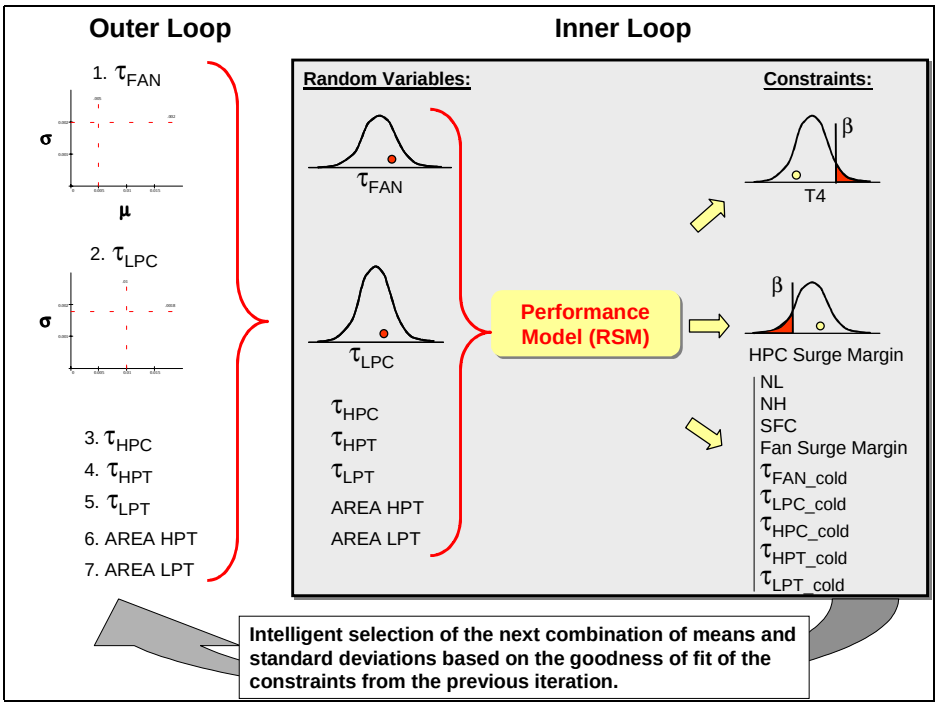

Figure 4-4: Simulation Algorithm Using an Optimization Plan

As mentioned in Chapter 2, there was more than one algorithm included in the optimization plan used to perform this work. The first was a genetic algorithm with a population size of 60 and a maximum trial limit of 2000. Again, the genetic algorithm served as the tool that explored the entire design space and found the regions of the design space that satisfied all of the constraints. The genetic algorithm was followed by NLPQL with a maximum iteration limit of 1000. The purpose of the NLPQL algorithm was to further explore the regions of feasible design space that were discovered by the genetic algorithm.

Although the objective function of the optimization plan is to maximize the standard deviations of multiple component geometries, the objective function in this case is not truly multi-disciplinary. Because all of the objective functions are of the same type, one objective function can be made by summing all of the standard deviations together and setting the objective function to maximize this sum. To sum the standard deviations in a manner that will allow this single objective function to be accurate, each standard deviation would have to be normalized over the range of its design space and

equally weighted with all other standard deviations. Although the optimization routines used in this work maintained each of the standard deviations as independent objective functions, the single objective function method may be simpler to execute with some optimization algorithms. The advantage of maintaining each of the component standard deviations as independent functions lies in the simplicity of changing the weighting functions. For example, if it was determined that maximizing the standard deviation of the HPT was twice as important as maximizing the standard deviation of the fan, changing the weighting function of the HPT to 2.0 is simpler and quicker if the functions are maintained independent rather than as a lump sum.

In addition to weighting functions that can be applied to optimization algorithms, it is also important to consider normalizing. Normalization accounts for the inherent difference in the magnitudes of some of the design variables. Assume that the minimum standard deviation of the LPC clearance that will be explored by the optimization algorithms is 0.00029" and the maximum standard deviation that will be explored for the LPT Nozzle Area is 0.32191". Since these design variables are on the order of 1000 times different, normalizing factors must be included to bring the two design variables within the same order of magnitude. It is important that all of the design variables in the same optimization plan be of the same magnitude because without it, one of the design variables could have a falsely significant impact on the overall objective function.

4.3.3 Comparison of Simulation Methods

There are advantages to both methods mentioned above but throughout this research, it has become clear that the method that employs the use of an optimization plan is far superior. The advantages of the purely stochastic method described in section 4.3.1 are that it is not dependent upon programming any optimization algorithm and therefore does not suffer from any of the algorithm's inherent weaknesses. Essentially, the stochastic method is simpler to set up and execute and requires no knowledge of optimization algorithms or additional software.

The variation simulation that uses an optimization plan is superior to the stochastic method because it makes intelligent decisions on where in the design space the next point should be selected. Using the purely stochastic method, 1 million combinations of means and standard deviations were attempted and only 105 unique solutions were found to be satisfactory. This implies that 99.9895% of the computational time was spent exploring regions of the design space that were infeasible. This fraction of feasible solutions also provides a sensible estimate for the percentage of the design space that can be expected to contain satisfactory solutions and how difficult it can be to find the space where all the constraints can co-exist. Although returning essentially the same results, the stochastic method took 3 days of computer time (Pentium III, 600 Mhz single processor) while the optimization plan method took only 8 hours on an even slower machine (Pentium II, 400 Mhz single processor).

The intelligent selection of the next set of design points based on the results of the previous iteration allows the optimization plan method to spend much less time exploring regions that do not satisfy all of the constraints. However, because part of the optimization plan contains a gradient-based algorithm (NLPQL), care must be taken to simulate a sufficiently large number of engines in order to ensure repeatability. The simulation of only 500 engines per iteration of the inner loop was done for the stochastic method (enough to obtain a good statistical average). In comparison, 50,000 engines were built per iteration of the inner loop for the optimization method. Although 50,000 far exceeds the amount required to get a good statistical average, repeatability is critical to a gradient-based optimization algorithm. As the number of iterations of the inner loop approaches infinity, the quantiles of the resulting distributions become constant. Since a gradient-based optimization algorithm will often hold all but one of the inputs constant in an attempt to quantify the change that a single input will have, it is important the algorithm be able to repeat this process consistently. If the number of iterations of the inner loop is not extremely high, the outputs of the RSM will not be consistent even though the inputs are identical. This is due to the inherent nature from which the inputs are randomly generated. The requirement for a large number of iterations of the inner

loop is not detrimental overall due to the fact that the number of outer loop iterations can be significantly reduced over that of the purely stochastic method.

It is important to note once again that the algorithms described above are not limited to Normal input distributions. Other distributions can easily be modeled with only small changes to the code. Moreover, the output distributions may not necessarily be of the same type as the input distributions. It is entirely possible that a Weibull input distribution could be skewed by the RSMs so that the output distribution is Normal instead of Weibull. This transformation of distributions was not observed in this research but if it were, the accuracy of the algorithm would be unaffected. The algorithm would be accurate in a case where the distribution is transformed because of the way the constraints are handled. When the algorithm takes the quantiles of the output parameters, a specific percentage of the outputs must pass regardless of the distribution governing them. In other words, a three-sigma pass rate means 9,974 out of 10,000 engines must be acceptable regardless of how the 10,000 engines are distributed.

4.4 Cost Modeling

Up to this point, all of the analysis has been based upon a three-sigma pass rate (level of reliability). By fixing the pass rate, all the unknown design variables can be numerically iterated until an optimum solution is found. However, there is a justifiable interest in wanting to explore other levels of reliability and their associated cost. The ultimate goal of performing cost modeling is to find the region of the design space that results in the least cost. As mentioned in Chapter 2, accurate cost modeling is quite involved and can introduce a significant amount of error. For this reason, cost modeling has not been the focus of this work but was performed in order to show qualitative results.

The cost modeling for this research was done using a simple cost-tolerance model of the form:

$$Cost_{part} = A + \frac{B}{T^2} \tag{4-2}$$

where A represents a fixed cost and B represents the cost of manufacturing a component to the tolerance T. This type of cost model was also explored by Ostwald [10] in an effort to illustrate the tradeoff between cost and tolerance. The two competing objective functions in this type of modeling are cost and reliability. As Figure 4-5 shows, an increased reliability (tight tolerance) results in a much higher cost of production but lower cost due to failure. As the tolerance is tightened, the amount of scrap is reduced due to the fact that more of the engines are passing all of the specified criteria. The qualitative illustration of competing objectives in Figure 4-5 lends itself well to multi-disciplinary optimization.

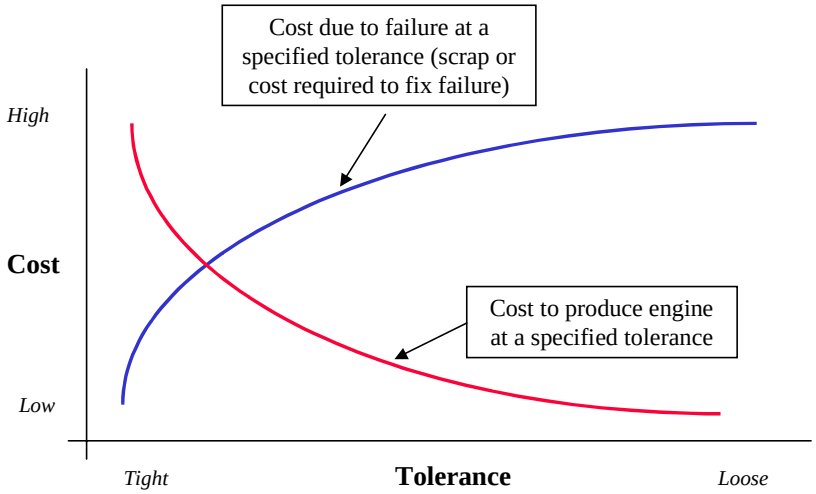

Figure 4-5: Cost-Tolerance Tradeoff

Although not including all of the associated costs, Equation 4-2 can be used to get a good estimate of the trend describing the cost of production as a function of the tolerance. However, a function must also be established that describes the cost of failure as a function of the tolerance. This function can be rather subjective due to the fact that it must account for several factors including the cost of having to scrap or fix parts that fail some criteria. A realistic failure function should also be able to account for less quantifiable factors such as customer satisfaction and timely delivery. Without having

detailed and specific knowledge of a product line, cost modeling of a turbine engine should be approached with caution.

Another factor that should be considered when modeling costs is the fact that the cost tolerance function is most likely discontinuous. In reality, a decrease in tolerance will not follow a smooth curve as shown in Figure 4-5. This is caused in part by the fact that certain tolerances can only be achieved with specific types of machining processes. For example, when the tolerance limits of a milling machine are not enough to meet the specified criteria, a more expensive yet more precise process such as grinding must be used. This step change in machine processes carries with it a new cost description and thus a discrete cost function. Figure 4-6 illustrates the difference between continuous and discrete cost functions.

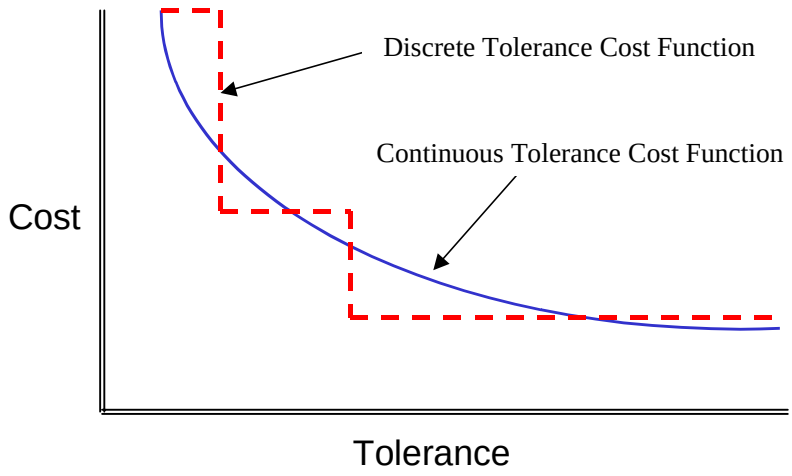

Figure 4-6: Discrete and Continuous Cost-Tolerance Functions

5 Results and Discussion

5.1 Variation Simulation Model Results

A number of methods exist to present the results of a variation simulation model that involves multiple dimensions. Because the results of this work exist in fourteen dimensions (two dimensions for each of the seven design variables), no single graphical method of presentation will entirely capture all of the results and their association within the design space. As a result, a number of graphical and tabular methods are presented that illustrate the results of the modeling. The results shown in this chapter were calculated on a 400 Mhz Pentium II single processor and required approximately 8 hours of CPU time.

