USING OPTIMIZATION AND LEAN PRINCIPLES TO DESIGN WORK CELLS AND MAKE CAPITAL PURCHASE DECISIONS FOR HOLE DRILLING OPERATIONS IN TURBINE AIRFOIL MANUFACTURING

By
Thomas E. Neal Jr.

B.S., Systems Engineering, United States Naval Academy, 1998

Submitted to the Sloan School of Management and the
Department of Electrical Engineering and Computer Science
in partial fulfillment of the requirements for the degrees of

Master of Business Administration
And
Master of Science in Electrical Engineering

In Conjunction with the Leaders For Manufacturing Program at the
Massachusetts Institute of Technology
May 2006

Signature of Author____________________
Sloan School of Management
Department of Electrical Engineering and Computer Science
May 5, 2006

Certified by____________________
Janice A. Klein
Senior Lecturer, Sloan School of Managment
Thesis Advisor

Certified by____________________
David H. Staelin
Professor, Department of Electrical Engineering and Computer Science
Thesis Advisor

Accepted by____________________
Debbie Berechman, School of Management, Executive Director, MBA Program

Accepted by____________________
A. C. Smith, Chairman, Committee on EECS Graduate Students
Department of Electrical Engineering and Computer Science

USING OPTIMIZATION AND LEAN PRINCIPLES TO DESIGN WORK CELLS AND MAKE CAPITAL PURCHASE DECISIONS FOR HOLE DRILLING OPERATIONS IN TURBINE AIRFOIL MANUFACTURING

By
Thomas E. Neal Jr.

Submitted to the Sloan School of Management and the Department of Electrical Engineering and Computer Science in partial fulfillment of the requirements for the degrees of Master of Business Administration and Master of Science in Electrical Engineering

Abstract

Classical manufacturing work cells have machines to perform each operation in the process, the number of each type of machine being chosen so that all machines would be equally busy. Although design of work cells for producing one product is straightforward, the design of a multi-product work cell is much more complex. Each product might require different machinery or require different processing times per operation, which complicates leveling the workload between machines.

The decision of what type of machinery to purchase for a particular operation can be complex in itself when comparing a number of technologically different alternatives. Although the machines' operating characteristics might be well known, it is often difficult to understand how each technology will affect the overall production system.

Any new work cell implementation is bound to cause friction within the organization. Change must be properly managed if it is to be done correctly. To implement any new structure for a work cell, both management and hourly employees must learn a new way of doing things. Often these new methods encounter significant resistance if not implemented in a way that takes into account their cultural impact.

This thesis describes a method for designing multi-product manufacturing work cells that utilizes optimization techniques to select capital machinery and uses Lean principles to assemble the machinery into an efficient and effective unit. It also shows how the best machinery for an operating environment can be chosen from two competing technologies. Finally, the thesis describes methodology for implementing new work cells in a change-adverse culture.

Although this thesis applies to the manufacture of almost any product, its example involves a factory manufacturing turbine airfoils and, more specifically, the design of work cells for drilling holes in turbine airfoils.

Thesis Supervisor: David H. Staelin
Professor, Department of Electrical Engineering and Computer Science

Thesis Supervisor: Janice A. Klein
Senior Lecturer, Sloan School of Management

Acknowledgements

First of all I must thank my wife for all of her support throughout this process. Although well prepared by my naval career she managed to weather yet another deployment with us living apart. She has also put up with all of the time I've spent writing this thesis, and has endured my seemingly endless grumbling through the difficult parts. I know that I could not have been successful without her.

Next I would like to thank my classmate and partner at Pratt & Whitney, Bret Awbrey. He endured the same successes and hardships that I faced at Pratt. He was a great sounding board for ideas, and he could always provide a unique perspective. I know that this process would have been immensely more difficult without his presence.

I would also like to thank Dave Pepe of the Turbine Module Center. As a technical lead for hole drilling he was a great source of information on technology, process control, and anything else that happened on the shop floor.

I would similarly like to thank my company supervisor Frank Gill. He was instrumental at getting me plugged in at Pratt and getting me started on this project. He was also a great source of information and a good sounding board for ideas about Lean manufacturing.

Additionally I would like to thank Steve Lindsey. He was my project champion and, more than anyone else, kept me motivated and focused on the project. He showed constant interest and enthusiasm for it, and this helped me to keep driving forward.

Next, I would like to thank George Joseph. His knowledge of the Turbine Module Center's financial system was what allowed me to relate what happened on the shop floor with the financials of the business and that was invaluable to this work.

I would also like to thank my thesis advisors, Jan Klein and Dave Staelin. They often steered me onto the right track when I was at the verge of loosing my direction. They also provided a constant barrage of insightful questions that kept my mind churning and helped me not to get too entrapped in the local Pratt & Whitney culture so I could maintain an outsider's perspective. Their advice in dealing with cultural issues was invaluable.

Finally I would like to thank Pratt & Whitney and in particular the Turbine Module Center for sponsoring this work and listening to my recommendations.

Table of Contents

Chapter 1: Introduction

Designing new multi-product work cells and implementing them in an existing manufacturing environment is a challenging task. It requires decisions about the technology to use in the machining operations, decisions regarding the number and types of machines to use in the cells to adequately machine the various products, and an implementation process that takes into account the cultural challenges inherent to change.

This thesis presents solutions to these problems in the context of building work cells for the hole drill operations of high pressure turbine vane manufacturing. It is based on research done at Pratt & Whitney's Turbine Module Center (TMC), located in East Hartford, Connecticut.

1.1 Description of the Problem and Objectives

Pratt & Whitney's Turbine Module Center is responsible for manufacturing high pressure turbine blades and vanes—collectively referred to as airfoils. Drilling small air cooling holes in these blades and vanes is one operation within the manufacturing scope of the TMC. Currently this machining is being done with equipment that is twenty to thirty years old. The age of the equipment makes its operation increasingly unreliable. Downtime on the equipment has been significant in recent years. Process variation is also significant with such old equipment. In addition the TMC has been facing constantly increasing demand for both military and commercial parts. These factors all challenge the TMC to meet its demand.

Pratt & Whitney management has determined that drilling holes in high pressure turbine airfoils should remain a core competency. Therefore, the problem that the TMC faces and the problem that this thesis proposes to solve is to determine a strategy to update the current hole-drill operations for turbine vanes into an improved system which can handle the expected increases in demand as efficiently and cost-effectively as possible.

The strategy behind updating the hole-drill operations for vanes is centered on the notion of procuring new capital hole-drill equipment and integrating it into work cells. There are such a large number of vane models that are drilled that it would be impossible to create a cell for each model. Therefore the cells to be designed must be capable of drilling multiple products. The first objective of this thesis is to find a method to choose sufficient capital equipment to meet the expected demand in the most cost effective way. The second, related, objective is to find a method to determine the most efficient way to deploy the equipment into cells and most efficient way to operate the cells.

In particular, the TMC is investigating two competing technologies for drilling round holes: using a laser and using a hollow spindle electro-discharge machining (HSEDM). The third objective of this thesis is to show how this determination can be made for the TMC.

Finally, the TMC has been drilling these holes using legacy equipment and methods for years. In addition to the challenge of determining what equipment to buy and how to best utilize it, the TMC will also face challenges in implementing anything new in the factory. Any culture that has been doing something one particular way for a long time will have a natural aversion to change, and the TMC will be no different. Therefore, the fourth objective of this thesis is to introduce techniques that will help these changes to be enacted and the new work cells to be brought online in a way that speeds cultural acceptance from both management and shop workers.

1.2 Thesis Approach

The approach of this thesis is to answer sequentially each of the above objectives. The first two objectives are intertwined. It would be impossible to choose the number of machines to purchase without knowing how they will be used in the cells, since this information affects the capacity of the system. However, how the equipment will be utilized within the cells actually has two components. One component is how many of each machine should be put in the cell and in what order, so that the workload is even. The other component is how to arrange the machines and how the operators will work in order to use them efficiently. Only the first component, dealing with the number and type of machines to put in each cell, directly relates to the first objective, which is to determine the number and type of machines to buy for the overall system.

The approach followed in this thesis is first to establish standard principles for how the cell should operate. This includes a general template for the standardized work involved and an idea of the configuration of the machines (e.g., aligned in a row, in a U-shape, etc.). For this thesis, the standard principles for how the cell should operate will follow standard Lean manufacturing principles. Once this standard is determined, the questions regarding the number and types of machines to include in each cell and therefore the number and type required throughout the system to meet demand can then be answered.

Optimization techniques are used to compute the number and type of machines required in each cell to meet demand. The optimization takes as inputs the demand volume and product model mix as well as information about production characteristics inherent to the machines and the layout of various cell types. As outputs it provides the number of cells required and the number and type of machinery in each cell. The thesis discusses the challenges inherent to the optimization as well as the pros and cons of various optimization techniques and describes in detail the optimization technique found to work best. An underlying part of the optimization and analysis is a cost model that helps to link the process aspects of the machines to the financial aspects of using them.

The thesis addresses the third objective—determining the best choice between two competing technologies—with the optimization model described above. Once an optimization is determined to meet the first two objectives, the same optimization can be used to determine which technology works better in the system. The optimization can be run utilizing each technology, and the most cost effective technology should be chosen. Additionally, the thesis examines the characteristics of each technology that make them more or less competitive.

The fourth objective is met with a recommendation for implementation of Lean principles based on the challenges that are likely to be faced at Pratt & Whitney's TMC. It is based on established principles for implementing Lean manufacturing and is customized to the author's observations of the specific nature of the TMC.

1.3 Summary of Findings

The numbers enclosed in this thesis have been altered to protect Pratt & Whitney's proprietary information. The numbers are nonetheless consistent with the results of the project at the TMC and they demonstrate how successfully the objectives listed above were completed.

In general, the optimization framework was found to be extremely beneficial. Figure 1-1 is a table that shows the potential value of implementing changes such as the ones described in this thesis. Figure 1-1 is based on using the methods described in this thesis to solve a hypothetical case, consistent with the one at the TMC. This hypothetical case and its solution are described in Section 4.2. Figure 1-1 is broken down into cases in which no optimization was performed (i.e., the traditional method was used), manual optimization was performed, or automatic optimization was performed. It should be apparent that the optimization framework allows for great improvements over traditional methods (by almost a factor of 4), even though the automatic optimization was not significantly better than manual optimization done using the framework (showing only a 25% improvement). This validates the methodology if not necessarily the automatic optimization method.

Figure 1-1 also shows the monetary savings derived if Lean principles are used and if they are not. This shows the significant advantage to using Lean principles in setting up the work cells. These savings are primarily derived from the added labor that is required to produce parts in a non-Lean cell. The labor costs add up quickly, especially in a labor-driven cost system like the one used by the TMC.

More specifically, this thesis shows the usefulness of looking at the difference between fixed and variable manual time in a machining operation. In this case, the differences between the fixed and variable manual time inherent to the laser and HSEDM processes demonstrated the supremacy of the laser process, despite individual attributes of the HSEDM process that suggested the contrary.

The thesis also finds that proper process control can have a significant effect on the controllability of a process and that the controllability of a process done on one machine can be the key to deciding between it and a similar process done with another machine utilizing a different technology. In this case, although HSEDM has superior controllability, process control measures allowed for sufficient controllability of the laser process to bring it into parity with HSEDM in this respect.

Figure 1-1

<table>
<tr><th colspan="4">Efficacy of Optimization and Lean Principles</th></tr>
<tr><td>Optimization Technique</td><td>No Optimization Using Lean</td><td>Manual Optimization Using Lean</td><td>Auto Optimization Using Lean</td></tr>
<tr><td>Utility Function:</td><td>2,400</td><td>8,800</td><td>11,000</td></tr>
<tr><td>5 Year NPV:</td><td>($2.4 Million)</td><td>$3.6 Million</td><td>$5.8 Million</td></tr>
<tr><td>Lean Principles</td><td colspan="2">Using Lean Principles</td><td>Not using Lean Principles</td></tr>
<tr><td>Yearly Savings</td><td colspan="2">$10 Million*</td><td>$4.5 Million*</td></tr>
</table>

* Assuming work cells implemented as determined by Manual Optimization

Chapter 2: Background

This chapter presents background information that is necessary to understand the objectives as they apply to the TMC. First the chapter will present the technical background information necessary to understand how holes are drilled in turbine vanes and the associated difficulties and challenges faced by these processes. Second, the chapter will discuss the current culture at the TMC, specifically looking at the challenges faced by the TMC, the history of the TMC, and the obstacles to change that exist at the TMC. Finally, the chapter will discuss relevant literature.

2.1 Technical Background of Drilling High Pressure Turbine Vanes

Drilling holes in turbine airfoils is not so straightforward a task as some might think. Although mankind has been drilling holes in metal for thousands of years, there are certain requirements for turbine airfoils that make this an especially difficult task. This section describes some of the strict criteria that exist for these holes and some of the different technical approaches that have evolved to meet these criteria.

Turbine airfoil holes are critical to the operation of jet engines and therefore have exacting tolerances. Turbine blades and vanes are located immediately behind the combustion chamber in jet engines. This is the hottest part of the engine and, in fact, the temperature of the exhaust gasses being pushed through these airfoils is actually higher than the melting point of the metals out of which the airfoils are created. The airfoils have two special properties that allow them to survive in this environment. First, they are coated with thin layers of temperature resistant materials which help to insulate them from the heat. Second, the airfoils are cast hollow, with many holes drilled from the outside of the airfoils to the inner cavities. Air is then forced through the stems of the airfoils, through the cavities, and out the holes. The circulation of this cooling air helps to remove the heat imparted by the jet exhaust. These cooling holes are therefore vital to allowing the movement of the air through the airfoils, distributing the cooling air evenly across the surfaces of the airfoils, and metering the flow of the air to maintain the desired temperatures within the engine.

The holes that are drilled in the airfoils must meet very stringent requirements in order to ensure that they can pass enough air to the right places to properly cool the airfoil. First, they must be positioned accurately so that the air traveling through the holes gets to where it needs to be to provide cooling. This means that the holes need to be in drilled in the correct spot on the vane, but it also means that they must be drilled at the proper angle to spread the air. Second, they need to be of the proper size. Holes that are too small obviously cannot pass enough air to ensure adequate cooling, and holes that are two large rob air from other holes. Third, the shape of the hole is important. The most basic hole is a cylindrical one, but holes are often conical or even rectangular conical, somewhat resembling an inverted pyramid. Cylindrical holes merely deliver air but conical and rectangular conical holes act as a diffuser which more evenly spreads the air across the airfoil. Fourth, the holes must be drilled in such a way that the material properties of the airfoil are untainted. If the hole is drilled poorly it might leave a very uneven surface on the walls of the hole. This uneven surface could have detrimental affects on the aerodynamics of the channel by affecting the location of the laminar-turbulent interface, which could affect the flow rate through the hole. Finally, a hole that is drilled too deep may result in "back wall strike," which occurs when the wall on the opposite side of the cavity is marred by the drilling process. At best this could affect the airflow through the cavity by upsetting the aerodynamics of the inner surfaces, and at worst this could degrade the structural integrity of the airfoil.

From a functional standpoint, the ideal would be to drill a hole in the airfoil that is located perfectly, of the perfect size, with the right shape, with good metallurgical quality, and without striking the back wall of the airfoil. From a manufacturing standpoint, the ideal would be to do all of this predictably, rapidly, and inexpensively. This is no small feat.

The two prevailing methods to drill these holes are with a laser and with an Electrode Discharge Machine (EDM). Both are very accurate and are capable of drilling very small holes (on the order of a few thousandths of an inch). Although there has been experimental research into trepanning shaped holes with lasers, there are currently no laser drilling machines that can machine shaped holes in materials as hard and thick as those used for turbine airfoils in a production environment. Therefore lasers can only be used to drill cylindrical and conical holes. For shaped holes, a comb-type EDM is used, because the hole shape follows the electrode shape and the electrode shape can be machined to virtually anything. Hollow Spindle EDM (HSEDM) is another method of drilling cylindrical and conical holes. This type of EDM drilling uses a single hollow, rotating electrode instead of comb-shaped electrodes in order to drill round holes very rapidly. Lasers drill at a much faster rate than EDMs. Lasers can drill holes on the order of one hole every two to three seconds. Comb EDMs can drill a row of holes (usually ten to twenty holes) in three to five minutes. Even HSEDMs, which are considered very fast, can only drill holes on the order of one every ten to fifteen seconds.

It is important to understand the sources of variation that cause imperfections in order to understand how equipment can be designed to meet these requirements. There are a wide variety of factors that affect each requirement. The following paragraphs discuss basics and details about the problems that have been known to cause manufacturing difficulties.

Most hole location variation is caused by tooling inaccuracies. For purposes of this thesis, "tooling" refers to a fixture that holds the turbine vanes in a precise position and that is mounted into the drilling machines. The tooling is removable from the machines and typically the vane is loaded into the tool prior to the tool being loaded into the machine for drilling. The tools are designed to grip the vanes in very few, small locations. Typically there are only three grip points where the tool contacts the vane. This is done to reduce the wear on the vane from the tooling and also to allow maximum access to the vane without the tool being in the way. However, since there are only a limited number of contact points, the vane may shift in the tooling, or not be mounted properly into it, which is one cause for hole location variation. One way to combat this is to use special mounting fixtures and gauges to load the vanes into the tools. This process typically involves loading a tool into a complex fixture and loading the vane into the complex fixture as well. This fixture precisely aligns the tool and the vane and the tool is then engaged to grip the vane. Finally, the vane and tool are checked using precise gauges to ensure that they are properly mounted together. This takes much of the tool mounting inaccuracy out of the process, but it is very time consuming. A second way to deal with this variation is what is known as "cast to size." Certain vanes have simple enough geometries that when they are cast, certain features can be guaranteed to be located with extreme precision. For these vanes, the points where they are mounted into the tools can be cast very accurately, negating the need to use complex fixtures and gauges. Unfortunately, most vanes cannot be cast to size because their geometry is too complex to support this. The third and latest way to deal with this sort of tooling error is to have the machine probe the part once it has been loaded. The machine determines the orientation of the part and compensates to drill the part properly, regardless of tool mounting error. This allows parts to be "soft loaded" into tools without special fixtures or gauges.

Machine variation can also cause hole location variation. While the tools contact the vanes in a complex way at various points, the machines contact the vanes only at one point – the point where the hole is drilled. For a laser, very few things that upset the position of the hole, since the drilling beam is inherently straight and therefore drills where it is aimed. The majority of positional inaccuracy – attributable to the machine – is in the servo mechanisms that rotate and position the part for drilling, but these inaccuracies are very small in comparison to the inaccuracies caused by the tooling. EDM drilling has more machine-dependent inaccuracies. Comb-type EDM drilling requires electrodes to be mounted into fixtures which the machine positions to the vane. In this process, there are a number of things that can cause positional errors. First the electrodes themselves are created with a manufacturing process that has inherent error, though very small. Next, the electrodes frequently are damaged so that one of the many teeth of the comb is out of position. Additionally, the electrodes frequently are improperly mounted into its fixture. Any of these errors would result in the hole drilled by that electrode being out of position. The current practice for reducing positional error in comb EDM drilling is to use an array of precisely aligned holes into which the electrode is dipped prior to the start of drilling with that electrode. If the electrode dips in and out of the holes without contact, the electrodes are positioned properly and drilling can commence. If the electrodes contact the block, an error is generated and the operator must realign the electrode. For HSEDM drilling, a single electrode is spun and fed through a guide to drill holes. The spinning action and guide help the electrode to be pushed straight into the vane where the hole should be drilled. Some variation, however, can still occur due to the electrode bending slightly after it leaves the guide. This is especially likely when the electrode is drilling at a high incidence angle. Yet, this still causes only a small amount of hole location variation compared to the errors seen in comb EDM drilling.

