A Capital Equipment Capacity Planning Methodology for Aerospace Parts Manufacturing in a High-Mix, Low Volume Environment

by

Matthew A. Reveley

B.S. Mechanical Engineering, Tufts University, 2006

Submitted to the MIT Sloan School of Management and the Department of Mechanical Engineering in Partial Fulfillment of the Requirements for the Degrees of

Master of Business Administration
and
Master of Science in Mechanical Engineering

In conjunction with the Leaders for Global Operations Program at the Massachusetts Institute of Technology

June 2012

Signature of Author ______________________________________
MIT Sloan School of Management, Department of Mechanical Engineering
May 11, 2012

Certified by ______________________________________
David Hardt, Thesis Supervisor
Ralph E. and Eloise F. Cross Professor of Mechanical Engineering

Certified by ______________________________________
Roy Welsch, Thesis Supervisor
Eastman Kodak Leaders for Global Operations Professor of Management

Accepted by ______________________________________
Maura Herson, Director of MIT Sloan MBA Program
MIT Sloan School of Management

Accepted by ______________________________________
David Hardt, Chairman of the Committee on Graduate Students
Department of Mechanical Engineering

This page intentionally left blank.

A Capital Equipment Capacity Planning Methodology for Aerospace Parts Manufacturing in a High-Mix, Low Volume Environment

by

Matthew A. Reveley

Submitted to the MIT Sloan School of Management and the Department of Mechanical Engineering on May 11, 2012 in Partial Fulfillment of the Requirements for the Degrees of Master of Business Administration and Master of Science in Mechanical Engineering

Abstract

A static capacity planning model was developed and tested following a four-phased framework. This model was developed for the purposes of capital planning for capacity requirements at a large aerospace parts manufacturing plant. Implications for capacity planning of the nature of the aerospace industry, as well as the company and plant being studied are discussed, as well as the current state of capacity planning.

In phase I of model development, an appropriate modeling solution is selected. In phase II, information is collected from the user base as to the desired user experience and functionality of the model, as well as the parameters that should be considered in it. Phase III involves assessment of the parameters' impact on capacity, and identification of appropriate data sources to feed the model. Additionally, phase III recommends changes to current data structures in order to optimize the balance of model accuracy with minimal incremental resource allocation. In phase IV, the mathematical model is explained, and the user interface is developed. With a working model, the results are validated with the shop floor, identifying gaps in data sources previously unobservable.

Following model development and validation, the model is applied to a subset of the shop, and used to develop recommendations for addressing predicted future capacity constraints. Application of the model reveals a blind spot in current heuristics-based planning, where high development loads can lead to immediate capacity constraints, but effects of the experience curve can actually cause this constraint to disappear on its own, without the need for excess equipment purchases.

Finally, extensions of the research and lessons learned are discussed, suggesting future project work within the plant studied, as well as elsewhere in the company and in other companies or plants.

Thesis Supervisor: David Hardt
Title: Ralph E. and Eloise F. Cross Professor of Mechanical Engineering

Thesis Supervisor: Roy Welsch
Title: Eastman Kodak Leaders for Global Operations Professor of Management

This page intentionally left blank.

Acknowledgments

The author would like to thank first and foremost his colleagues and advisors, as well as the staff and faculty at MIT, the Sloan School, and the Leaders for Global Operations program and alumni. Specifically, the author would like to recognize: Professors Roy Welsch and David Hardt for their guidance and for multiple treks up to Maine; Don Rosenfield, Jane Deutsch, Ted Equi, Jeff Shao, Patty Eames and the rest of the LGO administration for their endless patience and diligence; Jeremy Lieberman, Christina Williams and Jason Chen for their input, advice and support.