Figures 5-1 through 5-7 are three-dimensional drop-line scatter plots that show the results of the optimization performed by iSight™ V5.5. Each of the seven plots shows the design space of an individual component and the points in the design space where the most feasible results were found. The x-axis of each of the plots shows the range of means for that particular component over which the optimization algorithms where allowed to explore. Likewise, the y-axis is the range of standard deviations for that component to which the algorithm was limited. The z-axis indicates the level of feasibility that was determined for a specific iteration of design variables. The numbers contained in the boxes found on each of the drop-line scatter plots gives the optimum mean and tolerance as determined from the single most feasible point in the respective design space.

The quantifiable factor, feasibility, is a measure made by iSight™ as to how well the given combinations of design variables satisfy the constraints. A feasibility of 1.0 indicates that the selected combination of means and standard deviations did not pass one or more of the constraints. A feasibility of 7.0 or greater indicates that the selected combination of means and standard deviations did in fact pass all of the criteria. A feasibility of 8.0 or 9.0 (maximum possible feasibility) is a value assigned to

combinations of those design variables that passed all of the constraints by a greater margin than any feasibility of 7.0. Thus, a feasibility of 9.0 or any subsequent 8.0 is the best fit of all the constraints.

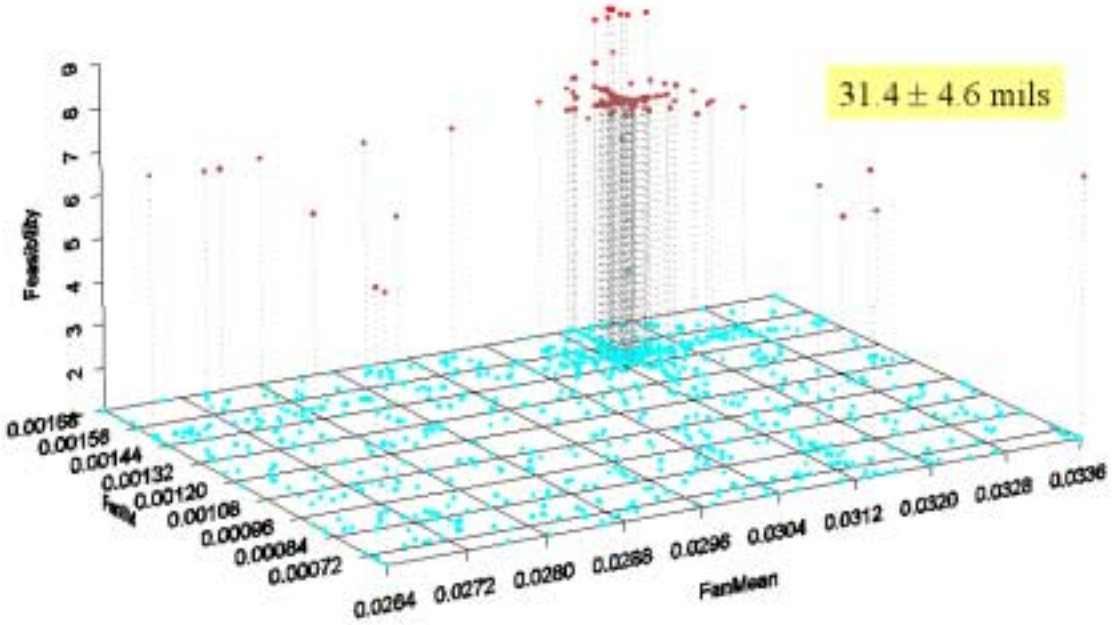

Figure 5-1: Fan Tip Clearance Design Space

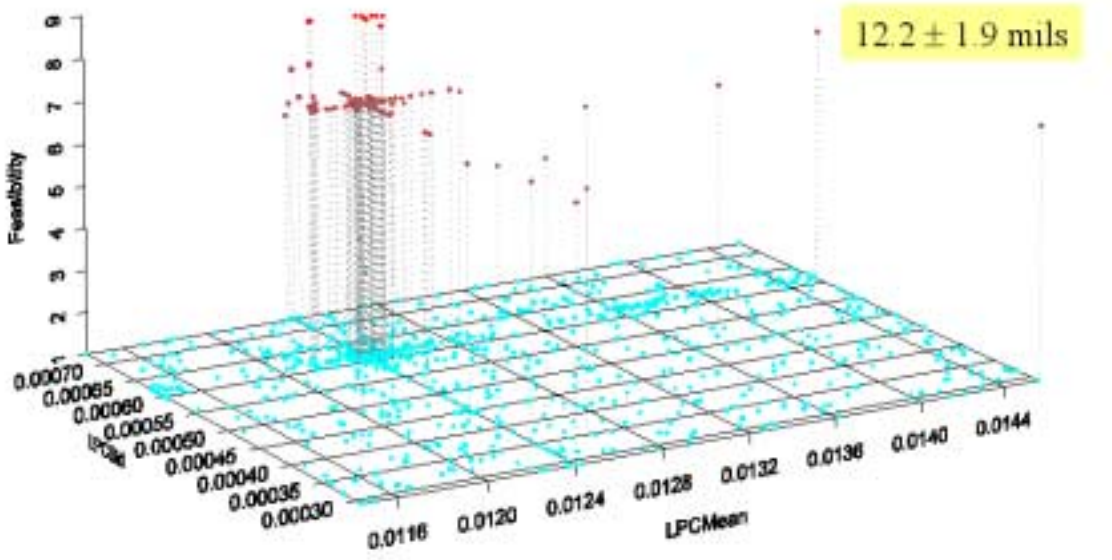

Figure 5-2: LPC Tip Clearance Design Space

The plots show several interesting phenomena. First, all of the plots show a distinct grouping of points. Additionally, among the clusters, each result shows a tightly clustered series of 9.0 feasibilities nested in and centered above a group of 7.0 feasibilities. This result is encouraging because it appears to indicate the presence of a single most highly feasible solution surrounded by a small region of acceptable, yet less feasible solutions.

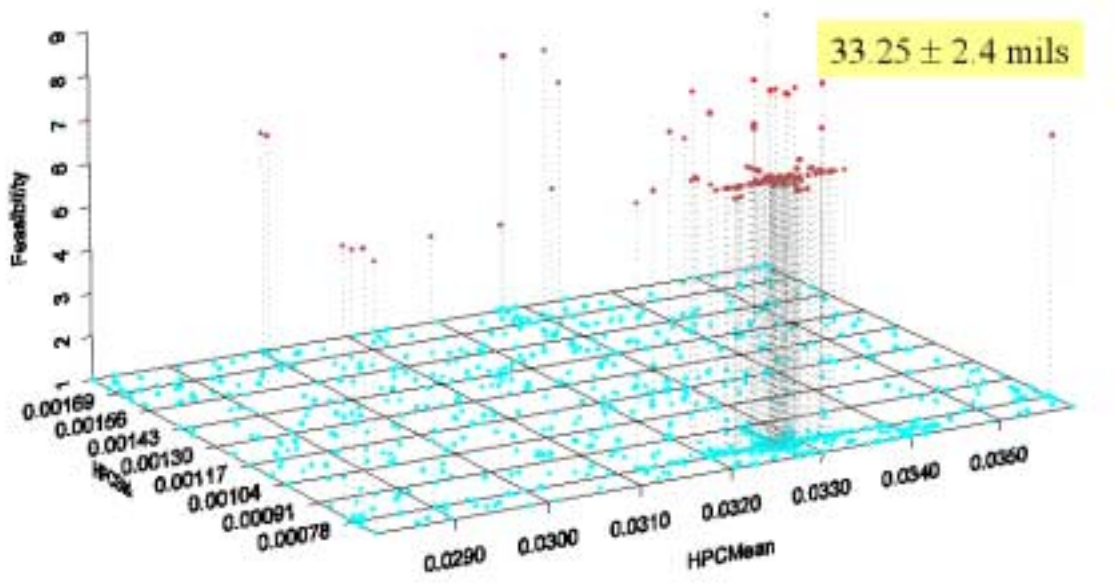

Figure 5-3: HPC Tip Clearance Design Space

Another commonality seen in all of the graphical results is that there are points throughout the design spaces that show acceptable solutions (Feasibility = 7.0) far outside of the tight cluster of points. These dispersed solutions occur because there exists feasible design space outside of the clustered areas that still satisfies all of the criteria. It is important to note that the solutions presented here are not the only solutions that satisfy all of the constraints. Because the objective function was to determine the maximum standard deviations with equal weights on all components, there can exist solutions that have a wide range of individual component tolerances, yet the normalized sum of the standard deviations is virtually the same. After comparing several different simulations

starting at different points in the design space, it became apparent that the dispersed solutions illustrate tradeoffs between certain components. By reducing the tolerance of one component, another component's tolerance can be increased. This tradeoff is expected, and examples can be seen in the maintenance of turbine engines.

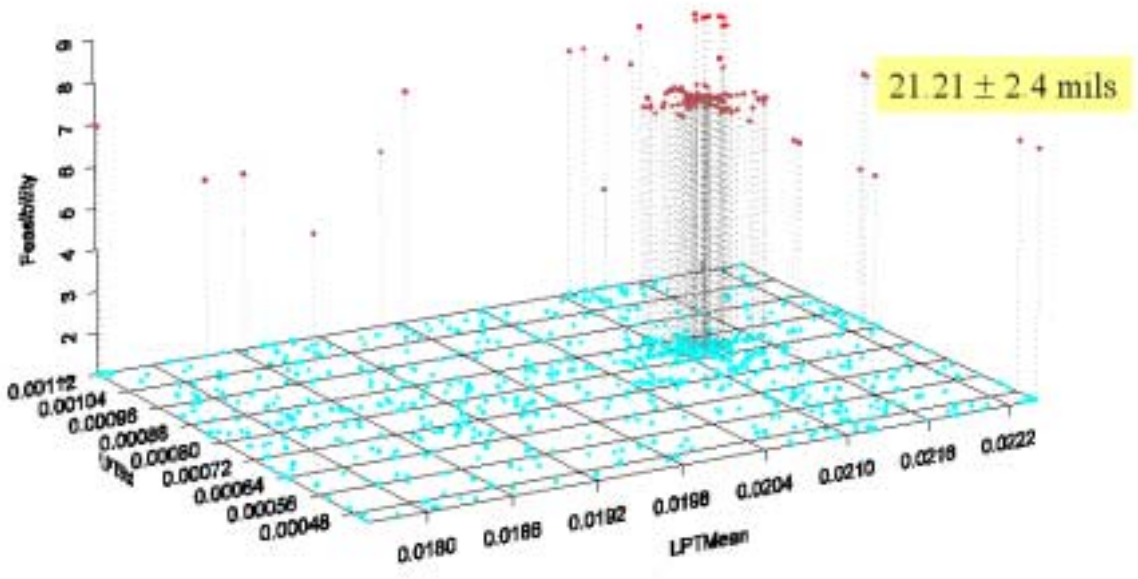

Figure 5-4: LPT Tip Clearance Design Space

It is common for an engine maintenance organization to extend the performance life of an engine by replacing a single component. An engine that is failing performance can sometimes be brought back to acceptable limits simply be replacing the HPC and not replacing any other components. The new HPC can help to overcome the performance losses resulting from a turbine that may have had a small tip clearance increase and is experiencing a loss in efficiency.