After drilling, hole location is checked through optical measurement. At Pratt & Whitney, two types of machines are used to do this. Both use a five-axis positioning mechanism to accurately position a part affixed to a tool. The machines then take a digital picture of where the hole should be and find how far the hole is from this location. The holes appear darker than the surrounding surface and the machines use an algorithm to determine the geometric center of the hole and its displacement from the expected hole position. As with the drilling machines, most of the error in the measurement process occurs during the tooling and a gauging process, since tooling is required to mount the vanes into the measurement machines.

Hole size variation is caused primarily by the machining process used. While the laser drilling process is the least susceptible to hole position variation, it is the most susceptible to hole size variation. Hole size is determined by the energy density that the laser delivers to the surface of the vane. The energy density that is delivered to the surface of the vane is determined by a very complex process with a huge number of variables, each of which has inherent variation. A brief description of the laser process will help to understand these factors.

The process begins with a flash lamp that is pulsed to "pump" a lasing medium. Although chrome-doped ruby has been used in the past, the typical lasing medium that is currently used is neodymium-doped yttrium-aluminum-garnet. This is commonly referred to as an Nd-YAG laser or simply a YAG laser. This pumping action excites electrons in the lasing medium, raising them to a high energy state. Once a population inversion is achieved there are sufficient electrons in upper energy states to produce sustainable photon emissions. Some electrons will spontaneously drop to a lower state, giving off photons. These photons will cause other electrons to drop energy states and give off even more photons. Since the lasing medium has a 100% reflective mirror at one end and a 99% reflective mirror at the other, most of the light travels back and forth, building up more energy as more photons are released, and the photons then go on to free even more photons. The 1% of the light that is released through the partially reflective mirror is the usable laser energy. In theory, the lasing process produces coherent light that is tightly collimated. However, because of the high power of the lasers used for drilling, a significant amount of heat builds up in the lasing medium. Because the heat builds up uniformly in the material (since it is generated by the lasing process) but is removed peripherally (due to conduction with the surroundings), the lasing medium actually warps, which distorts the resulting laser beam. This distortion adversely affects the quality of the beam because the distortions essentially cause the beam to act slightly out of focus. Therefore, the beam is then transmitted through a compensating telescope which re-focuses the beam to improve its quality. The beam is then sent through a system of mirrors that align it with the vane and a final focusing lens to focus the beam exactly onto the surface of the vane. The hole size is proportional to the energy density delivered to the vane, and therefore to change the hole size the flash lamps are driven at a higher voltage that excites more of the lasing medium, which gives off more photons and results in a higher energy laser beam.

Anything that causes the beam not to be properly focused onto the surface of the vane results in a change in the delivered energy density of the beam. The biggest source for variation in the focus of the beam comes from the setting of the compensating telescope. The temperature profile of the YAG rod changes whenever the laser power is changed, which slightly distorts the beam. With the beam distorted, the setting of the compensating telescope is off by some amount and therefore the quality of the beam is lowered. Also, over time the flash lamps degrade with use, making them less efficient and making them require higher voltages in order to get the same laser power output. With an open loop control system this means that the laser power will diminish slightly over time, which affects the energy density of the beam. Even if the voltage to the flash lamps is compensated to maintain a steady laser output power, the higher voltage will impart more heat into the YAG rods and therefore will have a slight effect on the temperature profile of the rod. This in turn will affect how well compensated the beam is.

The beam can also be disrupted by dust or slag (molten metal thrown off during the percussion drilling process of the laser). This material absorbs some of the energy from the laser beam and therefore lowers the final energy density.

Older lasers had a fixed compensating telescope that had to be adjusted manually. Newer lasers, though, have a compensating telescope that can be controlled electrically. This allows the compensating telescope to be adjusted every time the power is changed. Since power changes are a major source of the heat changes that bend the lasing rod, this helps to minimize the effects of changing laser power on the output hole size. Another feature that new lasers often posses is called variable spot monitoring (VSM). This feature allows hole size to be adjusted by a lens mechanism rather than by changing power to the flash lamps. VSM changes the diameter of the collimated laser beam prior to entering the final focusing lens. When the beam diameter changes, it is focused differently on the surface of the part. This changes the energy density at the drilling site and therefore changes the size of the hole. The benefit of VSM is that it removes many of the factors of variation listed above since it does not change anything in the long lasing process other than the position of one lens. Both of these features have greatly improved the performance of new lasers—as compared to the older lasers that currently reside in the TMC.

Part geometry can also cause hole size variation in the laser drilling process. When each vane is cast there are slight variations in the position of the internal cavity with respect to the external surfaces of the part. This is referred to as "core shift." Core shift results in wall thickness variation between parts, which to the laser drilling process is hole depth variation. Lasers do not drill perfectly cylindrical holes. Because the shots are focused on the surface the beam becomes slightly out of focus as each pulse cuts further into the part. This reduces the density of the energy that is delivered to the part further down the hole. Also, depending on the pulse repetition rate, there may still be a plume of microscopic slag pieces in the vicinity of the hole from a previous pulse when the next pulse is delivered. This plume can further absorb energy from the beam. Some of these effects are compensated by the fact that the laser beam reflects off of the side of the hole, which helps to concentrate it at the bottom and increase the energy density; however, all of these effects work together in an unpredictable fashion and contribute significant variation. Naturally, more of this variation is present as the hole gets deeper. In general, though, the holes end up being slightly conical with the bottom of the hole being smaller than the top, and a bigger difference being with a deeper hole. Therefore wall thickness variation is a significant contributor to hole size variation for the laser drilling process, and in fact is the leading contributor in most cases.

EDM drilling, on the other hand, has significantly less hole size variation than laser drilling. EDM drilled holes form in the same shape of the electrode used to drill them. They have the exact shape, but are slightly larger (on the order of a thousandth of an inch) due to overburn from the machining spark. The only sources of variation for EDM drilling are electrode size variation and overburn variation. Although the machining process for the electrodes has variation, it is extremely small and well within the tolerances for this process. Likewise, overburn variation is present due to the varying nature of the electrical generation and control inherent to the machine's power supply, but it is very small and well within the tolerances of the process. Since the holes take the shape of the electrodes they do not taper more or less depending on the thickness of the part and therefore are much less affected by part geometry variation.

Hole size is measured through a process called "airflow." To airflow a part, the part is mounted in a machine that forces air through the part. By measuring the pressure ratio between the supply pressure and the pressure in the vane, the amount of air flowing through the vane, which is mostly indicative of how large the holes are, can be determined. This measure also can be indicative of hole shape and material properties, as they too affect the airflow. This process can be done for all of the holes in the vane or for specific regions, although there are no procedures or specifications for individual hole airflows. This is a good overall measure for how well the holes have been drilled because it provides information on exactly how the vane will perform in service. Unfortunately, the measurement process is rife with variation. The largest source of variation is unknown air leaks. Although a calibrated standard is airflowed occasionally, this does not reveal leaks in the fixture for the vane or leaks through taped-off holes that are not intended to be airflowed. Another large source of variation is the measurement itself, which is an averaged reading over time. The readings fluctuate quite a bit and therefore it is difficult to get consistent results. The final source of variation is the operator. The results can fluctuate depending on how well or how consistently the operator mounts the vane in the fixture and tapes off the unmeasured holes.

There is relatively little hole shape variation since shaped holes are drilled with comb EDMs. As discussed above, this type of drilling is relatively free from hole size variation, and holes typically follow the shape of the EDM very well. If holes were to be attempted by trepanning with a laser or with successive cuts with a HSEDM, the potential for variation would be much higher since the shapes would actually be a conglomeration of small round holes. For the purposes of this thesis, however, the drilling of shaped holes will be assumed to be done with comb EDMs.

Material property variations are a direct result of the machining process. The two biggest concerns regarding material properties are whether the drilling process leaves micro-cracks along the ridges of the holes and whether it leaves slag particles in the holes. The micro-cracks form from the stress of the drilling process. Slag collects if the drilled material is not removed from the area well. In general, the laser drilling process causes more micro-cracks than does the EDM process because it involves a number of laser-induced explosions in the material to drill the hole instead of a steady etching away of material. Likewise, the laser process also leaves more slag since there is no removal mechanism other than circulating air. The EDM process must be done with the part and electrode submerged in some sort of dielectric fluid. Circulation of this fluid helps to remove material from the holes. Both the laser and EDM processes operate sufficiently within the required tolerances, though, and material property variations are not a significant factor in the quality of the parts.

Back wall strike can occur in both the laser and EDM process, but it has been virtually eliminated with current processes. In order to prevent a laser from striking the back wall, a filler material is injected into the vane's inner cavity. This filler material diffracts light and therefore can take a large number of pulses before a hole is drilled through to the back wall. Unfortunately, adding and removing this backing material is very time intensive. Some laser manufacturers have methods for breakthrough detection, but none of them have proven to work well enough to prevent damage to the back wall of the vane cavity. To prevent an EDM from drilling the back wall of a vane, the electrode feed rate is monitored and when it jumps up suddenly, this is a sign that the electrode no longer has resistance and therefore has broken through the wall. Even with very small cavities this provides sufficient indication of breakthrough that the process can be stopped before the back wall is damaged.

It is important to understand the exact manufacturing process at Pratt & Whitney in order to understand how this machinery is operated. There are two primary hole drilling operations: the drilling of shaped holes and the drilling of round holes. Each of these processes has a number of additional, auxiliary processes that depend on the technology used for drilling the holes.

As described above, shaped holes are always drilled with a comb EDM, so the process is relatively constant. The vane is first mounted into its tooling using a gauge fixture. This typically takes a few minutes. The tooling is then mounted into the machine. This takes roughly a half of a minute. Next the machine is started and automatically cycles through the drilling sequence. This assumes, though, that there are electrodes ready for the machine to use. Sometime during the cycle the electrodes must be replaced so that the machine has fresh electrodes to drill the part. The electrodes are loaded into the electrode tooling and then the electrode tooling is loaded into the machine. Loading electrodes requires significant manual time. The electrodes must be removed from their container, inspected, sanded to improve conductivity with the tooling, and mounted in the tools, all without bending any of the delicate comb teeth. The tooling is then placed into the machine. This process can take a couple of minutes per electrode. When the machine finishes drilling holes in the part, the tooling is removed from the machine and the vane is removed from the tooling. Since comb EDMs use an oil based dielectric fluid, the parts must then be either washed or burned out. Washing can take on the order of five minutes, while the burnout cycle takes a total of fifteen minutes although it is a two-stage process which indexes every seven and a half minutes, meaning that two parts can be in a burnout oven at any one time. After the washing or burning out is finished, the process for drilling shaped holes is complete.

Round holes can be drilled with a laser or with a HSEDM. The process for drilling round holes with a HSEDM is very similar to the process for drilling shaped holes. First the part is loaded into the machine. The part does not require a gauging tool to mount it into the tooling since the HSEDM has the capability to probe the part and compensate for any mounting inaccuracies; it is simply "soft-loaded" into the tooling, which is kept mounted in the machine. Replacing electrodes is also different for a HSEDM. Electrodes are removed from a storage container and inserted into a collet which is then tightened to grip the electrode. The collet is then inserted into a test chamber, which forces water through the electrode, and the operator checks for leakage. Assuming that there is no leakage, the collet and electrode can then be inserted into the machine. Once the vane has been drilled it is removed from the machine and put into a vibration sander. The sander removes any built-up slag from around the holes. Some additional manual sanding may be required to remove all of the slag. Since the HSEDM uses distilled water as the dielectric fluid, there is no need for further washing or baking.

An even more complex process must be followed when a laser is used for drilling round holes. The first step in this process is to inject backing material into each vane. For holes that are not drilled through the ends of the vanes, a machine can automatically do this, but it takes a few minutes. For holes that are drilled through the ends of the vanes, though, a different backing material must be used. The normal backing material is not adhesive enough to stay on the ends of the vanes. Typically this is added after the normal backing material is injected into the vane cavity. The vane can then be mounted into the laser. Currently, the vane must be mounted to tooling in a gauge fixture and the tooling must be mounted into the machine, but, as with the HSEDM, the newer lasers have probing capability which will allow the vanes to be soft-loaded into the machines. The laser then automatically cycles through the hole drilling sequence. Once the drilling is finished, the vane can be removed from the machine. The vane must then be placed in a burn-out oven to remove the backing material. After it is burned out, the vane is placed in a vibration sander to remove any slag build up. Some additional manual sanding may be required to remove all of the slag.

Turbine airfoil holes must be drilled to very exacting tolerances, yet in order to meet production requirements they must be drilled very rapidly as well. Laser drilling, EDM drilling, and HSEDM drilling are the prevalent methods used to drill these holes. While laser drilling is much faster than EDM and HSEDM drilling, laser drilling is generally more prone to variation, especially with respect to hole size as seen through the airflow measurement.

2.2 Background of the Culture at Pratt & Whitney's Turbine Module Center

Pratt & Whitney has been manufacturing aircraft engines since 1925 and in doing so has built up a considerable culture. The people at Pratt & Whitney pride themselves in their products. They firmly believe that they build the best aircraft engines in the world, and as such they feel a strong rivalry with GE Aircraft Engines and Rolls Royce. Pratt is an engineering-driven company that has recently been trying to improve its operations and manufacturing capabilities. Yet, although it designs cutting edge products, their manufacturing operations can be improved.

That Pratt & Whitney is an engineering-driven company can be seen in the fact that since the dawn of the jet age, it has held the majority of military engine contracts in the United States. Military engines are more complex than their commercial counterparts. They use more cutting-edge technology and they are built to be operated much more aggressively and in much harsher environments. Only the best engineered engines survive the military's grueling competitive selection process. Pratt & Whitney has triumphed over and over, and this is a tribute to the engineering focus of the company. As a testament to its strong engineering focus, Pratt & Whitney won the 1952 Collier Trophy for its development of the J57 engine and the 2001 Collier Trophy for the integrated Liftfan propulsion system for the Joint Strike Fighter. The Collier Trophy is awarded annually to those who have made significant achievements in advancing aviation. Pratt was also given special recognition in 1970 for its part in making the 747 functional. GE has won the award only once for aircraft engine design, and Rolls Royce has never won the award. These achievements further reveal Pratt's focus on engineering.

The fact that Pratt & Whitney is striving to improve its manufacturing can be seen in its recent history, when it has been working to adopt Lean manufacturing. During the 1980s, Pratt made efforts to "Lean out" its manufacturing. Operations managers had come to realize that they were wasting a lot of money building parts that differed in name only. Parts were constantly being re-engineered for minor gains in performance and each successive part type was often designed to use a totally new manufacturing technique. Pratt attempted to implement cross-functional teams, including Integrated Product Development cross-functional teams, to ensure that manufacturing and procurement, as well as engineering, had a voice in new design changes. Also, Total Quality Management was introduced as "Q-plus." Pratt tried to standardize its materials and its manufacturing processes and tried to physically align its manufacturing processes in the order of production. From 1983 to 1991, it reduced the average distance traveled by a jet engine part from eighteen miles to nine.[1]

In 1991, Pratt & Whitney found itselfs confronted with a substantial reduction in military engine sales due to the end of the Cold War and a boom of civilian orders that would soon subside. Art Byrne, the head of Wiremold, a nearby company that had recently gone through a Lean revolution with the help of the Japanese consulting firm Shingijutsu, acted as an example for the senior management at UTC. Management decided that these Lean principles were what were needed for them to improve their operations. Shortly after George David took over as head of UTC, he was able to hire Shingijutsu to consult and lead the Lean transformation at Pratt. George David also transferred Karl Krapek, the CEO of Carrier—another UTC company which Krapek was leading through its own Lean transformation—to head Pratt and to implement Lean manufacturing there. Throughout the 1990s, Krapek and his team at Pratt introduced a number of Lean improvements at the company. They began with the Middletown final assembly plant but soon branched out as they realized that they could obtain only limited gains in the final assembly facility until the plants feeding it parts were improved. The majority of Pratt's middle managers and shop floor workers had worked at Pratt for many years and were reluctant to change. Often, improvement activities would lead to dramatic improvements, but these improvements could not be sustained as managers and workers trended back toward their old ways as support was shifted elsewhere. Also, many improvements depend on related operations or facilities and the improvements could not be sustained without the related facilities being improved; however only a finite number of projects could be done.[2]

---

[1] Womack and Jones 1996, 164-167

[2] Womack and Jones 1996, 167-174

The Turbine Module Center was one of the many facilities where Lean manufacturing was implemented. Even in 1993, the TMC was facing the same challenge that it faces today. Turbine airfoils operate in the hottest part of a jet engine. They are subject to extreme conditions and they wear out frequently. The turbine airfoil business is actually focused on supplying replacement parts. In fact, jet engines are often sold at a loss, knowing that profits will be made through replacement parts such as turbine airfoils. Therefore a profitable turbine airfoil business is essential to the company. In 1993, the TMC was in the process of implementing Lean manufacturing, but this was causing shipping schedules to be missed and back orders to skyrocket, greatly impacting Pratt's cash flow. Ed Northern was brought in as general manager from Inter Turbine, a small turbine airfoil refurbishment company. Northern had experience with Lean manufacturing from his time at GE Aircraft Engines. After a one-time headcount reduction of 40% of his employees, he began to push Lean methods at the TMC. Flow lines were introduced in product oriented work cells, new Lean machines were procured, part design changes were limited, and single piece flow was introduced in as many places as possible. Between 1993 and 1995, the value of overdue parts fell from $80 million to zero, inventory dropped by half, manufacturing costs were reduced, and labor productivity rose dramatically. At the same time, the TMC, as with the rest of Pratt & Whitney, went through a quality revolution where Lean principles were applied to improve the quality as well as the throughput in the factories.[3]

After using Pratt & Whitney as a stepping stone Ed Northern left Pratt & Whitney shortly thereafter for a job elsewhere, and the TMC continued on its journey toward Lean manufacturing. At the same time Pratt & Whitney introduce a program called Achieving Competitive Excellence (ACE) which is its Lean manufacturing system. While Ed Northern was able to use a crisis to help enact change, after he left and the crisis abated it was much more difficult to keep up the improvement pace. Compared to the rapid achievements of the Ed Northern period, the Lean transformation was largely stagnant during the remainder of the 1990s and early 2000s, although certain improvements were implemented. When the factory was moved from New Haven to East Hartford the hole drill and coating facilities were organized into a cellular structure. High speed grind lines that work with single piece flow were introduced, and a new method to fixture turbine blades, which no longer required the blades to be encased in

---

[3] Womack and Jones 1996, 176-180

resin, was developed. In 2003, a new, highly sophisticated hole drill line was created utilizing HSEDM machines that incorporated significant automation and inter-machine communication. This line turned out to be a failure because of technical problems. It was essentially designed to do too much, and therefore ended up doing nothing well. Although the line failed, it was a clear sign that improvements were still being attempted in the TMC and that there was significant support for them.

Currently, the TMC is a conglomeration of Lean states. Some parts of the TMC, such as the high speed grind lines, show significant Lean manufacturing influence. Other parts, however, such as the vanes hole-drill cells and the coatings cells, are far from Lean. This thesis is concerned with vanes hole drilling and therefore will describe the current state of those cells in detail.

Drilling holes in turbine vanes takes place in five hole-drill cells. Each cell is part of a larger product cell as shown in Figure 2-1. There are five product cells. Product cells 1 and 4 are for processing commercial vanes while cells 2 and 3 are for military vanes. Product cell 5 is for a new military vane. Each product cell has a grind cell, a hole-drill cell, and a final inspection cell. These various parts are loosely coupled and, except for the grind cells, operate only in loose coordination.

