2007-03-28

Optimizing Product Variant Placement to Satisfy Market Demand

Jonathan Roger Parkinson
Brigham Young University - Provo

OPTIMIZING PRODUCT VARIANT PLACEMENT TO SATISFY MARKET DEMAND

by

Jonathan R. Parkinson

A thesis submitted to the faculty of

Brigham Young University

in partial fulfillment of the requirements for the degree of

Master of Science

Department of Mechanical Engineering

Brigham Young University

April 2007

BRIGHAM YOUNG UNIVERSITY

GRADUATE COMMITTEE APPROVAL

of a thesis submitted by

Jonathan R. Parkinson

This thesis has been read by each member of the following graduate committee and by majority vote has been found to be satisfactory.

Date	Jordan J. Cox, Chair
Date	Carl D. Sorensen
Date	Brian D. Jensen

BRIGHAM YOUNG UNIVERSITY

As chair of the candidate's graduate committee, I have read the thesis of Jonathan R. Parkinson in its final form and have found that (1) its format, citations, and bibliographical style are consistent and acceptable and fulfill university and department style requirements; (2) its illustrative materials including figures, tables, and charts are in place; and (3) the final manuscript is satisfactory to the graduate committee and is ready for submission to the university library.

Date

Jordan J. Cox
Chair, Graduate Committee

Accepted for the Department

Matthew R. Jones
Graduate Coordinator

Accepted for the College

Alan R. Parkinson
Dean, Ira A. Fulton College of
Engineering and Technology

ABSTRACT

OPTIMIZING PRODUCT VARIANT PLACEMENT TO SATISFY MARKET DEMAND

Jonathan R. Parkinson

Department of Mechanical Engineering

Master of Science

Many companies use product families in order to offer product variants that appeal to different market segments while minimizing costs. Because the market demand is generally not uniform for all possible product variants, during the design phase a decision must be made as to which variants will be offered and how many.

This thesis presents a new approach to solving this problem. The product is defined in terms of performance parameters. The market demand is captured in a preference model and applied to these parameters in order to represent the total potential market. The number and placement of the product variants are optimized in order to maximize percentage of the potential market that they span.

This method is applied to a family of mountain bikes and a family of flow-regulating disks used in industrial applications. These examples show that usage of this method can result in a significant increase in potential market and a significant reduction in production costs.

ACKNOWLEDGMENTS

The overarching theme of this work has been patience. I'm very appreciative of the patience of all those who've supported me in this endeavor.

To my family, thank you for the three years of praying that I would finish my thesis. I'm grateful for my parents who have been endlessly supportive of me. I'm also very grateful to my brother, Matt, for all of his help and for loaning me his enthusiasm for two years.

To my advisor, Jordan Cox, thank you for not giving up completely on me. I'm grateful for the countless hours of instruction and for the advice on my thesis and on life in general. Above all, I'm grateful for the friendship.

I'm appreciative of my committee members Carl Sorensen and Brian Jensen, for the time they've spent and their very helpful comments and corrections.

Finally, I'm grateful to my Father in Heaven for guiding me through this period of my life, for inspiring me, and for buoying me up during the difficult times. Without Him, nothing would be possible.

Table of Contents

List of Tables

List of Figures

Chapter 1

Introduction

The ultimate goal of most companies is to make money by offering products that satisfy their customers. The challenge that companies commonly face is how to hold on to their current customers while changing their products in a way that will appeal to potential customers. This is often done by taking a single product and offering several variations of it.

Products used to be offered as one size fits all. The classic example of this is the statement attributed to Henry Ford that people could have any color car they wanted, so long as it was black. More recently companies have expanded their product offering in order to capture new customers and to remain competitive. There are various methods that can be employed to change the product. For example, one common method is to offer different colors of the otherwise exact same product. Market research is done to determine which colors will be best received by the current customers and which colors might bring in new customers. A cost-benefit analysis is performed on this data and a decision is reached about which colors will be offered. The offered set of product variations will satisfy most, but rarely all, of the potential customers.

Another common but more complex method is to offer scaled or skewed configurations of a product. For a t-shirt this may again be a simple process of applying market research to determine what sizes to offer. For a bicycle, an automobile, or a jet engine it is much more complicated. In addition to the relative complexities of the products, there are differences in the kinds of configuration options that are available. Rather than discrete parameters like color, each of these will have many continuous parameters adding to the complexity of the problem. It is always in a company's best interest to balance the cost of offering products with the potential market for those products. If more than the necessary number of product variants are offered, or if some of the offered variants are too similar to each other, then money is lost manufacturing, shipping, and storing products that don't bring in any more customers or make the current customers more pleased. If too few product variants are offered, or if there are significant gaps in the set of variants offered, the company risks losing potential customers.

This thesis presents a method to determine which product variants to offer and how efficiently those variants are meeting the customer demand. It makes use of preference modeling, set theory, parametric design theory, and optimization. Given the necessary market information and product design information it will compare the efficacy of different numbers of variants and optimize the variants to maximize the potential market. This enables a company to more effectively meet their customers' demands while minimizing their product variant inventory. This method is illustrated through two examples.

1.1 Background

There are different parameter sets that can be used to describe a product. Of these many parameters, there are a few that are generally of interest to customers. This set of "performance parameters" align with customer requirements and will be used throughout this thesis in the quantification of customer satisfaction. Other sets (e.g., constant parameters, constraints, and design variables) are used as well during the design process to specify the product. By specifying new values for the design variables, the designer is able to generate new design variants. These values can be determined in a number of ways, including the experience of the designer or through the use of optimization. The effectiveness of a set, or family, of variants can be assessed by determining how well it fulfills the customer requirements.

A product's design space is the set of all possible variations of the design of that product. It is defined using design variables rather than performance parameters in order to handle infeasible regions of space and other design constraints.

Because market demand is a function of performance parameters, another space is needed to analyze and visualize how well a set of variants satisfies the market demand. De-

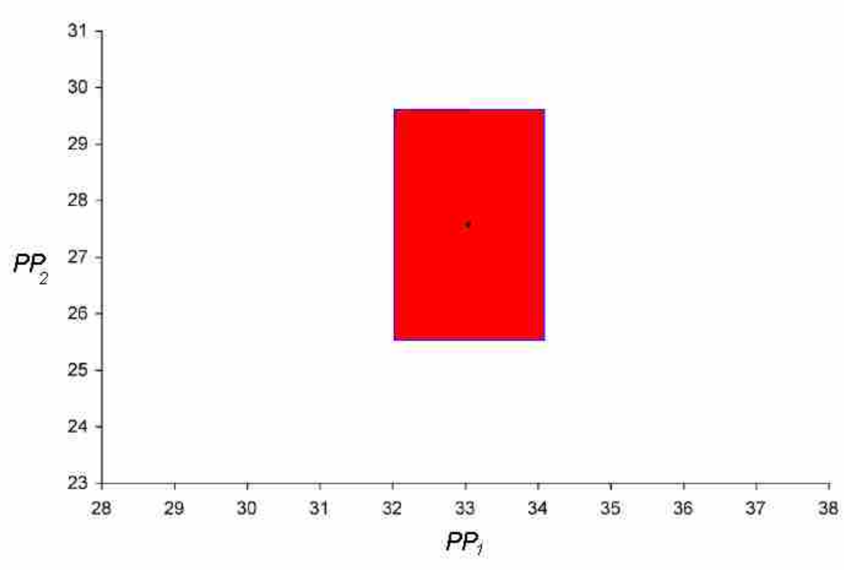

Figure 1.1: The area spanned by a product variant due to its adjustability

fined in terms of performance parameters, this space is called the Product Offering Design Space (PODS).

Every point in the PODS represents a set of customer requirements. A product variant that satisfies the customer requirements of a region of the PODS is said to span that region. For discrete variables this is rather simple: a vehicle could have 4, 5, or 6 tires, but it can't have 4.36 tires. Therefore a variant that has four tires spans all customers requiring 4 tires. It does not span customers requiring 3 or 5 tires. If that space is continuous, then because a variant is represented as a point in the PODS, no number of variants would be enough to span it. Fortunately the ability of many products to be modified or fine-tuned by the customer will allow a single product variant to span more than just the point of the PODS where it is located. Examples of this include adjustments to seats in cars, multiple holes in belts, and the adjustable width of the mouth on a crescent wrench. This inherent ability of a product to be adjusted is termed adjustability and is shown in Figure 1.1

The quantification of any metric can only be done to a certain degree of accuracy due to the resolution of the instrument used. Some instruments inherently have a finer degree of resolution than others. Because of this, a product variant only needs to meet the customer requirements with the accuracy discernible by the customer. If the customer is only able to measure to 0.001 mm, then there is no discernible difference between a variant that measures 0.00098 and one that measures 0.0011. Additionally, a customer may be able to measure something to a higher resolution than the customer is actually interested in. The customer may be able to measure to 0.001 mm, but if the customer only cares that it be accurate to 0.01 mm, then that takes precedence over the resolution of the instrument and a product that measured 0.0098 would be considered good enough. This resolution, either in the instrument or in the mind of the customer, is used to determine the increment that will be used for each customer requirement.

Though there are always exceptions, customer interest in a product is generally focused in a specific region of the PODS. Most customers don't want a monster truck nor a clown car, but rather prefer a size of car somewhere in between. In order to focus on this area of interest, a range to be explored is specified for each customer requirement. The region of the PODS delimited by these ranges is called the Product Offering Domain (POD).

The POD represents the region of the PODS in which a product variant might be desirable to a customer. However, not all variants within the POD are equally likely to satisfy a customer. Also, it is more important to offer a variant in a region of the POD that is of interest to many customers than it is for a region that only satisfies a few. This market interest is captured in a Preference Model as a function of the customer requirements. This function represents customer preference as a surface in the POD. For example, a product that has two customer requirements will have a two-dimensional POD. The Preference Model is a surface in the third dimension above the area identified by the POD (Figure 1.2) The volume under this surface represents the total market demand for the product.

1.2 Significance of Research

As customer needs vary, so do the desired product behavior and specifications. If these requirements were plotted on a graph, a single customer's requirement would be represented by a point on that graph. Taken as a whole, customer requirements can be represented with a continuous design space. Customer demand within that space varies and can be represented by a probability distribution. This facilitates the unequal weighting

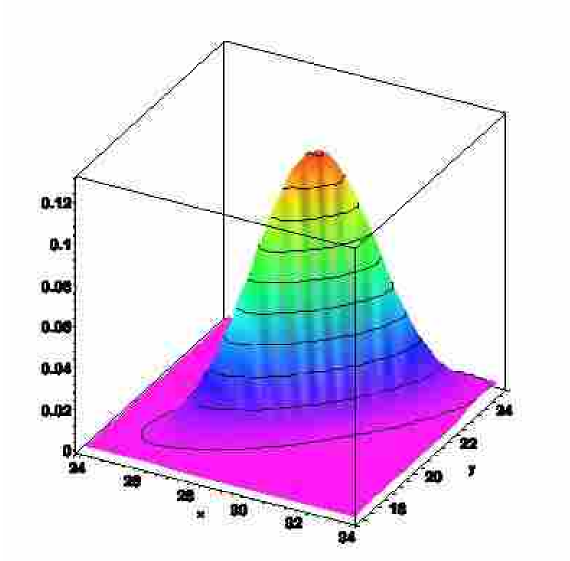

Figure 1.2: A preference model applied to a product defined by two parameters. In this case the preference model is a bivariate normal distribution.

of points within the design space—a point that satisfies many customers can be considered differently than a point that only satisfies a few.

It is often desirable to span a continuous PODS with a discrete set of product variants. Consider the problem faced by bicycle manufacturers: people from every segment of society ride bikes. These people come in a continuous array of sizes and proportions and have different riding position preferences. A similarly continous array of bike frame sizes is required in order to perfectly fit each of these potential customers. However, very few of these customers can tell if a bike fits them perfectly and are willing to sacrifice a little bit of that perfect fit in order to lower the final cost of the bike. Bike manufacturers are aware of this and, as a result, generally produce three or four frame sizes spaced evenly apart. For example, Trek's top of the line mountain bike, the Fuel 100, comes in 15.5", 17.5", 19.5" and 21.5" sizes. These four frames are expected to meet the needs of the vast majority of their customers. When a new bike model is being developed, the questions arise: "Over what range should the bike be offered?" "At what intervals within that range should the bike be offered?" "Is an even distribution the best way to satisfy our customers?" "Is our company's objective to satisfy as many customers as possible, or to please a smaller customer base but improve the experience for each of those customers?" Increasing the

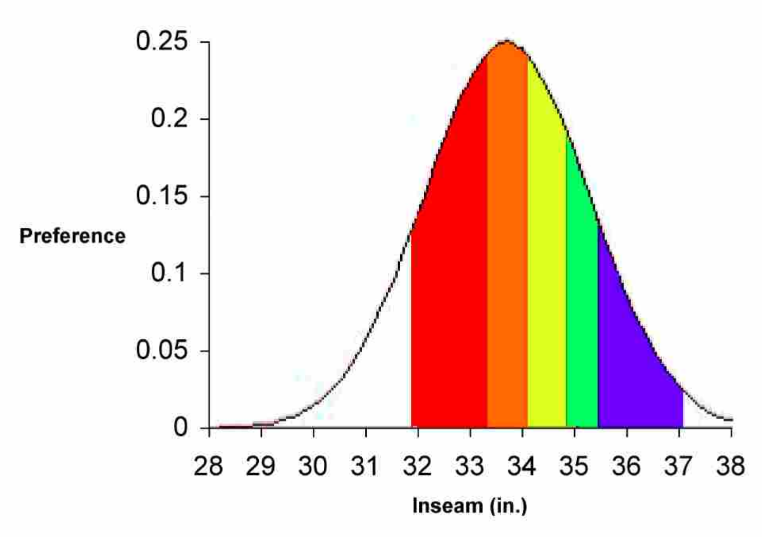

Figure 1.3: A set of three overlapping product variants.

number of sizes offered decreases the amount of adjustment necessary and improves the fit for the consumer. However, it also increases the costs associated with manufacturing and shipping and increases the burden on the retailers who must stock and store the various sizes.

Once the number of sizes to offer has been decided, the company must select which sizes to offer. If the sizes are too close together they may overlap, which provides no additional value (Figure 1.3). If they are too far apart, then gaps are left where potential customers will not be able to find a bike that fits them (Figure 1.4). The choice of sizes can be optimized to satisfy the highest percentage of the market demand (Figure 1.5).