Some of the characteristics of the optimization routines can also be seen in the graphical results. The genetic algorithm appears to have done a good job of covering the entirety of the design space searching for regions of high feasibility. This coverage is evident by the abundant generation of points that covers the bottom of the x-y plane (feasibility = 1.0) in each of the component design spaces. The only consistent solutions

were found near the singular cluster of points and this became the point where the algorithm began to focus most of its effort. After finding the space in which high levels of feasibly were consistently encountered, the genetic algorithm handed the task of finding the absolute best fit over to the NLPQL algorithm. The NLPQL algorithm, a gradient-based method, began to change each parameter incrementally to determine which combination best fit the constraints. Evidence of the incremental changes performed by the NLPQL algorithm can be seen in the plots. The cross or "+" pattern seen in each of the clustered areas of Figures 5-1 through 5-5 is the NLPQL algorithm incrementally stepping in a single dimension of the represented design space in order to determine if the next step returns better results than the previous step.

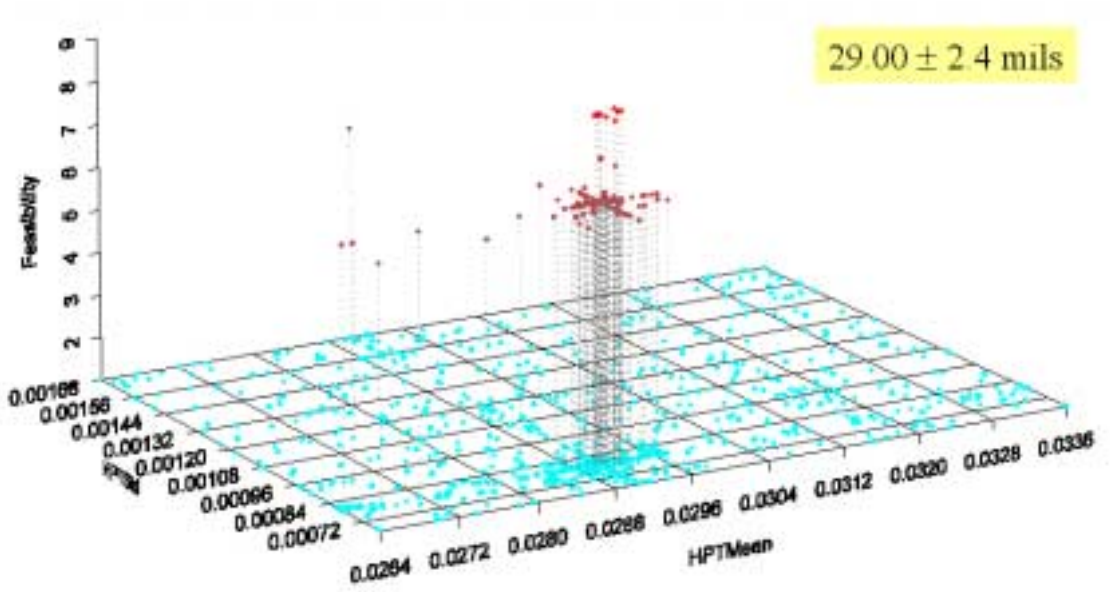

Figure 5-5: HPT Tip Clearance Design Space

Two plots show trends different from any of the others. These two unique plots are for the HPT and LPT turbine nozzle area, Figures 5-6 and 5-7 respectively. Given the fact the sensitivity analysis indicated that these two components were the most influential, it is not surprising that they show unique trends.

The uniqueness of Figure 5-6 arises because it does not have an equal number of feasible solutions on all sides of the most feasible solution. Unlike the other plots,

Figures 5-6 shows an “L” shaped pattern as opposed to the “+” shaped pattern. This distinct pattern indicates the presence of a constraint that is consistently failed because of exceeding a sharply definable limit in the design space. Notice also that there are no solutions found outside of these hard limits that ever satisfy all of the constraints. This result would appear to indicate that there is a limited availability for tradeoff when trying to further increase the mean or standard deviation of the HPT nozzle area. The inference is that by tightening the tolerances on all the other components, the variability of the HPT nozzle area can not be increased beyond the very specific limits seen in Figure 5-6. However, since there are an infinite number of combinations of design variables, it is not possible to prove with this method that no tradeoff exists for the turbine nozzle areas.

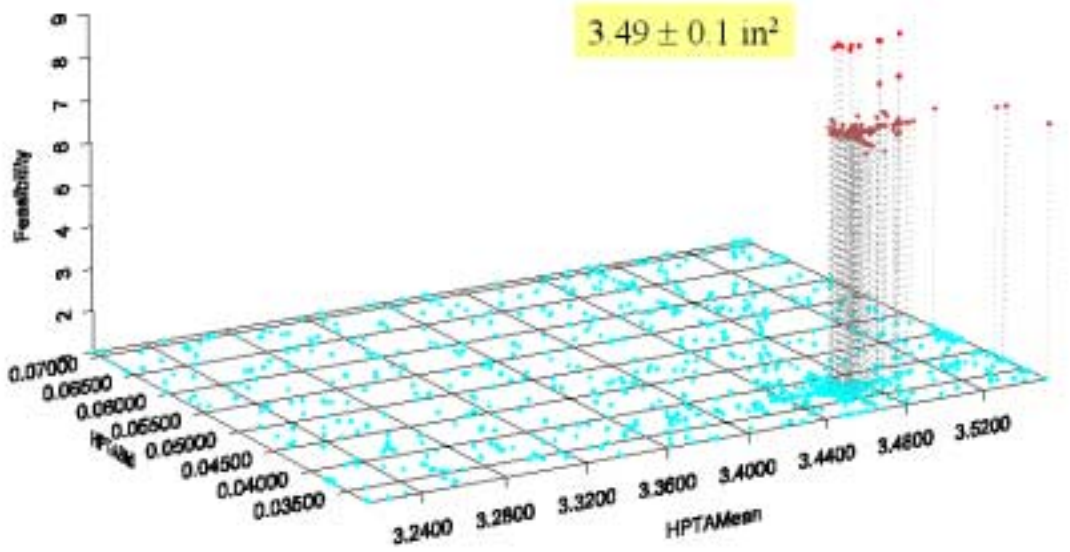

Figure 5-6: HPT Nozzle Area Design Space

Figure 5-7 shows the design space of the LPT nozzle area and is similar to the HPT nozzle area, but appears to be more forgiving with respect to standard deviations. As Figure 5-7 shows, there is a range of standard deviations that can satisfy all the constraints. However, there is a hard limit to the mean value that can apparently never be exceeded at the three-sigma level. For the LPT nozzle area, there are no solutions found past the hard limit indicating that exceeding the mean value of 14.94 in^2 consistently violates a constraint. It should be noted that this result does not imply that any single

value for a turbine nozzle area greater than 14.94 in^2 will cause an engine to fail performance. It only implies that the *mean* value of the distribution of turbine nozzle areas cannot exceed 14.94 in^2. In fact, with a mean value of 14.94 in^2 and a Normal distribution, approximately 50% of the turbine nozzle areas will exceed 14.94 in^2.

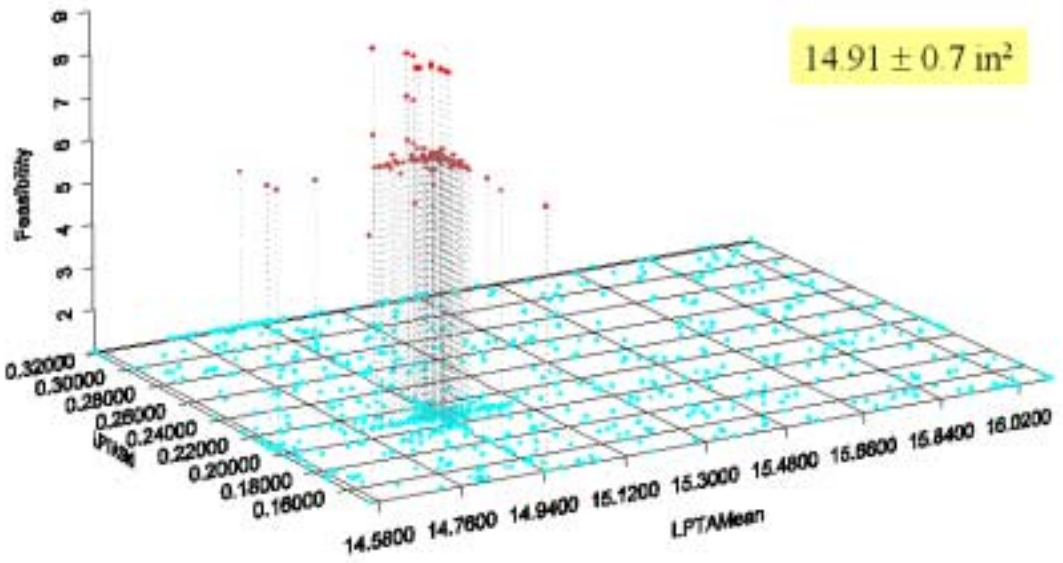

Figure 5-7: LPT Nozzle Area Design Space

In addition, all of the plots show solutions entirely contained in the design space, with no feasible solutions on the edge of the space. This result is important to observe since it indicates the design space is sufficiently large for each component. By not having a cluster of solutions in a corner or on an edge of the design space, it is evident that the algorithm was allowed enough latitude to adequately explore the design space and that the most feasible solution was not passed over because it exceeded the limits of the design space.

Before proceeding with the results, it is important to discuss how the tolerance of a part is determined from its standard deviation. As is common in probabilistic design, the "3σ rule" was used for this work. To illustrate the concept of the "3σ rule," an excerpt from Dai & Wang [1] is given:

"For tolerances, the standard deviations are estimated by setting 3σ equal to tolerances on the 99.74 percent limits. For example, if a cylindrical part has a diameter of 50 ± 0.5 mm, this means the diameter mean is 50 mm, and the standard deviation is $\sigma = 0.5 / 3 = 0.17$ mm. The number of cylindrical parts with diameter outside the limits [49.83, 50.17] mm is less than 27 in 10,000 parts produced."

While the 3σ rule is indeed an estimate, it is a simple and accurate tool that allows us to determine the tolerance from the standard deviation of a Normal distribution.

Table 5-1 lists the original mean value of all the components and the associated mean and tolerance that resulted from the optimization routine. The optimized mean value and optimized max tolerance shown in Table 5-1 were taken from the results shown in Figures 5-1 through 5-7. The last run of the optimization routine that returned a feasibility of 9.0 was taken from each of the design spaces and determined to be the best combination available to satisfy all of the constraints. At this point, it is important to reiterate that the objective function of the algorithm that produced the results seen in Table 5-1 was to maximize the standard deviation (tolerance) only, but have the ability to manipulate the mean value in order to do so. Furthermore, the weighting of the objective function was set to 1.0 for each component. In other words, maximizing the tolerance of one component was equally as important as any other component.

Table 5-1: Original Mean and Optimized Values

	Original Mean Value	Optimized Mean Value	Optimized Max Tolerance
Fan Tip Clearance (mils)	30.00	31.37	4.65
LPC Tip Clearance (mils)	13.00	12.20	1.88
HPC Tip Clearance (mils)	32.00	33.25	2.35
HPT Tip Clearance (mils)	30.00	29.00	2.35
LPT Tip Clearance (mils)	20.00	21.21	2.36
LPT Nozzle Area (in^2)	15.33	14.91	0.72
HPT Nozzle Area (in^2)	3.38	3.49	0.11

Figure 5-8 further quantifies the results of the variation simulation by showing the relative distance from the mean that is encompassed by the tolerance range. The optimized HPT Area mean was determined to be 3.49 in^2 with a tolerance of 0.11 in^2 and therefore has a tolerance variation of 0.11/3.49 = 3.15% around the nominal. A smaller percent variation from the nominal represents a tighter relative tolerance. Again, the trends seen in Figure 5-8 can be related to the component sensitivities. The tightest

relative tolerances belong to the two most sensitive components. This result confirms the results of the sensitivity analysis presented in Chapter 3 of this document.