Figure 2-1

Vanes Business Unit

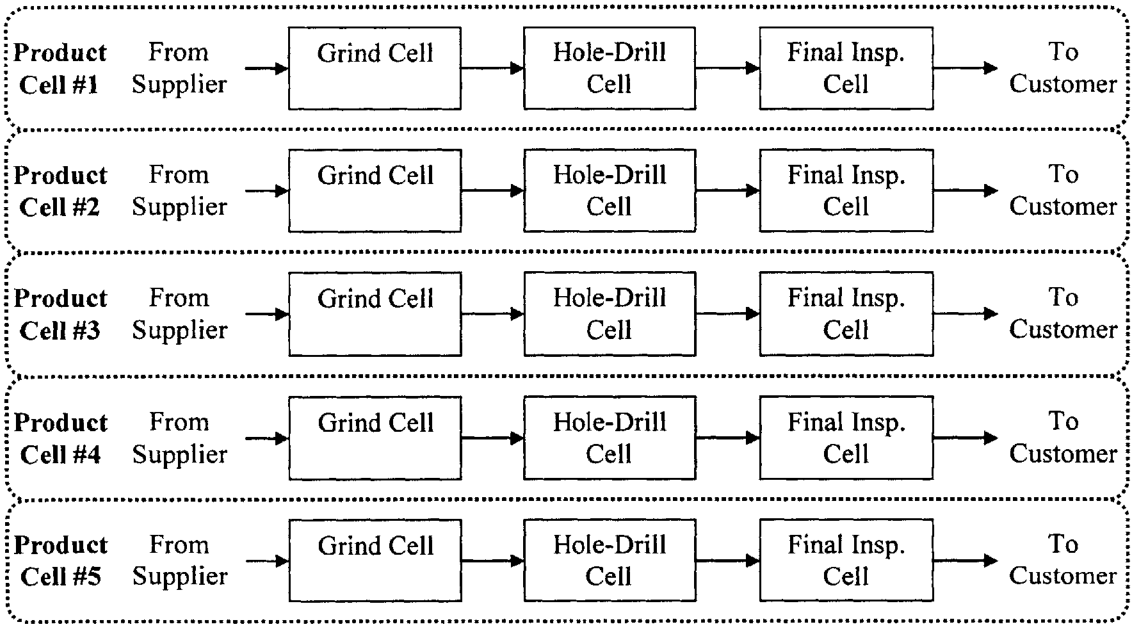

The general process for producing a turbine vane is that after being received from a casting vendor the vane is ground in the grind cells then taken into a different part of the factory to be coated. It is then returned to the hole-drill cell, where holes are dilled. After that, it may be subjected to further coatings or moved to the final inspection cells. After final inspection, it is shipped to the customer which is either the engine assembly facility or an outside customer if the part is to be a spare.

Parts are rarely drilled in a single product cell. All of the hole-drill cells, regardless of the product cell to which they belong, typically operate together as a quasi functional group. Also, while there are only five product cells, there are almost twenty different products that must be worked on the cells. A few products are scheduled to run at any one time and the products are scheduled sequentially in an attempt to meet the demand. There is a significant amount of downtime on the hole-drill machines. Additionally, the machines tend to be very idiosyncratic in that they process certain parts much better than other parts. This means that although there are quite a few hole drilling machines across all of the cells, only a few of them actually process parts well, which means that if one of them is down it significantly affects production. These machines are also quite old. They average about twenty five years in age which is quite old for laser or EDM drilling machines. Not only does this contribute to their downtime, but it also causes them to display a significant amount of operational variation. This drives up the rate of inspection required for parts and therefore hampers production. It also contributes to rework which further hampers throughput. Since the machines are operated more as a functional group than adhering to their product cells, each machine typically has its own operator, who spends most of the time merely watching the machine run.

The shop floor workers have all been with the company between twenty five and thirty years and are quite set in their ways. When asked about Kaizen events and continuous improvement projects, the typical reply is: "I used to take part in those, but they never resulted in anything changing, so I avoid them now." Likewise, during a recent ACE training event conducted by a Lean consultant no participants showed up event though a number of people were required to attend. Shop floor managers do not enforce ACE training requirements and do not promote ACE initiatives. This is primarily because their priority is to ship parts. As more than one manager at TMC has said, "You're not going to get fired by not doing continuous improvement, but you will get fired if you don't ship enough parts." So managers have their workers man machines or fill gaps in attendance rather than learn about Lean. Also, the union environment does not support Lean initiatives. Due to a falling out between Pratt and its union when ACE was first implemented, the union contract specifically states that no employee can be obligated to take part in a project that has the purpose of continuous improvement. There is such animosity between the union and the ACE program that recently during a project meeting a union representative exploded at a manager when the manager introduced himself as the "ACE manager" rather than as a "continuous improvement manager." Yet, other union members hold ACE in a higher regard and in fact have volunteered to become "ACE Pilots," which are ACE program representatives on the shop floor. These ACE Pilots can have a strong impact on how Lean is embraced by the workforce, but they are typically kept busy collecting ACE metrics data and participating in Kaizen events instead of mentoring shop floor workers and junior managers with regard to ACE initiatives.

The culture at Pratt & Whitney reflects its recent history of improvement initiatives. Through the 1980s, 1990s, and 2000s, Pratt has attempted to implement cross-functional, Lean thinking and has had varying success. The culture at the company, and especially at the TMC, therefore, is one of "been there, done that" and shows little enthusiasm for trying new ideas, especially Lean ideas which they feel have already been tried and proven to be failures. It is particularly difficult to implement new changes in this type of culture.

2.3 Literature Review

A considerable amount has been written on topics related to the objectives of this thesis. The first two objectives are concerned with designing work cells and purchasing the appropriate equipment for the work cells. Lean manufacturing texts discuss the basic principles behind how to design lines of flow and a number of previous LFM theses describe different ways to schedule high mix product lines. The third objective is an extension of the first two, although Lean manufacturing texts also describe the attributes of Lean tooling, which are important when selecting one machining process over another. The fourth objective, which is concerned with the implementation of the Lean methodology, is also covered in detail in Lean manufacturing texts.

The primary Lean manufacturing texts that were valuable to this thesis were *Lean Thinking* and *The Machine that Changed the World*, both by James P. Womack and Daniel T. Jones, and *Lean Production Simplified,* by Pascal Dennis. While the books by Womack provided excellent background information, theory, and ample examples of Lean implementation, *Lean Production Simplified* provided much more specific instructions that were directly applicable to creating Lean work cells. The examples of Lean implementation at various facilities that were covered in *Lean Thinking* were especially useful in shaping the plan for implementation at the TMC.

An understanding of the Theory of Constraints is also vital to managing flow through a work cell, especially when variation exists in the manufacturing processes. Eliyahu M. Goldratt's books *The Goal*, *Critical Chain*, and *Theory of Constraints* were all used to understand the Theory of Constraints.

Stephen Muir's (LFM '97) thesis on optimal allocation of product loadings and equipment sharing helped to guide the efforts of this thesis and provided a great example for optimal part loading in a slightly simpler problem. Alison Page's (LFM '01) thesis on forecasting mix-sensitive tool requirements under demand uncertainties also provided some guidance as to how to meet the first two objectives of this thesis.

This thesis combines the ideas of Muir and Page into one strategy for acquiring tooling under demand uncertainty for a system where product loading is optimized. However, it is much more an extension of Muir's ideas than of Page's, since the uncertainty of demand that Page dealt with was left to be dealt with through running multiple repetitions of the optimization for various potential demand scenarios in this thesis. Although based on the Lean methods and design theories of Womack, Jones, and Dennis, this thesis goes one step further to better address the difficulties inherent to designing cells that must process a high mixture of products.

Chapter 3: Optimal Work Cell Design

This chapter describes the methodology that was used to design the work cells responsible for drilling turbine airfoil holes. The general methodology was to define how the work cells should function, build a mathematical model to describe this functioning, and finally to optimize the layout of each cell and the number of cells using mathematical programming.

3.1 Design Problem

Muir studied the capacity of various semiconductor factories that produced the same product lines and found that because of slight equipment variations between the factories, each had a different capacity which depended on the product mix that it was running. Muir found that the capacity of a factory could be changed by varying its assigned product mix. Therefore, when the factories were considered together—what he called a virtual factory—the overall capacity of the virtual factory could be maximized. He maximized the overall capacity by dividing up the overall demand so that each factory was assigned a product mix that was as similar as possible to its optimal product mix. Muir used Linear Programming to divide up the overall demand in order to maximize the capacity of the virtual factory.[4]

Muir provided the foundation for how this thesis looks at the capacity of the hole drilling cells at the TMC. Each hole-drill cell is like a factory in Muir's analysis. Each one is capable of processing at least some of the products that another cell produces, but because of differences in equipment between the cells, each cell has a different capacity for a particular product. Muir's optimization could be performed on these cells to optimize their capacity by properly allocating product mix to each cell. However, the objectives of this thesis are to inform decision making about equipment to purchase, and therefore how to design and implement the hole-drill cells.

This problem goes one step beyond Muir's optimization. Instead of making a decision on the assigned product mix for each cell to optimize the capacity, the objective is to manipulate the equipment in each cell so that sufficient capacity can be attained to meet demand for the minimum monetary cost. Essentially this is a double optimization. The objective is to optimize the number of cells and design of each cell assuming that the resultant configuration will have a product mix that is optimally determined by a Linear Program similar to Muir's.

---

[4] Muir 1997

One might question whether this is all necessary. One might argue that if Muir had found that each factory worked exactly the same way, he would have had nothing to optimize since the product mix would have affected each factory the same and the simple answer would therefore have been to allocate the same product mix to each. Yet this is not the same thing as saying that had Muir been able to design the factories that he would have wanted to make them operate identically. Neither would this thesis desire to make all of the hole-drill cells in the TMC identical. An example will help to illustrate why.

Assume there are two vane products that need to be run through a hole-drill factory consisting of two cells. The demand for vane A is 20 units per month and the demand for vane B is 80 units per month. Vane A differs from vane B in that it requires that it be drilled with machine Y and machine Z where vane B only needs to be drilled by machine Z. Machine Y has a capacity of 20 units per month and machine Z has a capacity of 20 units per month. If the cells were identical, both cells would have one Y machine and three Z machines and the product mix would be the same on both cells. The utilization of Z machines in each cell would be (10 A + 40 B) / (20 x 3) = 5/6 = 83.33%. The utilization of each Y machine in each cell would be 10 A / (20 x 1) = ½ = 50%. However, if cell #1 had four Z machines and no Y machines and cell #2 had one Z machine and one Y machine, the utilization would be different. In cell #1 the Z machines would be utilized 80 B / (20 x 4) = 1/1 = 100%. In cell #2 the Z machine would be utilized 20 A / (20 x 1) = 1/1 = 100% and the Y machine would be utilized 20 A / (20 x 1) = 1/1 = 100%. This demonstrates how having different cell types can be more efficient than having identical cell types, depending on the demand of each product and the capacity of each machine.

One might then argue that each product should have its own cell as in the example above. Depending on the characteristics of the products and the machines, this might not be the best solution either. If a single machine has a capacity far in excess of the demand for one of the products and this machine is applicable to more than one product, it would not make sense to purchase two of these machines, one for each cell running only one product. There is also a concern about redundancy. If only one cell can produce a given product, then should that cell become inoperable, a backlog for that product would form. In the TMC, another problem exists with this potential solution. There are almost twenty different product types that all have different processing characteristics for hole drilling. It would be impractical to have twenty different cells, one for each product. Because of the available machinery each cell would have a capacity far in excess of demand and the machine utilization would be extremely low.

The current methodology for designing work cells to handle multiple product lines is to consolidate those product lines into as few product families as possible. This is certainly a reasonable solution. By examining the characteristics of each product, products can be grouped together by characteristic. Typically, a product attribute chart is drawn to compare the products. Along the top of this chart are listed all of the processing steps required by all of the products. Along the side are listed all of the products. X's are entered in the boxes at the intersection of the products with the steps that are applicable to them. By looking at the chart one can visually identify similarities between products and thus form product families.

As an example of this, assume there are four products A through D. Assume between all four of these products, seven processing steps, T through Z, are required. The product attribute chart might look like the one below.

Figure 3-1

Steps:	T	U	V	W	X	Y	Z
Product A	X	X	X				X
Product B		X	X	X	X	X	
Product C	X		X				X
Product D		X	X	X	X	X	

The chart makes it clear that products B and D should be grouped into one family. Likewise, although not a perfect match, the chart shows that products A and C can most likely be grouped together as a family, understanding that product A requires one additional step. It might make sense in this case to have two cells. One cell could perform processing steps U through Y and would process parts B and D. The other cell could perform processing steps T through V and Z and would process parts A and C.

Yet this methodology misses one very important factor. It assumes that every product requires an identical amount of processing time at each step, and by doing so it misses load balancing each product. Replacing the X's in the above product attribute chart with processing times reveals that the products are not as similar as initially thought.

Figure 3-2

Steps:	T	U	V	W	X	Y	Z
Product A	1	2	2				3
Product B		2	2	3	3	5	
Product C	2		4				6
Product D		2	2	2	1	1	

The significance of the difference in processing times is that each product requires a different number of machines to balance the cell. Any imbalance in the cell causes inefficiency and poor machine or worker utilization. For the example above, products A and C would still work well as a family of products since they require the same ratio of equipment for each step. To balance the cell there would need to be three times as many machines working step Z as working step T and twice as many machines working steps U and V as working step T. Since this is consistent between product A and C, those products would still make a good family. Products B and D, on the other hand, would not constitute a very good family of products. To balance the cell running product B, the cell would have five machines doing step Y and three machines doing steps W and X for every two machines doing steps U and V. But for product D there would only need to be one machine doing steps X and Y and two machines doing step W for every two machines doing steps U and V. This means that if a cell were built to balance the load running product B, while this cell was running product D many of the machines doing steps W, X, and Y would not be utilized and therefore utilization would drop.

In most circumstances, and certainly at the TMC, this problem is unavoidable. Since each turbine vane product has a different number of holes, and since the drilling times are dependent on the number of holes in each part, the products are inherently going to have different processing times. So the question is which parts should be grouped together, taking into account their processing times, to maximize the utilization and therefore minimize inefficiency. This is where optimization comes in. By comparing different products' processing times, an optimization routine could combine the products together into a family to be processed in the same cell while minimizing the amount of equipment purchased to create these cells. This is the methodology followed by this thesis.

3.2 Lean Standards

A mathematical model must be built before an optimization can be performed, but before a mathematical model can be built, there must be very strict rules regarding how the work cells operate. These rules make the cells operate predictably, which is necessary if they are to be described with mathematics.

If there must be strict rules regarding how the work cells operate it only makes sense for the rules to promote the efficient operation of the cells. Lean manufacturing and the Theory of Constraints are two tools promoted by modern operations thinkers. Lean manufacturing has taken root at Pratt & Whitney in the form of the Achieving Competitive Excellence (ACE) operating system. Therefore, the operating rules for the cells being modeled for this thesis were designed using Lean manufacturing principles and following the fundamentals of the Theory of Constraints.

In its most basic form, Lean manufacturing is the systematic identification and removal of any waste associated with manufacturing activities. Lean manufacturing strives to remove all non-value-added activities and to simplify and make efficient all value-added activities. Taiichi Ohno of Toyota identified the basic types of waste. Motion, delay, conveyance, correction, overprocessing, inventory, overproduction, and knowledge disconnection are all forms of waste that can be readily seen in any factory. Lean manufacturing employs various tools to reduce these wastes. For example, single piece flow and pull systems are used to eliminate excess inventory and delay while Lean work cell layout helps to eliminate excess motion and conveyance.[5] It is far beyond the scope of this thesis to describe Lean manufacturing in detail but a familiarity of it is necessary to best understand the scope of this work. For a more complete description, *The Machine that Changed the World* and *Lean Thinking* by Womack and Jones are excellent references.

The aspects of Lean production that are most relevant to designing work cells to drill turbine airfoil holes are cycle times, standardized work, rapid changeover, and cellular machine layout. Defined cycle times are what links together each part of the cell. Cycle times provide an interval at which parts are moved from one operation to the next. By adhering to the cycle times, no inventory builds up between operations since even if one operation is faster than another, the part does not get indexed until every operation in the cell is complete and every part can be

---

[5] Dennis 2002

indexed to the next operation. The cycle time is what determines the throughput of the cell. It is defined by the slowest process in the cell, but this can be minimized by leveling the loading in the cell. For instance, if operation A takes one minute and operation B takes two and a half minutes, then one would want sufficient resources in the cell to do two to two and a half B operations for each A operation. This would effectively cut in half the cycle time of operation B. In this example, the overall cycle time of the cell could be something on the order of one minute and fifteen seconds. Strict cycle times and load leveling are what link Lean manufacturing to the Theory of Constraints. The Theory of Constraints shows how a manufacturing system can only produce as fast as its slowest—bottleneck—operation, and dictates that to increase throughput the bottleneck resources must be expanded.[6] Operating with a cycle time limits a cell's throughput to the operation with the slowest cycle time (the bottleneck), and load leveling helps to maximize the efficiency of the system by ensuring that there is not a great disparity between the resources of the bottleneck and the resources of the other operations.

Standardized work is another key ingredient to Lean manufacturing. It helps to enable a predictable and stable cycle time by defining the exact manipulations that each operator must perform each cycle. If the operator deviates, it is likely that an operation will not be completed in the cycle time. Standard work also allows load leveling of the manual work required in a cell. It can be used to divide up the manual time so that each worker works the same amount of time during each cycle. This helps to maximize labor utilization as well as to prevent a worker from being overworked. As long as the assigned cycle time for the cell is greater than the standard work time for any operator, the operators will have sufficient time to complete the required tasks. Of course the standard work and cycle time must take into account potential problems so that the line is not constantly being stopped because a worker is not given sufficient time to complete her tasks in light of common problems. At the same time, the cycle time should be close to the standard work time so that if there are problems in the cell, they do cause cell stoppages. Although stoppages decrease throughput and productivity, they help to bring attention to the problems so that the problems can be solved permanently. Work stoppages incentivize managers to commit the proper resources to solving problems that cause reoccurring stoppages.

Rapid changeover allows the line quickly to be switched from machining one product to machining the next. This is crucial for work cells which must machine more than one product,

---

[6] Goldratt 1984

such as the hole-drill cells in the TMC. Reducing the changeover time from one product to the next reduces the optimal batch size for running each product. This allows the line to be switched over more often, which allows greater responsiveness to changes in demand or disruptions in the part supply to the cell.

Cellular machine layout helps to reduce the cycle times of the cells. Having each machine in close proximity to the rest and having the machines lined up in order of operations allows parts to be taken from one operation to the next without wasting time. It also makes managing the cell easier because the parts travel in a defined path and thus the cell's state of production is obvious.

These principles helped to provide the structure for the mathematical model. Most importantly to the mathematical model were the assumptions that the number of machines and number of workers define the cycle time of the cell. This is enforced through standardized work, which ensures that personnel are utilized well, and through strict adherence to cycle times, which ensures that machines are utilized well since the cycle times are defined primarily by the automatic time of the machines.

3.3 Work Cell Mathematical Model

The mathematical model for the hole drilling work cells is actually two separate models. The first model is an optimization framework that allows the determination of cycle times, capacities, and cost for various cell types and facilitates comparing them to find the best combination of appropriately sized cells. The second model is a cost model that uses the results of the optimization model to explain the financial impact of the modeled state. This section only deals with the optimization framework. The cost model is described in Section 4.1.