A similar scenario occurs among companies producing custom products. Many companies produce customized products per customer specifications in order to fill a single order. While this approach meets customer requirements, it can lead to enormous inventories or catalogues of parts and products due to legacies of individual custom projects and higher production costs. These design inventories are expensive to maintain and are often so cumbersome to navigate that it offsets the benefit of reusing an existing design. Rel-

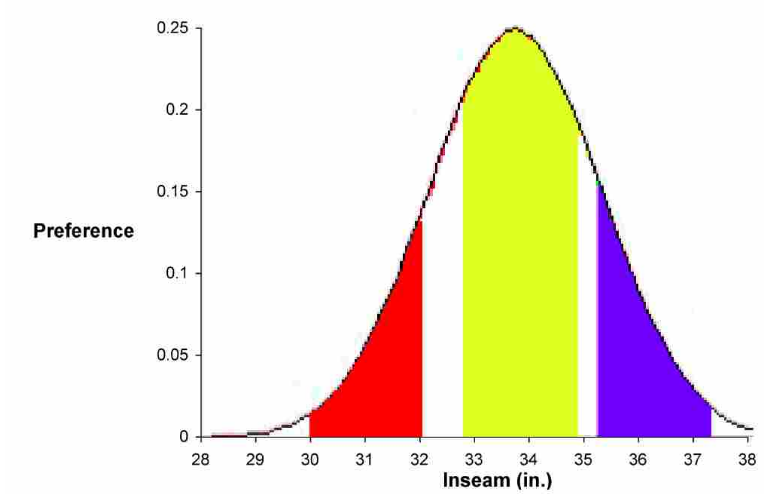

Figure 1.4: Gaps, or regions of the POD between product variants that are not spanned by any of the variants.

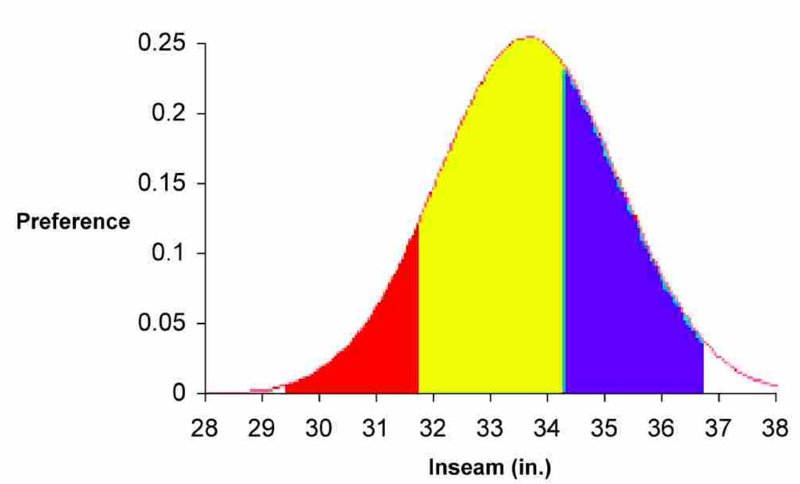

Figure 1.5: Product variants that are well distributed.

evant questions are: Which legacy product variants should be removed? Where should a new product be offered?

This thesis presents one answer to those questions. It uses preference modeling and optimization to identify gaps and redundancies in a product offering design space, and to identify the set of designs that will best span the POD.

1.3 Research Limitations

This thesis makes no attempt to determine the needs of a customer or of a market in general, or to develop a method of visually representing the design space. The developed method will only consider a single product family or technology at a time. Also, it is assumed that the models used accurately represent the product, cost, and customer needs. There is no differentiation between customer requirements that are satisfied by a product with adjustability and customer requirements that are satisfied by a tailor-made product that fits with no adjustments.

1.4 Outline of Thesis

This thesis presents a method of determining a spanning set of product variants. It also lays the groundwork for research into several other proposed approaches to this problem. The remainder of this thesis is organized as follows. Chapter 2 summarizes related literature and provides background information concerning the product offering design space (PODS), set theory, and the statistics involved. Chapter 3 explains the methodology and illustrates it using a generic example. Chapter 4 applies this methodology to a line of mountain bikes. Chapter 5 applies the methodology to a flow control valve example. Chapter 6 summarizes the work completed and suggests possible directions for future work.

Chapter 2

Background and Literature Review

This chapter presents an overview of the research pertaining to this thesis. Areas of interest are preference modeling, methods of representing design space, and product families.

2.1 Preference Modeling

Öztürk et al. [1] provide a good explanation of the basic concepts involved in preference modeling. They cite examples of preference modeling being used in the fields of economy, sociology, psychology, political science, artificial intelligence, computer science, temporal logic and the interval satisfiability problem, mathematical programming, electronic business, medicine and biology, and archaeology as well as decision analysis. Various types of preference are discussed including "better than", "bigger than," and order preference. They also discuss various ways of measuring superiority and the concept of an object being "near" another in preference.

There are a variety of different methods of developing a preference model. Herrera-Viedma et al. [2] present what they terms a consensus model. This model captures differences between individual opinions and group opinions. Petiot and Grognet [3]state that often subjective assessments are as important to a customer as technical functions. They present a method to capture the subjective portion of preference and use this information to better position future products in a high preference area.

Benferhat et al. [4] develop what they term a bipolar view in preference modeling. This categorizes preference into positive and negative preferences, both of which can vary by degree. It handles preferences separately which can lead to inconsistencies between the two when they are brought together.

Michalek et al. [5] address the problem of developing marketing and engineering design goals in isolation from each other. They apply analytical target cascading to coordinate both disciplines in order to determine an optimum solution to both. This method is applied to the design of analog household scales. The results show that the optimum product generally falls short of marketing predictions but surpasses what engineering would have developed alone.

For the purposes of this thesis, the preference model is a summation of binary preferences. A customer either does or does not want the product at a given point in the design space. The preference model is equal to the scaled sum of customers that would buy a product variant at that point in the design space. It does not take into account the degree to which the customer desires the variant at that point. This could certainly be done in future work by applying a bias function to the preference model.

2.2 Design Space Representation and Exploration

In order to identify a spanning set of designs, a model of the design space must first be developed. Many methods exist for exploring design space and for handling the constraints and regions of infeasibility. Bates and Wynn [6] model feasible regions as well as proximity to infeasible regions which may impact the robustness of the design. They accomplish this using Hilbert bases and defining positive definite kernels on hyperspheres. Kuchcinski [7] uses a method called constraint logic programming to explore design space while still satisfying constraints. This method can be used in common optimization routines. Miramond and Delosme [8] focus on efficiency in exploring design space. During the exploration or optimization process, this process simultaneously explores the available computer resources and assigns tasks in such a way as to maximize the usage of system resources.

Siddique and Rosen [9] research representations of combinatorial design spaces and the effects of constraints on these spaces. They also examine how to represent feasible regions and how to determine the size of these feasible regions. The methodology is illustrated with a coffee maker product family example. Jiang and Yan [10] also examine constraints from a product family (building block) point of view. Corbett and Rosen [11] demonstrate a method of finding feasible regions of design space in platform commonization by using discrete design spaces. Their approach focuses on "defining configuration design spaces for engineering design, with an emphasis on...their combinations into larger spaces that more completely capture design requirements." It summarizes design spaces that model physical connectivity, functionality, and assembly considerations. They introduce a new design space that models flows among components. All of these bodies of work provide methods of handling constraints and infeasible space that will likely exist in the design examples used in this thesis.

A Product Offering Domain Space (PODS) as defined by Hunsaker [12] is a three dimensional box bounded by the minimum and maximum levels of performance, cost, and quality. These metrics were chosen in order to best estimate the capacity to make a product (pg.9). The PODS provides a representation of the desired and actual envelopes. The end goal was to compare the demand for a product with a company's ability to make that product. A similar approach will be used in this thesis. The PODS research was focused on visualizing the design space and therefore limited to three dimensional space. However, for this work the visualization is less important than the ability to analyze a product's complete performance. Therefore, multiple performance metrics will be used to define the design space as determined by the customers' specifications. These metrics may include cost and quality as well if they are determined to be of interest to the customer and to be beneficial to capture in the model.

2.3 Product Families

Various attempts have been made by companies to cut costs in the development of products with similarities. Product families came about as a result of these attempts. The most commonly discussed method of developing a family of products is a modular approach. A family of products often share common components thus cutting manufacturing, inventory, and maintenance costs. These shared components are called modules and together they make up a product platform. The products created from a product platform make up a product family. Many research projects have explored different facets of the design and usage of product families.

Simpson [13] provides an extensive overview of the previous research on product platform theory and design. He discusses the advantages and disadvantages of using product platforms, their applications, various methods of developing platforms and designing with reuse in mind. He then discusses product family topics that would likely benefit from future work. Jose and Tollenaere [14] provide a more recent survey of research in product families. This work is more focused on efficient product family design.

Nanda et al. [15] attempt to formalize the product family definition and development process by introducing the Product Family Ontology Development Methodology (PFODM). They then apply this ontological approach to disposable cameras. Complimenting this ontological approach, Nanda et al. [16] also propose a method of capturing and storing this information.

There are several different optimization techniques that have been applied to the problem in order to evaluate a product platform. Thevenot et al. [17] use a genetic algorithm to asses the level of commonality within a product family in order to select the modules to be used in the product platform. Rai and Allada [18] use a multi-agent framework in a multi-objective optimization in order to determine the Pareto-design solutions for a set of modules. Wang et al. [19] use a simulated annealing algorithm in order to maximize the similarities among components.

The use of product platforms has many benefits, but it can also cause a product's performance to be altered. Nelson et al. [20] focus on the effect of product platforms on product performance. They use multicriteria optimization techniques to analyze the performance of two nail gun products using a product platform as compared to their performance when designed independently. This facilitates a decision as to whether or not a product platform should be used for several products, or if they should be developed independently.

The majority of research has been focused on the development and improvement of product families. A minority of researchers have focused instead on a methodology to optimize the product variants used to span market segments. Nidamarthi et al. [21] illustrate a method of optimizing a set of variants on performance versus cost. Michalek [22] develops another performance versus cost method that makes use of preference modeling. De Weck et al. [23] demonstrate a method to optimize the number of product families and the number of variants within each family in order to find the maximum profit. This method defines products in terms of a few key performance parameters rather than in terms of their design variables. By mapping the industry-wide cost and sales for these products it is able to determine the optimum number of families and variants to offer. They apply this method to an automotive product platform. De Weck et al. discuss the number of product families and variants to offer; however, they do not discuss their selection. This thesis will use this approach of using the performance parameters for the product definition. It is only applied within individual product families rather than across multiple families.

Chapter 3

Methodology

This chapter outlines the methodology used to define and evaluate different sets of product variants. A product definition is developed based on the performance parameters and another definition is developed using the product's design variables. The concept of adjustability is described and applied to each parameter. The incremental value to be used along each parameter is selected. The product offering domain (POD) applies a range for each parameter and represents all the possible product variants to be examined.

There are several desirable characteristics that a method of evaluating the spanning sets of product variants should have. It should take customer requirements as an input so that customer preference data can be used. It should also take into account regions of design space that are infeasible, such as locations where designs can not be created because of physical, financial, or manufacturing constraints. The method should be applicable to a wide array of products. The application of this method should result in a metric that can be used to quantify how well a set of product variants spans the potential market. There should be a way to visualize these results so that they can be easily understood and presented. Additional metrics should evaluate the missed opportunity of potential market that was not spanned, as well as market that is overlapped by multiple product variants. The proposed methodology incorporates all of these characteristics. An example using generic parameters will be given in order to fully illustrate this methodology.

3.1 Product Definition

A product's design space is a representation of possible single-point designs of that product. The process of spanning a design space begins with developing a definition for it. Several possible methods of defining a product were considered. Two parameter-based methods were found to be useful for this methodology. The first method defines the product in terms of parameters that describe its performance characteristics. These are generally the parameters that are of interest to the customer and are termed *performance parameters*. The second method defines the product in terms of a set of *design variables* that, along with constants and relationships between the variables, fully constrain the design of the product. For a given set of design variables there is only one possible product variant that can result. The first step in developing a definition of the product is identifying both the performance parameters and the design variables. Once these have been selected a range and the amount of adjustability are identified for each parameter. Lastly, the increment to be used for spacing candidate product variants is identified.

3.1.1 Performance Parameters

Customers shopping for a bicycle are more likely interested in how well the bike fits than in the geometry that results in that fit. Parameters of interest to that customer would include the inseam size for which the bike is a good fit. For an aircraft engine the parameters would include thrust and weight. These parameters describe the performance or attributes of the product and are called performance parameters. The performance parameter definition of the product consists of the parameters specified by the customer or the parameters used by the customer to make a purchase decision. When the spanning is calculated the customer requirements will be used to represent the set of product variants.

For example, Product A has three performance parameters that are used by customers when spec-ing the product or making purchasing decisions. These are PP_1, PP_2, and PP_3. These will be used as the example continues in the following sections.

3.1.2 Design Variables

During the design process, feasible design space is typically defined, as described earlier, by the parameters, constraints, and design variables. This definition facilitates design decisions and the application of computer-aided engineering (CAE) tools like solid modeling, optimization, and finite-elements analysis, which are used to calculate the feasible design space. Mapping the performance parameters to the design variables can be non-trivial and it is possible that the performance parameters of a candidate variant appear feasible when the design is not. There is some inherent understanding of feasibility in the performance parameters, but not an exact mathematical one. For example, a customer generally knows not to ask for a Hummer that gets 35 miles to the gallon. However, that customer doesn't know exactly what mileage the company is capable of designing the Hummer to get. Also, a customer may specify multiple performance parameters that map to only one driving parameter. Based solely on these parameters there is no way to determine if the design is over or under-constrained. Additionally, a given set of performance parameters can be met with a variety of different design configurations.

For these reasons, the set of performance parameters cannot be used as the design definition of the product. For example, specifying the inseam length of the intended customer does not provide adequate information to define the height of the bike. Other information is needed relating the inseam length to the bike height. Additionally, it is possible to satisfy the inseam length performance metric with various different bicycle configurations. Therefore, in order to adequately define the bicycle, a set of design variables must be used. Design variables are the parameters necessary to fully constrain the product. Relations between the performance parameters and the design variables will be used later to map information back and forth between the two. The design variable portion of the product definition consists of the driving parameters, or the parameters that are required in order to generate all the parameters necessary to create the product. This definition of the product is used in a black box to generate potential product variants, which will then be mapped back to the performance parameters in order to calculate their spanning fitness.