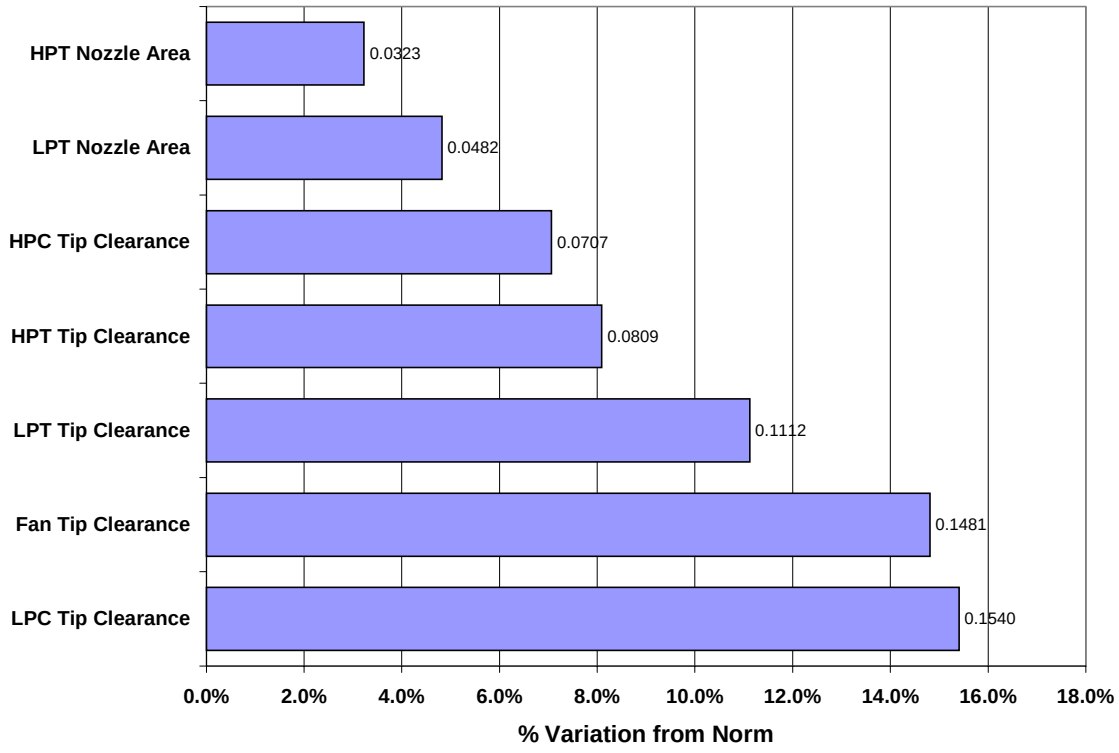

Figure 5-8: Percent Variation from Norm for each Component

5.1.1 Multidimensional Considerations

It is important to note that the optimization algorithm varies fourteen independent variables while any single three-dimensional plot shown in Figure 5-1 through 5-7 contains spatial information on only two dimensions of the design space. This lack of dimensional representation can be deceiving if viewed without regard to the complexity of multi-dimensional space. An example of this deception can be seen in Figure 5-1. This plot shows a cluster of infeasible design combinations directly under the cluster of feasible points. While this is seemingly contradictory, one must remember that only two of the fourteen dimensions are shown. As an example, imagine that a feasible solution is found that contains the fan clearance mean, *a*, and the fan clearance standard deviation, *b*. Now imagine that the next step of the optimization algorithm involves changing only the mean value of the HPT tip clearance and that this new combination fails to pass all of the

criteria. This new design point is not feasible, but since changes in the HPT tip clearance are not reflected on the fan design space plot, the new infeasible point now lies directly beneath the last feasible design point (*a, b*) in the fan design space. There are methods available, such as the computation of a Euclidian Distance between solutions, that aid in further illustrating these multi-dimensional concepts.

A normalized Euclidian Distance is useful in determining how close two points are in multi-dimensional space. Unlike the plots in Figures 5-1 through 5-7 that show information in only two of the fourteen dimensions, a normalized Euclidian Distance contains information on the spatial relation of all fourteen dimensions and provide a much more accurate view of the relationship between any two points. Equation 5-1 describes a normalized Euclidian Distance.

$$E = \sqrt{\sum_{i=1}^{n} \frac{\left(x_i - x_{i_o}\right)^2}{\left(\Delta_{range}\right)^2}} \tag{5-1}$$

where

n is the number of dimensions in the design space
$\mathbf{x_i}$ represents the value of any single dimension in the design space
$\mathbf{x_{io}}$ is the reference value for the same variable described by x_i
$\mathbf{\Delta_{range}}$ is the difference between the maximum and minimum values for x_i in the design space

It is important to normalize the Euclidian Distance over the range of each variable to ensure that any potential difference in the magnitudes of the variables does not skew the result. The result of a normalized Euclidian Distance is a dimensionless number that quantifies the distance between any two points in n-dimensional space.

Applying Equation 5-1 to the results of the optimization routine reveals some interesting results. Again, the most feasible solution (the results shown in Table 5-1) from the optimization routine was determined to be the last run in which a feasibility of 9.0 was attained, and it was this value that was used as the reference value. A normalized Euclidian Distance from this reference solutions was then calculated for each of the other points attempted by the optimization algorithm whether the point was feasible or not.

The goal of this type of analysis is to quantify, in fourteen dimensions, the distance between the optimum (reference) point and all the other points that were attempted by the algorithm. The expected outcome was that all of the results that had a feasibility of 9.0 would be closely clustered and be a good distance from the points where the feasibility was 1.0. This, however, was not found to be the case as shown in the plot of Euclidian Distance and Feasibility shown in Figure 5-9.

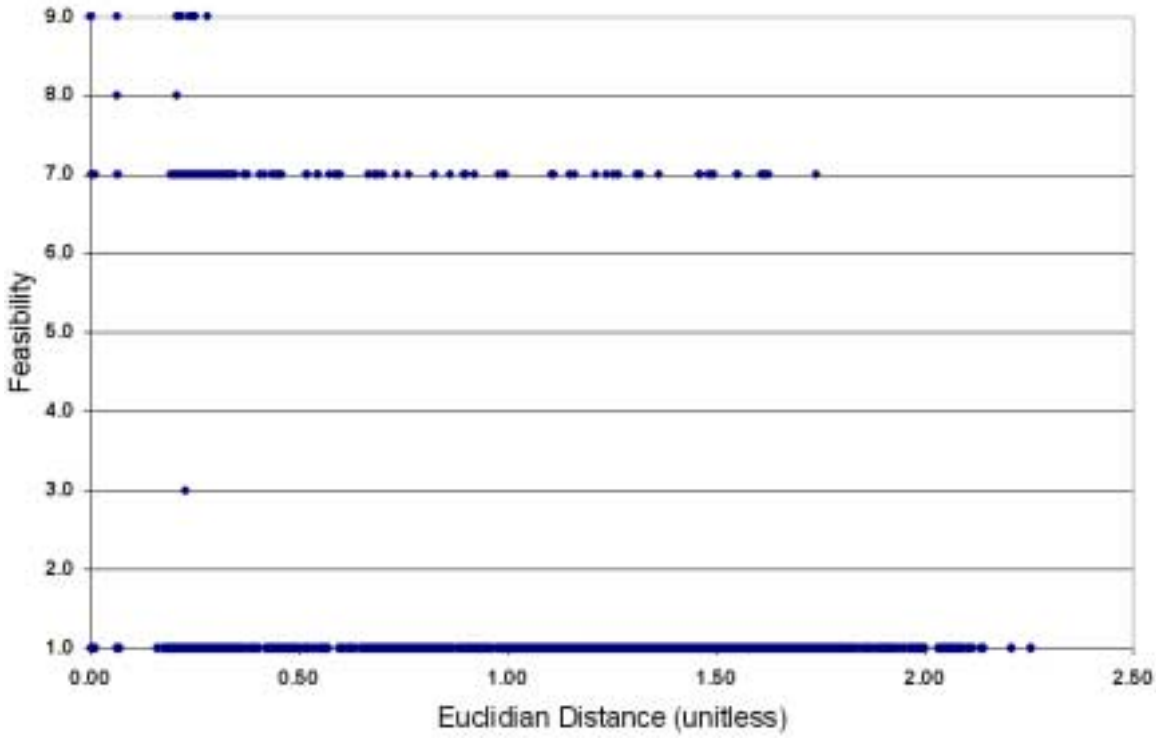

Figure 5-9: Plot of Normalized Euclidian Distance

Figure 5-9 shows there are solutions that are not feasible very close to solutions that are entirely feasible within the fourteen-dimensional space. Again, the smaller the Euclidian Distance, the closer that point is to the optimum solution. This figure shows solutions with a feasibility of 1.0 only a very short distance from solutions that have a feasibility of 9.0. While this result was not expected, further analysis reveals a potential explanation. It is the means and standard deviations of both turbine nozzle areas that are partly responsible for infeasible results being very close to the feasible results. As Figures 5.6 and 5.7 show, there appear to be sharply definable limits where the design space goes from being highly feasible to highly non-feasible over a very short distance.

Small changes in the means and standard deviations of the turbine nozzle areas will result in only small changes in Euclidian Distance, yet a large change in the performance of the engine. Once again, this result relates to the conclusion that of all the components included in this work, the turbine nozzle areas are the most influential to the overall engine performance.

Another point to consider when examining Figure 5-9 is that the information contained in the Euclidian Distance does not account for the direction taken in any dimension, only the absolute distance from the reference point itself. As an example, the maximum standard deviation for the HPT Nozzle Area was found to be 0.0375 in^2, and it is this value that was set as the reference value for the Euclidian Distance calculations. During the optimization, an HPT Nozzle Area standard deviation was tried at 0.04 in^2 and another point at 0.035 in^2. Both points are equidistant from the reference value of 0.0375 in^2 yet the 0.035 in^2 standard deviation yields a feasible result while the 0.04 in^2 fails some or all of the constraints. This limitation of directional information in the design space is also partly responsible for the unusual results seen in Figure 5-9.

In order to confirm that the results presented above were not simply a function of the dimensions of the design space or the initial starting position of the optimization algorithms, several simulations were conducted. Since a genetic algorithm does have some stochastic characteristics, other simulations were performed that had different initial starting positions to ensure that the solutions presented above were not by happenstance. Furthermore, selected regions of the design space were reduced to show that the original solutions were not a function of the unique geometry of the initial design space. The design space was altered in such a manner that spaces of highly infeasible results were eliminated, but not the area in which the original, most highly feasible solutions were found. The results from the optimization routines that had different initial starting positions and altered design spaces were similar to the results presented above showing that the initial starting position and the initial values were not factors that determined the outcome of the results. There were, however, other equally feasible

results discovered at the same time. Again, the additional solutions simply illustrated the potential tradeoffs between certain component tolerances.

5.1.2 Confirmation of Results with a Performance Model

Because all of the results discussed thus far were generated using the response surface models, it is important to run the results through the actual performance code to ensure that the output is the same as predicted by the RSM. To do this confirmation, the engine model was run 500 times with the distributions of the tip clearances and turbine nozzle areas defined by the results found with the optimization routine. By confirming the results with the performance code rather than the RSM, any modeling error or boundary limitation inherent to the RSM can be removed and the results presented with even greater confidence. Again, due to the high R^2 values of all the RSMs, very little difference between the prediction made by the RSM and the performance code is expected.

After using the engine model to simulating the manufacturing and assembly of 500 engines under the optimized distribution, Figure 5-10 was generated. Figure 5-10 shows the distribution of each performance parameter included as a constraint. Accompanying Figure 5-10 is Table 5-2 that lists the constraint limits, the three-sigma quantiles of the performance parameter distributions, and the margin by which performance parameters passed their respective constraints. The bold black line in Table 5-2 separates the performance constraints and the geometric constraints.