Section 3.1 describes how optimization can be used to choose the number of machines to purchase and how to deploy the equipment in work cells. Although the two primary hole drilling operations are drilling shaped holes and drilling round holes, only about half of the vanes have shaped holes and require being processed by both of these operations. Optimization is not required to level the loading in a cell that has only one operation, but it can help identify the number of machines to use in that cell to meet demand and whether that cell should be one big cell or two smaller ones, and it can still determine how round hole only cells fit in with the rest of the cells, if at all.

For this thesis, the optimization model is constructed as a complex spreadsheet that performs two functions. First, it provides a model for cycle time and manual time for every part family for various cell configurations. Second it provides a framework for optimizing the number and type of cells to be designed into the factory. The optimization function is a result of the ability to change the model while observing a utility function until the maximum utility is achieved. This model will be described from the micro level to the macro level.

The foundation of the model is a set of calculations that define the automatic and manual time associated with various individual operations that take place in hole-drill. For instance, there are calculations for how long it takes to drill every part with comb EDMs, HSEDMs, and lasers. The calculations are essentially based on multiplying an average time per hole for each process by the number of holes in each type of part. For comb EDMs, though, since they drill whole rows of holes at a time the calculation is the product of the number of rows by the average time to drill a row of holes. Manual time for various operations is calculated in a similar fashion.

The next level of the optimization model combines the times for these operations into total times for the entire hole drilling process for each part. Determining the overall cycle time inherent to the cell requires knowledge about the design of the cell. Yet one of the purposes of the ensuing optimization is to choose the appropriate cell designs. The principles described in Section 3.2 set the ground rules for how the machines are operated and how they are linked together, but the performance of the cell still depends on the number of each type of machine in the cell and the number of operators working in the cell. Although there are almost limitless combinations of equipment and machinery, the number of parameters was limited in order to allow optimization in a reasonable timeframe. Three parameters were varied.

The first of the three parameters was the number of comb EDMs in the cell. The second was the number of round hole-drillers in each cell (either lasers or HSEDMs). The third was the number of operators required to operate each cell. These variables were limited to four values in order to keep this part of the model a reasonable size and in order to limit the computations necessary for the optimization function. Cells could have 2, 4, 6, or 8 comb EDMs. Cells utilizing HSEDMs could have 4, 10, 12, or 14 HSEDMs. Cells utilizing lasers could have 2, 4, 6, or 8 lasers. Cells could have 1, 2, 3, or 4 operators. These values were determined through trial and error while analyzing various trial cell configurations and are constrained by practical factors such as the space available for a cell. For vanes without shaped holes, the cells were considered to have no comb EDMs, and only the numbers of round hole-drillers and operators were considered. This resulted in 128 combinations of cycle times being calculated for parts with both round and shaped holes (64 combinations utilizing lasers and 64 utilizing HSEDMs) and 16 cycle time combinations were calculated for parts with only round holes. This part of the model would continuously update all of these cycle times for any change in the automatic or manual time associated with a given operation.

Each of these combinations of equipment represent "mini-models" of potential cell types for each product. The optimization then subsequently chooses the best combination of these "mini-models" instead of having to manipulate the number of machines and workers during the optimization. This method was chosen due to the highly non-linear effect that equipment and manning has on the cells. The process of optimization through choosing from the mini-models is challenging, but much simpler and faster than if the optimization also had to calculate the cycle times for each combination of machinery and workers that it was comparing.

Besides calculating cycle times, this level of the model also calculates a utility value for each mini-model. The utility value is calculated in the same way as the one used for optimization, which is detailed in subsequent paragraphs. This utility value shows which mini-model is preferred by each product as if only that product were produced on the cell type that the mini-model represents. The model then takes the average of these preferences weighted over the expected demand of each product to determine the best overall mini-model as if all products had to be produced on a single cell type. These results help to give the user of the model an idea of which cell types are preferable to each product and which types would likely be chosen by the subsequent optimization.

The next level of the model is where the optimization is implemented. It allows for the optimization to choose up to eight hole-drill cells as if there were that many product cells and to assign production volumes of products to each. Although there are only four grind cells, the coating process decouples the grind cells from the hole-drill cells. The number of final inspection cells could easily be manipulated to match the optimized number of hole drill cells. A product cell then would likely contain only hole-drill and final inspection cells with the grind cells operating as a separate entity.

The optimization works as follows. Demand is one input. The decision variable is the volume of each product assigned to each cell. By assigning a volume of a certain product to each cell, the remainder (demand minus assigned) is reduced. Assigning parts to a cell also causes the utilization of the cell to rise. Keeping the utilization below a nominal value and driving the remainder to zero means that the cells chosen have sufficient capacity to satisfy the demand. The utilization of the cell depends on the parameters of the cell (principally cycle time and manual time), which come from the mini-models. These parameters are also fed into the cost model, which is discussed in Section 4.1. The optimization part of the model happens in coordination with the modeling function described above. The optimization essentially works to maximize a utility function:

$$Utility = \frac{TotalproductionTimeAvailable}{CycleTime \times CellCost^2}$$

This utility function essentially expresses the amount of times that a part can be produced in the given available production time (*TotalproductionTimeAvailable/CycleTIme*) divided by the cost of the cell squared. Expressing how many times a part can be produced in a given time gives an idea of how well that cell type produces the part and therefore how efficient it is. This is a good quality to have as utility when the overall objective is to find the most efficient combination of work cells. Dividing this quantity by cost ensures that cost is also taken into consideration. The cell cost is squared to properly scale the function to the importance of cost. The square function was determined through trial and error. Multiplying the cell cost by a factor would not have any effect on the optimization since this is a ratio, so the only way to increase the effect of the cell cost is to raise it to a power. The square function was chosen as its results most closely align the utility value to the savings determined by the cost model. Yet it is important to understand that the objective is not only to reduce costs, but to find savings given highly efficient work cells.

The actual optimization, however, is slightly more complex than the utility function described above since the cycle time varies for every product. The actual optimization is:

$$Max\left(\frac{\sum_{ChosenCells}\left(\sum_{PartsAssigned}\left(\frac{TotalproductionTimeAvailable}{CycleTime \times PartDemand}\right)\right)}{\left(\sum_{ChosenCells} CellCost\right)^2 \times Factor}\right)$$

$$\text{S.T.}\quad \left\{\begin{array}{l} PartDemand - PartsAssigned = 0 \\ PartsAssigned = Integers \\ \sum\limits_{PartsAssigned} CycleTime < TotalproductionTimeAvailable \end{array}\right\} \qquad \begin{array}{c} Decision \\ Variables \end{array} = \left\{\begin{array}{l} PartsAssigned \\ ChosenCells \end{array}\right\}$$

Assigning parts to chosen cell types is a simple linear problem. The solution is to assign parts to the cells in the order of how efficiently each cell processes the part. Unfortunately, the task of choosing cell types is highly non-linear. It is non-linear because the cycle times for each product vary disjointedly with the number of machines and operators in a cell. The disjointedness is due to the fact that the bottleneck shifts between different equipment types or personnel as the number of each type of equipment and the number of operators is varied. Traditional non-linear solvers (e.g., Excel Solver) do a very poor job of finding the global optima for this problem and instead tend to get stuck on local optima. A number of sub-optimization routines were tried with varied success. Typically, a user of the model can do almost as well as these routines by doing “manual optimization.”

The model assists the user in doing this “manual optimization” by displaying the utility of the operator’s choices and showing which cells most efficiently drill each part. This information helps the operator to pick the best cells. Although this in no way provides a truly optimal solution, the solutions can be quite good. The automatic optimization routines are discussed in detail in the next section.

3.4 Optimization Routines

To recap, the optimization routine must choose between zero and eight cells from the possibilities modeled with the mini-models. It must then assign production volume for each product to these cells so that demand is met. It must work through these tasks in such a way as to find the number and types of cells to choose that allow the maximum utility, which essentially meets demand with the least monetary expenditure. As discussed above, this is a highly non-linear problem due to the fact that adding machines or workers shifts the bottleneck around in the cell and results in the cycle and manual times changing non-linearly. Since the optimization must assign whole numbers of parts to each cell, the optimization is an integer problem in addition to being non-linear.

Standard non-linear programming is similar to linear programming in that it can be thought of geometrically as finding the highest or lowest point in a region bounded by lines, surfaces, or hyper-surfaces—depending on the dimensionality of the problem—which are determined by the constraints of the program. For a linear program those lines are always straight, the surfaces are always planes, and the hyper-surfaces are always hyper-planes, which is why it is called linear. For non-linear programs, though, the lines are curves, and the surfaces and hyper-surfaces are irregular and curved. For linear programming it can be shown that the optimal point will always be at the intersection of the constraint surfaces. This guarantees that an optimal solution can be determined very quickly with a linear program since only the intersections need be checked. For non-linear programming though, the optimal point can be anywhere along the surface. This makes finding the optima more difficult and time consuming.[7]

Typically, Non-Linear Programming (NLP) involves algorithms that trace the constraint surfaces and follow the path with the highest gradient to a point where the highest gradient is zero. Unfortunately, the point on which the algorithm ends may be the optimal point, or only a local optimal point.[8] The easiest way to visualize this is to think of the problem as finding the highest peak in a mountain range. Using this type of algorithm, one would find the direction of steepest ascent from the starting point and work up the mountain until one could go up no further (i.e., the steepest gradient becomes zero). At this point, an optima has been found. But this might not be the global optima. The next mountain over could be taller, but since the process started on this mountain, it only ended at this peak. Most optimization routines then use bounds to help rule out local optima. For example, if one decided that the tallest mountain had to be at least 1000 feet higher than the starting point, one would ignore the first optima if it were only 700 feet higher. One would then continue searching until one found a peak that was at least 1000 feet higher than the starting point. This still does not ensure that the global optima is achieved, but it does decrease the probability of choosing a local optima.

These algorithms further increase their chances of finding the global optima by adjusting the step size for determining a gradient. Assume again that you want to find the tallest peak in a mountain range. Instead of looking for the steepest slope at your exact point and following that up, you could look at the elevations at every point exactly a mile from you. Then you would go

---

[7] Ragsdale 2004, 354-356

[8] Ragsdale 2004, 356-360

to the point that was the highest and again look at all points around you for a mile and choose the highest. You would continue to do this until you were at the highest point compared to all the points a mile around you. You could then cut your looking distance to a half mile. You would look at the elevations of all points around you at a radius of a half mile, choose the highest elevation and go there. You would continue this process, continue cutting down the look distance until you were actually at the peak. It is highly likely that this would be the global optima, but it is still not certain. This method could still miss a very skinny peak close to your starting point. The only way to ensure that you find the highest peak would be to check the elevation at every point in that mountain range, but that would be prohibitively time consuming. This example only covers two constraint dimensions (North—South, and East—West). The complexity of other problems with more dimensions would increase by a factor of the dimensions, as would the time it would take to look at all of those points and compare them.

Mathematical programming problems become even more difficult when there are integer constraints. Even linear programs become more time consuming. For a linear program the optima is always at the junction of constraints, but if it is an integer problem it is unlikely that the junction of constraints will occur exactly at an integer solution. Therefore a linear program must find all of the intersections and then compare all the feasible integer solutions near each junction to find the best one. The same issues exist with a non-linear programming problem, but it is even worse because of the additional number of potential solutions in an NLP.[9]

For the optimization problem at hand, traditional NLP techniques would be unlikely to work. Given the problem for choosing just one cell the optimization would have to look at all the feasible mini-models, and for each trial it would need to use a Linear Program (LP) to optimally distribute production volume for each product to the cell. Since each answer is based on the mini-model, and since the mini-models differ in such a non-linear fashion, the NLP algorithm would almost certainly find a local optima. In trials using Excel Solver's Standard GRG Nonlinear Solver, the solutions showed zero improvement over manual solutions. To pick a single cell it is possible to actually run through all of the options (approximately 210 total) and pick the best one. However, when the problem is expanded to choose more than one cell the time it takes to do so increases by a factor to the power of the number of cells, and the problem becomes unfeasible for practical purposes. Rudimentary calculations show that assuming each

---

[9] Ragsdale 2004, 244-251

calculation (including the running of the embedded linear program) takes one second, it would take approximately 1.4 x 10^{182} years to check all of the possible combinations for eight cells.

Since determining the optima through brute force calculations would take billions of billions of years, and since traditional NLP routines proved essentially worthless, a different sort of optimization had to be considered. Dynamic Programming (DP) is another type of non-linear programming. Dynamic Programming determines an optimal decision strategy for decisions made in stages. The problem can be restated as a Dynamic Programming problem if choosing the cells is considered to take place sequentially. Basically, this means that first the cell type is chosen for cell #1, then the cell type is chosen for cell #2, and on until all cells have been chosen. Likewise, if the problem was solved sequentially running through the different products, it could also be formulated as a DP problem. The general idea of DP is that subsequent decisions "learn" from earlier decisions, which helps to remove the necessity of working through all possible combinations, and therefore can determine the optimal solution in much less time. Mathematically, a deterministic DP algorithm looks like the following[10]:

$$x_{k+1} = f_k(x_k, u_k)$$
$$k = 0,1,\ldots,N-1$$

Where:

k indexes discrete time (or stages).
x_k is the state of the system and summarizes past information that is relevant for future optimization.
u_k is the control or decision variable to be selected at time k.

And the utility function is additive and represented in the form:

$$g_N(x_N) + \sum_{k=0}^{N-1} g_k(x_k, u_k)$$

Where $g_N(x_N)$ is the terminal utility for the last stage and $g_k(x_k,u_k)$ is the utility at each stage.

If this problem were to be structured as a DP it could be done in the following fashion. The problem would begin by looking at all the possible combinations of how the first cell could be solved. Unfortunately, there would be a very large number of combinations that would need to be considered. Assuming that there were 200 possible cell types to consider, that there were ten products, that the volume of each product was 100, and even conservatively assuming that

---

[10] Bertsekas 1995, 2-3

the cell had the capacity to run only 100 parts of any product, then there would be combinations on the order of 200^(100^10) = an incredibly large number. This number would be so large that it would take far longer to solve than solving all of the combinations as mentioned in the paragraph above. Why doesn't DP provide an advantage like it's supposed to? This is because in the combinations scenario only the combination of potential cell types and eight potential cells is considered. The allocation of demand for each product is taken care of by the LP that runs for each scenario. Since each combination has all eight cell types chosen, the LP can optimize across all cells. Therefore, the LP deals with how the demand is allocated between cells. For the DP problem, though, since each cell is considered sequentially, an LP cannot be used to allocate demand between cells. An LP could assign the allocation of each product demand to the cell, but since cells are considered sequentially, the resulting solution would not be optimal. The same issue would be encountered if the DP was formulated to take each product sequentially instead of each cell.

Dynamic Programming was considered because in addition to being an optimization methodology, there are a number of sub-optimization routines that are closely related to Dynamic Programming. The above paragraphs should make it clear that finding the optimal solution to this sort of problem is nearly impossible in a practical time frame. Standard non-linear programming techniques and dynamic programming techniques were shown to require too much computation to be of any use. Yet in most practical applications a suboptimal solution is sufficient if an optimal solution is too difficult to obtain. For this problem a suboptimal solution should still provide significant advantage over the traditional method of cell design, even if there could be a theoretically higher gain with an optimal solution.

Two popular sub-optimization concepts that are related to Dynamic Programming are rollout algorithms and policy improvement algorithms. Since they are related to DP, they approach problems as if there were sequential steps to solve. A rollout algorithm is a suboptimal policy based on a one-step lookahead policy, with the optimal utility-to-go approximated by the utility-to-go of the base policy. The base policy is derived from some heuristic or suboptimal policy. In simple terms the rollout algorithm determines the best choice for the current stage by comparing all of the current choices given that it knows what will happen during the next stage given each possible choice. Notice that this is a one-step lookahead. It only looks at the potential effects of each of the current choices on the next stage. It is possible for the rollout algorithm to be an *l*-step lookahead, where *l* is the number of subsequent stages that are examine for each potential choice in the current stage. This method is suboptimal because it uses a heuristic or suboptimal method (the base policy) to determine what will happen in the subsequent steps given the choice in the current stage. Mathematically, a rollout sub-optimization looks like this:[11]

$$Max\left\lfloor g_k(x_k,u_k) + \hat{J}_{K+1}(f_k(x_k,u_k)) \right\rfloor \text{ to find the best policy } u \text{ for each } k = 1,2,3,\ldots,N$$

Where:

$\hat{J}_{k+1}$ represents a utility-to-go for choices made according to the base policy.
$f_k(x_k,u_k)$ is the effect of u_k as the choice for the next stage given the state x_k.
$g_k(x_k,u_k)$ is the utility gained up to and including stage k.

Essentially for each stage the algorithm finds the policy that maximizes possible utility given an assumption of how the subsequent stages would be chosen.

This type of algorithm can be used to solve the problem at hand because the base policy can be substituted for the unknown choices of subsequent cells, allowing the algorithm to use an LP to assign product demand across all cells. Choosing cells could be considered the sequential stages. The base policy would give an assumption of the subsequent cells that have not yet been chosen. The algorithm would choose the best first cell assuming that the rest of the cells would be chosen according to the base policy and then run the LP to distribute demand across the cells. The algorithm would then choose the best second cell knowing the choice of the first and assuming that the rest would be chosen according to the base policy. This would continue through all eight cells. For the cell allocation problem the sub-optimization would be set up as follows.

$$\text{For } k = 1,2,\ldots,8 \text{ choose } u_k \text{ to } Max\left[\frac{g_k(x_k,u_k) + \hat{J}_{k+1}(f_k(x_k,u_k))}{\left(\sum_{i=1}^{8} Cost\right)^2} \right] \text{ where:}$$

$$g_k(x_k,u_k) = \left(\sum_{i=1}^{k} \sum_{PartsAssigned} \frac{Total\,\Pr oductionTimeAvailable}{CycleTime \times PartDemand} \right)$$

$$\hat{J}_{k+1}(f_k(x_k,u_k)) = \left(\sum_{i=k}^{N} \sum_{PartsAssigned} \frac{Total\,\Pr oductionTimeAvailable}{CycleTime \times PartDemand} \right)$$

---

[11] Bertsekas 1995, 315-334

Where:

x_k is the state of cell k. The state of the cell is its type.
u_k is the choice of the cell type for cell k.
PartsAssigned is determined with a Linear Program that distributes parts across all cells.
Cost is the cost of cell k.
CycleTime is a function of x which is the state of the cell, meaning the type.
$\overline{\mathit{CycleTime}}$ is a function of $\bar{x}$ which is the assumed state of cells not yet chosen according to the base policy.
TotalproductionTimeAvailable is a constant.
PartDemand is a constant.

For this problem formulation, x_k and u_k are essentially the same thing. However, it helps to think of x as the actual state of the cell where u is the policy that chooses that state. Therefore the point of this algorithm is that by knowing the x's and $\bar{x}$'s, the best u's can be chosen.

This methodology greatly simplifies the computations involved in finding a solution, but the quality of the solution depends on how good the base policy is. The simplest base policy would be to assume that all subsequent cells will be of a certain, constant type. For example,e it could be assumed that all remaining cells would be of a type that contained four EDMs and two lasers and was operated by three workers. This cell layout can process both ceramic and metallic coated parts as well as parts with both shaped and round holes so it would be a good "middle-of-the-road" cell. Or the remaining cells could be assumed to be of the type that the model shows drills all of them the best: the type that would have been chosen if all cells had to be the same. Likewise, the remaining cells could be assumed to be those found through manual optimization (as described above). The base policy could also factor in the types of cells chosen up to that point. It could assume that certain cell types would favor certain product lines and that the remaining cells would be those that favor other product lines. There are a huge number of ways that the base policy could work. Most of them have advantages and disadvantages. The more complex policies may be slightly more accurate at guessing what cells to choose, but they may require too much computation or even be too burdensome to program. On the other hand, the simpler base policies could be easily programmed and quickly computed, but may not provide sufficient accuracy. For the analysis in this thesis the base policy that was chosen was that the remaining cells would be those determined through manual optimization. The best solution derived manually was used to determine what subsequent cell types would likely be.