Additionally, gratitude is due to United Technologies, Pratt and Whitney, and the North Berwick Parts Center for their support of the LGO program and for sponsoring this internship and being gracious hosts. Specifically, thanks go to: Mike Newsky, Mike Papp, Steve Ruggiero and Chris Karcher for their support of the project, and for their endless wealth of knowledge; Ron Boone, Gino Veneroni, Kerry Thompson, Pat McKenna, Wayne Sevigny, Ken Allen and Colin Claffey for their levity, knowledge, and patience; Mariam Correa for her overwhelming support and enthusiasm; Wade Maynard, Pat Reagan, Dick McDonald, Ty Cryer, Keith Pooler, Mac Mann, Ken Dickie, Steve Mongiat, Mike Macmahon, Rene Thibeault and Dave Cote for their time, effort, and humor.

Finally, the author would like to thank his friends and family for their support.

This page intentionally left blank.

Table of Contents

List of Figures & Tables

1 Chapter 1 – Introduction and Thesis Overview

Virtually all large manufacturing organizations are faced with difficult decisions regarding the investment of capital in manufacturing equipment. For capital-intensive industries such as aerospace, understanding and quantifying the risks associated with these investments, and the demands driving the need for investment are critical to a firm's success. (Vranakis & Chatzoglou, 2012) These challenges are amplified as the manufacturing environment becomes more complex.

Aircraft engine part manufacturing represents an archetypal example of a high-complexity heavy manufacturing environment, characterized by an extremely diverse product portfolio with relatively low part-by-part volumes. This complexity vastly increases the amount of data that is necessary to fully characterize the manufacturing system. Unfortunately, not all of this information is relevant to the business' day-to-day decision making, so there is a diminishing and eventually negative marginal benefit to monitoring and controlling each subsequent parameter. To deal with this problem, large manufacturing organizations pick and choose the degree to which the myriad parameters of the system are monitored and controlled based upon their relative utility and perceived impact on relevant decisions.

This thesis seeks to apply a structured approach to developing an extensible tool for understanding capital equipment capacity for the purposes of long-term capital planning, while operating under the constraints of a high-complexity manufacturing environment. Through this process, organizational "blind spots" are identified upon which poor investment decisions can be made without the aid of such a tool.

1.1 Strategic Capacity Planning in Manufacturing Organizations

Capacity planning has been extensively studied as it relates to investment decision-making in manufacturing industries. Specifically, there has been extensive research into the application of capacity planning methodologies in the semiconductor manufacturing industry, where capacity expansion can cost firms billions of dollars. (Geng & Jiang, 2009) Geng et. al. evaluated a host of capacity planning methodologies for the semiconductor industry, including static capacity modeling, the "neighborhood search method", and applications of linear and stochastic programming. A static capacity model is aggregated, and, while fairly easy to use, may be subject to substantial inaccuracy due to this aggregation. Neighborhood search methods attempt, through trial and error within a simulated environment, to identify the optimal investment strategy by evaluating small incremental changes around the current strategy. By contrast, mathematical programming uses mathematical modeling to systematize the optimization process in order to evaluate the full investment space.

Tenhiälä et. al. assess various planning methodologies' efficacy depending upon the type of process under consideration. Four types of processes are assessed, "job shops", "batch processes", "batch processes with bottleneck control", and "production lines", organized in decreasing order of complexity. Planning methodologies assessed are "rough-cut capacity planning", "capacity requirements planning", "finite loading with capacity leveling" and "finite loading with optimization", in order of increasing sophistication. It is suggested that rough-cut capacity planning, or looking at high-level aggregate data, is most effective for job shops characterized by extremely high mix, low-volume manufacturing with pooled resources. Secondly, it is suggested that batch processes, characterized by high-mix low-volume production with some task interdependence as seen in mixed-model cellular shops, are best suited for capacity requirements planning, enabling planning for individual machine or aggregate labor capacities. Finally, Tenhiälä et. al. suggest that the most sophisticated methods of finite loading with capacity leveling or optimization (which enable planning for shift schedules and setup time reductions as well as jig and tool capacities) are best suited to the least complex environments such as batch processes with bottleneck control where overall rate is determined by a single asset, and production lines, where tasks are completely sequential. (Tenhiälä, 2011)