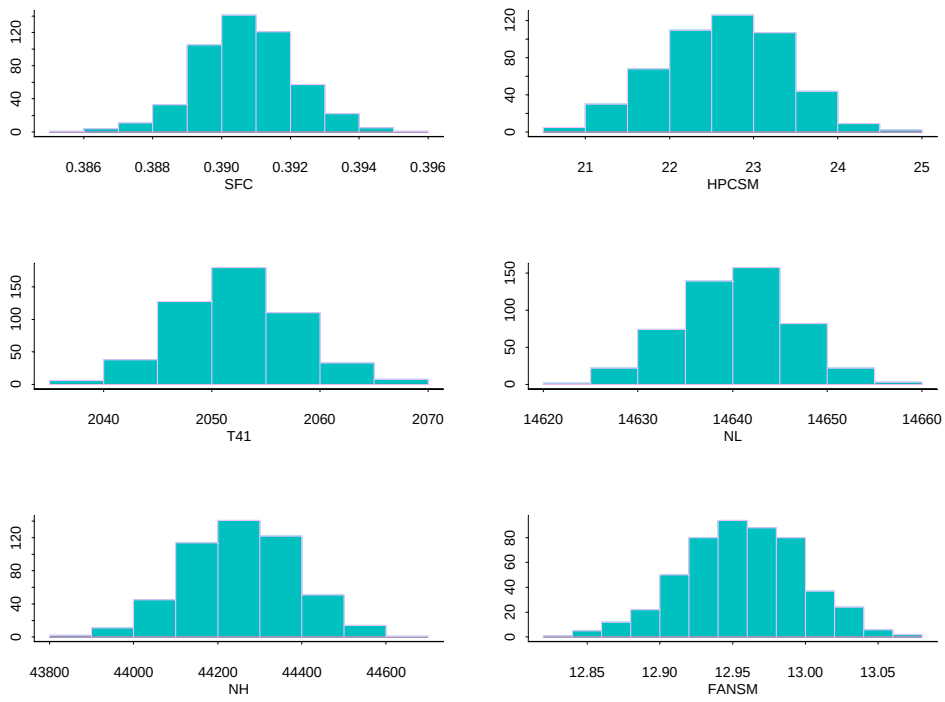

Figure 5-10: Performance Parameter Histograms

Table 5-2: Constraints and Optimized Results

Constraint	*Lower Limit*	*Nominal*	*Upper Limit*	*3 Sigma Quantile*	*Pass Margin*
Fan Tip Clearance (in)	**0.020**	0.030		**0.0357**	**1.7843**
LPC Tip Clearance (in)	**0.010**	0.013		**0.0138**	**1.3829**
HPC Tip Clearance (in)	**0.014**	0.032		**0.0351**	**2.5058**
HPT Tip Clearance (in)	**0.016**	0.030		**0.0311**	**1.9426**
LPT Tip Clearance (in)	**0.009**	0.020		**0.0236**	**2.6170**
LPT Nozzle Area (in^2)	**NA**	15.3292		**15.5797**	**NA**
HPT Nozzle Area (in^2)	**NA**	3.3832		**3.5909**	**NA**
HPC Surge Margin (%)	**20.000**	20.840		**20.7342**	**1.0367**
Fan Surge Margin (%)	**12.500**	13.0047		**12.8415**	**1.0273**
SFC (lbm/hr/lbf)		0.3915	**0.3950**	**0.3946**	**1.0011**
T41 (Deg F)		2052.60	**2090.0**	**2068.4**	**1.0104**
NL (RPM)		14630.94	**14665.0**	**14655.7**	**1.0006**
NH (RPM)		44519.58	**44595.0**	**44578.6**	**1.0004**

The results in Table 5-2 are virtually identical to the results predicted by the RSM and confirm that the RSMs were accurate over the range in which they were used. The three-sigma quantile represents the limits of the distribution at which 99.74% of the data points are contained within. Since 500 engines were simulated in the generation of Figure 5-10 and Table 5-2, a failure of one engine still constitutes a three-sigma pass rate

since 499 / 500 = 99.8%. However, if two or more engines out of 500 fail, then the results would be unacceptable at a three-sigma pass rate.

The pass margin, which is the last column shown in Table 5-2, is a factor quantifying the magnitude by which the three-sigma quantile passed the respective constraint. For example, the three-sigma constraint for NH was 44595 RPM and the resulting distribution had 99.74% of the engines at or below an NH of 44578.6 RPM. Since 44595 / 44578.6 = 1.0004, the NH constraint was passed by a margin of only 1.0004. As Table 5-2 shows, all of the pass margins are above 1.0 which indicates that all of the constraints were in fact met or exceeded. Furthermore, it is evident that the constraints limiting the standard deviations from being increased further were the performance parameters and not the geometric parameters. All of the geometric parameters had pass margins above 1.3 while the performance parameters had pass margins no greater than 1.0367 (HPC Surge Margin). The closer the pass margin is to 1.0, the closer that particular constraint came to failing its criteria. This being the case, NL and NH appeared to be the performance parameters responsible for limiting a further increase to any of the component standard deviations.

Since the optimization algorithm was allowed to freely manipulate the mean values within the design space in an attempt to maximize the standard deviations, it is of interest to compare the resulting mean values with the original mean values used in the performance code of the engine. Table 5-2 includes a column titled Nominal that lists the original nominal performance values that result from running the engine model with the nominal geometric values found in the actual engine. The dark vertical line associated with each component in Figure 5-11 shows the tolerance band for that component over a range that has been normalized by the nominal value of that component. The horizontal line at y = 1.00 indicates the relative position of the original mean value of that component. The height of the black or white bar shows the relative position of the optimum mean result presented in the results of Table 5-1. A white bar indicates that the optimum mean was found to be higher than the original mean and a black bar indicates that the optimum mean was found to be lower than the original mean. By normalizing

each of the tolerance ranges by its respective nominal value, all of the components can be clearly presented in one plot. From Figure 5-11 it is evident that again the LPT and HPT nozzle areas must have been issues of concern during the assembly and manufacturing of these engines.

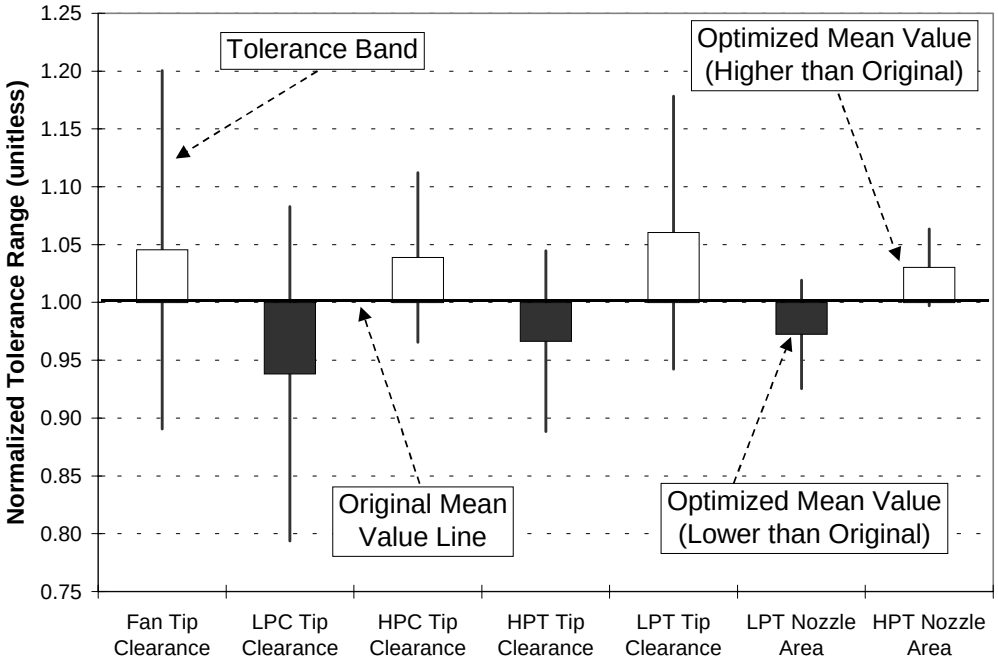

Figure 5-11: Normalized Tolerance Range Results

The HPT nozzle area in particular is an issue worth exploring further. The results of this research indicate that the build specification for the HPT nozzle area that will allow for the most variation in all of the components at a three-sigma pass rate is 3.48 ± 0.11 in^2 [3.59 to 3.37 in^2]. The original build specification for the engine was an HPT nozzle area of 3.38 in^2. Essentially, the extreme limits of the allowable tolerance appear to be at virtually the same point as the mean value of the actual build specification. This result implies that even very small variations in the original HPT nozzle area will force the engine to fail one or more performance specifications at a rate considerably greater than 1 / 99.74 = 0.010027%. This is an important point to consider because if the mean value of the HPT nozzle area were set to 3.48 in^2 rather than 3.38 in^2, any associated tolerance of the nozzle area would be much more forgiving regarding the performance of the engine. Furthermore, adjusting the mean value of a component is a relatively

inexpensive task when compared with adjusting the tolerance of a component, but can be equally beneficial.

It is difficult to quantify where the best mean value lies without including performance variation. If only single point nominal values are considered, the engine model indicates that all the performance parameters will be acceptable when the original set of mean values is used. However, when variation is included, the mean values are forced to change in order to maintain acceptable performance over a wider range. By using nominal values without variation, which is common practice for engine manufacturers today, there is no quantification of how often a given set of input geometries will yield an acceptable product or how much variation can be introduced before a failure is realized.

Although not unreasonable, the optimized tolerances that are presented in Table 5-1 are tighter than those found on many gas turbine engines. At the same time, the tolerances are based on a three-sigma pass rate, which is a pass rate rarely achieved by engine manufacturers. If the simulation were performed at a two-sigma level, the tolerance ranges would be much wider. Although variation simulation modeling of the type illustrated here is a powerful tool, it should only be used with respect to actual manufacturability. Like many designs that work well on paper, it is possible to discover results that are simply not practical regarding manufacturability. Results from an optimized variation simulation model should always be checked against the limits of real world manufacturing capability.

5.2 Cost Modeling Results

As mentioned before, all of the work to this point has been done assuming that three-sigma is the optimum performance pass rate for an engine. This, however, is most likely not the case. It has long been accepted that tighter tolerances are more expensive to manufacture and assemble than wider tolerances. What is not so largely accepted is the method to model the cost of a part as a function of its tolerance. Although not the

focus of this work, some basic cost modeling was performed to illustrate the concept and the tools required to model the cost of producing an engine at different levels of reliability.

Using Equation 4-2, simulations were performed at different levels of reliability in order to illustrate the method by which the least cost reliability can be determined. To establish a baseline for the cost modeling, the values shown in Table 5-3 were arbitrarily selected. Table 5-3 shows that to manufacture and assemble a fan disc with tip clearances that have a tolerance of ± 3 mils will cost $1.50. To manufacture an LPC with tip clearances having a tolerance of ± 1.3 mils will cost $2.00, and so on. The total cost to assemble and manufacture the engine under the reference tolerances is $15.50. Although the costs are impossibly low and for illustrative purposes only, the different costs per part assignments do carry with them an aspect of proportionality. For example, an HPT disc is assumed to be twice as expensive as a fan disc.