Each parameter has a range over which it is feasible. Inside this range there may be pockets of infeasibility. This could be due to constraints in manufacturing or physical limitations in the actual geometry of the part. These regions of infeasibility may be a function of the parameter itself or of any of the other parameters used in the definition. The range of feasible values and any regions of infeasibility are identified and will be used in the POD definition and spanning calculation.

Continuing the example, Product A is defined by two design variables, DV_1 and DV_2. By reversing the functions used to predict the performance parameters, it is determined that DV_1 is a function of PP_1 and PP_2. DV_2 is a function of PP_3 only (Equation 3.1,3.2).

$$DV_1 = F(PP_1, PP_2) \tag{3.1}$$

$$DV_2 = G(PP_3) \tag{3.2}$$

3.2 Product Offering Domain

In order to measure how well a set of variants spans a region of design space, that region must first be defined. The Product Offering Design Space (PODS) can be considered to extend off to infinity along each axis. This methodology is only concerned with the region of the PODS deemed significant to the market. Because this region represents the domain over which a product will be offered, it is termed the Product Offering Domain (POD). Many aspects of the product definitions will be used to build the POD.

The first step in creating the POD is to define its axes. Each performance parameter is used for an axis. An additional axis represents the customer's preference, and will be used when the preference model is applied. Therefore, a product defined by two performance parameters will require a three-dimensional POD. Using the Product A example, PP_1 and PP_2 become the 'x' and 'y' axes in cartesian space. The third axis will be used by the preference function.

A two dimensional product definition can easily be represented in cartesian space.

3.2.1 Range of Offering

The next step in creating the POD is to define its bounds along each of the performance parameter axes. These bounds form the region in the design space in which the company is interested in creating a product line. Values for these bounds can be found through market research or, if the company is modeling an existing product line, from analyzing that product's portfolio. A company's product portfolio consists of the many legacy designs that are still offered. Each of these designs can be categorized by the performance parameters customers use when selecting from among them in order to determine the currently offered range. Market research can show the range that potential customers would be interested in. However, it is not always advisable to incorporate all of the design space in which a customer may desire a product. Using the truck example, a customer may be okay with getting 14 mpg, but would be more than happy to have a truck get 30 mpg. In order to design such a truck, the company would have to sacrifice size, weight, and horse power. Therefore, rather than waste iterations on designs that are infeasible or that do not satisfy the target market, the company may choose to set the range maximum lower than the (in this case infinite) value a customer might be interested in.

A range must be specified for each of the performance parameters. For example, a customer purchasing an engine may do so based on its speed, torque, and weight. These are the performance parameters for this product line. Because several models are offered within the line, there exists some range of performance capability exhibited by the legacy designs. If the weights of the five engines in the product line are 5, 15, 17, 22, and 40 kg, the range is from 5 to 40 kg. Using set notation the offerings can be represented as: A = $\{5,15,17,22,40\}$ and the range as A_range = [5,40]. Because these ranges are applied to the performance parameters, if possible they are also converted to ranges for the design variables in order to focus the perturbations in search of candidate designs. Alternatively, for an existing product line the portfolio of legacy designs can be analyzed in terms of the product's design variables, in which case no conversion is necessary. If the company is launching a new product, the range is selected based on market research and the business strategy of the company.

Manufacturing feasibility must also play a large role in the selection of these ranges. Each design variable should be looked at separately from the other variables to ensure that its range is feasible. There are likely to be interactions between the different design variables causing infeasible regions in the design space. These are not allowed to alter the ranges of the design variables and will be handled during the spanning calculation phase. In other words, the range for each parameter should incorporate the largest and smallest possible values for that parameter, regardless of the infeasible regions of space in between. For Product A, a range is chosen for DV_1 from 25 to 45 and for DV_2 from 180 to 320. It is noted that there is a region of infeasible space from 30 to 34 for DV_1 for any value of DV_2. This will be used in a later step.

$$DV_1 range = [25, 45]$$
$$DV_2 range = [180, 320]$$

Now that ranges have been found for each axis, the POD is defined by the intersection of these ranges. Therefore, a design problem with two performance parameters will create a rectangular shaped POD before the application of the preference function.

3.2.2 Adjustability

Once the performance parameters and design variables have been determined, the amount of adjustability is identified for each performance parameter.

Each product variant is represented by a set of performance parameters defining a point within the POD. Each variant is able to satisfy the customers who desire a product with those same performance specifications. However, since the performance parameters and customer requirements must be exactly the same for the customer to be strictly satisfied, taken together these discrete points are unable to span all or even a significant portion of the POD. In order to create a spanning set of designs a new way of looking at design instantiations must be developed that allows for the possibility of satisfying a customer with a product that does not exactly meet the performance specifications they requested. Rather than design and manufacture custom products tailored to perfectly fit each customer, many companies design products so they can be adjusted after-market to better fit the customer. This adjustability can be achieved through modular components or through modifiable products. Modular adjustability generally comes in the form of components swapped in and out by the company at the customer's request (i.e. Dell), or in the form of after-market modifications made by the customer. Modifiable products have adjustability designed into them. Examples of this include adjustable seat height on a bike to accommodate various stand-over heights and adjustable stem lengths, and fore/aft seat position to accommodate various torso and arm length ratios. This methodology is only concerned with adjustability derived from the ability of a product produced at a specific point in design space to please customers within a specified range of that point. Therefore, only the second form of adjustability will be addressed.

The amount of adjustability is determined for each performance parameter. This value could be zero, it could be a constant, or it could be a function of the parameter itself and any of the other parameters. In the example of Product A, it is found that PP_1 is able to adjust +2 and -1. PP_2 has no adjustability in either direction (Table 3.1).

3.2.3 Parameter Increments

It is uncommon to see a car advertised as having a 2.19 liter engine, with two decimal places of accuracy. However, it is quite common to see a car marketed as a 2.2 liter or 3.0 liter engine, with one decimal place of accuracy. This is because a decision has been made by the company to not offer, or at least to not market, an increment smaller than one decimal place. This could be due to limits in testing, measuring or manufacturing accuracy. The measurement process may be unable to accurately measure smaller than .1 liters (unlikely in this example) or it may not be cost effective to ensure a manufacturing tolerance of less than .1 liters (also unlikely). Using the automobile engine example, it is more likely that this one decimal place increment is due to limits in the perception or interest of the customer. A customer simply does not care about a .01 or even a .05 difference in engine size. The automobile manufacturers are aware of this and therefore use a one decimal place increment to market their engine sizes, though it is quite likely that the actual size of the engine is measured more accurately. For this methodology, the increment to be used for each performance parameter is identified. In order to identify the increment to be used, both the market perception and the manufacturing increments are analyzed.

Table 3.1: Adjustability of each performance parameter

	$+$	$-$
PP_1	2	-1
PP_2	0	0

The perception of the market relates to the performance parameters rather than the design variables. Testing often applies to performance parameters as well. Since the candidate product variants are generated using design variables, any increment derived from the market perception and applied to the performance parameters is also converted and applied to the design variables. This will determine the size of the perturbations applied to each of the design variables when exploring the design space for candidate variants. Continuing the example, it is found that the market only perceives a difference of 10 for PP_1, 2 for PP_2, and 200 for PP_3. Testing on PP_2 is only accurate to 1. Because the market perceivable difference is larger than the testing increments (which is usually the case), its value will be used for PP_2. These performance parameters are converted to values for the design variables which are 1 for DV_1 and 10 for DV_2 (Table 3.2).

Limits in manufacturing are generally related directly to the design variables. Therefore, the parameter increments resulting from these limits can easily be applied to the design variables but will need to be converted and applied to the performance parameters. Additionally, any limits in testing are also applied and mapped accordingly. However, as with market perception, testing is rarely the limiting factor. For Product A, limits in manufacturing result in an increment of 2 for DV_1 and 1 for DV_2. No testing is done on the design variables for Product A.

Each design variable is analyzed independently using both methods. The method resulting in the larger of the two increments is used. These parameter increments will be used in the implementation of the spanning calculation. From the Product A example, for $DV_1$the limit on manufacturing yields a larger increment of 2 and will therefore be used. For DV_2 the market perception yields the larger increment of 10 and will be used.

DV_1 Increment: 2

DV_2 Increment: 10

Table 3.2: Mapping performance parameters to design variables

PP_1	10	DV_1	1
PP_2	2		
PP_3	200	DV_2	10

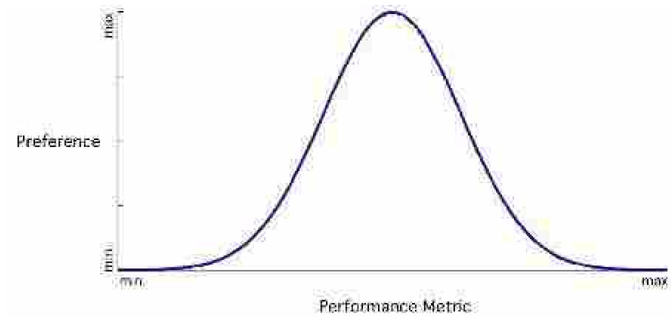

Figure 3.1: The market demand model, applied to a product defined by a single performance parameter.

3.2.4 Market Demand Model

The next step is to develop a model of the market demand as a function of the performance parameters. This is termed the Market Demand Model and it finally makes use of the extra dimension in the POD. Because it represents the influence of the market and will be used to calculate the market opportunity that can be spanned by a design, it is defined in terms of the performance parameters.

Thus far the performance parameters of interest have been identified and a POD has been chosen. However, within that domain there will be regions of greater or lesser market demand. The market demand model is a function of the performance metrics. This model weights the POD according to market demand and gives some understanding of how well a product at a specified location in the design space will be received (Figure 3.1).

If this is being applied to an existing product, the legacy designs can be used to develop this function. Each legacy design is plotted in the design space. For a problem with two performance metrics, these are plotted on a planar graph (Figure 3.2). Regions of space with more designs are assumed to be more highly desirable and those with fewer designs are conversely less desirable. The preference is calculated by using the parameter increments to create a grid and summing the number of legacy designs within each cell of the grid. The results are then fitted with a curve (Figure 3.3).

If there are no legacy designs or if the legacy designs are assumed to not accurately represent the market demand, this function needs to come from marketing data or from customer requests. For the purpose of this thesis, it is assumed that the data needed to de-

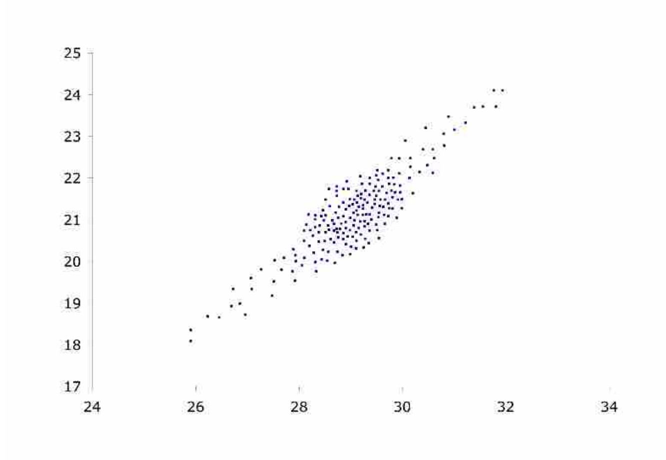

Figure 3.2: Legacy designs plotted on an xy plane. More points indicates a higher market preference.

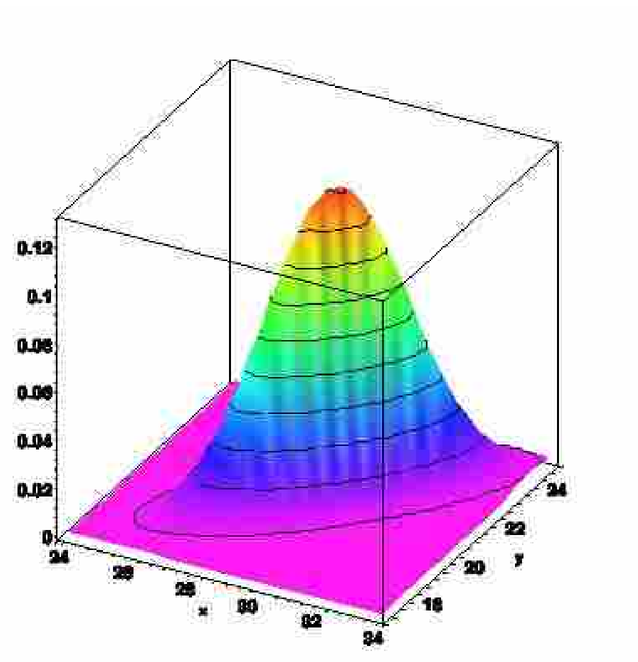

Figure 3.3: A preference model created from the legacy designs from Figure 3.2.

velop the market demand model already exist. Several methods are available to develop the market demand model from this information. Because the preference function represents market demand, it can be directly applied to the performance parameters. The first case study will do this using statistical models from market research. The second case study will cut up the POD into cells, sum the number of legacy designs in each cell and thereby get a value for the popularity of each cell. By applying either of these methods to a 2D problem, a contour plot is generated with peaks in the 3rd dimension.

3.3 Spanning Analysis

Thus far the performance parameters and design variables are identified and ranges and increments selected for each, and a model of the market demand has been applied to that region of the design space. All these components together make up the POD. With the POD defined, it is now time to analyze sets of product variants and find an optimal set. A preliminary analysis will first be conducted, followed by calculations for the total market spanned, the amount overlap, and the minimum, maximum, and average percentage of the total market spanned by each variant.

3.3.1 Preliminary Analysis

If desired, a preliminary analysis can be performed on each of the parameters in order to have a rough understanding of the percent spanned along one axis of the POD by a set of designs. This is a one dimensional calculation and is performed for each parameter individually. It is calculated for each parameter by summing the spans of each product variant in the candidate set, taking care to take into consideration any overlap, and dividing it by the total range of that parameter. Equation 3.3 shows this calculation for PP_1 for a given set of designs.