Earlier, the concept of policy improvement was mentioned as a popular concept in Dynamic Programming. Policy improvement is a type of policy iteration. Policy iteration was designed as a method for solving infinite horizon DP problems where there are a finite number of states and the optimal policy is a stationary one (i.e., determining the state to which to go so that staying there will minimize the cost over a large number of iterations). Policy iteration begins with a single stationary policy. The policy is subjected to a policy evaluation step which uses the single policy to solve a set of linear equations which represent the effect of following that policy (the state to which to go) from every possible initial state. The system of equations has as many equations as there are states, so by solving them the cost function for each state can be estimated. This cost function is not known, since only one policy is being evaluated, which means that only the cost-to-go values of the nodes from each state to that single state are being evaluated. Next, the policy is subjected to a policy improvement step. The policy improvement step determines which policy is the best (the state to which to go) and then uses that as the new policy. The process then begins over and this new policy is used to solve another set of linear equations and another policy improvement step is performed. This cycle continues until the best policy found in the policy improvement step is the same as that found in the previous cycle. Essentially, it keeps going until it cannot improve the policy any further.[12]

The cell picking problem that this thesis considers does not meet the strict requirements for policy improvement. It is not an infinite horizon problem, although it does have a stationary solution in the sense that there are a certain set of cell types that will most efficiently process the given demand. Technically, the problem could be set up to be an infinite horizon problem with a stationary optimal solution. However, to do this, the stages would have to be a choice of the combination of all eight cells types that would be chosen over and over until the best solution was found. However, as was discussed above, the sheer number of combinations of stages makes it impossible to solve the set of linear equations that would result. Instead, looking at the problem as successively picking cell types through eight cells, this is no longer an infinite horizon problem, and since the stages are the choice of a single cell, the optimal policy is no longer stationary.

Yet, the basic idea of policy improvement can be applied to choosing cell types. Earlier this thesis discussed of how to pick the best base policy for a rollout algorithm. The choice of a

---

[12] Bertsekas 1995, 374-376

good base policy becomes less important if a type of policy improvement is implemented along with the rollout algorithm. The rollout algorithm, as discussed above, would be implemented sequentially by picking the best cells (from one to eight) assuming that the cells not yet picked would be of the types that were found through manual optimization. The quality of the result depends largely on the quality of the base policy and how close it is to the final result. If this process were to be repeated with a new base policy that is the old result, then base policy improves and becomes stronger. Like traditional policy improvement, the process would stop when the new solution is not improved over the solution from the previous cycle, or if the solution improves by a small amount within some defined margin, or if the solutions begin to oscillate back and forth between two solutions. Again, this method does not guarantee an optimal solution; it is only meant to find a good, suboptimal solution with a limited number of calculations.

The policy improvement approach to the algorithm makes it less important which policy is the initial one, but the algorithm still provides an answer that is strongly correlated to the base policy. Since the algorithm has a bias toward cost reduction (due to the squared total cost term), the algorithm tends to pair down the capacity of the cells whenever possible.

One problem with this policy improvement method is that it does not account for the case where one cell might be chosen that allows another cell to be negated. For example, if there were four cells currently and the algorithm was iterating through cell types for the third cell and found that all of the parts could be allocated to the first three cells, the algorithm might not chose this cell type because the cost of the fourth cell (even if it is not being used) is still counted against the utility function. Therefore, the heuristic algorithm must check for this case and adjust the utility calculation to check to see if the resulting utility is an improvement over the current best.

A non-integer linear program was used for the actual optimization in Excel. Excel's integer linear program algorithms are quite unstable and it was determined that they would be unsuitable to run in a heuristic that greatly varied the initial state of the model for each iteration involving the Solver. Additionally, when an integer linear program was run for this model, it took on the order of a minute for Solver to work through all of the sub-problems necessary to come to a solution. That would make the heuristic (which runs the Solver thousands of times) far too time consuming for practical implementation. Instead of using the integer program, the heuristic solves non-integer linear programs and rounds the decision variables down. The danger in doing this is that it does not guarantee an optimal solution. However, since the LP is already being run for a sub-optimization heuristic, it makes sense to get a non-optimal answer quickly rather than waste time for the optimal answer. The rounded integer LP was never more than 1% off of the non-integer LP and usually its error was around 0.05%. This means that even if the answer is non-optimal, it is assuredly within the bounds of the heuristic. Another danger in rounding the decision variables is that it may cause the constraints to be violated. However, in this case the decision variables are the number of parts to run in each cell for each product type. If the numbers are always rounded down, this moves the function further from the constraint. One of the constraints is the maximum time that the cell is available to work. If the number of parts is rounded down, then the time used will drop. The other constraint is that the number of parts run must be non-negative. Since the non-integer LP results in positive numbers, even rounding them down will never cause the result to be below zero.

The heuristic also requires that the remainder (the number of parts left over after all parts are worked) equal zero. This is not added as a constraint in the LP since it complicates the LP and causes the LP to fail to find feasible solutions. Since the objective function of the LP is maximized by working as many parts as possible, the LP will always give an answer with the lowest remainder possible. Therefore, the LP can be run without the restriction that the remainder equal zero, but the heuristic algorithm will discard any solution from the LP that has a remainder. The benefit here is that the Excel Solver is more reliable without that constraint.

To summarize, the algorithm that this thesis has developed as an automatic sub-optimization routine works as follows, and was implemented using Excel:

1) Start with a base policy that is the type of each of the eight cells derived from manually using the model and utility function to determine a good solution.
2) Iteratively look at each possible cell type for the first cell.
    a. Each iteration consists of trying a cell type for the first cell, assuming the rest of the cells are according to the base policy, and running an LP to distribute demand across the cells optimally.
    b. Each iteration also checks to see if any other cell has no parts assigned to it and can therefore be removed. If another cell has no parts assigned to it the

algorithm adjusts the utility function to remove the cost of that cell and checks again to see if the new utility value is the best. If so, the other cell is removed.

3) Pick the type of the first cell that results in the highest utility, as long as the remainder is zero.
4) Iteratively look at each possible cell type for the next cell.
    a. Each iteration consists of trying a cell type for the next cell, with the earlier cells determined by the previous iterations and the later cells according to the base policy.
    b. Each iteration also checks to see if any other cell has no parts assigned to it and can therefore be removed. If another cell has no parts assigned to it the algorithm adjusts the utility function to remove the cost of that cell and checks again to see if the new utility value is the best. If so, the other cell is removed. Pick the type of the next cell which results in the highest utility, as long as the remainder is zero.
5) Repeat steps 4 and 5 for each remaining cell until all eight are picked.
6) Repeat steps 2 through 6 using the result from step 6 as the new base policy. Continue this cycle until there is no improvement in the result of step 6.

Chapter 4: Demonstrated Benefit

The discussion of mathematical models and optimization routines in Chapter 3 has all been highly theoretical. This chapter takes those concepts and demonstrates how they were used to derive actual benefit to Pratt & Whitney. This chapter first shows how new processes and equipment can be monetarily quantified. It then describes an example case and shows the benefit attributable to Lean processes and design optimization.

4.1 Modeling Cost

Building a mathematical model and optimization routine is useful from a theoretical design standpoint, but in order to justify the implementation of any factory design there must be a comprehensive cost analysis to show the perceived benefit in monetary terms. One of the problems with cost justification for policy and layout improvements such as these is relating the changes on the shop floor to cost. This is especially difficult in a company that does labor based cost accounting. The following is a description of the cost model that is used in this thesis to relate the recommended changes derived from the optimization model to actual cost savings.

In general, the cost model determines the yearly savings for process changes by allowing the future state to be changed and comparing it to the current costs. The savings are determined by:

$$[(\text{Cost/Hole} \times \text{Hole Volume}) + \text{Scrap/Rework} + \text{TVA}]_{\text{Current}}$$
$$- [(\text{Cost/Hole} \times \text{Hole Volume}) + \text{Scrap/Rework} + \text{TVA}]_{\text{Future}}$$

Where:

Scrap is the product lost due to poor quality.
Rework is the product that is processed multiple times due to poor quality.
TVA is Temporary Vendor Assist which is the cost for production capacity given to a vendor because it cannot be manufactured in-house due to capacity limitations.

The specifics of each component of the savings are detailed below. The costs for scrap, rework, and TVA are all straightforward and easy to compute. Computing the cost per hole, however, is more difficult. The model can predict how future costs would be affected by changes to equipment and processes only if the processes are somehow related directly to cost. Since the TMC utilizes a direct labor cost basis, this provided a useful framework for relating processes to cost. In direct labor cost accounting, costs are determined by the following equation:

$$[(MT + WT) \times LR + RT \times LR] \times (1 + FR) \times (1 + OR)] = Cost$$

Where:

MT = Manual Time
WT = Walk Time
LR = Labor Rate
RT = Realization Time
FR = Fringe Rate
OR = Overhead Rate

This equation begins with manual and walk times as well as realization time and then applies the labor, fringe, and overhead rates. Since the manual, walk, and realization times are direct consequences of the process, this method accomplishes the task of relating process to cost.

Manual time is the time the worker must spend processing each part. The walk time is time the operator spends traveling between machines. Walk time is necessary to the operation, although it provides no inherent value itself. Manual and walk time are determined through time studies in which an engineer monitors how long a worker takes to do a specific operation.

The labor rate is the average rate that workers are paid. This is determined by the payroll and for the purposes of the model is fixed.

Realization time is the amount of production time that is not accounted for by manual time. Essentially, realization time can be thought of as follows. If, on average, 1,000 parts can be drilled by one cell in a month, the total time spent drilling parts would be, for example, forty hours per week for two shifts over four weeks. This would be 40 x 2 x 4 = 320 hours. Therefore, each part is processed for 0.32 hours. However, if the manual and walk time associated with manufacturing these parts is only 0.25 hours per part, then there is a gap. The realization time is the 0.07 hours in between. Realization time accounts for time wasted through things like machine downtime, parts shortages, changeover time, and any other process inefficiency. Realization is determined by monitoring the actual throughput over a long time and comparing it with the expected manual and walk time determined from time studies. Once realization time is determined, it is typically converted to a factor (i.e., realization rate) which can be factored with labor time to get the cost of the sum of labor time and realization time.

Fringe rate is an additional rate to account for workers' fringe benefits. Fringe rate is also determined from payroll data and is constant.

Overhead rate is a factor that relates the direct labor expense to the overhead expense. This is calculated on a periodic basis by comparing actual overhead expenses to actual labor expense. The overhead rate then becomes a constant for the subsequent time period.

In order to bill a customer using this system, the number of parts produced is accounted for and this is turned into a number of "hours" which is equivalent to the labor time (i.e., manual plus walk time). The labor, realization, fringe, and overhead rates are applied to the "hours" to determine the cost. Therefore, the accounting method is an iterative process. It determines the various rates by analyzing past performance and then actually accounts for costs by applying these rates to the number of parts produced in a given time period.

The cost model built for this thesis uses the same methodology as direct labor based cost accounting. It takes manual, walk, and realization times as inputs, applies the fixed labor and fringe rates, and calculates realization and overhead rates to determine the cost of each operation. It then adds up the pertinent operations to find the cost of drilling a certain part. Dividing that cost by the number of holes in each part gives the cost per hole to drill a specific part.

To use this information for capital equipment purchases, the data was refined into three part categories. The vanes have round holes drilled into a metallic or non-coated surface, round holes drilled into a ceramic coated surface, or shaped holes, and vanes may have both round and shaped holes. All the parts with a certain type of hole are listed with the operations associated with drilling that type of hole. The yearly volume of each part is applied so that the resulting cost per hole for the specified hole-type ends up being the following function:

$$\frac{\sum_{PNs} PC \times PV}{\sum_{PNs} NH \times PV} = Cost \; / \; Hole$$

PC = Cost to drill a single product
PV = Yearly volume for that product
NH = Number of holes in that product
PN = All relevant products

The realization rate is calculated separately. Realization rate is calculated per product per work cell. This accounts for time when parts are not being worked such as during changeovers or equipment downtime. The model takes as an input the number of minutes of downtime expected for each cell per shift and the number of minutes of downtime expected per month as well as the number of shifts that the cell will be operated per month. The model also estimates the changeover time as a part of realization. The model assumes that a cell conducts five changeovers (approximately once per week) for each product that is processed in the cell. The realization time is then divided by the labor time to get the realization rate. In this way, a hypothetical future state can affect cost through its shortening of realization time.

The overhead rate is also calculated separately. In essence, a factor is applied to the current overhead rate to get the future overhead rate. The overhead adjustment factor is determined based solely on capacity changes. If capacity is increased and is still within the demand, then the number of parts produced for the cost (which should stay essentially the same) should go up, thus reducing the overhead rate. The increase in capacity for each product is applied to a weighting factor to determine its effect on the overall overhead for hole drilling. Then the aggregate effect for hole-drill is applied to another factor that weighs hole-drill's effect on the overall business unit and results in the overhead adjustment factor which goes into the cost per hole calculations above. In this way, hypothetical future states that result in increased capacity affect the cost through the lowering of the overhead rate.

Scrap, rework, and Temporary Vendor Assist (TVA) are all costs that go directly to the bottom line. Future scrap/rework cost is determined by factoring the current scrap/rework cost from cost-for-quality data. The factor is determined subjectively through an understanding of the drivers for the original scrap and rework and an understanding of how new machines and processes will reduce this in the future state. For simplicity's sake it was assumed that new machines would cut scrap and rework by one half. This seems significant, but the total scrap and rework loss is only about $100,000 per year. Half of that is insignificant compared to the savings from other improvements.

TVA is calculated simply as the difference between demand and capacity, although it is capped to the current TVA amount. The model calculates the TVA required for each part number for the future state and compares it to the current TVA to determine the savings.

The model also performs a simple Return on Investment (ROI) analysis. It takes the total yearly savings and applies it to a learning curve to determine projected future savings. From this it subtracts depreciation and the initial investment costs to determine a pseudo-cash-flow projection. It then performs five- and ten-year NPV and IRR calculations on this pseudo-cash-flow.

4.2 Optimization Example Case

Section 3.4 describes the automatic optimization routine that was derived for use in this thesis. This section uses the mathematical models described in Sections 3.3 and 4.1 and the optimization routine described in Section 3.4 to solve an example problem based on the one encountered at Pratt & Whitney's TMC. Since information from Pratt is proprietary, this fictitious example case is given in lieu of actual results from the project at the TMC. This example, however, does conform very well to the situation and results found for Pratt. The number of products, the physical characteristics of each product (numbers of round and shaped holes), and some of the processing characteristics (times to drill certain holes), as well as cost data, have been changed to protect proprietary information. The example problem allocates a monthly demand of 10,000 parts in seventeen product lines between up to eight different cells. Of the seventeen product lines there are example products with round and shaped holes, products with round holes only drilled in metallic coated parts, and products with round holes only drilled in ceramic coated parts. This mimics the full product line of the TMC.

As Figure 4-3 depicts, the example case shows that there is a significant advantage to using the optimization framework and doing manual "optimization," while there is only minimal benefit from the automatic sub-optimization.

As a baseline a set of cells was designed using a traditional factory design approach. This is likely how the TMC would have designed its hole drilling facilities given the setup of the example problem. All products are aligned in an attribute chart. The attributes look something like this:

Figure 4-1

Production Steps:	Drill Shaped Holes	Drill Round Holes in Metal	Drill Round Holes in Ceramic
Military Parts	Yes	Yes	
Civilian Parts			Yes

This indicates that only two types of cells are needed: one type that is capable of drilling shaped holes and round holes, and another type that is capable of drilling round holes through ceramic. One recently introduced military product differs greatly from the older military parts. This product (Product A) has many more shaped holes than the other military parts and therefore would likely have been given its own work cell. Since the TMC has been very interested in utilizing HSEDM for drilling, it is likely that all of the military parts would have been drilled with the HSEDM.

After knowing the order in which the equipment should be aligned in a cell, the traditional method calculates the loading on each type of equipment using the most restrictive product. Assume the most restrictive military part (non-Product A) is called Product B. Product B has the following cycle times:

Figure 4-2

Step	EDM	HSEDM
Step 1	20 min	155 min
Step 2	20 min	41 min

Therefore, the best equipment ratio for running this part would be 2 EDMs to 10 HSEDMs (8 for step 1 and 2 for step 2). A cell loading calculation determines that two of such cells are needed to support the expected demand.

A similar process is done for Product A and for the civilian parts. The results are that the cell for Product A should have 4 EDMs and 8 HSEDMs and that there should be two cells for civilian parts, each containing 4 lasers. The decision of cell size for the civilian parts, since they have essentially one type of operation, was weighed heavily toward the manageability of the work cell as judged by the first line supervisors. Also, these five cells are a convenient mix since there are currently five cells in the TMC and the management would likely appreciate having it remain that way.

When the model is manipulated for manual optimization, it derives a much different cell configuration. Manual optimization finds that having one cell with 8 EDMs and 4 lasers, one cell with 4 EDMs and 2 lasers, one cell with only 4 lasers, and one cell with 4 HSEDMs is best. The automatic sub-optimization routine finds that one cell with 4 EDMs and 8 lasers, one cell with 4 EDMs and 2 lasers, and one cell with 4 HSEDMs is best. These three configurations are depicted below, along with the cost of each, assuming that EDMs cost $750,000, HSEDMs cost $340,000, lasers cost $550,000, each depreciates over ten years, and the average salary of a worker is $81,000.

Figure 4-3

	No Optimization	Manual Optimization	Automatic Optimization
Cell 1	4 EDMs 8 HSEDMs 4 workers	8 EDMs 4 lasers 4 workers	4 EDMs 8 lasers 4 workers
Cell 2	2 EDMs 10 HSEDMs 4 workers	4 EDMs 2 lasers 4 workers	4 EDMs 2 lasers 4 workers
Cell 3	2 EDMs 10 HSEDMs 4 workers	3 lasers 4 workers	6 lasers 4 workers
Cell 4	4 lasers 3 workers	4 HSEDMs 3 workers	4 HSEDMs 2 workers
Cell 5	4 lasers 2 workers	None	None
Yearly Cost:	**$3,537,000**	**$2,746,000**	**$2,530,000**
Yearly Savings:	**$9,649,025**	**$9,987,196**	**$10,059,962**
5 Year NPV:	**($2,380,977)**	**$3,641,020**	**$5,835,952**

When the configuration found using the traditional approach is plugged into the model its utility value is only 2,357.95, compared to that determined by manual optimization of 8,762.31 and that determined by automatic sub-optimization of 11,187.59. However, because of the squared cost factor, these utilities are not linear. Yet, the important factor is not utility, but instead the overall savings, although the utility function acts as a proxy for savings for purposes of running the optimization. Assuming that the same Lean methods were used to govern the running of the cells and that the only difference is in the optimization framework used to size the cells, the numbers still show a vast benefit to using some sort of optimization.

The yearly savings figures are relatively consistent between the three configurations. This is primarily due to the fact that each one is sized for a specific demand, and the capacity is a large factor in driving the savings. All these factors also assume that the cells operate in a Lean fashion, and as will be discussed below, the Lean operations of the cells also contribute significantly to the savings. The NPVs, though, vary significantly between the different cases. The NPV is due not only to the savings, but also to the capital expenditure. Therefore the cases with the lower yearly cost will have higher NPV. Again, this illustrates the desire to improve efficiency (i.e., reduce cost per hole) with the lowest capital expenditures (i.e., low yearly cost), which is exactly what the utility function forces the optimization to do.