Table 5-3: Cost Modeling Initial Conditions

	Fan	LPC	HPC	HPT	LPT	HPT Area	LPT Area	TOTAL
Cost/Part	$1.50	$2.00	$2.00	$3.00	$2.00	$3.00	$2.00	$15.50
Reference Tolerance	+ - 3.0 mils	+ - 1.3 mils	+ - 3.2 mils	+ - 3.0 mils	+ - 2.0 mils	+ - 0.13 in^2	+ - 0.6 in^2	

The baseline cost will become a reference point by which the cost modeling simulations will be compared. Consider a solution that passes the required level of reliability with a fan tolerance that is ± 4 mils rather than ± 3 mils. The cost to manufacture the fan will be reduced from $1.50 to [$1.50 / (4 mils / 3 mils)^2] = $0.84 since the cost is proportional to the inverse of the square of the tolerance. These simple cost modeling calculations were done at levels of reliably ranging from 1.0 (66.26% pass rate) to 3.0 (99.74% pass rate) where the least expensive combination of tolerances was selected for each level of reliability tested.

The other cost model that must be included in order to determine total minimum cost is a function that describes the penalty for failing to meet all of the criteria. Again,

this function can be rather subjective, but should incorporate the cost associated with having to scrap an engine because it failed to meet some criteria, or the cost to fix the engine that failed performance. The penalty function used in this cost modeling is the product of twice the cost of the baseline engine ($15.50 x 2 = $31.00) and the failure rate associated with a specific level of reliability. For example, at a one-sigma pass rate, the failure rate is 1.0 - 0.6826 = 0.3174, which gives a penalty function of $31.00 x 0.3174 = $9.84. Twice the cost of the baseline engine is used to simulate the cost associated with scraping a defective engine and having to build a new one, resulting in one engine being built for the cost of two. This is a gross simplification of what would happen in an actually engine manufacturing environment, but does provide a qualitative trend of cost versus failure.

Figure 5-12 shows the results of the cost minimization. The dashed line represents the least cost that can be achieved when manufacturing the engine at a given reliability. As the dashed line indicates, higher levels of reliability require a much greater cost to produce since the tolerances have to be much tighter. There is a unique trend seen in the dashed line in that it appears to approach a constant cost value once the pass rate goes below a level of 1.5. This same asymptotic relation was seen in cost modeling performed by J. R. He [14], where he noted it was common to see in increase in cost as the tolerance was decreased and a cost that essentially remained constant as the tolerance is continually increased. The dotted line shows the cost associated with an engine failing some or all of its performance criteria. As the pass rate is increased, fewer engines are failing performance so there is less need to spend money scraping or fixing these engines. The solid line is the sum of the production cost curve and the failure curve and shows the reliability where the least total cost of production can be realized. Given the assumptions noted above, a two-sigma (95.44%) pass rate will result in the least total cost of production.

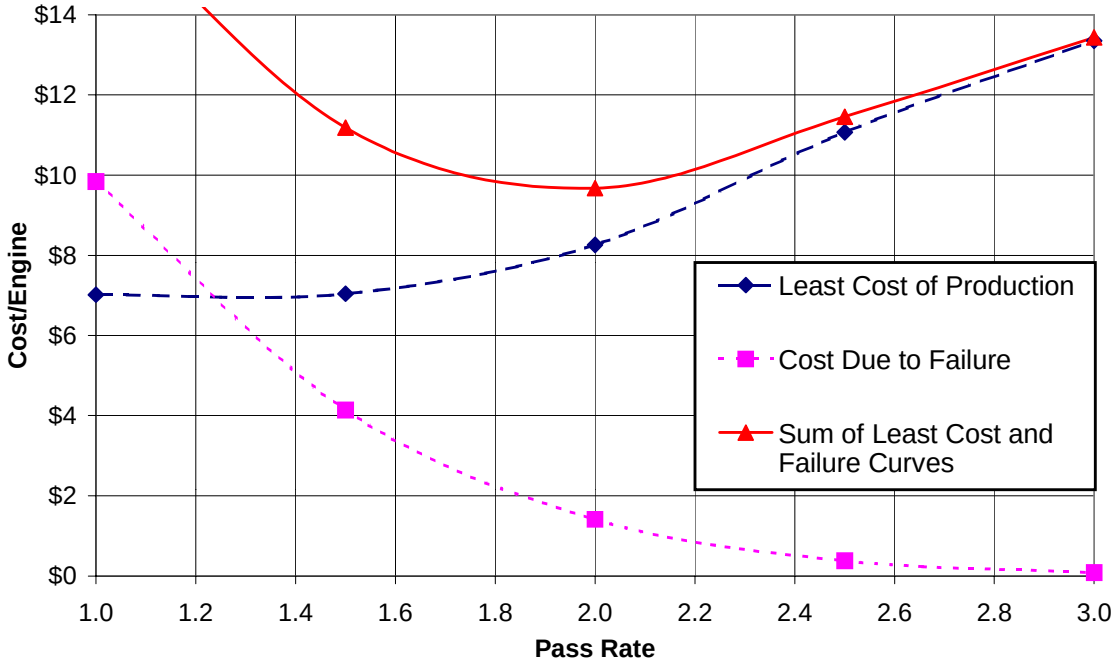

Figure 5-12: Cost Modeling Results

Further examination of Figure 5-12 shows that the solid line illustrating the total cost of production has a much steeper slope towards lower pass rates than towards higher pass rates. By moving equidistant in each direction along the Pass Rate axis from the optimum pass rate of 2.0, it is evident that the cost associated with producing an engine at a lower than optimum pass rate is more expensive that producing an engine at a higher than optimum pass rate. Although purely qualitative, the implication is that an error on the higher side of the optimum pass rate is better than an error on the lower side of the optimum pass rate.

6 Conclusions

The goal of this work was to show the tools and procedures required to determine the maximum tolerances of critical turbine engine components that allow for a predetermined pass rate of a number of performance or geometric constraints. This project is an advancement over previous work because it is a performance-based method for determining tolerances, and adequately captures the interactions of all the tolerances included in the analysis. Previous methods of determining tolerances focused only on individual components and were not determined with respect to the tolerances of other components. Specifically, this work was able to determine the maximum tolerances of seven critical turbofan engine geometries that allowed a three-sigma pass rate of five geometric and six performance parameters. The number of components modeled and the number of constraints is limited only by the computational power available to the user. While this work was done for a moderate bypass military turbofan engine, the algorithms are applicable to a wide range of turbine engines. Land-based gas turbines, auxiliary power units, micro-turbines, turbofans, and turbojets could all be analyzed at any desired pass rate with the methods presented in this research.

6.1 Component Modeling

There are two basic steps to performing a variation simulation model on a gas turbine engine. The first step is to create a model that can predict changes in performance as a function of the geometry. The second step is to generate an algorithm that simulates building engines under different conditions. The performance parameters selected as constraints should be those parameters that are guaranteed to a customer or critical minimums that define the line between the product being able to perform the required task or not. A good criterion for selecting which geometries should be included in the simulation can be done with a sensitivity analysis. Critical geometries that affect the product performance and have unknown tolerance ranges are good candidates to include in the model.

To model changes in performance as functions of geometry, there is no substitute for experimental data. Unfortunately, experimental data that involves changing tip clearances and turbine nozzle areas is difficult and expensive to attain. Destructive testing would be necessary since the tip clearances of the components in question must be filed or ground down. In the absence of experimental data, correlations can be used to simulate the effect of changing tip clearances. Accurate and detailed correlations will help ensure a more precise final solution.

The sensitivity analysis done in this work was found to be essential in establishing which parameters to include in the model and explaining the trends seen in the results. As expected, the highly influential components were found to consistently be the factors that had the most effect in shaping the results. Of all the components included in this analysis, the turbine nozzle areas were the most influential. Although the sensitivity study performed in this work was a simple parametric analysis and did not capture any non-linear responses, it was sufficient to use as a qualitative guide in explaining the engine performance trends. Although not as accurate as performing a full variation simulation, a sensitivity analysis alone can be used to give engine manufacturers a good idea which components are potentially responsible for causing performance problems.

6.2 Variation Simulations

Once the performance as a function of the geometry is adequately modeled, the rest of the process is essentially done entirely with linear or non-linear programming and statistical analysis. The generation of a response surface model is critical to this type of work due to the sheer number of calculations that need to be performed. The response surface models used in this work were found to be incredibly accurate over the range in which they were design to predict with no adjusted R^2 being less than 0.99965. Generally speaking, the wider the range and the more points that can be used to generate the RSM, the better the result will be. A large number of points (relative to the number of independent variables) are advantageous when generating an RSM since it will allow a greater chance to capture any special features throughout the surface. An RSM that has a

wide range of validity when compared with the design space that will be explored is also important to ensure that any associated standard deviation will allow for all of the predictions to be made within the range of model validity.

The definition of the design space is also crucial when programming a variation simulation model. Larger design spaces will take more computational power and time to explore, but offer the advantage of finding solutions that will not be possible in smaller design spaces. By limiting the design space, the user could potentially be ignoring solution space that may be optimal. A good indicator that the design space is adequately large is the presence of clusters of solutions freely standing in the design space and not touching one or more of the design space borders.

The fact that commonly known optimization algorithms can be effectively applied to the optimization of a standard deviation is concept well illustrated by this research. The selection of an optimization algorithm that will be used to optimize a stochastic process is important because not all algorithms are well suited to this application. A large number (50,000) of simulations at each selected standard deviation are required to effectively use a gradient-based optimization algorithm. The larger the number of simulations performed at each selected standard deviation, the more consistent the quantiles of the results will be and the more accurate the prediction of the next step by the gradient based algorithm. Furthermore, a genetic algorithm with a population of approximately 4 times the number of independent variables coupled with the NLPQL algorithm appears to work well for optimization applications like those presented here. It was found that searching for solutions without an optimization algorithm resulted in spending a tremendous amount of computational time searching infeasible design space. Although computationally intensive, the method that does not involve an optimization plan can be used to preliminarily explore the design space but it is ineffective for finding an optimal solution.

The capability to trade tolerances on one component for the tolerances on another component is an important conclusion realized after performing several different

simulations. Some components offer considerable tradeoff potential while others, such as the turbine nozzle areas, appear to have limits that can never be successfully exceeded. This can be useful information for an engine manufacturer who is looking to reduce costs on an expensive component, and can do so by tightening the tolerances on a different, less costly component. Although tradeoffs between the tolerances of some components exist, there appears to be a limit of the total variation that can be introduced. A single simulation can contain several highly feasible solutions, each with unique component tolerances, yet the normalized sum of those tolerances is nearly constant for each simulation. In essence, the optimization algorithm is maximizing the sum of the variances and can achieve this maximum through a variety of different combinations of individual tolerances. In other words, no additional variance can be added to the system but the variation that does exist in the system can be traded between certain components.

Analyzing the results of a multi-dimensional optimization can be a complex task. Since there is no single graphical method available to effectively illustrate fourteen dimensions, a combination of plots and tables were generated. A normalized Euclidian Distance between points was shown to be a practical tool for determining the relation between a reference point and any other known point in the multi-dimensional space. By quantifying the multi-dimensional distance between the optimum solution and all other solutions, a true representation of the distance between feasible and infeasible solutions can be generated and explored.

6.3 Cost Modeling

Modeling and optimizing the costs of turbine engine manufacturing is a difficult task. A very detailed cost model and penalty function must be available if there is to be any confidence in the results. If these cost models exist, this work was shown to be a viable platform by which an optimization of manufacturing costs can be performed. The results from a cost modeling simulation will tell an engine manufacturer what the optimum pass rate of their engines should be to maximize profits and the tolerances that are associated with the optimum pass rate.

If detailed cost models are not available, this research still provides some insight into cost minimization. If a pre-determined pass rate is assigned, the algorithm will find the maximum tolerances that allow the performance constraints to be satisfied. By maximizing the tolerance, the cost is essentially being reduced since it is commonly accepted that the cost is inversely proportionally to the tolerance. Furthermore, a cost engineer can apply some basic quantitative experience and refine the model even more by incorporating weighting functions. A cost engineer who is aware of the fact that an HPT disc costs nearly twice as much as a fan disc may choose to make maximizing the standard deviation of the HPT twice as important as maximizing the standard deviation of the fan. This being the case, the weight of the objective function responsible for maximizing the standard deviation of the HPT can be set to 2.0, while the weighting function of the fan kept at 1.0. This method will not reveal what the lowest cost will be or help in determining if the pre-determined pass rate is an optimum, but it will provide the tolerances that will result in the least cost, given the assumptions that were made.