This approach is simple, yet yields information that can be helpful to test existing sets of variants or to do a cursory test to yield a group of candidate sets. However, its usefulness is limited because it ignores any relationships between the parameters, assumes a uniform market demand across the entire range, and does not subtract out any pockets of infeasible values for the parameter. The results of this analysis for the example are shown in Table 3.3.

$$\%PP_1Spanned = \frac{29-26+35-30+44-36}{45-25} = 0.8 \tag{3.3}$$

3.3.2 Total Volume: Market Opportunity

Using the simple example of a product with one performance parameter, on a 2D plot the market demand model is represented as a curve. The area under this curve bounded by the POD represents the total area of the POD and the potential market for that product. This is termed the Market Opportunity and is shown in Figure 3.4. For an example with two performance parameters, the market demand model would generate a surface and the volume under it would represent the market opportunity, as shown in Figure 3.5. In order to standardize terms across any dimensionality, the terminology from the 3D example is used. Therefore, for a POD of any dimensionality the market demand model is considered to generate a geometric entity termed a surface, and the space under that will be termed the volume, which represents the market opportunity.

If the market demand model is a mathematical function, the volume underneath it can be calculated by integrating the function with respect to each of the variables. If the model is composed of values at discrete points in the POD, the volume can be calculated by summing the values at each point used to develop the model. The former approach will be used in the first case study and the latter will be used in the second.

3.3.3 Volume Divided into Cells

Using the POD and the market demand model in their current states would allow a variant to be placed anywhere along the continuous design space within the bounds of the POD. This is generally not what a company would want to do. For example, offering an automobile engine with 1.931 liters would not make much sense, from an engineering

Table 3.3: Preliminary analysis shows spans for all performance parameters.

Design#	PP_1	$PP_1 Range$
1	27	26-29
2	31	30-33
3	33	32-35
4	37	36-39
5	40	39-42
6	42	41-44

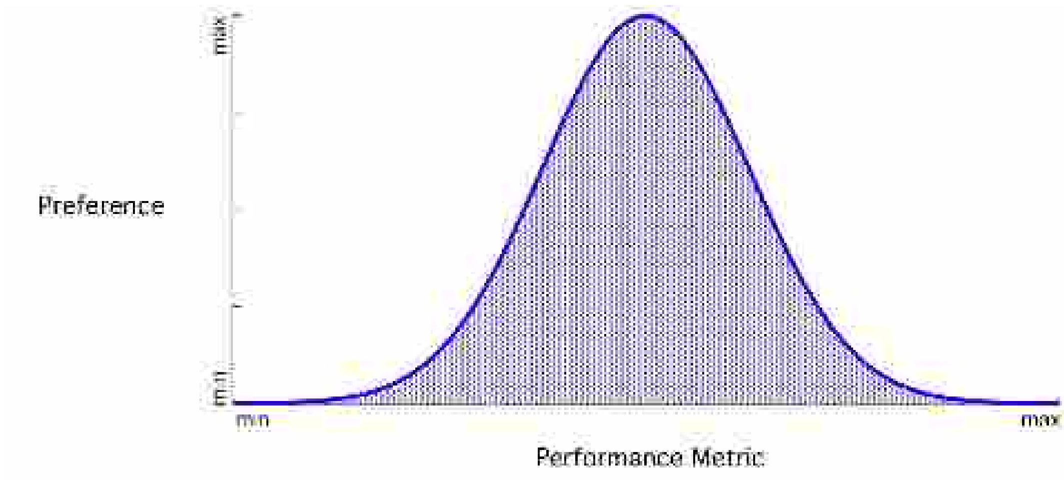

Figure 3.4: The area under the curve represents the market opportunity for a product defined by one performance parameter with a normal distribution preference function

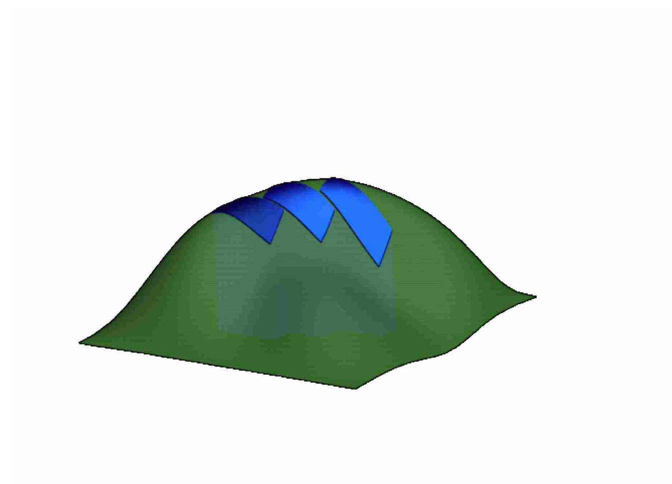

Figure 3.5: The market demand model, shown as a surface. Three product variants are spanning regions in the volume underneath the surface.

or marketing standpoint. The parameter increments developed during the product definition stage will now be applied to the POD. The intersection of these increments in each dimension divide the POD into cells. Each cell represents a possible location for a candidate product variant in the optimization routine. The space covered by a given cell also represents area that would be spanned by a variant at that location.

During the POD definition stage, any infeasible design regions were noted but nothing was done about them. The cells in these infeasible regions cannot simply be discarded because a feasible variant may be able to span into an infeasible region due to its adjustability. Therefore it is necessary to handle all tests for feasibility when generating the candidate variants using the design variables. If a specific variant is infeasible it should not be considered a candidate and should not be evaluated in the spanning calculations.

3.3.4 Spanning Due to Adjustability

Thus far, the ability for a candidate design to span a region of design space is controlled entirely by the size of the cell, which is a function of the parameter increments. As discussed earlier, some parameters also have an inherent amount of adjustability that allows them to span a wider range of space. The adjustability identified in the product definition stage is now applied to each parameter. If this adjustability is larger than the respective increment for that parameter, it will allow a single design to span multiple cells.

3.4 Effectiveness of Spanning Set

A basic analysis of how well a set of designs spans a POD has already been shown. As discussed, that method has several drawbacks which necessitate a more accurate calculation. This section presents several metrics that, when used together, give a much clearer understanding of the effectiveness of a candidate set of designs.

3.4.1 Total Market Spanned

The most obvious and possibly most important metric is the Total Market Spanned. This metric is the union of all the cells spanned by a set of designs and is given as a percentage of the market opportunity that is spanned. It is calculated by taking each candidate in the set and marking where in the design space that candidate is located. Then the cells that are spanned by that design are flagged. This is done for each of the candidates. The total market spanned is the sum of the volumes of all the flagged cells, divided by the total market opportunity. Using this approach ensures that cells which are spanned by more than one design (overlap) are only counted once.

3.4.2 Overlap

The percentage of the total market spanned gives a good idea of if there are enough designs being offered. In addition to this it would be helpful to know if there are too many designs or if the designs are being used inefficiently. The overlap metric does this. It is calculated at the same time as the previous metric. When marking the spanned cells, each design is compared with each cell and if the cell is spanned, its volume is added to the total volume spanned by that design. Once all of the designs have been analyzed, the volume spanned by each design is summed. This volume includes any volume overlapped by multiple designs. This volume divided by the total market spanned gives the overlap metric. A value of 1 means there is no overlap. A value of 2 means that on average any region of the POD spanned by a design is spanned by two designs. Because of the way it is calculated, it can never be less than 1.

3.4.3 Min, Max, Average Cell Volume

Several smaller calculations can be performed if they are deemed beneficial to the company. The minimum, maximum, and average volumes spanned by each of the designs can all be beneficial. Each of these may include some overlap and should therefore be used carefully. Comparing the minimum to the average volume spanned by an individual design can indicate the total percentage spanned that would be lost if one fewer designs were used in the set. Similarly, comparing the maximum to the average can indicate how much of the total is being spanned by the most significant design. Additionally, it can be beneficial to look at the statistics for the individual cells. Comparing the minimum and maximum cell volumes with the total can provide information about the boundaries used in the POD definition.

3.5 Cost vs. Coverage Trade-off

One of the most useful applications of these tools is to perform a trade-off analysis between maximizing the coverage of the potential market and minimizing the cost. In the first case study (Chapter 4) there is no information available for manufacturing costs, so this analysis was not performed. In the second case study (Chapter 5) some rough cost estimates were available. These were used to analyze how the cost is affected when the number of variants offered is minimized and their distribution is optimized.

3.6 Summary

This chapter presented the methodology that will be used in the examples in the following two chapters. A rudimentary example was used to illustrate the concepts. The following chapter uses a real world example of a mountain bike product family which will further explain how the method is applied and how real world obstacles are overcome.

Chapter 4

Case Study: Mountain Bike

Mountain biking began as a fringe sport and experienced much of its early growth in Northern California. There the first documented competitions were held in 1976. The first company to exclusively manufacture mountain bikes started in 1979. In the twenty-seven years since then mountain biking has experienced astonishing growth and as of 1996 is even an Olympic event. The design of these bikes has changed radically since their inception, and most companies have spent much of their research and development efforts in this area. However, during that time the market for mountain bikes has changed as well. The mountain bike has long since overtaken the road bike as the do-it-all bike of the masses. It is used by the young and old of both males and females. However, for many years bike manufacturers did not recognize this trend and continued to design mountain bikes exclusively for males in their 'teens and twenties. The female demographic has different needs from its male counterpart because of the difference in average height, but until the 1990's females were largely overlooked as potential customers. Since then, companies have begun offering a wider range of frame sizes, and many offer an entire line of bikes designed specifically for women.

This chapter applies the methodology developed in the previous chapter to select mountain bike product variants of the Novara Ponderosa FSL mountain bike product family. The company's "fitting guide" will provide the parameters used in the product definition. Population distribution data will be used to represent the market. Their current product offering will be analyzed and discussed. Optimization will be performed in order to improve the product offering and an optimal set of variants will be determined, analyzed and discussed.

4.1 Product Definition

As discussed in the methodology, the product needs to be defined in terms of both performance parameters and design variables. Bicycle design is deceptively complex and can include a variety of performance parameters. Many professional riders are so highly attuned to the way their bike fits that they only use bikes that are custom built for their bodies and riding styles. However, for this case study the problem has been greatly simplified. Assumptions and simplifications have been made in order to match the fitting process used by an average customer buying an off-the-shelf but still competitive mountain bike.

4.1.1 Performance Parameters

When a customer is shopping for a mountain bike, the salesperson uses two performance parameters to determine if a bike fits correctly: the customer's reach and inseam length. Reach is measured as a combination of torso length and arm length. The actual function used to calculate reach can vary from company to company. According to the REI Bike Fit Basics[24], torso and arm length are measured as follows:

> Torso: Seat a coffee-table book against your crotch with the spine parallel to the floor. Measure from the spine of the book to the little V in your throat just above your sternum. This is your torso length. Arm: Find the end of your collarbone. Grip a pencil in your fist and measure from the pencil to the tip of the collarbone. This is your arm length.

$$Reach = \frac{Length_{torso} + Length_{arm}}{2} \tag{4.1}$$

Inseam is a much more familiar measurement, but few people know how to measure it accurately. The inseam, again defined in the REI Bike Fit Basics(Figure 4.1) , is measured as follows:

> ...stand against a wall with your feet about 8" apart. If your shoes are on, take them off. Put the book between your legs, comfortably up against your crotch. Make sure the spine is parallel with the floor. With the tape measure, determine the length from the spine of the book to the floor. This is your inseam measurement.

These definitions and equations yield the two performance parameters that will be used throughout this case study.

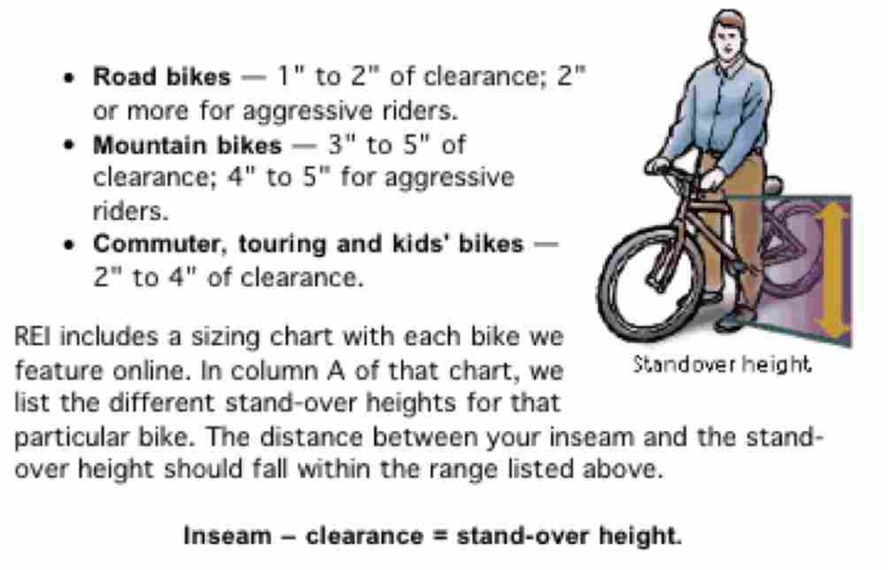

Figure 4.1: REI Bike Fit Basics: stand over height. [24]

4.1.2 Design Variables

The bike is also defined in terms of design variables in order to generate candidate designs that can be analyzed using the performance parameter definition. The design variable product definition has already been done by REI. This definition is not readily accessible to the public. However, by analyzing the geometry of the existing line it is possible to determine some of the design rules used in the definition. These rules will be used to ensure that no infeasible designs are considered as candidate variants. However, it is possible that these "rules" are more flexible than they are being made for this study. This will be taken into consideration during the analysis.