Determining the benefits of the optimization was merely a matter of determining what cell types would have been utilized without optimization, plugging them into the model, and comparing the savings with the savings achieved through the optimized configurations. Determining the benefits of Lean methods used in the cells, however, is much more difficult. Non-Lean work cells are difficult to model because they lack the structure required for proper modeling. This thesis uses a financial cost model to quantify the effects of the operational recommendations. The model calculates the financial results of operational changes and compares them to the initial state. Therefore, the financial model contains information about both the initial and future states. Yet the initial state is one with non-Lean work cells operating legacy equipment and the future state is one with Lean work cells operating new equipment. To use this information to determine the value of the Lean methods alone requires the separation of these two effects. The best way to separate the effects is to look at the differences between the finances of the initial and final states. Once the major differences are determined, they can be analyzed to determine the cause for each change. Once determined, the causes can be classified as being due to the new equipment or to the adoption of Lean methods. Although the two causes are interdependent, this method provides a reasonable approximation.

Section 4.1 shows that cost savings at Pratt & Whitney can be modeled as the sum of TVA savings, scrap/rework savings, and savings from improvements from hole-drilling operations as seen in the cost per hole calculations. TVA and scrap savings are independent portions of this sum because they directly affect the bottom line. Cost per hole savings are based on Pratt's labor costing model and can be described with the following formula:

$$\Delta\ [(\text{Labor Time} + \text{Realization Time}) \times \text{Labor Rate} \times (1 + \text{Fringe Rate}) \times (1 + \text{Overhead Rate})] = \Delta\ \text{Cost per hole} = \text{Savings}$$

Therefore the cost savings from Lean efforts or new equipment can be determined individually by analyzing which of these terms are affected by either Lean methods or new equipment and how each term individually affects the overall savings. To give structure to this analysis a few key assumptions were made. First, it was assumed that capacity improvements are primarily a function of machine automatic time and that any improvements in capacity were therefore due to new equipment which has faster automatic cycle times than the old equipment. This is not entirely true, but it simplifies this approximation. The second key assumption was that improvements in scrap and rework would be primarily a result of new equipment since the new equipment has less variation leading to less scrap and rework. Again, this is not strictly true, but it simplifies the approximation.

The initial state is compared to the future state assuming the configuration found through manual optimization as described above. The cost model shows that there is $9.6 Million of annual savings due to this plan prior to accounting for scrap and TVA. Since TVA savings are based on taking production back in-house and away from vendors through improvements in capacity, it is assumed that these savings are due solely to new equipment. Likewise, as stated above, scrap and rework savings are assumed to be due solely to new equipment. That means that the $9.6 Million savings is due to the factors of labor time, realization time, labor rate, fringe rate, and overhead rate. Since the labor rate and fringe rate do not vary with improved operations, they do not contribute to the savings. Therefore the savings is really a function of labor time, realization time, and overhead rate.

According to the cost model, the change in overhead rate is due to increased capacity in the plant. Increased capacity means that the same overhead is spread across more parts, which effectively lowers the overhead rate. Therefore, since overhead is a function of capacity, it is assumed that the savings due to the change in overhead rate is due solely to new equipment. Substituting the original overhead rate in the cost model for the future state, the total savings is decreased from $9.6 Million to $9.4 Million. Therefore, at most, only $9.4 Million in savings can be attributed to Lean methods.

Realization time consists primarily of four elements: machine downtime, setup time, expected excess time (i.e., the excess time that will exist regardless of how many improvements are made in the system), and wasted time. The operations model assumes that the future state will have the expected excess time and therefore the future state already accounts for this. Thus, none of the change in realization time is due to expected excess time. Setup time is also accounted for in the future state model, but it may not be the same as initial state setup time. Setup time can be affected by both new equipment (that allows for faster changeover) and Lean methods (which also help to promote faster changeovers). The savings due to reduction in setup time is therefore assumed to be evenly split between new equipment and Lean methods. In order to use the model to determine quantitatively what the value of downtime and setup time is, it was assumed that in the initial state downtime would be doubled and setup time would be doubled. This assumption is made since the actual downtime for the current state is twice that used in the model for the future state. Likewise, the assumption that the future state cuts setup time in half is well supported by the fact that current setup times take on the order of a shift and future setup time was modeled to take less than half a shift given the new equipment and standardized work for changeovers. Since setup time is only half due to new equipment, it was only multiplied by 1.5 instead of by 2. By doubling the downtime and factoring setup time by 1.5, the model is essentially adjusted for the affects of new equipment. These changes resulted in savings dropping from $9.4 Million to $9.2 Million. This is a relatively small change, so even if the above assumptions do not hold true, the effect would be small. The only part of realization time not accounted for so far is waste. It is assumed that all reduction in wasted time is due to Lean practices. This is rightly so since the whole concept of Lean manufacturing is the systematic removal of waste. To determine the overall effect of realization time reductions on the overall savings, the old realization rate was substituted for the new one in the cost model. This results in the savings dropping from $9.2 Million to $3.8 Million: a difference of $5.4 Million, of which is all attributable to Lean methods.

The only other factor in the cost equation that has not been accounted for is labor time. Labor time is the time that is actually spent doing manual work by the operators. Some of this time was undoubtedly removed through Lean practices that reduce wasted movement, but much of it is also undoubtedly the result of more automated machines. Most notably, the incorporation of soft-load capability in all of the machines would have a significant effect on lowering labor time. Therefore, conservatively, the remaining $3.8 Million savings due to labor time reductions is assumed to be due solely to new equipment.

Figure 4-4

Total Savings:	$9,987,196.17	New Equipment	Lean Methods
- TVA	$9,656,311.17	$330,885.00	$0
- Scrap/Rework	$9,609,849.53	$46,461.64	$0
- Overhead Savings	$9,430,748.41	$179,101.12	$0
Savings except Realization Time and Labor Time	**$9,430,748.41**	**$556,447.76**	**$0**
- Downtime & Setup Time	$9,215,472.06	$215,276.35	$0
- Expected Overhead Time	$9,215,472.06	$0	$0

- Wasted Time	$3,772,874.50	$0	$5,442,597.56
-Labor Time	$0	$3,772,874.50	$0
Total:	**$0**	**$4,544,598.61**	**$5,442,597.56**

In summary, out of a total savings of $10 Million (including TVA and scrap/rework savings), $5.4 Million can conservatively be attributed to Lean methods being incorporated in the operation of the new work cells. This means that for the work cells designed by the manual optimization case, Lean methods accounted for at least 54.5% of the savings.

Another concern in the design of a factory is how robust the factory is to changes in demand. The model does not directly take the probability of the demand into account. Instead, the idea is to run the model with various demands to determine the best cell designs to cover the likely range of demand and product mix. This was not done for this example, although it was done for the project at the TMC.

This optimization example demonstrates a number of important findings. First it demonstrates the importance of using a systematic, utility-based framework for determining the correct design for work cells in a factory. In this case, the manual optimization approach shows a higher annual savings with only 80% of the yearly expenditure. This results in the difference between a 5 year NPV loss of over $2 Million and a gain of $3.5 Million. Second, this example demonstrates that automatic sub-optimization heuristics can find further savings but that the problem is so highly nonlinear that the savings are only marginally superior to those found using manual optimization. Finally, the example shows the important role that Lean methods play in the implementation of the new work cells. In this case, the Lean practices are conservatively shown to account for at least 54% of the savings that the new work cells demonstrate.

4.3 Findings Delivered to Pratt & Whitney

So far, this thesis has described in general the methodology that was used at Pratt & Whitney to deliver a design for new hole drilling work cells. Yet the deliverable given to the TMC was much more specific than what has been described here. This section sheds more light on what type of information was determined from the analysis done at the TMC and delivered to Pratt & Whitney.

Pratt was given a simplified explanation of the theory behind the optimization framework as well as a complete description of the mathematical and cost models, including a comprehensive explanation of the assumptions used. It also received an explanation of how its value stream would change given the proposed changes.

Most importantly, though, Pratt was given specific recommendations for how many of what types of machines should be bought and where exactly each machine should be placed in the factory. Although standardized work for the new cells was left up to the ACE group, an example of standardized work was provided, which detailed exactly how the cells should operate for maximum personnel and machine utilization. The TMC was also given a detailed list of which parts should be run through which cells, as well as backup configurations should one cell be down for maintenance. It was also given analysis similar to that found here regarding the expected financial gains to be had should the recommendations be implemented, as well as sensitivity analysis for variation in machine downtime and cycle times. Finally, the TMC was given a description, similar to Chapter 5, which described the pros and cons of choosing HSEDM or lasers for drilling round holes. These pros and cons included a sensitivity analysis for the variation of each machine cost.

Additionally, the report given to Pratt included a section that detailed how equipment which Pratt was committed to purchase could be utilized in the case that Pratt was unable to purchase more new equipment beyond that to which it had already committed. This "low cost" plan included a description of which cells should operate which parts, a description of how the cells should be arranged, an example of standardized work, and the expected benefit of implementation. Finally, the report included a description of how best to implement the recommendations, including a sample time line that detailed how different aspects of the plan, such as enforcing a strict cycle time that couples together all of the machines in a cell, could be phased in to make the adjustment as easy as possible for the workers and supervisors.

In summary, the report to Pratt was similar to this thesis in that it described in some detail the optimization framework and in greater detail the required Lean methods to be used. Unlike this thesis, however, it also included much more specific recommendations and cost justifications, including sensitivity analysis for various factors.

Chapter 5: Deciding Between Two Manufacturing Methods

One of the questions in which Pratt & Whitney was particularly interested, was whether or not to pursue drilling round holes with HSEDM versus the traditional laser method. Section 1.2 stated that this thesis would use the results of the optimization to answer this question. Figure 4-3 showed that the factory designs using manual and automatic optimization both relied primarily on lasers for round holes, although each of the results had one small cell of HSEDMs. This might seem counterintuitive, since it was expected that there would be a clear winner. Naturally there would need to be lasers to drill ceramic coated parts, but if the HSEDM was superior one might expect the results to look somewhat like the configuration of no optimization in Figure 4-3, with HSEDM drilling all non-ceramic coated parts and lasers drilling ceramic coated parts. If the laser was superior, it would be expected that all round holes would be drilled with lasers. Yet the results of the optimization did not match either of these cases. This chapter compares each method in more detail, including analyzing some of the key assumptions behind how these two techniques were modeled, and then will explain why the optimization turned out the way it did.

5.1 Detailed Comparison of HSEDM and Laser Drilling

Section 2.1 provides an overview of how both lasers and HSEDMs drill holes and includes an analysis of the problems to which each method is prone. This section will distill that down into a clearer picture of the differences between laser and HSEDM drilling.

In Section 2.1, it was noted that laser drilling is faster than HSEDM drilling but more prone to variation. The complex auxiliary processes inherent to laser drilling were also listed. But is laser drilling really faster if it requires so many auxiliary processes? Typically the speeds of machines are compared on the basis of their overall cycle time which is based largely on the automatic time of each machine. However this line of thinking no longer applies when the task changes from choosing between one of two different machines to choosing a number of machines from two different types. From a process standpoint, the speed of a machine is not necessarily of primary concern. For example, if machine #1 is half the speed of machine #2, but one operator can operate twice as many machine #1's than machine #2's, then the potential capacity of the machines is equal. Granted, this does not take into consideration footprint or machine cost. This is based solely on capacity based on labor, which is a good perspective to have considering the financial model used in the TMC (i.e., labor based cost). It can be shown that the only factor that can really be used to compare the capacity of two machines is the amount of manual time each machine requires per part. The machine with the smallest manual time per part has the best theoretical capacity.

Mathematically the number of machines that a worker can operate is defined as:

$$Machines = \frac{CycleTime}{ManualTime} = \frac{AutoTime}{ManualTime}$$

The Cycle Time is essentially the longest automatic time (plus load time). One process is faster than another by a factor relating each machine's automatic time. The factor that the laser process is faster than the HSEDM process is essentially:

$$Factor = \frac{AutoTime_{EDM}}{AutoTime_{Laser}}$$

Therefore, to determine which process actually allows a single worker to process the greatest number of parts, the above two equations must be compared. For a part to prefer HSEDM over laser:

$$\frac{\frac{AutoTime_{EDM}}{ManualTime_{EDM}}}{\frac{AutoTime_{Laser}}{ManualTime_{Laser}}} > \frac{AutoTime_{EDM}}{AutoTime_{Laser}} \Rightarrow \frac{ManualTime_{Laser}}{ManutalTime_{EDM}} > 1 \Rightarrow ManualTime_{Laser} > ManualTime_{EDM}$$

Since manual time is what actually determines the process speed of a machine, the extra manual time in the laser process from the auxiliary processes is very significant when comparing it to HSEDM in terms of speed. It seems that in terms of manual time, the HSEDM is actually faster than the laser. Also, from a management standpoint, the option with the fewest auxiliary processes is easier to manage and has less risk of work stoppages due to equipment breakdown.

It is easy to quantify and compare how much faster one process is than another. However, it is not so easy to understand and deal with the variation of a process. Section 2.1 described in detail what causes variation in the laser process. Yet the important question is how this variation plays a role in production. First of all, high variation can cause scrap and rework, which obviously lowers throughput and hurts production. High variation also requires a lot of monitoring to keep the process in control in order to minimize scrap and rework. For hole size variation, that means that excessive airflow measurements are being made. Airflow measurements are very time consuming. This means that they require additional workers, which add to cost. It also means that airflow can become the bottleneck operation and directly slow down production. Finally, high variation means that measurements should be done more often in order to catch problems before they cause scrap. For instance, if a part requires two successive operations in a laser, it would be advantageous to airflow the parts between the trips to the laser so that if the part was incorrectly drilled, it could be scrapped or reworked early in the process so as not to waste time running a bad part through later operations. Ideally, though, if the process were stable enough, the parts could be airflowed only once, at the very end of the drilling process. This is dangerous, since not only are a number of operations being inspected with one airflow measurement, but there is also a lag time between when the part is drilled and when it is inspected, and during that time other parts would have been run on the same machine. This risks scrapping not only one part, but many parts; therefore, to support doing airflow measurements in this way, the process must be very stable.

Yet the real question is not which process is more stable, because the HSEDM process is the obvious answer. Instead, the question is whether the laser process can become stable enough to support implementation of the desired process flow. The work cells that were modeled in Section 4.2 all assume that process variability is low enough to allow the placement of the airflow inspection at the end of the drilling operations and that in most cases not every part is inspected. This means that there is a sampling plan, and that there might be six or seven parts in production between the machine that drilled the hole and the airflow inspection that would reveal a problem. With a HSEDM, the process easily meets these requirements, but the laser process is much less stable.

To support the cycle times defined in the model, the inspection rate must be lowered to only 1 out of 24 parts inspected. Add to this a six-part lag time between the first hole drilling operation and the airflow measurement, and there must be an average run length (ARL) greater than 30. The ARL is the number of parts that can be run before an out-of-specification part is produced. Essentially, ARL = $1/\rho$, where ρ is the probability of having an out-of-specification part. ρ can be determined from statistics about the process. Assuming the process variation is distributed normally the standard deviation of the process is determined and then a normal distribution table is used to find the probability that a part will land outside half the specification limit given that the mean is zero.

For airflow, the specification is typically 1.1 – 1.3. This means that the specification limit is 0.2, and half of that is 0.1. Excel [=1-2 x NORMDIST(0.2/2,0,S.D.,1)] can be used to find ρ. As an extra safety measure, ρ was multiplied by a factor of 2, which means that by sampling at the rate equal to the ARL the process should only be halfway to out-of-specification by the time it is sampled. Although not mentioned specifically ARL is directly related to Cp. The table below is based on a recent run of parts on an old laser at the TMC and some preliminary data from one of the new lasers that the TMC was considering which has recently been delivered. Note that the sampling data used to determine the standard deviations for both the old and new laser systems were not rigorous. There was not sufficient time during the author's internship after the new laser was delivered and run in to conduct a proper sampling experiment. These numbers, however, provide ballpark figures for comparing the two systems.

Figure 5-1

Parameter	**Old Laser**	**New Laser**
S.D.	0.052	0.020
ρ	5.447%	0.0000403%
ρ x 2	10.894%	0.0000807%
ARL	9.2 = 9	1,238,521
Sufficient?	No	Yes

Although the old process was not sufficient to support the proposed plan the new process is more than sufficient to support the new plan. Again, this is very preliminary data, but it does indicate some improvement in process stability and after even more strenuous testing it is probable that the new lasers will have sufficient stability to negate the stability advantages of the HSEDM.

These statistics assume that the process stays in control. If the process goes out of control, the lag between the drilling step and the inspection step may cause a number of parts to be scrapped out. Before this system is put into place, an analysis should be done on the major contributors to the process going out of control in order to determine how large the errors could be. If they cause the process to go out of control yet allow the parameters to stay within the tolerances, a delay can still be implemented. If, however, the common problems cause grand shifts in the airflow and there is no other way quickly to determine whether these problems occurred, a lag between drilling and inspection is infeasible. For the TMC, most problems that cause lasers to go out of control either do not shift the airflow measurements a significant amount each repetition or are easily identifiable through simple visual inspection and would therefore likely be caught immediately.

In order to take advantage of this reduced variation to decrease the sampling rate, a very strict process control plan must be put into place. A strict sampling plan would be followed and all results would be statistically analyzed with the use of common control charts. The results of the analysis could immediately be fed back to the appropriate machines through an automatic network. With all of the sampling done at one station, preferably by one operator, and following a sampling plan, much of the variation in the measurement process would also be reduced.

Overall, although the HSEDM has far less inherent process variability, the laser process can be brought into a state of sufficiently low variation to support the Lean work cell plan and therefore lasers can compete with HSEDM with regard to variation.

Section 2.1 did not discuss one important factor in comparing these two drilling methods: the economic ramifications. The most significant economic factor is the purchase price of each machine. Currently a laser costs approximately $550,000 while a HSEDM costs approximately $340,000. However, the laser is three to five times faster than the HSEDM, so three to five times as many HSEDMs would need to be purchased in order to acquire the same capacity as with lasers. It is immediately obvious that the economics of the situation do not favor HSEDM. Yet, as discussed earlier, it is fallacious to compare capacity on a machine-to-machine basis. Capacity must be viewed from a system standpoint. Comparing a fully modeled system, as was done in Section 4.2, is the only way truly to understand the economic ramifications of using either method, but it would seem probable that HSEDM would be less favored in most circumstances.

Although lasers have faster cycle times, they are not necessarily faster overall since it is the manual time of a machine that actually determines its capacity in a labor-limited environment. Similarly, although the HSEDM process has lower variation than the laser process, it is likely that the new laser processes have decreased variation to a point that they are on par with HSEDM for practical purposes.

5.2 Explaining the Results of the Model for HSEDM and Laser Choices

Knowing all of the advantages of HSEDM over laser, it would be expected that the optimization would result in the choice of HSEDMs to drill all parts except those that are ceramic coated. Yet the model showed that the best solution was involved using all lasers except for one small cell of HSEDMs. The explanation for this mystery lies not in what cells were chosen, but rather in the details of the optimization. The optimization not only chooses the best cell types; the LP part of the optimization also distributes the allocation of the demand to each of the chosen cells. Understanding why one small cell of HSEDMs was chosen comes from looking at what product demand was assigned to that cell and what the utilization of the cell was.

The best solutions all had a small cell of HSEDMs which processed almost exclusively one product. This product will be referred to as Product A. In fact the utilization of the cell was only about 60%. So what is so different about this product and why should it prefer HSEDMs over lasers (which are preferred by all the other products)? The only noticeable difference between this product and others is that it has very few holes. In fact, while most parts have a few hundred holes, this product has only about a dozen.