7 Summary and Recommendations

7.1 Applications of Current Work

The algorithms presented in this thesis offer a tool that can be used to bring the experience of individual component designers together to ensure a consistently satisfactory product. Without a tool such as this, component designers are forced to set tolerances on their designs based solely on the requirements of the individual component and only guess at the total effect that the tolerance they assign will have on the overall product. This research is a performance-based method that introduces a means by which interacting tolerances can be accurately determined without guesswork or expensive redesign.

A time in which this work would be particularly advantageous would be during the design phase of a new engine or new engine components being fit to an existing engine. By identifying the maximum possible tolerances that satisfy an acceptable level of reliability early in the design phase, the cost of manufacturing an engine can be reduced. The lower production cost can be realized for a number of reasons. First, this tool will aid the engineer in identifying less critical dimensions which may allow for the use of simpler and less costly manufacturing processes. Furthermore, this engineering tool can identify potential manufacturability problems prior to production, and corrections can be incorporated that will lower the performance failure rate and result in less post-production rework.

Variation simulation modeling also provides a means by which to perform a trade study of different tolerances. Using a VSM such as the one presented here, an engineer would be able to determine how tight the tolerances have to be on a newly fit compressor given the tolerances of an existing turbine, in order that the engine passes a selected level of reliability. If, perhaps, no solutions were found, the engineer would have an indication that there is no such tolerance on the compressor that will satisfy the performance criteria and that the tolerances of the turbine would also have to be altered (i.e. the turbine must

also be replaced). Taking this idea one step further, simulations can be done that identify the least expensive method by which an engine can be brought from a condition where performance is not met to a condition where performance is met. An engineer would be able to explore the possibility of replacing the LPC and HPC in order to avoid having to replace the HPT. Perhaps by changing the LPC and HPC, an acceptable performance can be met while avoiding the more costly replacement of the HPT.

Engine modeling that includes variation can also be used well after the design phase when an engine is already in production. After identifying which component is causing a performance failure, a VSM can be run to determine the tolerances that have to be met on that component in order to allow for acceptable performance. This type of modeling could eliminate the expensive practice of guessing at the tolerances required on newly serviced parts and having to replace them if the initial guess was wrong. A few incorrect guesses can result in hours of maintenance and testing before acceptable performance is achieved.

A real world example of the potential uses of a VSM can be taken from one of the largest energy producers in the United States. The energy company in question operates large industrial gas turbines for power generation. When performing maintenance on the HPT and LPT sections, a variety of turbine nozzle vanes are ordered, each with different flow areas. When the engines are reassembled and tested after the maintenance work, they sometimes fail to meet the performance specifications. By having a variety of different turbine nozzle vanes on hand, the technicians can continue to replace the vanes until the engine passes performance. A VSM would help quantify the proper mean area and tolerance of the turbine vanes that could result in the company not having to purchase additional vanes that will not be used.

7.2 Recommendations for Future Work

While this work represents a unique improvement to turbine engine modeling technology, it stands to benefit even further from continuing improvements. Two major

areas in which this work can be improved upon are the number of components modeled and the loss models that accompany each of them. There are only seven components included for analysis in this work, and they were selected based on the belief that they are some of the most influential geometries found in a turbine engine. Some of the components were found to be truly influential, while others were eventually determined to be less important than originally anticipated. By including other geometries such as bypass area, casing diameter, nozzle area, and a number of others, a more complete and accurate look at the engine system can be made.

Again, all of the geometries that are included for analysis must be accompanied by models that accurately predict the change in performance as a function of the change in geometry. These component models are the second major area where improvements can be incorporated into the existing work. By performing experiments that are aimed at determining changes in performance as a function of changes in geometry, data that are more accurate can be incorporated into the VSM. If experimental data are not practical, CFD may provide an acceptable alternative. Either CFD or experimental work will represent an improvement over the simple geometric correlations currently used in this work.

Another area of improvement includes performing the simulations with distributions other than Normal. A Normal distribution represents an ideal case of variability where as a Weibull distribution may be more representative of an actual manufacturing environment. A Weibull distribution will help capture phenomena such as machine tool wear and worker fatigue. This is a somewhat minor change to the algorithm since virtually any distribution can be simulated given the fact that there is data to support the use of such distribution.

Another area that may be worth researching was theorized by the author during the development of these algorithms but never explored. It may be possible to work the problem in reverse with respect to the way in which it was presented in this thesis. Rather than start with a large design space and look for feasible solutions that result in

acceptable distributions, it may be possible to start with acceptable distributions and reverse the process to "back out" the feasible design space. While this method may prove to be mathematically difficult, it could result in a significant time savings and a much more accurate depiction of the feasible design options.

The last topic that will be mentioned as a recommendation for future work will be the focus of the remainder of this chapter. This topic is the influence of time dependent performance degradation on turbine engines. This area of research is of unique interest to organizations responsible for the maintenance of large fleets of aircraft such as commercial carriers and military forces. While this work predicted the reliability of newly manufactured engines, time dependent modeling could be used to predict the reliably of turbine engines over their lifecycle.

7.3 Applications of Future Work

Time-variant performance predictions offer a tool by which turbine engine performance can be predicted over the lifecycle of the engine. To perform this type of analysis, accurate models of how time degrades the performance of each component are necessary. Furthermore, a large amount of data are required in order to get a good probabilistic model that governs the occurrence of a specific type of failure. However, the implications of being able to accurately predict the performance of an engine over time are significant.

As a turbine engine is operated over time, some of the components undergo performance degradation because of blade tips and seals rubbing, combustion deposition, and a number of other phenomena. The result of these detrimental phenomena is that the engine will deliver less performance as time continues. A common factor used to gauge performance degradation is the turbine inlet temperature. A higher corrected turbine inlet temperature is considered degradation in performance for the fact that the engine must work at a hotter temperature to deliver the same performance. Figure 7-1 shows the degradation of corrected turbine inlet temperature as a function of time for a single

military turbofan engine. Each dot represents a snapshot in time taken by the engine when recording its performance at a predetermined point in the flight envelope. The time on the x-axis of Figure 7-1 is shown increasing over the span of several months and a notable increase in the corrected turbine inlet temperature is highlighted. Specifically, on the date shown by the arrow, something caused a dramatic increase in the corrected turbine inlet temperature. This increase was actually caused by a high g maneuver that exceeded the limits of the engine. The result was a seal and/or tip clearance rub that permanently affected the performance of the engine. This example will be used to illustrate the method by which a time-variant performance prediction can be made.

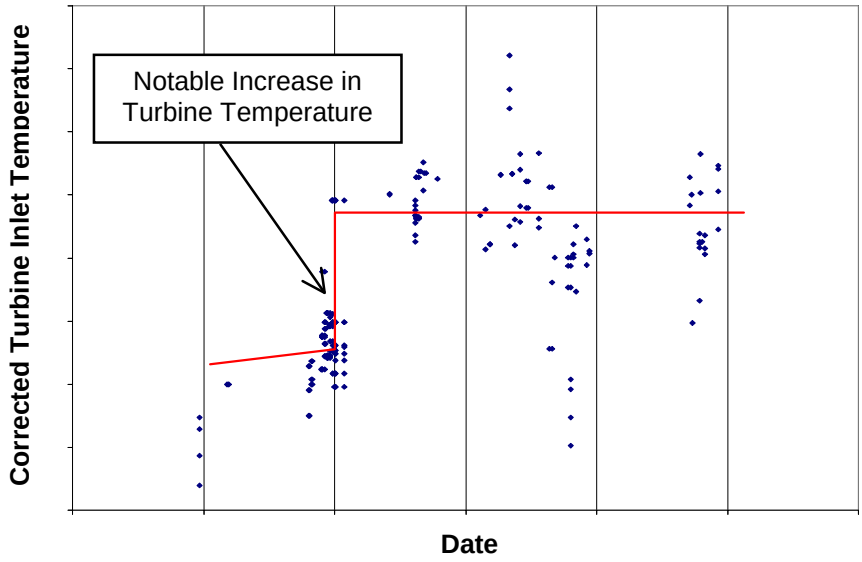

Figure 7-1: Time-Variant Corrected Turbine Inlet Temperature

The first step in generating a time-variant engine performance prediction is to establish a fault tree. Turbine engines suffer from performance degradation over time for many different reasons. Some engines begin to suffer from combustion deposits, while others like the example given in Figure 7-1, suffer from tip clearance rubs. Whatever the reason for reduced performance, there is an associated probability that any specific fault will be the reason for causing an engine to fail performance. Moreover, it is quite possible that any single fault is not the entire reason for a performance failure but rather an interaction of different faults. The purpose of a fault tree is to assign probabilities to specific faults and fault interactions that govern the occurrence of those faults. As an

example, given enough time-variant profiles of engine similar to those shown in Figure 7-1, it is possible to establish a probability that a tip rub will occur as a result of a high g maneuver.

The fault tree must be followed by models that quantify the effect on performance as a function of experiencing one of the faults. A degradation model may be as simple as associating any tip rub fault with a corrected turbine temperature increase of 3.0% (as seen in Figure 7-1), or as complex as modeling all the interactions associated with an increase in tip clearance. The more inclusive the model, the more accurate the prediction will be.

Having access to large amounts of data is the key to generating time-variant models. In essence, predictions of younger engines are being made based on the statistics associated with older engines. Furthermore, since not all engines are used in the same manner, there are further categories that must be defined. For example, military fighter engines used as trainers have a different mission and suffer from different faults than fighter engines used in an operational theater. This further breakdown of categorical usage may help narrow the variance of the individual fault trees and degradation models.

Since large amounts of detailed engine performance data are generally not available, there are methods to help fill in the missing pieces. It is common for engine monitoring systems to record a few key performance parameters during each flight, but not to record all the parameters that will be need to produce detailed time-variant models. In the absence of detailed information, an accurate performance model can be used along with an optimization algorithm to help predict what the unknown performance parameters must have been based on the known performance parameters. As an example, say that the atmospheric conditions, the turbine inlet temperature, and the engine speed are known and that the SFC is not known. Using an accurate performance model and an iterative algorithm, a good prediction of the SFC can be made based on the parameters that were known and input to the performance model. By employing this method,

additional data can be extracted from known data points and used in the generation of time-variant models.

There are a number of notable advantages that result from being able to accurately predict the performance of an engine over time. First, it can provide a knowledge based, forward looking method for depot level organizations to forecast maintenance budgets. This is an issue of particular concern to military organizations that are often required to predict the budget for the next fiscal year based on the budgets required from years past. While projecting budgets based on the previous year is not necessarily inaccurate, it does not reflect the additional requirements that will be needed based on an aging engine inventory that is not degrading in a linear fashion. Time-variant predictions can be used to illustrate the fact that an entire family of engines may fall out of a performance specification over a very short period where simple extrapolation of engine trend data will not reveal any such occurrence. Moreover, this tool can be used to determine the most effective timing and level of maintenance in order to prevent a surge of engine performance failures. By having the ability to know when a large percentage of the engines in an inventory are most likely going to fail, appropriate fixes can be incorporated into regularly scheduled maintenance. This forward-looking maintenance methodology will significantly reduce the resources required to perform unscheduled maintenance.

The union of the work presented in this research and the potential future applications illustrated in this section begins with the component performance loss models. The loss models used in this work are essentially at time equal to zero, for a newly manufactured engine. By transforming the loss models used in this work into models containing time-variant information, time can be included in the optimization algorithm. Figure 7-2 shows an example of how the time-variant optimization might be adapted from the current work. In addition to having data describing the distribution of tip clearances at time equal to zero TACs (Total Accumulated Cycles), data would also have to be available at other points in time. By having data at 800 and 1200 TACs, it is possible to see how time affects the distribution of tip clearances for an engine under

known operating conditions. With this new time-dependant description of tip clearance distributions, optimization routines can be run that include time as an independent variable, constraint, or objective function.