In order to illustrate how performance parameters are mapped to design variables, two variables commonly used in the design of a bike were chosen. The first, stand over height, is directly related to the leg length, and therefore the inseam length (Figure 4.2)[25]. The stand over height is most likely not used as a design variable when designing a new bike. Generally the seat tube length would be used. Stand over height is a function of the seat tube length, the seat tube angle, and the height of the crank set. These latter two are constrained by the desired bike geometry and are generally held constant between different bike sizes of the same model. However, the seat tube length will vary from size to size

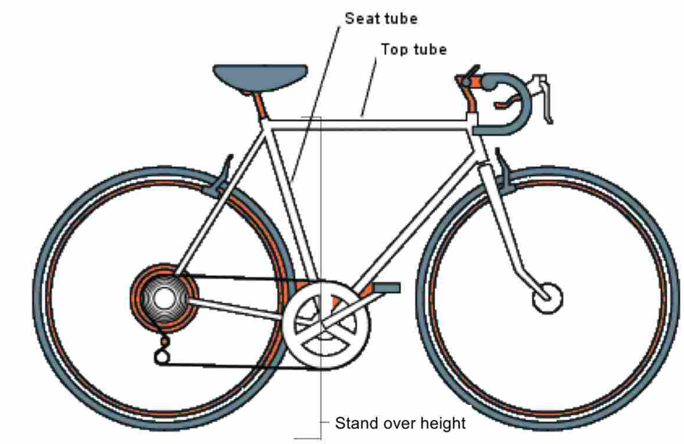

Figure 4.2: Diagram of a bike illustrating the design variables stand over height and top tube length. [25]

(Figure 4.3). This measurement is often used to market bicycles. The 15", 17" and 19" sizes that are commonly offered are referring to the seat tube length. However, these sizes no longer correspond consistently with the seat tube length, especially in higher end, full-suspension bikes. On many of these, the geometry no longer requires a seat tube, but companies persist in using this measurement to market their bikes. Because of the arbitrary nature of the seat tube length measurement, and because the stand over height is a readily available variable that maps directly to the seat tube length, the stand over height will be used as the design variable. The stand over height is given as a function of inseam length in Equation 4.2.

$$height_{standOver} = length_{inseam} - 4\text{"} \tag{4.2}$$

The second design variable to be used is top tube length. This is is a function of reach as shown in Equation 4.3.

$$length_{topTube} = length_{reach} - 6\text{"} \tag{4.3}$$

Novara Mountain Bikes

Model	Ponderosa				Ponderosa FS & FSL		
Size	13"	15"	17"	19"	15"	17"	19"
Seat Tube Center-Top	15.0	17.0	19.0	21.0	19.0	19.0	21.0
Seat Tube Center-Center	13.0	15.0	17.0	19.0	15	17.1	19.0
Top Tube	20.7	21.6	22.2	22.9	21.6	22.2	22.9
Standover	28.6	29.0	30.4	31.7	29	30.4	31.7
Rear Center	16.9	16.9	16.9	16.9	16.9	16.9	16.9
Front Center	23.7	24.7	25.2	25.7	24.8	25.4	25.6
Wheelbase	40.6	41.5	42.0	42.5	41.7	42.3	42.5
Offset	2.0	2.0	2.0	2.0	2.0	2.0	2.0
Head Tube Angle	71	71	71	71	71	71	71
Seat Tube Angle	72	72	72	72	73	73	73

Figure 4.3: Novara Bike Sizing Chart, illustrating which variables could be selected as design variables.

The values used for these equations come from REI Bike Fit Basics. This document specifies a range of values for fitting both reach and inseam. For this case study, the average value is used. The range specified will be used to define adjustability. The 4" clearance used for the inseam is the recommended clearance for a typical rider. More clearance is recommended for more aggressive riders.

4.2 Product Offering Domain

The mountain bike POD exists in a three dimensional space defined by the two performance parameters, *inseam* and *reach* in the x and y directions, and by the market preference in the z direction.

4.2.1 Range of Offering

The range for each performance parameter is obtained from the market distribution data. A minimum and maximum are chosen for each parameter so that most of the potential market is captured within that range. In order to include as much of the market as possible, the far extremes of the population distribution data were used for each parameter.

This results in a range that covers very nearly 100% of the population, but because both parameters are normally distributed, it also results in a range that is much wider than will probably be spanned by the set of designs. These ranges for the performance parameters must also be converted to ranges for the design variables. This is done using the same equations used to initially derive the design variables (Equations 4.3, 4.2). The calculations for both parameters are shown below:

$$\begin{aligned} standoverHeight_{min} &= inseam_{min} - 4 \\ &= 28 - 4 \\ &= 24 \end{aligned}$$

$$\begin{aligned} standoverHeight_{max} &= inseam_{max} - 4 \\ &= 38 - 4 \\ &= 34 \end{aligned}$$

$$\begin{aligned} toptubeLength_{min} &= reach_{min} - 6 \\ &= 23 - 6 \\ &= 17 \end{aligned}$$

$$\begin{aligned} toptubeLength_{max} &= reach_{max} - 6 \\ &= 31 - 6 \\ &= 25 \end{aligned}$$

4.2.2 Adjustability

As stated previously the values for adjustability come from the ranges specified in REI Fit Basics for reach and inseam. It specifies ”3” to 5” of clearance” for the stand-over height for a typical rider and 4” to 5” for a more aggressive rider. The value for the typical rider is used which results in adjustability of ± 1” that is applied to the $height_{standOver}$ design variable. This is a fair assumption because the seat post can be raised or lowered in order to fine tune the distance between the seat and the pedals to fit a specific customer. This maps to an adjustability of ± 1” for the *inseam* performance parameter.

The fitting guide also specifies that “you can go about 2” longer or shorter on the top-tube length with this formula.” Adjusting the seat and the handlebars generally compensates for the 2”. This adjustment comes from sliding the seat forward or backward, or by increasing or decreasing the height of the handlebar stem. This gives ± 2” of adjustability to the $length_{tobTube}$ design variable, which maps to an adjustability of ± 2” for the *reach* performance parameter.

4.2.3 Parameter Increments

Good values to use for the parameter increments are more difficult to come by, as they are based more on a trade-off with cost and are generally chosen by a company. Companies generally offer bikes in 2” increments for the stand-over height. In order to determine if any improvement could be made by using a smaller increment, the value of .5” was chosen. The value for top-tube length is a function of the rest of the geometry of the bike. Its length is used more as a check to see if a bike with the correct stand-over height will be a good fit and to help decide between two sizes that could both be a good fit. Therefore, the value selected for a specific design will likely be used only as a guideline when creating the final design. Since it is generally not used to market the bikes, it does not need to be a nice round number like the stand-over height. As such, the increment value can be fairly arbitrary as long as it is small enough to be useful. A value of .5” was deemed to be small enough but not so small that it would unnecessarily slow down the optimization routine.

A summary of the values chosen for the mountain bike product definition is provided to use as a reference for the rest of the example (Table 4.1).

4.2.4 Market Demand Model

A population height distribution was used to synthesize the market demand model. The distribution data was divided into categories of gender, race, and age, but no data was available for either males or females under the age of twenty. The bike model chosen for the analysis is the top of the line Novara mountain bike: the Ponderosa FSL. Novara sells a separate line of women's bikes but not a top of the line bike. Therefore, it is assumed that they are targeting a mostly male demographic with this bike, but that it will also attract some female riders. Novara offers four sizes for most of their models but only three for this one. Because it was the smallest size that was eliminated, it is also assumed that their target market is an older market.

Using these assumptions, it was decided to use the population information for all males over the age of twenty. If the female population were included, it would skew the population to a shorter height. This will be taken into consideration in the conclusion. The population data gave only information on height. In order to populate the market demand model for inseam as well, a common ratio of .485 * height was used [26]. A normal distribution of the inseam with a standard deviation of 1.56 was applied using the result of this calculation as the mean.

These two normal distributions were combined into a bivariate normal distribution (Figure 4.4). This was used as a mathematical model representing the market demand. By using all ethnicities and age groups, it is assumed that the market demand is uniform across all these groups.

Many assumptions were made in the creation of this market demand model. These were deemed acceptable because this case study is a proof of concept for the methodology and should not and will not be used as a final recommendation on mountain bike design.

Table 4.1: Mountain bike product definition

Performance Parameter	Design Variable	Adjustability	Increment
inseam	$height_{standOver}$	±1”	.5”
reach	$length_{topTube}$	±2”	.5”

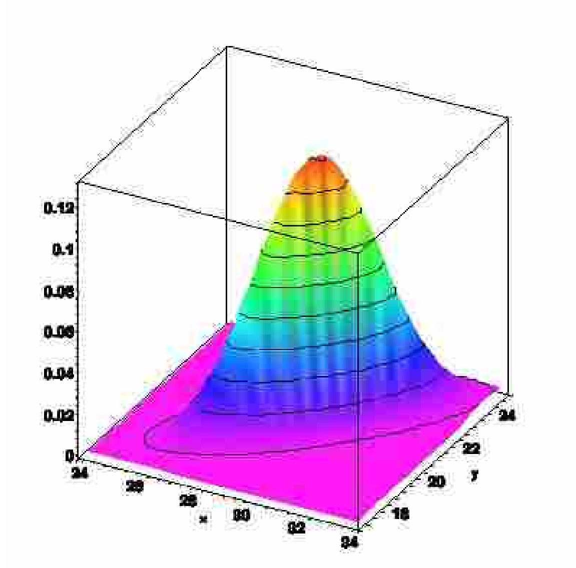

Figure 4.4: Market demand model and market opportunity

4.3 Spanning Analysis

4.3.1 Preliminary Analysis

As discussed earlier, the POD is using much broader ranges than will likely be covered by the chosen set of designs. Using these ranges will result in a deceptively low value for the percent of POD spanned. Because of this, the smallest and biggest sizes of the current product offering will be used, with adjustability applied, for the ranges in this analysis. This results in a range of 32.2 to 36.8 for the inseam and 25.6 to 30.9 for the reach. By inspection it can be seen that the current product offering (Table 4.2) will span these ranges completely.

This analysis is only useful for illustrating that the placement of the middle size spans the interior well and that there is overlap which could be used to cover the upper and lower regions of the POD.

Table 4.2: Current product offering

Parameter	15”	17”	19”	Range
inseam	33.2	34.6	35.8	32.2 - 36.8
reach	27.6	28.2	28.9	25.6 - 30.9

4.3.2 Total Volume: Market Opportunity

As discussed previously, in this case study the market opportunity consisted of every male over the age of twenty. When plotted on the design space, the market demand model was represented as a surface. The volume beneath this surface down to the x-y plane represented the total market opportunity (Figure 4.4).

4.3.3 Volume Divided into Cells

This volume was then divided into cells based on the increment value specified in the product definition. The cell boundaries were projected up onto the surface. The volume under the surface bounded by each cell was calculated and tabulated. Each cell represented the range spanned by one possible location for a design.

4.3.4 Spanning Due to Adjustability

The adjustability allowed each candidate design to span multiple cells. Because the increment in both the length and height directions was a constant .5", the number of cells spanned by a design was also constant. The product definition stated an adjustability of $\pm$1" in the height and of $\pm$2" in the length. Using these values for the adjustability, any given design could span its own cell, plus two cells on each side in the height direction, and plus four cells on each side in the length direction. Therefore, each candidate design spanned a five by nine patch of cells in the POD (Figure 4.5).

4.4 Optimization

Because of the way the calculations are structured, the optimization is not able to investigate a different number of designs within a single routine. It was decided to investigate how well the POD could be spanned by three designs, four designs, and five designs, and then to compare the results. For each scenario the optimization objective was to minimize the percent of the POD that was not spanned by the set of candidate designs. Several algorithms were tried and the Hooke-Jeeves algorithm was found to do the best job of exploring the possible combinations and finding an optimum set of designs. Each

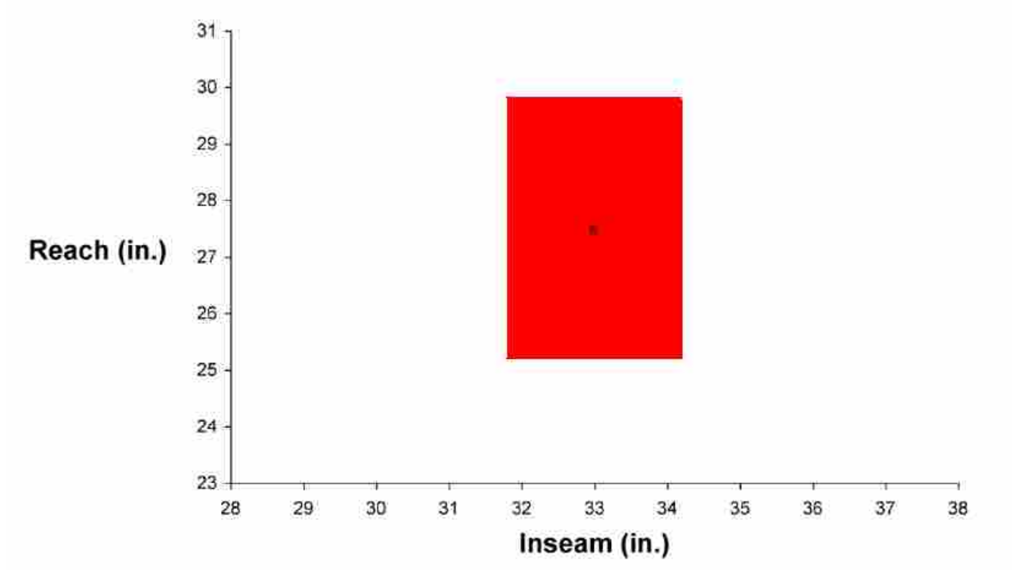

Figure 4.5: Spanning due to adjustability of ±1” inseam and ±2” reach for a variant with an inseam of 33” and a reach of 27.5”.

routine ran two hundred evaluations. The results of each are discussed below, as are the results from the set of three designs currently being used by REI.

4.5 Effectiveness of Set of 3 Variants

The initial set of three variants was analyzed to determine how well it spanned the POD. After, the distribution of these three variants was optimized in order to maximize the percent of the potential market spanned.

4.5.1 Current Product Offering

All of the calculations were done using the parameter values from the three sizes currently offered by REI (Table 4.2). The results of these calculations are shown in Table 4.3. A plot showing the location of each size and the amount of space it covered was generated (Figure 4.6).

Figure 4.6 shows that there is substantial overlap in the *x* or *height* direction that could be used to increase the percentage of POD spanned. Figure 4.7 shows the location of the designs in relation to the height parameter of the market demand model. This shows

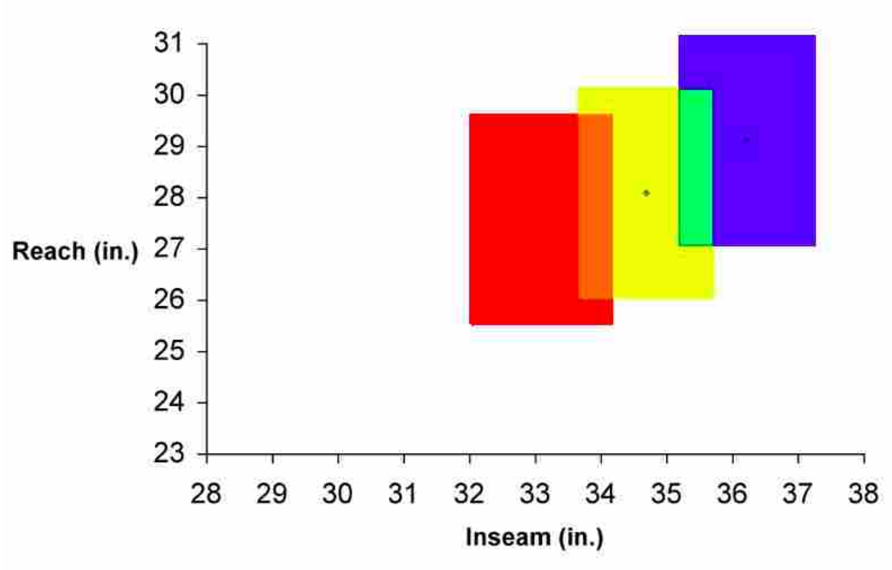

Figure 4.6: Spanning of current product offering

that if the overlap is ignored and the same spacing between sizes is used, they could benefit by shifting the sizes down .5” or more. Optimizing the placement of the three sizes should find a set of designs that are spaced further apart and possibly shifted down.