This means that lasers or HSEDMs must be sensitive to the number of holes in a part. Naturally, the more holes a part has, the longer it is going to take to drill. However, this only affects that automatic time of the machine, and, as Section 5.1 demonstrates, it is actually the manual time that determines capacity. Thus one of the two must have manual time that varies with respect to the number of holes. As discussed in Section 5.1, lasers require more manual time per part due to the added steps of filling the part with backing and burning out the backing, but this is completely independent of the number of holes in each part. HSEDM, however, requires very little manual time in dealing with each part. Yet HSEDM does require a significant amount of manual time for electrode replacement. Electrodes must be replaced more often for parts with more holes since they are used more. Therefore, it is the HSEDM process that is sensitive to the number of holes in a part.

For a product with a dozen holes, electrodes only need to be replaced every few parts, but for a product with hundreds of holes, electrodes must be replaced a few times per part. Therefore, HSEDM has more capacity than laser drilling only when the manual time required for electrode replacement per part is less than that required to fill backing into a laser-drilled part and burn it out. Figure 5-2 shows manual time as a function of the number of holes in a part. The laser process is represented by a straight, level line while the HSEDM process is represented with a line with definite slope.

The immediate benefit of this analysis is that it shows that investing in HSEDM may not be the best choice since it is optimal for only one part. However, a more long term benefit of this analysis is that it shows how the process can be improved to favor more products. The HSEDM that is used in the model and used to produce the graph below is a particular design that is to be produced for Pratt. However, a number of levers affect the manual time per hole, such as the electrode wear rate, which is determined by electrical characteristics of the EDM power supply, the length of the electrodes, and the time it takes to load electrodes into the machine. All of these levers can be changed to some degree through modifications to the design of the machine. Pratt's understanding of these characteristics and the importance of looking at fixed and variable manual time instead of just overall manual time will allow it to work with its vendor potentially to produce a future HSEDM which fits better with the number of holes per part common to its product line.

Figure 5-2

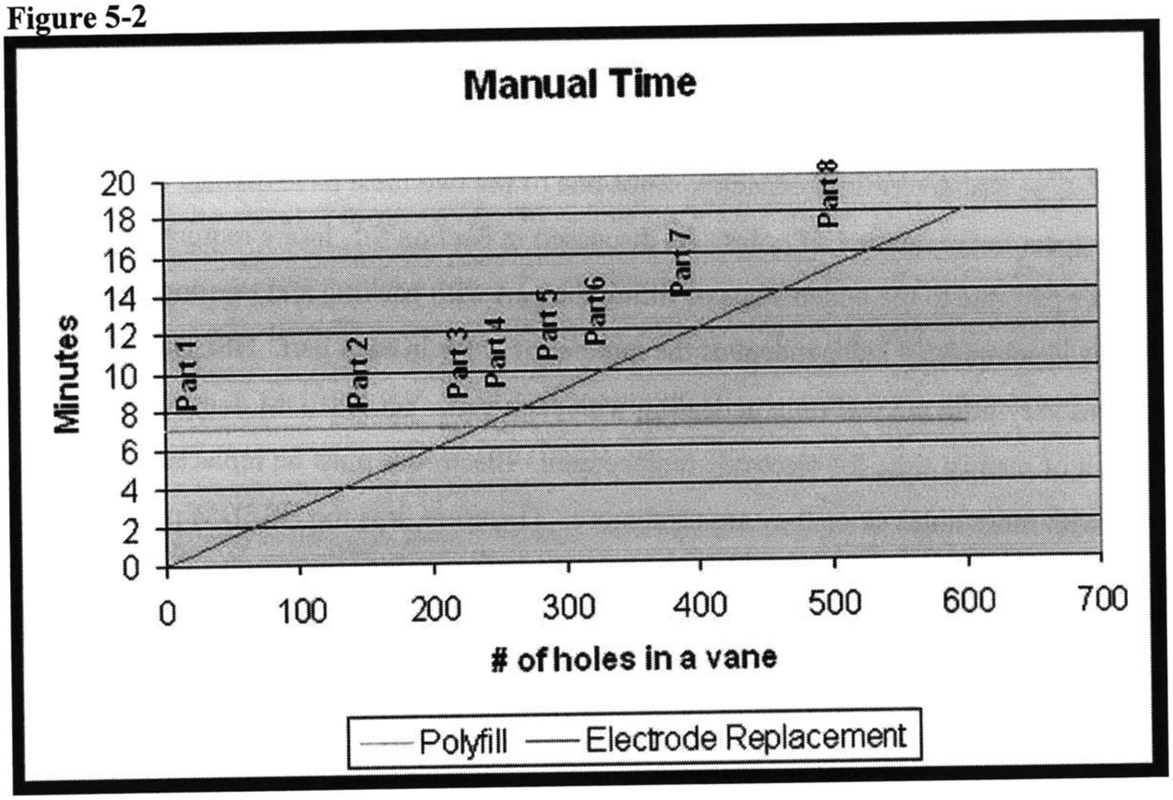

Chapter 6: Roadblocks to Changing a Culture

The previous chapters discussed the potential for increased production and reduced cost that exists through the combination of Lean operating principles and new equipment organized in an optimized fashion. Mathematically, this all makes perfect sense. Unfortunately, the world of manufacturing is not as mathematical as many would like. This is especially true when it comes to dealing with people. Be it the shop floor workers, the first line supervisors, the support staff, or even upper management, people naturally resist change. But in order to make any of these possibilities a reality, Pratt must begin to do things in a way that is very different from how things are currently done. This chapter looks at the roadblocks to change at Pratt & Whitney's Turbine Module Center. The following analysis may seem like a list of problems, and overly pessimistic, but every organization resists change, and before a plan can be enacted to help accomplish change—as is provided in Chapter 7—there must be a serious attempt made to understand why change is difficult.

This thesis uses the Lenses model to analyze the TMC in order to determine what roadblocks stand in the way of change. The Lenses model analyzes an organization from three different perspectives, or "lenses." The first perspective is the strategic design lens. This lens looks at an organization from the perspective of how it is structured to fulfill a purpose and accomplish tasks. This lens might look at what organizations within the company are designed to fulfill what roles and who reports to whom. The second perspective is the political lens. This lens looks at an organization from the point of view of the people who are in it. This lens looks at things such as the struggles for power between different people and the interests of various stakeholders. Instead of viewing a company as a cohesive unit working toward a common goal, as is the case with the strategic design lens, it looks at a company as a group of individuals interacting in a variety of ways for a variety of reasons. The final perspective is the cultural lens. This lens looks at an organization from the perspective of the common attitudes and beliefs that are shared by the members of the organization and their collective history. This lens looks at things such as attitudes and symbols. Unlike the political lens, which looks at people as being individuals, the cultural lens looks at how the organization has instilled some commonality in its members. These lenses will help to reveal the roadblocks to change that are present in the TMC.[13]

6.1 Roadblocks Seen Through the Strategic Design Lens

The strategic design lens is concerned with how the TMC is organized. The TMC performs four basic tasks for each of its principle products: turbine vanes and turbine blades. It grinds blades and vanes, it drills holes in blades and vanes, it coats blades and vanes, and it performs final inspection on and ships blades and vanes. These are the goals for the TMC. The direct manufacturing part of the TMC is split into three sub-organizations, the blades business unit, the vanes business unit and the coatings business unit. Grinding, drilling, and final inspection for blades and vanes are done in their respective business units, but coating for both blades and vanes are done in the coatings business units. There are also auxiliary organizations at the TMC such as the engineering, quality, and ACE organizations, which support manufacturing.

The TMC is managed by a vice president. Below him are the heads of manufacturing, finance, engineering, etc. Below the manufacturing manager are the heads of the vanes, blades, and coatings business units as well as TMC's ACE manager and others. As with many contemporary organizations, most people at Pratt belong to at least one chain of command in a classical matrix organizational structure. There are product-focused teams, which are the business units, and functional teams, such as engineering and maintenance. For instance, a manufacturing engineer who is a hole-drill expert might be assigned to assist vanes hole drill engineering efforts. His product team would therefore be the turbine vanes business unit and he might work for the head of the vanes business unit. He would also be a member of the manufacturing engineering team and report to the head manufacturing engineer in his functional team. This is a typical matrix organizational structure.

Matrix organizations are appropriate for organizations like Pratt that are driven significantly by technology. The functional groups allow the members to stay current in their fields while the product groups help them to stay focused on a particular task. For instance, in the TMC, it is very important for the vanes hole-drill experts to stay in contact with the blades hole-drill experts since the processes are virtually identical. Yet it is important for both the

---

[13] Carroll 2002

blades and vanes hole-drill experts to be accountable to the managers of each business unit so that they can keep their work focused on the particular needs of that business unit.

The maintenance group is another matrix organization. It is responsible to the heads of the blades, vanes, and coatings business units for keeping production machinery running, and it is also responsible to the head of the maintenance department.

In theory this matrix structure is appropriate for Pratt & Whitney, yet in practice it is not fulfilling its purpose. There are incessant complaints from the business unit head and those directly responsible for manufacturing that they do not receive the help that they need from the functional groups. The most significant complaint is about the maintenance organization. When a machine breaks down, a trouble call is placed. The maintenance organization is supposed to come and fix the machine. Yet the metrics of the maintenance organization are such that it is only the time to respond that matters, and not how long the machine is down. If a machine has a problem and the maintenance technician arrives immediately, but cannot figure out the problem or fix it immediately, his responsibility is essentially complete since he responded to the call. But for the manufacturing organization, every minute the machine is down is a minute of lost production. Therefore, although there is some interaction between the maintenance people and the business unit heads, there is no real alignment in their incentives. Likewise, a different maintenance person might show up each time and therefore each maintenance person has no direct responsibility to a particular machine or production method. Essentially what this means is that although the organization is a matrix, the functional dimension of the matrix is much stronger than the product dimension, and this disparity causes problems.

A similar problem can be seen with the scheduling of manufacturing engineers. Often the projects on which they are assigned to work are not in line with the true goals of the organization as defined by the business unit managers. Instead, the projects are assigned by the functional head of manufacturing engineering. The head of manufacturing engineering attempts to work with the business unit leaders in order to make an appropriate schedule, but this breaks down when there are constant changes in production, machine status, and priorities. If the business unit managers had more control the direction of the manufacturing engineers they could better ensure that the engineers were working on the appropriate priority for the current situation, despite the rapidly changing nature of the business.

The ACE program is one matrix organization that is aligned more toward the product dimension than toward the functional dimension. The corporate ACE group is relatively small and does not have the resources to work with every sub-organization. This leaves the local ACE organizations to be much more directly managed by the manufacturing organization. As stated above, in the TMC, the ACE manager reports directly to the head of manufacturing for the TMC. Fortunately, the TMC's ACE group is sufficiently knowledgeable in continuous improvement that it is not hampered by its loose association with the company wide ACE program. However, they suffer from a separate organizational problem. Currently the ACE team helps to further continuous improvement by running Kaizen events. It is used by the manufacturing organization as extra labor to complete improvement projects. Although this has been helpful, the team is too small to drive significant change and way too small to follow up and enforce all of the initiatives. Having them run Kaizen events also has the effect of taking any responsibility for continuous improvement away from the manufacturing managers. Therefore, the manufacturing managers do not have buy-in for the initiatives, do not understand how to use the initiatives, and why the initiatives are important, since they are not involved. Additionally the workers see the managers as taking no interest in continuous improvement and therefore take away that it can't be very important. The problem with the structure of the ACE organization is that it is being improperly utilized. When a Japanese sensei comes to a company he does not lead Kaizen events. Instead, he mentors the manufacturing managers to lead, take ownership of, and follow up on initiatives. It is an advisory role that improves the efficiency of manufacturing by training those responsible for manufacturing—the managers, instead of taking the responsibility away from them.

The strategic design issues described above reveal some significant roadblocks to change at the TMC. Many matrix organizations are driven far more forcefully by the functional dimension than the product dimension. This makes it very difficult to cause change from the product, or manufacturing, side. Any change initiative that is attempted could be drowned through a lack of support from the auxiliary organizations. Additionally, the principle change agent for the TMC—the ACE group—is hamstrung since it is utilized as a group of change agents instead of as a group of mentors. Although these are significant roadblocks to change, they are ones that can be remedied through some organizational changes, such as those described in Chapter 7.

6.2 Roadblocks Seen Through the Political Lens

The political lens is concerned with the power bases and interests of various stakeholders in an organization. The interests and powers that these stakeholders possess have significant effects on the survivability of any change movement. A number of key stakeholders at the TMC would be involved in any changes as described in this thesis, including the management of TMC, and the corporate management—most significantly the finance element, the shop floor workers, the union, the ACE group, the manufacturing engineers, the maintenance group, outside vendors, the facilities group, and many more. Fortunately for the TMC, though, most of these groups are accepting of change. This is not to say that they are advocates of change, only that most of them have no incentive to maintain the status quo other than human nature. Two groups, however, have significant power and have the incentive to maintain the status quo.

One group that has some incentive to maintain the status quo is the union. First of all, the union wants to protect jobs. It will assume that any change is a threat to jobs. Second, there are a number of personalities at the TMC that, frankly, enjoy "winning" against management. They view management as the bad guy and see themselves as the good guy. They will bring up trivial problems because they view any win against management as victory for their side. Any recommendation for change that arises from management will be resisted wholeheartedly. Additionally, the union has a special disdain for the ACE program. Any change initiative that is linked to ACE or Lean is likely to be viewed in a very bad light.

The union has significant power to support their interests as well. Currently the union contract provides that union members are not required to do any action that is solely for the purpose of continuous improvement. The union also has significant power to bring up grievances. Yet, it is not perfectly aligned with the workforce. Many workers on the shop floor see the union in as bad of a light as they see management. Although there are often jokes about how poorly managed the TMC is, there are an equal number of jokes told about how poorly led the union is. However, in the end, the workforce is likely to back the union over management if it comes down to it.

Even though the union may oppose change, it should be understood that it is not doing such merely to cause problems. There is a difficult history at Pratt between the union and management and the union representatives truly feel that they are doing what is required in the interests of both the workers and management. The issue at Pratt, as with union issues at many companies, is more with the communication between the union and management than with each of the parties. In order to enact changes that the union might oppose, this communications barrier must be breached.

Beyond the union, Pratt's senior management is the next group of stakeholders who have an incentive to maintain the status quo. The Turbine Module Center provides over one seventh of the revenue for Pratt & Whitney as a whole. Any deviation from the status quo puts that revenue at risk. Although management would likely be willing to support new equipment and minor operating changes, the change would need to be implemented transparently, doing one-to-one swapping out of equipment, and process improvements would have to be gradual so that no significant part of the value chain would ever be at risk. This is a perfectly legitimate position for management to hold. However, no significant or lasting change can occur in this manner.

The political perspective reveals even more roadblocks to change at the TMC. First, there is the union, which will likely oppose any major improvement program directed by management because there is distrust between them due to a long history of poor communication. Secondly, senior management themselves will also likely hamper a wide-scale improvement program since it may put a large revenue stream at risk. Like the strategic design roadblocks, these roadblocks are quite significant. Unfortunately, though, these roadblocks cannot be overcome through simple reorganization. Overcoming these roadblocks requires the collusion of the union and the ability to separate any improvement programs from what higher management perceives as the current state value stream. This is certainly possible if a way can be found to begin this change in a way that minimizes the union's concerns and also minimizes the amount of revenue that management perceives as being put at risk.

6.3 Roadblocks Seen Through the Cultural Lens

The cultural lens is concerned with the culture of the organization: employees' perceptions and the meanings that they assign to situations. The TMC is an old organization and many of its employees have been working there for over twenty years. The TMC has been through a number of changes. It went through various improvement programs, such as the one Ed Northern directed, and it has been moved a number of times, most recently from North Haven to East Hartford. This history has grown a strong culture at the TMC, and it is one that will provide a number of roadblocks to change.

The first part of the culture that is relevant is a "been there, done that" mentality regarding improvement programs. Most of the workers have participated in Kaizen and similar events only to be disappointed with the lack of results or the lack of management follow-through to keep the changes in place. Workers and management no longer have any confidence that improvement programs will actually get things accomplished. Workers are reluctant to take part in any improvement program because they fear it to be a waste of time, and shop floor management feels that it is more important for workers to be pushing through parts rather than spending time on improvement programs that are perceived likely to fail. Likewise, the history of Lean at the TMC has shown that any improvement programs that have been directed from upper management have eventually worn off. This is essentially what happened to Ed Northern's initiatives. The only impetus for change that he provided was himself and once he left improvements slowed. This has left the workforce believing that anything that comes from upper management will fade away and that if they do not like some new change, all they have to do is wait and it will go away. As described in Section 2.2, workers have seen attempts at continuous improvement changes before and have come away from those experiences with disdain. Workers are likely to avoid wanting to make any changes simply because they feel like they will be wasting their time. Their experience with management driven change initiatives has left them with the feeling that management doesn't have a clue as to how to change for the better. They feel that they know the best way to drill turbine airfoils and are resistant to change since they think they know better than management. This mentality is an obvious roadblock since it ensures that any future changes will be met with skepticism and reluctance. It also means that it will be difficult to get people to be motivated about a new change and to buy into it.

This "been there, done that" mentality has bred a class of workers who have little interest in their jobs other that punching the clock. They see no future and they are quietly awaiting retirement (which is not more than five years away for most of the workforce). The types of changes proposed for the TMC would require significantly more work from these workers. Currently, most of the workers operate a single machine and spend most of their time watching it cycle automatically. With the changes they would have to operate many machines and would be actively working much more of the time. Currently they may be actively conducting manual operations around 20% of the time. With the changes that would increase to close to 85% of the time. Some workers will simply not want to work more, but most of the workers will be reluctant to do this because they have an ingrained distrust for the equipment such that they feel they need to constantly be monitoring it. It will also be a significant change to have the workers working together as cell teams, with each operation dependent on the others. Each operator currently operates his machine, and there is a buffer insulating him from the other steps. Operators can take a break when they want and go at whatever speed they want. With the proposed changes, though, operators will have to be coordinated in everything they do and will lose much of their autonomy. This will be a significant change for them. Many will feel micromanaged. Most workers find some satisfaction in their jobs, either through relationships to coworkers or with their skill as craftsmen, but no one sees a real opportunity to do something they think is important and therefore everyone is content with mediocrity. No one is willing to work especially hard, and no one is willing to change, because operators do not perceive any past changes as having resulted in much good. Again, this is an obvious roadblock to future change.

A second aspect of the culture which impedes change is the mentality of the importance of meeting demand and shipping parts. It permeates form management to the shop floor workers. Although the shop floor workers are not motivated to work very hard, as described above, they know that they are judged by how many parts they put through a machine in a day. One might think, therefore, that the proposed changes would actually motivate them to make small changes to increase their efficiency, yet this is not the case. As described above, there is a culture of mediocrity and no one wants to be seen as the superstar. A superstar would not be viewed kindly by coworkers. Also, the equipment eccentricities drive workers to be inefficient. Although it is highly inefficient for an operator to sit and stare at a machine that is running automatically for a long period of time, it is understandable that the operator should desire to do so when his experience tells him that if something goes wrong, he needs to be there. Most of the operators have been "burned" in the past when they tried to be efficient and do multiple things at once and some equipment went wrong, and they ended up further behind than if they had just been paying attention to one operation at a time. If there wasn't such emphasis on throughput, this might not be the case.