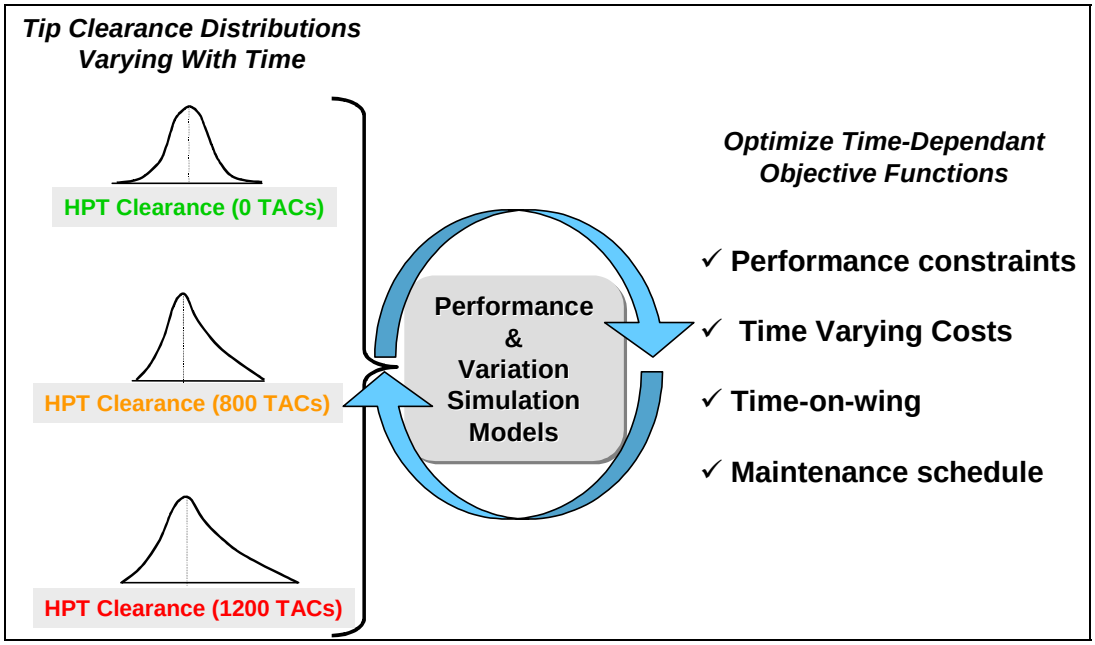

Figure 7-2: Time-Variant Optimization Concept

As an example, simulations can be run to provide answers to a number of questions:

1. Given a known initial starting condition, how long will the current engine last before it fails any performance specification?
2. When is the best time to replace the engine components to allow for a maximum reliability at the least cost (e.g. time-on-wing objective)?
3. What performance specifications does an engine have to meet today to ensure it will meet the minimum requirements 200 hours (or any time increment) from now?

Appendix: Response Surface Model Summary

```
Response Surface Model

Model input parameters:
      x01 = "fanclr"
      x02 = "lpcclr"
      x03 = "hpcclr"
      x04 = "hptclr"
      x05 = "hptarea"
      x06 = "lptclr"
      x07 = "lptarea"

Model output parameters:
      y01 = "SFC"
      y02 = "HPCSM"
      y03 = "T41"
      y04 = "NL"
      y05 = "NH"

MODEL DATA

Number of inputs:                   7
Number of outputs:                  5
Number of recorded designs:         400
Number of polynomial terms:         36
Rank of least-squares matrix:       36

Matrix with design outputs:

R^2 coefficients of output functions:

  output         R^2         adjusted R^2
-----------------------------------------------
   y01  |     0.99968     |     0.99965     |
   y02  |     0.99997     |     0.99997     |
   y03  |     0.99977     |     0.99975     |
   y04  |     0.99978     |     0.99975     |
   y05  |     0.99980     |     0.99978     |
```

Table A-1: Response Surface Model Coefficients

	const	fanclr	hpcclr	hptclr	hptarea	lptclr	lptarea	hptarea**2	lptarea**2	fanclr*hptarea
SFC	0.4683	-0.0380	0.4760	2.5489	-0.1107	0.3219	-0.0004	0.0194	0.0001	0.0138
HPCSM	-209.5714	2.4655	-91.3182	-359.1963	97.7466	8.1440	4.5294	-11.8250	-0.1479	-2.0343
T41	3716.7174	-289.9293	3436.9804	15885.8398	-816.1940	602.9603	-82.9784	128.4392	3.0186	80.0435
NL	15294.1312	381.0410	-285.0919	-2042.7487	71.6267	-1140.9521	-70.6918	-13.5718	1.5387	-17.4636
NH	33892.0976	-3418.2778	6126.1098	-34361.2148	-995.1602	12899.2061	1630.0854	-133.5575	-41.7548	87.6227

	fanclr*lptarea	lpcclr*hptarea	lpcclr*lptarea	hpcclr*hptarea	hpcclr*lptarea	hptclr*hptarea	hptclr*lptarea	hptarea*lptclr	hptarea*lptarea	lptclr*lptarea
SFC	0.0001	0.0063	0.0098	-0.0672	-0.0104	-0.2430	-0.0270	-0.0007	0.0000	-0.0036
HPCSM	0.2842	-9.0034	-0.8579	4.2011	3.4552	8.5654	5.7010	-11.7318	0.1755	1.7932
T41	3.2587	51.9864	66.8319	-399.0970	-95.9802	-1103.3164	-399.9930	1.9083	-0.7412	-1.8418
NL	-13.4287	12.2267	-11.1858	57.4338	1.9225	258.2036	25.4086	103.5422	0.2334	-2.6623
NH	229.8134	1869.5464	308.5193	582.9582	-90.8623	3154.8179	-1827.2394	-2088.1345	27.0984	15.2248

References

[1] Dai, S. and Wang, M., *Reliability Analysis in Engineering Applications* (New York: Van Nostrand Reinhold, 1992), pp. 72-73.

[2] Bury K., *Statistical Distributions in Engineering* (New York: Cambridge University Press, 1999), p.137.

[3] Early, J., and Thompson, J., "Variation Simulation Modeling – Variation Analysis Using Monte Carlo Simulation," *Failure Prevention and Reliability,* vol. 16 (September 1989), pp. 139-144.

[4] FAST Users Guide, Honeywell Engines and System, Phoenix, Arizona, 1995.

[5] Booth, T. C., "Importance of Tip Clearance Flows in Turbine Design," *von Karman Institute for Fluid Dynamics, Tip Clearance Effects in Axial Turbomachines Lecture Series 1985-05* (April 1985).

[6] Lakshminarayana, B., "Compressor Loss Correlations and Analysis and Effects on Compressor Performance," *von Karman Institute for Fluid Dynamics, Tip Clearance Effects in Axial Turbomachines Lecture Series 1985-05* (April 1985).

[7] Tallman, J. and Lakshminarayana, B., "Numerical Simulation of Tip Leakage Flows in Axial Flow Turbines, with Emphasis in Flow Physics," *Proceedings of the ASME Turbo Expo 2000* (May 2000), pp. 1-15.

[8] Copenhaver W. W., et al., "The Effect of Tip Clearance on a Swept Transonic Compressor Rotor," *Transactions of the ASME Journal of Turbomachinery*, vol. 118 (April 1996), pp. 230-239.

[9] Craig, M., "Managing Variation by Design Using Simulation Methods," *Failure Prevention and Reliability*, vol. 16 (September 1989), pp. 153-163.

[10] Ostwald, P. F. and Blake, M. O., "Estimating Cost Associated with Dimensional Tolerance," *Manufacturing Review*, vol. 2, no. 4 (1989), pp. 277-282.

[11] Rao, S. S., *Engineering Optimization-Theory and Practice* (New York: John Wiley and Sons, 1996).

[12] Schittkowski K., "NLPQL: A FORTRAN subroutine solving constrained nonlinear programming problems," *Annals of Operations Research*, vol. 5 (1985), pp.485-500.

[13] iSight Developers Guide, 1999, Document Version 2.0, Engineous Software, Inc., North Carolina

[14] He, J. R., "Tolerancing for Manufacturing Via Cost Minimization," *International Journal of Machine Tools and Manufacturing*, vol. 31, no. 4 (February 1991), pp. 455-470.

[15] McDonald, P. W., "Transient Model Applications," *von Karman Institute for Fluid Dynamics, Gas Turbine Engine Transient Behavior Lecture Series 1993-06* (1993).

[16] Horlock, J. H., *Axial Flow Turbines* (Florida: Krieger Publishing Company, 1966), pp. 107-108.

[17] Chase, K. W., et al., 1989, "Least Cost Tolerance Allocation for Mechanical Assemblies with Automated Process Selection," *Failure Prevention and Reliability*, vol. 16 (September 1989), pp. 165-171.

[18] Wilhelm, R. G. and Lu, S. C.-Y., *Computer Methods for Tolerance Design* (New Jersey: World Scientific 1992).

[19] Brook R. H. W., *Reliability Concepts in Engineering Manufacture* (London: Butterworths, 1972).

[20] O'Connor, P. D., *Reliability Engineering* (Washington: Hemisphere Publishing Corporation, 1988).

[21] Munoz R., "Optimization Strategies for the Synthesis / Design of Highly Coupled, Highly Dynamic Energy Systems," Ph.D. Dissertation, Virginia Polytechnic Institute and State University, Blacksburg, Virginia (2000), pp. 15-22.

[22] Ghiocel, D. M. and Neville, R. F., "Specific Probabilistic Modeling Issues for Gas Turbine Engine High Cycle Fatigue (HCF) Life Prediction," FAA/AFWL Application of Probabilistic Design Methodologies to Gas Turbine Rotating Components Workshop (1998).

[23] Arnulfi, G. L. and Massardo, A. F., "The Effect of Axial Flow Compressor Deterioration on Gas Turbines and Combined Cycle Power Plants," *IGTI-ASME Cogen-Turbo Journal*, vol. 8 (1993), pp. 607-614.

[24] Sues, R. H., et al., "Reliability Based Optimization for Design of Aeropropulsion Components," *American Institute of Aeronautics and Astronautics*, AIAA-96-1609-CP (1996), pp. 2511-2518.

[25] Basson, A. and Lakshminarayana, B., "Numerical Simulation of Tip Clearance Effects in Turbomachinery," *American Society of Mechanical Engineers*, 93-GT-316 (1993), pp. 1-16.

[26] Chima, R. V., "Calculation of Tip Clearance Effects in a Transonic Compressor Rotor," *American Society of Mechanical Engineers*, 96-GT-114 (1996), pp. 1-11.

[27] Baghdadi, S., “Modeling Tip Clearance Effects in Multi-Stage Axial Compressor,” *American Society of Mechanical Engineers,* 95-GT-291 (1995), pp. 1-17.

[28] Computational Science Education Project, *Introduction to Monte Carlo Methods* (1995).

Vita

The author, son of Lee Howard Sheldon and Peggy Jean Eckhardt, was born in Portland, Oregon on 20 July 1972. He attended Sunset High School in Beaverton, Oregon before proceeding to Oregon State University (OSU) where he received his B.S. in Mechanical Engineering in 1995. During his senior year, Karl served as the president of Pi Tau Sigma (Omega Chapter) and was voted as the OSU Mechanical Engineering Department's Outstanding Senior. Having attended OSU on a 4-year army ROTC scholarship, Karl was commissioned as a Second Lieutenant in the US Army Medical Service Corps upon graduation from college. During his 4 years in the army, he was stationed in Germany and served a year as member of both NATO & UN forces operating in the Balkans. After earning the rank of Captain, Karl receiving an honorable discharge from military service and chose to attend Virginia Tech to earn a Master of Science in Mechanical Engineering. Upon graduation from Virginia Tech, he will begin work for General Electric Corporate Research and Development in Niskayuna, New York.