4.5.2 Optimized Product Offering

The optimization routine was run using the current product offering as a starting point. The resulting set (Table 4.4) increased the percent POD spanned by 19.2%. All overlap was eliminated and the sizes were shifted down, with the largest size being .5”

Table 4.3: Current product offering of three variants

Size	Height	Length	Spanned Vol.
15”	33.0	27.5	.567
17”	34.5	28.0	.568
19”	36.0	29.0	.567
Overlap Ratio	1.7		
%Spanned	75.7		
%Open	24.3		

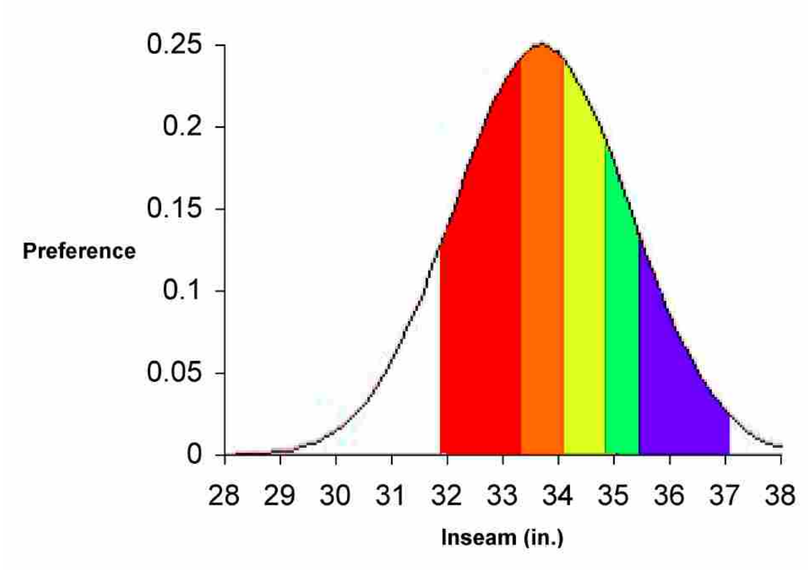

Figure 4.7: Spanning of current product offering inseam

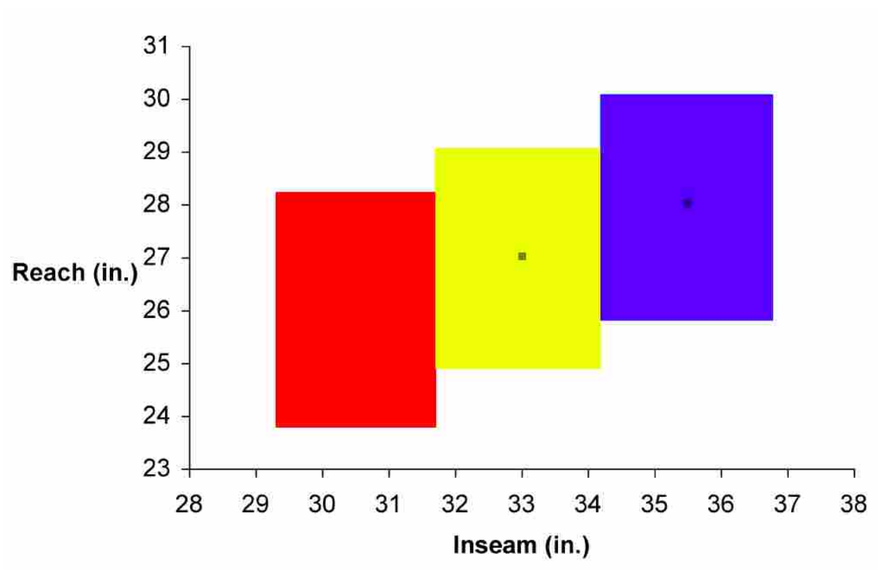

Figure 4.8: Spanning of optimized set of three variants

shorter than the current product offering (Figure 4.8 , Figure 4.9). However, by eliminating all overlap, any sort of safety factor on the adjustability has also been eliminated.

4.6 Effectiveness of Set of 4 Variants

The optimization routine was setup to optimize the placement of four mountain bike designs with the same constraints and function as before. As can be seen in Table 4.5,

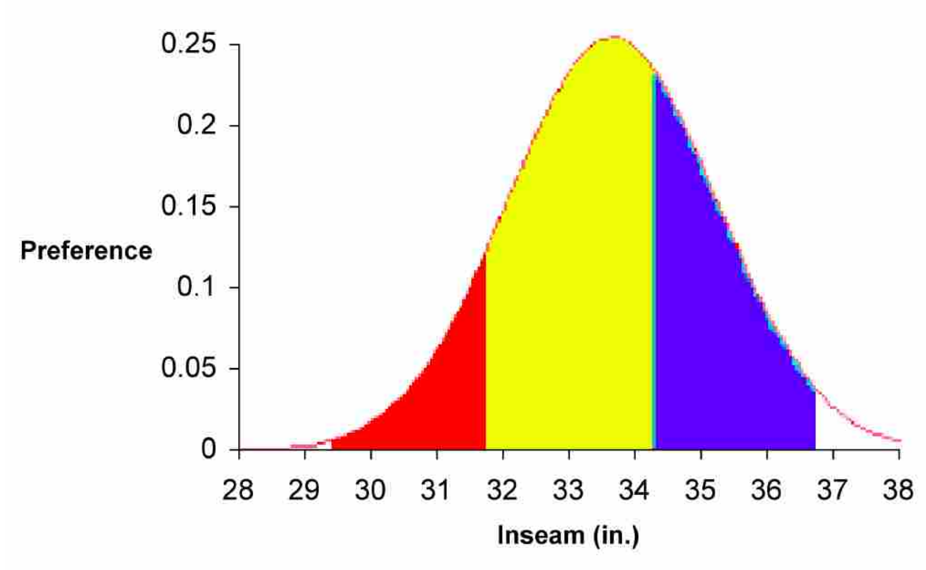

Figure 4.9: Spanning of optimized set of four variants: Inseam

there was little improvement from the additional design. The set of four designs was able to span 94.9% of the POD, an increase of only 1.5% over the set of three designs. Additional improvement is difficult to come by because the unspanned regions of the POD are all along the outside edge (Figures 4.10, 4.11).

4.7 Effectiveness of Set of 5 Variants

After viewing the results from the set of four designs, it was decided that there was little to be gained by using an additional design. If the remaining 3.6% of the market were deemed to be worth the investment, the same method could easily be applied to a set of

Table 4.4: Optimized product offering of three variants

Size	Height	Length	Spanned Vol.
Small	30.5	26	19.4%
Medium	33	27	61.2%
Large	35.5	28	19.4%
Overlap Ratio	1.0		
%Spanned	94.9		
%Open	5.1		

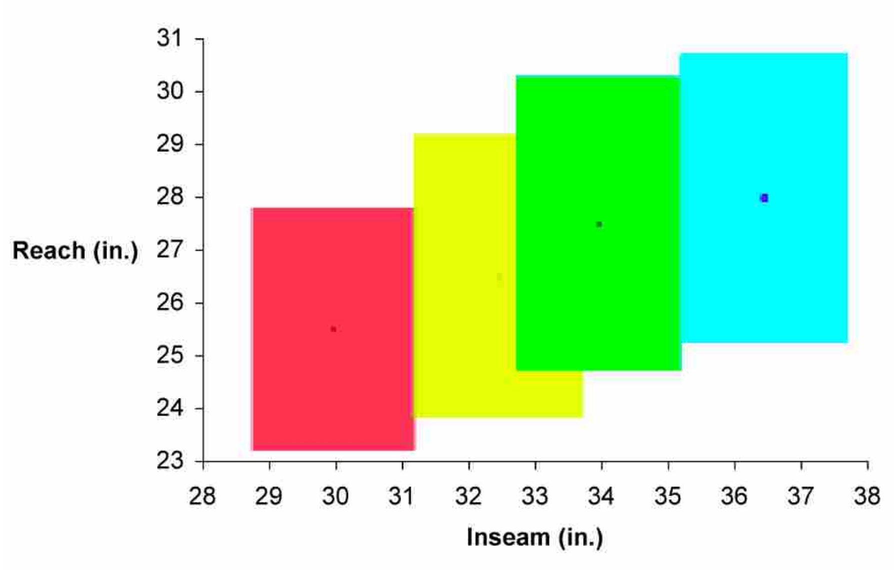

Figure 4.10: Spanning of optimized set of four variants

five variants. However, increasing the number of variants from three to four yielded only a 1.5% increase. Part of this reason is because the set of three variants , once optimized, had zero overlap. With the existing overlap in the set of four designs, it is possible that adding an additional variant would have slightly more substantial results.

4.8 Results

At this point it must be decided whether the increased potential market spanned by four designs over three is worth the added design, manufacturing, and inventory costs associated with the extra design. With the appropriate data a cost-to-benefit analysis could be

Table 4.5: Optimized product offering of four variants

Size	Inseam	Reach	Spanned Vol.
Small	30	25.5	10.7%
Medium	32.5	26.5	54.9%
Large	34	27.5	48.2%
XLarge	36.5	28	6.0%
Overlap Ratio	1.2		
%Spanned	96.4		
%Open	3.6		

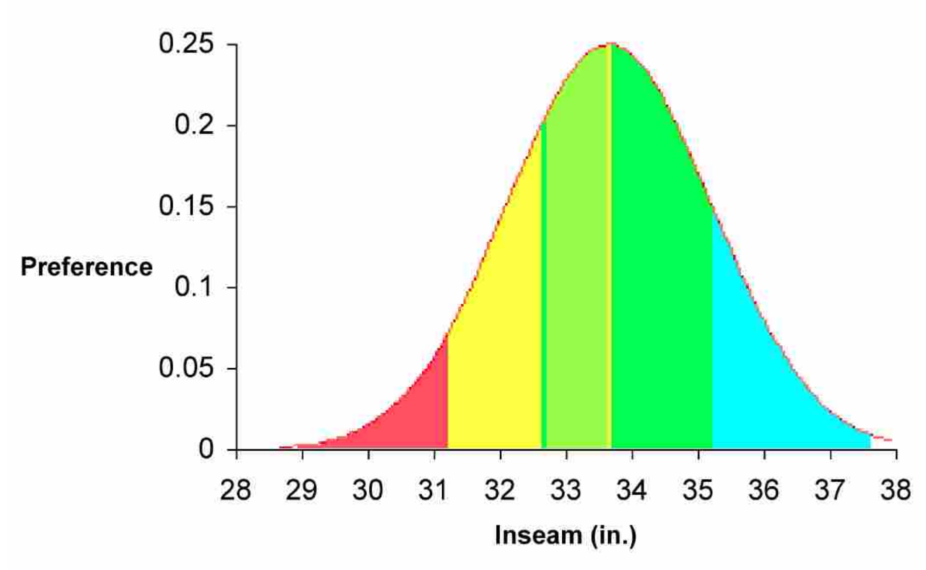

Figure 4.11: Spanning of optimized set of four variants: Inseam

performed, but such an analysis is outside the scope of this thesis. However, by observation it seems safe to conclude that, if chosen carefully, three designs are adequate to span the design space and that a fourth design is not worth the investment.

4.9 Conclusion

The current product offering is skewed too tall for the population distribution used in this case study. This is more surprising considering that had females been included in the population distribution, the current product offering would do an even poorer job spanning the market.

There are several possible explanations for why REI would use this set as their product offering.

- Mountain bikers are taller than average.
- REI has not researched the requirements of their intended market, or their intended market does not match the demographics used for this example.

- REI is designing to the casual, ride around town use of mountain bikes. Just as few SUV's are used off road, few mountain bikes are used for trail riding. A larger frame allows for a more efficient riding position.

- Many customers purchase a larger bike than what optimally fits them. This is partially due to the "bigger is better" ideology, and partially to the salespeople not educating the customer. This could inflate the sales of the larger frames and lead REI to shift the product offering larger.

Some of these explanations are more probable than others, but none of them emerges as the obvious answer.

The author is 6'1" and has tried the two larger sizes of the mountain bike used in this case study. Based on the REI Fitting Guide, the medium size frame was barely small enough to be a good fit. After trying out the medium size frame, the author felt that it was still about 1/2" too big for serious mountain biking. This experience lends credibility to the conclusion that REI's line of mountain bikes are skewed too tall. It should also be noted that this case study used data from the 2004 model. The 2005 model specifications show a decrease in height of about 1/4" on all three sizes. This may be considered a validation of these results. It is also important to note that they are working their way toward offering a competitive bike sized specifically for women. They do not as yet offer one comparable, but the top of the line women's bike is now much improved over the offering of previous years.

Chapter 5

Case Study: Flow Regulating Disks

There are various industrial applications that require controlling the amount of noise created when venting steam to the atmosphere. The most common method of reducing the noise is through silencers. These work similarly to a muffler on a car, reducing the transmission of sound to the air. The initial cost of silencers is relatively inexpensive. However, a traditional silencer can measure 12 feet in diameter and 20 high and weigh in excess of 32,000 pounds. This increases the structural and installation costs. Additionally, silencers don't hold up well to the wear and tear they are subjected to through normal operation, thus increasing the long-term cost to this solution.

Another method of approaching this problem is to control the venting of the steam in order to control the noise it creates. One company does this by using a stack of disks with channels etched into them through which the steam must pass. Turns in these channels slow the release of the steam, and forks in the channels control its expansion (Figure 5.1)[27]. These atmospheric resistors are more expensive up front, but with a longer life and measuring less than half the size of an equivalent silencer, overall they can be a less expensive solution.