In fact, this concern for the workings of the machines is its own part of the culture. Recently a new laser was brought into the TMC. One of its features was a highly customizable and programmable user interface. The workers immediately requested that it be set up to mimic the user interface on twenty year old machines. The author and the technician from the laser manufacturer spent nearly half an hour trying to describe ways that the laser could be set up so the interface was superior to the old machines, such as having one button set to do operations that used to take five minutes to program and having displays that showed more relevant information than did the old machines. Yet the operators could not be convinced. At first, it appeared that they were simply being stubborn or might have been afraid that if the machines were too easy to use, they lose job security. However, it seems that the real reason for their hesitation was that they knew that the old way worked and they did not want to trust the machine to operate correctly with operations being automated. As in this example, this culture of using unreliable equipment is another roadblock to change since it gives workers another reason to fear change.

The "been there, done that" mentality, the emphasis on throughput, and the distrust of equipment all work together to create a culture where workers would rather do things the old way than do things a new way that could be better. The culture has created a "comfort zone" for them from which they will be unwilling to deviate from without a lot of help, and this will be a significant roadblock.

The above paragraphs deal with the culture of the workers. Most of the observations also apply to the first- and second-level supervisors as well, since they spend enough time on the shop floor to have seen what the workers have seen. Additionally managers and supervisors will be reluctant to change since they will go from being "fire fighters," tackling exciting new problems every day, to being responsible for longer term improvements. This is a complete shift in perspective and will affect them on a daily basis. Most of the supervisors and managers currently at the TMC are there because they thrive in that sort of environment. Although many of them claim to desire a simpler life and a system that is easier to manage, in reality they are energized by the myriad of problems that they face every day and enjoy triumphing over them. It will not be easy to get them to turn into the type of manager that looks long term and, instead of throwing quick fixes at problems, investigates the details of each problem to fix them permanently.

Even the Lean advocates and experts at the TMC have a mentality that will be a roadblock to change. That mentality is that the factory must be stable before it can flow. Traditional Lean thinking uses what is commonly called the "House of Lean" to implement Lean manufacturing in a business. The House of Lean has stability as the lowest foundation and standardization as the next level of the foundation. The pillars of Lean that support the roof are just-in-time and Jidoka (i.e., load leveling). The roof is customer focus. Typically when implementing Lean it is advisable to build from the ground up. Only after stability and standardization have been secured can just-in-time and Jidoka be implemented. Essentially, this means that before a plant can try to implement Lean flow, it must first have stability and standardization. This is the philosophy at the TMC. The TMC is stuck with inadequate stability and has therefore been unable to go forward with flow in many areas, including the hole-drill cells within the vanes product cells. As this thesis has described, the old hole drilling equipment had such poor operating variability and such downtime issues that there could be no production stability.

Unfortunately, though, flow is the visible manifestation of Lean. A factory can be superbly stable but without higher throughput, low inventory levels, and responsive production no one will notice. The key to driving change is to show that change is possible and that its affects are good. When there is no visible sign that things have changed, it is nearly impossible to keep people motivated to continue their efforts to change. The TMC has exactly this problem. Although strides are being made toward stabilizing and standardizing production, these efforts are often short-lived because they alone have no visible impact. Every front line manager has been heard uttering the words “If we could only get parts to flow!” at least once. Even these cries of despair were against the lack of stability that was preventing the parts to flow consistently, but they illustrate that the focus is on the flow. Managers were not crying out, “If only the system variation had a lower standard deviation!”

Yet, flow cannot happen without stability and standardization, so it is actually a balancing act to match work on stability with just enough flow to keep the work force interested. Too much concentration on flow will make increasing flow impossible since flow cannot happen without stability. On the other hand, too much concentration on stability will result in no flow, which will starve the perception of accomplishment that is necessary for motivation.

This mentality of “stability first” is a roadblock to change, since the drive for stability starves people from getting the positive feedback that changes are having an effect.

Although it is very difficult to accomplish change in the shadow of a culture that opposes it, it can be done if that culture and its effects are taken into account. The plan in Chapter 7 takes the culture into account and is one way not only to accomplish change despite the culture at the TMC, but actually to change the culture for the better.

Chapter 7: Overcoming the Roadblocks to Change

Chapter 6 used the lenses perspective to analyze the roadblocks that are faced in making changes in the TMC. This chapter describes two approaches to overcoming these roadblocks. The first approach is one that could be adopted within the TMC itself, and the second is an approach to force the TMC to change through outside forces.

7.1 Accomplishing Change from Within

The TMC must take extraordinary measures in order to make the types of changes described in this thesis and recommended to the TMC by the author. The following paragraphs describe one way of overcoming the organizational, political, and cultural roadblocks in order to accomplish change within the TMC.

Most change advocates and consultants recommend that a crisis be used to drive change. Yet, there is no crisis within the TMC. Demand is high and business is good. Therefore, something other than a crisis must be used if change is going to be accomplished from within the TMC.

This method relies on the ability to build an autonomous plant-within-a-plant (PWP) at the TMC. This would be an area set aside from the rest of the operations at the TMC. Ideally it would have some sort of walls or curtain separating it from the rest of the factory and it would be painted differently within the area. Within the area there would initially be room for one single hole-drill cell. The equipment in this cell would consist as much as possible of new equipment, including all new hole drilling machines as well as a new airflow bench.

PWPs are widely used in manufacturing operations, although they are typically created to allow a product line to have a focused business strategy consistent with what is required by the product line.[14] The proposed use of a PWP at the TMC goes a step further, however, since it is designed to affect workers' perceptions as well as their processes.

Within the area, the cell would be arranged in a Lean line of flow as alluded to in this thesis and as described in the recommendations to the TMC. Although eventually the equipment may be moved through continuous improvement efforts, it must start out being aligned pursuant to the best design available. The user interfaces on the machines would all be different from their contemporary counterparts. All of the equipment would be painted or marked such that any

---

[14] Anupindi, Chopra, Deshmukh, Van Mieghem, and Zemel 1999, 24-25

outside equipment (e.g., carts, tables, lockers, etc.) that ends up in the area would stand out. 5S markings would delineate the location of all movable equipment. Essentially, the PWP would be a poster-perfect Lean cell. It would also contain all of the process control computing and networking equipment alluded to in Section 5.1 to maintain tight process control.

Besides the look, feel, and equipment in the cell being completely new, a new job classification would also be required to work in the cell. The job classification would be a worker cross-trained for hole-drilling operations, maintenance, and quality control. Only those workers with this classification would be allowed to work in the cell. To get this classification, workers would have to bid for the jobs in the new cell, be chosen, and then undergo special training. The workers chosen would be chosen irregardless of seniority. The workers in the classification would also be immune to being transferred or moved to different jobs for a period of at least a year. Note that one cell only requires three to four workers per shift, so this is not a significant demand for workers: less than 10% of the workforce would be initially used in the PWP. In addition, the cell leader, who would be specifically chosen for the job and a dedicated continuous improvement leader from the TMC's ACE staff, would be permanently and exclusively assigned to the work cell. The cell would also have one or more manufacturing engineers assigned to it and responsible to the cell leader.

The cell leader would not attend production meetings for a period of three months once the cell has begun production. The cell leader would coordinate directly with the hole-drilling materials representative (who is responsible for scheduling production through hole-drilling). There would be no production requirements for the first two months of operation, although the cell would receive priority for incoming parts from the coating operations. Prior to the cell commencing production, every effort would be made to build up an input buffer. For the first two months, that buffer would include scrap parts. The cell would also have a number of training parts on hand that can be re-used. There would be drillable training parts which would already have holes, but which could be passed from machine to machine and operated on by the machines (even though they won't be drilling into anything). There would also be airflow training parts, with various values, that would be substituted for the drilling training parts at the airflow operation in order to provide appropriate airflow training and allow for feedback to be sent to the machines to adjust for airflow. This training cycle would allow the workers in the cell to develop their skills and to correct the cell's problems even if there is a lack of incoming parts.

From day one, the cell would begin production as a flow of parts. Any problems at a station would cause the entire cell to stop until the problem is corrected, and the entire cell team would be involved in the corrective actions. There would be no blame for the operator should a stoppage occur or a bad part be produced. Instead, the team as a whole would deal with problems so that the problems could be eliminated permanently. The cycle time would initially be set very long and as these problems are resolved the cycle time would gradually be reduced. Ideas for improvement would be elicited from the entire team, and improvement events would include the entire team and be lead by the team leader with the assistance of the ACE representative. As cycle times fall, worker and machine productivity would rise, soon surpassing the contemporary cells in the plant, and eventually having productivity two to three times that found in the contemporary cells. The cell leader would be responsible for social rewards for the team for passing milestones and for having new improvement ideas. These social rewards would be recognition based, ensuring that people are appreciated for their contribution. It is imperative that the first cell be implemented correctly, as it would set the precedent for later cells.

The concept begins with a single cell, but once that cell is up and running and exceeding the production rates and quality of the current cells, another cell can be built as more contemporary equipment and cells can be retired. This will continue until the entire current hole-drilling operation is replaced by the new cells.

The above paragraphs explain how to accomplish change in the TMC, but it is important to understand why things should be done in this manner and how these things help to overcome the roadblocks described in Chapter 6.

Establishing a separate plant-within-a-plant is vital to pulling workers out of their comfort zone. Section 6.3 described how the “been there, done that” mentality, emphasis on throughput, and distrust of equipment all force the workers into a comfort zone involving doing things per the status quo. Making a PWP that is as different (visually as well as operationally) as possible from the current factory forces the workers out of their comfort zones. Emphasizing that things are different, as soon as the workers walk into the area, will help them to accept further changes in the way they do things within that area. The fewer ways in which they can relate the new cell to their previous experiences, the less likely they are to be captured by the old ways. Having new equipment and new user interfaces will help the workers to learn how the machines work from a new perspective. After they have operated the new equipment for a while, the efficiencies gained should boost their confidence in the equipment.

The PWP also helps to insulate the cell from the emphasis on meeting demand. To mitigate upper management's roadblock of not wanting to affect the current production, the PWP will have the appearance of being something that is set aside and not affecting the current state, which will continue to run almost exactly as it currently does. This is a sort of "out of sight, out of mind" policy that will help keep the heat off of the cell for a while. Also, it was mentioned above that the cell would not be required to produce any parts for two months nor would the cell leader even be present at production meetings. This does two things. First, it helps to relieve the pressure on the workers and cell leader in the new cell. Second, it sets low expectations up front for upper management. The low expectations will make upper management less likely to interfere with the cell and more likely to ignore it for a while. It might not be possible to get management to accept two months of no production. Instead, it might agree to one month of no production and one month of best effort before production goals are implemented. The point, though, is to set the expectations as low as possible in order to relieve the pressure from above. This single cell will eventually capture a significant percentage of the production—enough to amaze the harshest critics—if it is given sufficient time to ramp up.

The new worker classification helps to isolate the cell on a personal level. It makes it easy to control who is allowed to work in the cell so that only those that are more likely to be willing to accept the changes can be recruited. The workers will be incentivized to bid for the positions because it will involve getting additional training. It will also appeal to the workers that still have hope that things can change for the better. Also, the guarantee that the workers will be there for at least a year is a strong incentive to join, since they will not need to worry about being bumped out of the position by a more senior person for at least that long. Since the bids are taken irregardless of seniority it is also a way to build a classification of motivated workers who have no fear of being bumped from their positions by senior workers. The new job classification also helps to deal with the roadblock of the matrixed organization not being focused enough toward the product side. By qualifying the workers for operation as well as maintenance and by assigning manufacturing engineers and the ACE representative to the cell leader, this aligns the power back to the cell leader. Yet the workers should not be socially isolated from the rest of the workers in the plant. Should this happen they may be considered outcasts, which is undesirable since they need to act as ambassadors for the new system to the rest of the workforce. In the beginning lunch periods should be spent with the team together conducting training and socializing so as to build the team socially. Later, though, after the cell has begun production and started to overcome issues, it would be best for the workers to be able to eat lunch and take breaks with the rest of the plant.

Unfortunately, the union is not likely immediately to support a new job classification. The idea here is to get pull from those workers who want to change things and try something new. For example, the current ACE Pilots would be good candidates. Since this new classification would be completely voluntary, since there should be candidates wanting the classification (especially with some marketing effort on the part of management), and since the first set of jobs should only involve three or four workers, the union is likely to acquiesce. The fact that the new classification also includes training on quality control will also serve as an additional justification for the classification. The union would not be likely to support a job classification that included ACE or Lean training, given its history with ACE, but the same things can be taught under the guise of quality. The union is also not likely to want the bids for this position to be seniority-independent. However, most of the very senior people are not likely to want this job classification. It can be marketed as something different and something that will require a lot of work. Most of the very senior workers will be just as happy not to volunteer for the new classification, and therefore eventually the union will go along with that part of the deal as well.

The point here is not to force the union to accept the change. Rather the point is to build a classification of workers from volunteers and then to show through the example of this cell that not only can production be done more efficiently, but it can be done in a way that takes care of the workers involved and does not undermine union authority. Section 6.2 asserts that the difficulties between union and management are largely due to communication problems, and therefore by using this new work cell and its workers as an example, management can communicate to the union what is possible and that the new state of affairs is not a threat. The union will likely see this new cell as a "speed-up" cell, denoting what is to come. Therefore, the cell has to demonstrate to the union that what is to come actually works and is not a threat to the union or the workforce. The union must understand that the workers participating in this new venture are attempting something new for the company and the union and are not traitors to the union. If it does not work, the experiment would be small enough that no harm would be done. If it does, it will demonstrate that the needs of the workers and the union can be met.

How the operations in the cell are developed and ramped up, and how continuous improvement is done within the cell is also very important to overcoming roadblocks. By starting from day one with a flow mentality, this will help move the TMC's Lean experts away from the "stability first" mentality, as well as provide positive feedback for the workers in the cell that their efforts have an effect. The way that the cell deals with continuous improvement—from the organization of the Kaizen events to the fact that everyone in the cell is involved, including the cell leader—will help to change the "been there, done that" mentality once workers see that change can happen and can be sustained.

By developing a plant-within-a-plant, populated by cross-trained workers and motivated supervisors, changes such as the ones described in this thesis can be accomplished in the TMC and rolled out one cell at a time.

7.2 Accomplishing Change from without

Section 7.1 discusses a method for accomplishing change at the TMC without a crisis. If this is not possible, a crisis, such as those suggested by Womack and Jones must be created.[15] The easiest way to cause a crisis at the TMC is not on the inside, but rather from the outside.

Typically when demand outstrips capacity at a factory the managers and workers at the factory are in no danger of losing anything. This overlooks the obvious fact that the business cannot continue forever not meeting demand. Drastic measures may be needed to compensate for the capacity shortcomings. Since the factory itself would not be running efficiently, it is unlikely that the business would expand its own operations. This is exactly the position in which the TMC is now. It must find a way to increase its capacity to meet demand. One way would be to implement changes such as those described in this thesis, but if there is too much reluctance to change, that might not be possible. This, combined with a history of poor union relations—as is the case at the TMC—may force the company to go elsewhere for capacity.

For Pratt, there is ample precedent for this. Currently the TMC has a number of operations performed by vendors. This is called temporary vendor assist (TVA), and has been discussed previously in this thesis. If TVA prices get low enough, it is within Pratt's ability to

---

[15] Womack and Jones 1996, 250-251

completely outsource an operation or even a whole product. This has happened in recent years with less advanced turbine vanes. TVA could be used as one type of crisis that might inspire some willingness to change.

However, Pratt has stated that it wishes to maintain hole drilling capability and advanced coating capability. These are the technologies that allow the turbine airfoils to survive the extreme temperatures of the engine and are seen by Pratt as core competencies, and ones that provide a strategic advantage as long as they are kept in house.

The other crisis that Pratt could cause would be to open up a new, small, hole-drilling factory somewhere outside of Connecticut. It could be a non-union factory and a Greenfield operation. Pratt could take the concepts described in this thesis and much more easily enact them in a Greenfield environment with non-union workers. Although it would be risky to attempt to move all hole drilling to a new site, this would not be necessary. The operation would have to be large enough to be self-sustaining, but the real purpose of the new plant would be to compete with the drilling operations that stay within the TMC. This would put pressure on the managers and workers in the TMC to change in order to stay competitive. The mere threat that all drilling could be moved to the new factory should be enough of a crisis to motivate significant changes.

Although this method may work, it does so by forcing the union and the workforce to change. It will undoubtedly be better for Pratt, the union, and the workforce if everyone tries to cooperate to bring about change rather than having one party force the others into it. While this plan would likely have success, it should only be undertaken if there is little hope for getting the union and workforce to cooperate, as suggested in Section 7.1.

It is often easier to cause a crisis from outside the company than to attempt change without a crisis. In the case of the TMC, capacity deficiencies can be used as a lever since they can give rise to outside entities evolving to compete with the TMC. This competition can be the driving force that is needed to bring about change in the culture.

Chapter 8: Conclusion

Designing new multi-product work cells and implementing them in an existing manufacturing environment is a challenging task. It requires decisions about the technology to use in the machining operations, decisions regarding the number and types of machines to use in the cells to adequately machine the various products, and an implementation process that takes into account the strategic design, political, and cultural challenges inherent to change.

This thesis has presented solutions to these problems in the context of building work cells for the hole drilling operations of high pressure turbine vane manufacturing. The thesis has discussed and quantified the benefits of using an optimization framework for making decisions about how to set up work cells and assign production to each. It has also discussed and quantified the benefits of using Lean practices in the work cells in order to operate them most efficiently. The thesis has also shown how important the difference between fixed and variable manual time is to comparing two competing technologies in the case of HSEDM versus laser drilling. Finally this thesis has described techniques for promoting change in a change-adverse environment like the one at Pratt & Whitney's Turbine Module Center.

Taking a system-level view of the design problem is one key to building work cells that are truly efficient and cost effective. The other key is operating these cells in an efficient, standardized manner, as outlined by Lean manufacturing.

Bibliography

Anupindi, Ravi; Chopra, Sunil; Deshmukh, Sudhakar D.; Van Miegham, Jan A.; and Zemel, Eitan. 1999. *Managing Business Process Flows.* 1st Ed. Upper Saddle River, NJ: Prentice Hall.

Bertsekas, Dimitri P. 2000. *Dynamic Programming and Optimal Control.* Vol 1. 2nd Ed. Belmont, MA: Athena Scientific.

Carroll, John S. 2002. *Introduction to Organizational Analysis: The Three Lenses.* MIT Sloan School of Management Working Paper 14, 1-13. Cambridge, MA. Sloan Communications Office.

Dennis, Pascal. 2002. *Lean Production Simplified.* New York: Productivity Press.

Goldratt, Eliyahu M. 1990. *Theory of Constraints.* Great Barrington, MA: North River Press.

Goldratt, Eliyahu M. 1992. *The Goal: A Process of Ongoing Improvement.* 2nd Revised Ed. Great Barrington, MA: North River Press.

Goldratt, Eliyahu M. 1997. *Critical Chain.* Great Barrington, MA: North River Press.

Muir, Stephen H. 1997. Semiconductor Fabrication Capacity Gains Through Optimal Allocation of Product Loadings and Equipment Sharing Across Process Technologies. S.M. Thesis, Massachusetts Institute of Technology.

Page, Alison L. 2001. Forecasting Mix-Sensitive Semiconductor Fabrication Tool Set Requirements under Demand Uncertainty. S.M. and MBA Thesis, Massachusetts Institute of Technology.

Ragsdale, Cliff T. 2004. *Spreadsheet Modeling & Decision Analysis.* 4th Ed. Mason, OH: South-Western.

Womack, James P. and Jones Daniel T. 1996. *Lean Thinking.* New York: Free Press.

Womack, James P.; Jones, Daniel T.; Roos, Daniel; and Carpenter, Donna Sammons. 1991. *The Machine That Changed the World.* New York: Harper Perennial.