The current business process is to custom assemble each atmospheric resistor from a catalog of standard disks in order to meet the customer's requirements. Because of the wide range of applications and conditions, this necessitates a large catalog of disks. This chapter examines the current catalog and how it was developed. Various methods of reducing the number of disks while minimizing costs are demonstrated. The resulting sets of disks are compared in a trade-off study and a single best set of disks is recommended. In this application of the method, the objective is to reduce costs by reducing the number of disks

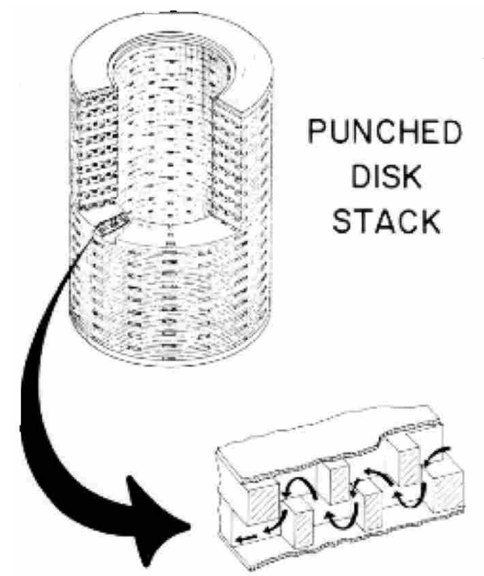

Figure 5.1: The disk stack from an atmospheric resistor. [27]

used in the catalog, and by using the lowest cost disks that will still satisfy the customer requirements.

5.1 Flow Control Disks

There are several characteristics of this case study that make it unique and more of a challenge than the mountain bike example of the previous chapter. First, the previous example was interested in the end product, delivered to the customer. It was easy to determine if the customer requirements were satisfied by the delivered performance parameters. However, this case study is interested not in the end product, the atmospheric resistor, but in the catalog of disks used in the assembly of that product. A given resistor is made up of an even number of disks of equal dimensions. Generally, the higher the flow rate of steam, the more disks will be used which will result in a taller stack and a higher capacity for controlling the steam. However, the same disk could be used to assemble stacks of different heights. Using ten of one disk for a stack would result in a resistor with a relatively low flow. Using forty of that same disk in a stack would result in a resistor that allowed a much higher flow. Thus, the same disk could span multiple discontinuous regions of the resistor POD. Additionally, rather than use more disks to handle a greater flow rate, a shorter and wider stack with more turns in each disk could conceivably be used as well. This means that there are multiple disks that could be used to satisfy the same customer requirements, but only if the design rules were changed to accommodate such flexibility.

The existing catalog of disks consists of 746 different disks. Most of these are not kept in inventory, but instead are available to the engineer to select when he or she is designing a new custom atmospheric resistor. These disks are evenly spaced along each of the design variables. From a design point of view it makes sense to use design variables to populate the catalog because it is easier to adhere to standard sizes this way, it appears to cover all regions of the design space, and it's easier to do with design variables than with performance parameters. However, these design variables do not map straight across to performance parameters and, as discussed earlier, because the disks are being used in stacks of different sizes one disk may map to multiple regions of performance parameters. Using a set of disks that are uniform in terms of design variables does not in any way guarantee that the resulting product family will be uniform in terms of performance parameters. In this specific case study, it is almost guaranteed to leave gaps and to create overlapping designs. Anywhere there is a gap, the engineer is unable to select a disk from the catalog for the resistor, and ends up having to design a custom disk instead. Anytime a custom disk is used it greatly raises the cost of the atmospheric resistor both in terms of development costs and manufacturing costs. Additionally, with the catalog defined in terms of design variables it is very difficult to determine how well it spans the POD, which is defined in terms of performance parameters. The black box that maps these parameters takes performance parameters as inputs and generates the design variables. In order to analyze the existing catalog, this black box would have to be used recursively for each of the 746 disks to determine which performance parameters would result in the given design variables. This was deemed to be an unnecessary increase in entropy and was therefore not performed.

5.2 Definition

The atmospheric resistor is the end product that is delivered to the customer. However, this case study is only interested in the disks used to create the resistors. Therefore, the product definition will be of the disks rather than the resistor. Additionally, there are three different materials that can be used to manufacture the disk. The material to be used is specified by the customer. There was no data available for this thesis concerning the demand for the different materials, so it was assumed that demand was uniformly distributed amongst the three. None of the other parameters was dependent upon material, so it was eliminated as a customer requirement in order to reduce the problem. This reduction on the disks in the catalog results in the 746 different disks. Assuming a uniform distribution, the selected set of disks would apply to all three materials.

5.2.1 Performance Parameters

The customer requirements are given to the company in terms of how the atmospheric resistor should perform. It is possible to map these requirements to requirements for the disks. However, because any disk can be used to create multiple sizes of stacks it is useful to keep the customer requirements as they are given and define the performance parameters in terms of the completed disk stack. Therefore, the performance parameters are inlet pressure, inlet temperature, flow, and noise. The inlet pressure and temperature are measured at the entrance to the resistor. Pressure is measured in terms of psia and temperature in terms of Fahrenheit. Flow is a mass flow rate metric defined by the company. Noise is measured in decibels 10 feet up and 10 feet over from the outside of the resistor. All of these parameters can be plugged directly into the black box to generate the design variables.

As part of the product definition a range was also found for each of the parameters. Data from actual customer orders is unavailable, therefore a best guess was used for each. This estimate was aided by product pamphlets available to the public on the company's website. The performance parameters and their respective units and ranges are shown in Table 5.1.

5.2.2 Design Variables

The design variables are the variables that are modified in order to generate a new design. Using the disk catalog as a starting point, it was determined that disk type and inner diameter were two of these. The third variable could have been either the number of turns in the grooves on the disk, or the outer diameter of the disk. The number of turns more directly affects the performance of the disk whereas the outer diameter is calculated from the required wall thickness for a given number of turns. For these reasons the number of turns was chosen as the third design variable.

The disk type is a discrete variable with three possibilities to choose from. The disk types vary in size and in whether they come in sets of two or four. The first two types come in sets of two and are either 1/16" or 1/8" in thickness. The third is 1/16" and comes in a set of four. The inner diameter is self-explanatory and is measured in inches. In the original catalog of disks the inner diameter ranged from 3" to 16" with several gaps in between. The range for the new catalog will be determined from the range set on the performance parameters. However, 3" is the minimum feasible inner diameter, and any candidate disk with a diameter smaller than that will be thrown out. In the original catalog the number of turns ranged from 2 to 30. Again the performance parameters will determine the range for the number of turns. However as each turn is a 90 degree turn alternating right and left, in order for the grove to exit perpendicular to the outside circumference of the disk only an even number of turns should be used. Because of this design rule, all candidates with an odd number of turns need to be rounded to an even number. In order to still satisfy the customer requirements, the number of turns will always be rounded up. Table 5.2 summarizes the design variables and their associated constraints.

5.2.3 Adjustability

The design of the disks does not incorporate any inherent adjustability. Generally adjustability comes from an adjustment that can be made by the customer after the product has been manufactured and assembled, but in this case there is no adjustment that can

Table 5.1: Performance parameters and their respective ranges

Parameter	Name	Min	Max
Inlet Pressure (psia)	P_i	30	450
Inlet Temperature (degF)	T_i	300	1,000
Flow	$Flow$	20,000	600,000
Noise (dB)	$Noise$	85	112

be made to the resistor after assembly. However, there is another form of adjustability that can be used by the company in selecting which disks to use. The customer gives the company specifications for the resistor they need, and the company is expected to deliver a product that meets these specifications or beats them. Assuming no size constraints, the company could conceivably use their top of the line resistor to satisfy all of their customers' requirements. This form of adjustability will be applied to the inlet pressure, temperature, flow, and noise. The higher performing the atmospheric resistor is, generally the more expensive it is to manufacture. Because of this it is generally in the company's best interest to deliver a resistor that is as close to the customer's specifications as possible. However, the company also desires to minimize the number of different disks they manufacture in order to minimize costs and production time. A trade-off study will be done to analyze these two competing objectives.

5.2.4 Parameter Increments

This is a very different kind of market than that of the mountain bike example. When ordering an atmospheric resistor the customer will specify exactly what their performance requirements are. As a result temperature and pressure specifications can be any value within the specified range and are rarely nice, round, easily-marketable numbers. From the company's point of view, there is no driving factor for the location of these increments. However the spacing of the increments should be driven by the cost trade-off discussed above. Rather than fix a somewhat arbitrary value for the temperature and pressure increments, a cost analysis will be performed illustrating the benefits of different possible increments. The remaining two parameter increments are more straight forward. An increment of 1 dB will be used for Noise. The Type is already set at three discrete values.

Table 5.2: Design variables and their associated constraints

Design Variable	Name	Constraint
Disk Type	*Type*	3 discrete possibilities
Inner Diameter)	D_i	3" minimum
Number of Turns	*Turns*	Must be an even number

5.2.5 Market Demand Model

In the previous example, the market demand model was developed using statistics for human metrics which were represented using a bivariate normal distribution. In this example there are four performance parameters with different modes of distribution. Actual previous sales records and forecasts for future sales are unavailable, so some estimates must be made regarding the demand for different regions of the design space, which are representative of different disks. First, it is assumed that the pressure, temperature, and flow are all uniformly distributed across their respective ranges. Because OSHA standards limit noise pollution it was decided that a normal distribution would be appropriate for the noise parameter. It was normally distributed about 100 dB with a standard deviation of 7 dB. Using these four distributions, a Monte Carlo simulation was run in order to generate 2,000 designs in terms of performance parameters. These are representative of the market demand.

5.3 Spanning Analysis

Because of the various complexities of this case study, the analysis was performed differently from the mountain bike case study. This problem has four performance parameters all of which have a much higher number of possible values. Together there are several hundreds of millions of possible combinations per disk. If the disks are combined into sets, the problem becomes very unwieldy indeed.

5.3.1 Set Selection Process

The market demand was given in terms of performance parameters. However, disk sets are created using the design variables to define the disks. In order to convert these performance parameters to their corresponding design variables, the 2,000 customer requirements generated by the Monte Carlo were fed into the black box. The ranges for the resulting design variables are shown in Table **??**. As stated previously, a single disk can potentially span multiple regions of the design space. Because of this, the 2,000 customer requirements generated only 867 different disk designs. Upon closer inspection it was determined that some of those designs were infeasible because they had an odd number of turns, and company design practice dictates that only an even number of turns can be used. All of the disks with an odd number of turns were rounded up to the nearest even number. As a result, only 566 disks were necessary to satisfy those 2,000 customer requirements. These 566 were the disks that would have been designed and delivered to the customers using the current company practice, without allowing for a higher performing disk to be substituted for a lower performing one. It is significant to note that each disk was used an average of roughly four times. This overlap indicates that the initial population of 2,000 customer requirements was sufficient to adequately represent the full range of customer requirements that the company can expect to see.

Of those 566, there were some that were used more often than others, ranging from eleven times for the most popular disk to only once for the least popular. The 566 disks were ranked in order of their popularity (Table 5.4). An analysis was performed on the cost of using all 566 disks as a percentage of the cost of using the current 746 disks. This cost was then compared with the cost of using only the disks that were used 11 times, then those used 10 or more, then 9 or more, and so on. For each analysis, the 566 disks were separated into two categories, those that would be used (Group A) and those that would not be (Group B). For each disk that would not be used, the cheapest disk from Group A that matched or beat the disk's performance was selected to replace it. The average cost of all the disks in Group A was calculated to give an idea of how much extra it cost to use higher performance disks in place of the lower performance ones. Invariably there were some disks left over in Group B that could not be replaced by any of the disks in Group A. These were noted as well and their average cost was also calculated. This gives an idea of what it would cost to make these additional disks on an as needed basis (Table 5.5).

Table 5.3: Ranges for each of the design variables, resulting from 2000 customer specifactions.

Design Variable	Min	Max
Type	1	3
D_i (in.)	3	18
Turns	2	48

5.4 Results

The differences between this case study and the mountain bike case study (complexity of the problem, number of dimensions, the lack of an actual preference model, and the uni-directional nature of the black box) made it impractical to do the same types of analyses on the results. Because the black box takes customer requirements and generates the design variables to satisfy those requirements, it is difficult to analyze how well the current catalog of disks spans the design space. It is assumed that the company is currently able to satisfy all customer requirements (within the chosen range) with the existing catalog of disks.

By assuming that the results of the Monte Carlo simulation accurately represent the preference model, it is then possible to calculate how well a new, optimized set will span the POD. Any set of disks that is able to satisfy all the requirements from the Monte Carlo simulation is assumed to completely span the POD.

5.4.1 Cost Trade-off

Once the disks were organized into sets based on their popularity, a trade-off analysis was performed on the number of disks and the cost. The same process as before was used to reduce the number of designs used to span the POD. However, this time the total number of disks s was used (the preferred set of commonly used disks and the disks that

Table 5.4: Disks ranked according to number of times used.

Popularity	Number of Disks
11	1
10	2
9	5
8	6
7	13
6	27
5	47
4	84
3	109
2	114
1	158

couldn't be replaced by the commonly used disks). These results are shown in Table 5.6. This was done for each set of commonly used disks, summing them with the leftover disks each time. The cost is a total average cost of all those disks and is shown as a percentage of the cost of using the 746 original disks. This is shown in Figure 5.2. It was determined that using the 101 most commonly used disks would leave only 68 disks that were not spanned by the 101. If the set of 101 commonly used ones and the set of 68 leftovers were combined, this would still not only reduce the inventory considerably (to 169 disks total) but would also reduce the average cost per disk used by 20%. If an additional 70 disks were used (or 239 total) the average cost would be reduced by 23%.

5.5 Conclusion

Using 2,000 system requirements generated by a Monte Carlo Simulation to represent the market demand, it was determined that the number of disks in the catalog could be drastically reduced. The existing catalog consisted of 746 disks. If the main objective was to minimize the number of disks used, the optimal set consists of 169 disks with an average price 20% lower than the existing catalog. If the objective is to minimize the cost, the optimal set has 239 disks and saves 24%.

Many assumptions were made in this study, but the overall concept was proved and could possibly lead to very significant cost reductions for the company. It should

Table 5.5: Results of disk redistribution and cost analysis

# of Times Used	Disks Used	Avg. Cost	Disks Remaining	Avg. Cost
11	1	$4,526	558	$9,068
10	3	$7,478	521	$8,976
9	8	$4,550	323	$9,009
8	14	$6,457	295	$9,274
7	27	$6,723	229	$8,838
6	54	$5,734	138	$8,973
5	101	$5,293	68	$10,046
4	185	$5,858	54	$10,452
3	294	$6,657	33	$10,802
2	408	$7,461	14	$9,819
1	566	$8,982	0	$0

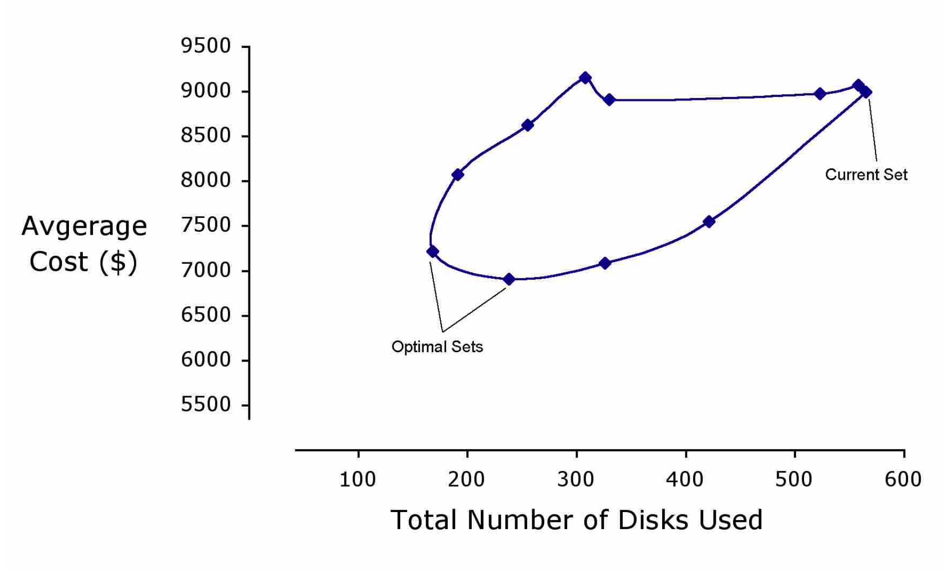

Figure 5.2: Trade off between number of disks used and the average cost of the disks.

be remembered that the disk choice was constrained by the current design process. If the process were altered so that in place of a tall, small diameter stack a shorter, larger diameter stack could be used (or vice-versa) and still meet the customer requirements, the total number of disks used could be further reduced.

Table 5.6: Results of disk number/cost trade off ranked by popularity

# of Times Used	Disks Used	Disks Left Over	Total Disks	Avg Cost	% Original Avg Cost
1	566	0	566	\$8,982	103%
2	408	14	422	\$7,539	86%
3	294	33	327	\$7,075	81%
4	185	54	239	\$6,896	79%
5	101	68	169	\$7,205	82%
6	54	138	192	\$8,062	92%
7	27	229	256	\$8,615	99%
8	14	295	309	\$9,146	105%
9	8	323	331	\$8,901	102%
10	3	521	524	\$8,968	103%
11	1	558	559	\$9,060	104%

Chapter 6

Conclusion

6.1 Summary

This thesis has presented and demonstrated a methodology for selecting an optimal set of product variants in order to maximize the percentage of potential market spanned with a minimal number of variants.

The product was defined both in terms of performance parameters and design variables. The design variable definition was used to suggest new candidate variants, to apply design standards, and to ensure that candidates were feasible designs. The performance parameter definition was used for the actual analysis of candidate sets because the performance parameters can be compared directly to potential customer requirements. All of the performance parameters in the product definition were used as primary axes in a Product Offering Domain (POD).

Each point within the POD represented a potential customer requirement. A candidate product variant was said to span that point if it could satisfy that customer requirement. The region that a single variant could span was determined by the product's adjustability and by the increment value used along each axis. Adjustability was considered to be any ability inherent in the product that allowed it to satisfy customer requirements beyond the single point in the POD where it is located. Common examples cited were waist belts and crescent wrenches. Adjustability allowed a single product variant to span a larger region of the POD. The increment was used to determine at what values a variant would be considered along each of the performance parameter axes. The increment value was based on manufacturing and testing limits as well as limits in the customer's ability to discern a difference in performance. The chosen increment values were used to create a grid of cells, the centers of which marked the location of a candidate product variant.

The market demand was captured using a preference model which was a function of the performance parameters. This function was applied to the POD as one additional primary axis representing market demand. The volume under the preference model represented the total market opportunity.

During the analysis, the cell of every selected variant was activated, as well as any adjacent cells that were spanned by a variant. The volumes of all of the activated cells were summed in order to calculate the potential market spanned as well as the percentage of the total potential market that was spanned. Optimization was performed in order to maximize this value while using the minimum number of product variants. The trade-off between number of variants and percent potential market spanned was analyzed. Overlap (regions that were spanned by more than one variant) was also calculated.

This method was first applied to a family of mountain bikes. The performance parameters were inseam length and reach. The design variables used were seat-tube length and top-tube length. The market preference model was created assuming the market consisted of a normal distribution of age 20 and older males. The existing product family of three variants was analyzed and it was found that there was overlap in the designs as well as significant regions of the POD that were not spanned at all. The percent of the potential market spanned was calculated to be 75.7% The variants were optimized, again using three variants. This produced a much better spread of the variants with zero overlap and 94.9% of the market spanned. The optimization was performed again, this time using four product variants. This only increased the percent spanned by 1.5%. It was concluded that three product variants were sufficient if they were efficiently placed.

Next, the method was applied to a family of flow-regulating disks used in the assembly of atmospheric resistors. This was a more complex problem that provided some difficult challenges to this method. The atmospheric resistor was the actual end product delivered to the customer. However, since the main component of the resistor is a stack of disks, and since all of the disks in that stack are identical, it was decided to apply the method to the disks rather than the entire resistor. Using the existing process to design a new atmospheric resistor, the engineer had a catalog of 746 different disks to choose from. By plugging the customer requirements into a black box the engineer could generate the design variables for the disks required for the stack. This illustrated a new aspect of the problem: the same disk could be used to assemble stacks of different heights. Using ten of one disk for a stack would result in a resistor with a relatively low flow. Using forty of that same disk in a stack would result in a resistor that allowed a much higher flow. Thus, the same disk could span multiple discontinuous regions of the resistor POD.

The performance parameters used were inlet temperature, inlet pressure, flow, and noise. The design variables were disk type, inner diameter, and number of turns. With four performance parameters, each with fairly wide ranges, there were an unwieldy number of possible permutations. Rather than exhaustively search them all, a Monte Carlo simulation was performed to generate 2,000 representative customer requirements. These were fed into the black box and the result was the set of 566 different disks that were used to satisfy the 2,000 requirements. These disks were ranked according to how many times they were used. Then an analysis was performed to see how cost would be affected by replacing the less often used disks disks with more commonly used ones that still met the customer requirements. It was determined that using the 101 most commonly used disks would leave only 68 disks that were not spanned by the 101. Doing this would not only reduce the inventory considerably (to 169 disks if it was decided to stock all of them) but would also reduce the average cost per disk used by 20%. If an additional 70 disks were used (239 total) the average cost would be reduced by 24%. Many assumptions were made in this study, but the overall concept was proved and could possibly lead to very significant cost reductions for the company.

6.2 Future Work

There are several areas exposed during the course of this work that would benefit from further work.

1. In the bike problem, it was assumed that all variants would yield the same profit per variant sold. This may not be true. In the disk problem the production cost was taken into account, but it was assumed that the sale price would be marked up equivalently, so that, as in the bike example, the profit is the same for any product sold.

Future work could incorporate the production cost and the sale price as performance parameters in the product definition.

2. The number of disks used in each stack was not taken into account when determining how often each of the disks was used. The results can't be used directly to determine which disks should be kept on hand in ready supply because it could be that one of the very popular disks was used exclusively in shorter stacks while one of the less popular disks was mostly used in taller stacks. In this case, the actual number of disks used could be roughly the same. Further work should take this into consideration when performing the analysis.

3. Adjustability was assumed to be continuous for this thesis. However, many product families make use of modules. These modules could be used as a form of discrete adjustability.

References

[1] M. Öztürk, A. Tsoukiàs, and P. Vincke, "Preference modelling," DIMACS, Tech. Rep., 2003. 9

[2] E. Herrera-Viedma, F. Herrera, and F. Chiclana, "A consensus model for multiperson decision making with different preference structures," vol. 32, no. 3, pp. 394–402, 2002. [Online]. Available: http://search.ebscohost.com/login.aspx?direct=true&db=axh&AN=BAST02157775&site=ehost-live 9

[3] J.-F. Petiot and S. Grognet, "A multidimensional scaling approach for product design and preference modeling," vol. 4. UMR CNRS 6597, Inst. Rech. Commun./Cybernetique N., Ecole Centrale de Nantes, 44321 Nantes, France: Institute of Electrical and Electronics Engineers Inc., 2002, pp. 54–59. 9

[4] S. Benferhat, D. Dubois, S. Kaci, and H. Prade, "Bipolar possibility theory in preference modeling: Representation, fusion and optimal solutions," *Information Fusion*, vol. 7, no. 1, pp. 135–150, 2006. [Online]. Available: http://www.sciencedirect.com/science/article/B6W76-4G4XBXY-1/2/978c42bdcab1ed2acb375f43061962a3 9

[5] J. J. Michalek, F. M. Feinberg, and P. Y. Papalambros, "Linking marketing and engineering product design decisions via analytical target cascading," *Journal of Product Innovation Management*, vol. 22, no. 1, pp. 42–62, 2005. [Online]. Available: http://www.blackwell-synergy.com/doi/abs/10.1111/j.0737-6782.2005.00102.x 10

[6] R. A. Bates and H. P. Wynn, "Modelling feasible design regions using lattice-based kernel methods," *Quality and Reliability Engineering International*, vol. 20, no. 2, pp. 135–142, 2004. 10

[7] K. Kuchcinski, "Constraints-driven design space exploration for distributed embedded systems," *Journal of Systems Architecture*, vol. 47, no. 3-4, pp. 241–261, 2001. 10

[8] B. Miramond and J.-M. Delosme, "Design space exploration for dynamically reconfigurable architectures," vol. I. LaMI, Universite d'Evry Val d'Essonne, UMR 8042 CNRS: Institute of Electrical and Electronics Engineers Inc., Piscataway, NJ 08855-1331, United States, 2005, pp. 366–371. 10

[9] Z. Siddique and D. W. Rosen, "On combinatorial design spaces for the configuration design of product families," *Artificial Intelligence for Engineering Design, Analysis and Manufacturing*, 2001. 10

[10] Z. Jiang and J. Yan, "Research and development on constraint-based product family design and assembly simulation," *Journal of Meterials Processing Technology*, 2003. 10

[11] B. P. Corbett and D. W. Rosen, "Platform commonization with discrete design spaces: Introduction of the flow design space," in *Proceedings of DETC'03*, 2003. 10

[12] D. Hunsaker, "A framework for modeling mass customization envelopes," Master's thesis, Brigham Young University, 1999. 11

[13] T. W. Simpson, "Product platform design and optimization: Status and promise," vol. 2 A. Department of Mechanical Engineering, Pennsylvania State University, University Park, PA 16802, United States: American Society of Mechanical Engineers, New York, NY 10016-5990, United States, 2003, pp. 131–142. 12

[14] A. Jose and M. Tollenaere, "Modular and platform methods for product family design: Literature analysis," *Journal of Intelligent Manufacturing*, vol. 16, no. 3, pp. 371–390, 2005. 12

[15] J. Nanda, T. W. Simpson, S. R. T. Kumara, and S. B. Shooter, "A methodology for product family ontology development using formal concept analysis and web ontology language," *Journal of Computing and Information Science in Engineering*, vol. 6, no. 2, pp. 103–113, 2006. 12

[16] J. Nanda, H. J. Thevenot, and T. W. Simpson, "Product family design knowledge representation, integration, and reuse," vol. 2005. Department of Industrial and Manufacturing Engineering, Pennsylvania State University: Institute of Electrical and Electronics Engineers Computer Society, Piscataway, NJ 08855-1331, United States, 2005, pp. 32–37. 12

[17] H. J. Thevenot, J. Nanda, and T. W. Simpson, "A methodology to support product family redesign using a genetic algorithm and commonality indices," vol. 2 B. Harold and Inge Marcus Department of Industrial and Manufacturing Engineering, Pennsylvania State University, University Park, PA 16802, United States: American Society of Mechanical Engineers, New York, NY 10016-5990, United States, 2005, pp. 1009–1018. 12

[18] R. Rai and V. Allada, "Modular product family design: Agent-based pareto-optimization and quality loss function-based post-optimal analysis," *International Journal of Production Research*, vol. 41, no. 17, pp. 4075–4098, 2003. 12

[19] H.-J. Wang, Q. Zhang, and T.-T. Zhao, "A sa-based method for developing modular product family." Liaoning Key Lab. of Intelligent Information Processing, Dalian University, Dalian 116622, China: Institute of Electrical and Electronics Engineers Computer Society, Piscataway, NJ 08855-1331, United States, 2005, pp. 2483–2488. 12

[20] S. A. I. Nelson, M. B. Parkinson, and P. Y. Papalambros, “Multicriteria optimization in product platform design,” in *Proceedings of DETC’03*, 1999. 12

[21] S. Nidamarthi, G. Mechler, and H. Karandikar, “A systematic method for designing profitable product families,” vol. 3. ABB Corporate Research Center, 68526 Ladenburg, Germany: American Society of Mechanical Engineers, New York, NY 10016-5990, United States, 2003, pp. 67–75. 12

[22] J. J. Michalek, “Preference coordination in engineering design decision-making,” Ph.D. dissertation, University of Michigan, 2005. 12

[23] O. L. De Weck, E. S. Suh, and D. Chang, “Product family and platform portfolio optimization,” in *Proceedings of DETC’03*, vol. 2 A. Engineering Systems Division, Massachusetts Inst. of Technology, Cambridge, MA 02139, United States: American Society of Mechanical Engineers, New York, NY 10016-5990, United States, 2003, pp. 175–185. 12, 13

[24] (2006, November). [Online]. Available: http://www.rei.com/online/store/LearnShareDetailArticlesList?categoryId=Cycling&url=rei/learn/cycle/clfitf.jsp 32

[25] (2006, November). [Online]. Available: http://www.bicycletheory.com/website/about/index.php?strWebAction=the_theory 33

[26] *National Health and Nutrition Examiniation Survey*. National Center for Health Statistics, 1988-1994. 38

[27] (2006, 11). [Online]. Available: http://www.ccivalve.com/pdf/210.pdf 49