For aerospace parts manufacturing at Pratt and Whitney's North Berwick Parts Center, a static model is built and implemented using the methodology proposed by Ozturk et. al. This model is built for Tenhiälä's capacity requirements planning capability, since the environment is best described as a "batch process". From the static model, a basic capacity planning method is applied using the neighborhood search methodology studied by Geng et. al. with very basic stochastic modeling used to guide manual scenario testing and analysis. This methodology was selected because of the ease of use of a static model, combined with the existing shop floor planning process, built upon local specialized process knowledge.

This approach also carries the smallest technical hurdles for implementation. Due to the high-mix, low-volume manufacturing environment at the North Berwick Parts Center and the high degree of flexibility inherent in the system for process adjustments, keeping the process of capacity planning primarily manual enables incremental improvement over the existing framework for decision-making. Implementing a fully automated system optimization model would be not only an extremely difficult undertaking owing to the high level of complexity and extremely large number of decision variables; it would also represent a substantial organizational change, significantly increasing the barriers to adoption.

1.2 Static Capacity Modeling

Ozturk et. al. propose a four-phased approach to developing a static capacity model within a manufacturing organization. These phases are as follows:

I. Selection of Modeling Solution
II. User Requirements Specification
III. Design of Data Structure and Information Flow
IV. Implementation and Validation (Ozturk, Coburn, & Kitterman, 2003)

This approach is applied through model validation, and tested within a subset of the plant's manufacturing organization.

1.2.1 Phase I – Selection of Modeling Solution

During the modeling solution selection phase, the proper mode of implementation must be selected. For example, in some situations a model can be implemented within existing software, or an off-the-shelf solution can be purchased. In others, however, situational considerations may dictate that a standalone model should be developed.

If a standalone solution is deemed to be the best option, then a platform must be selected, as well as a mode of deployment. Depending upon the structure of the organization, deployment may need to happen at a different level, or at multiple levels simultaneously. For example, if deployment is to happen within a production analytics department, then the solution may be able to push more of the analytical design onto the user, opting for a more flexible solution with a higher degree of customization. Alternatively, if day-to-day use is to be left to a capital planning department, or to upper management, the operational analytics may need to be more constrained in order to simplify the user experience.

1.2.2 Phase II – User Requirements Specification

Once a solution and platform have been selected, the target users' specific requirements must be determined. This specification should include not only critical parameters for the model, but also the degree of simplicity and range of capability required, as well as aspects of the user experience. User requirements can be broken down into two primary sub-tasks, determination of user requirements for functionality, and determination of expectations for user experience. It is of critical importance that any tool be designed based upon a target user experience, and that this step be continuously revisited throughout the course of model development in order to maximize the chances of adoption.

1.2.3 Phase III – Design of Data Structure and Information Flow

Once the user specifications have been determined, they must be distilled into a mathematical model, mapping the process of the conversion of the required inputs to the required outputs. In this phase, discrepancies between existing data sources and sources required to meet user specifications are identified, and must be addressed either through the inception of tandem projects to create or update the sources available, or by adjustment of the user specifications.

1.2.4 Phase IV – Implementation and Validation

The first three phases encompass the design and specification of a capacity model. This final phase involves the actual creation of the model, validation of its outputs, and the improvement process of using shop-floor feedback to improve the accuracy of the model. Until the model is validated, it cannot reasonably be used as a decision aid. Additionally, the user interface can be developed here, as well as the format in which data will be presented. The process of validating and updating the model based upon shop-floor feedback can force iteration of all three previous phases of model development. However, with careful execution of the first three phases, this may not be necessary.

1.3 Thesis Overview

The research for this thesis was conducted by applying this four-phased framework to the development of a capacity planning model for Pratt and Whitney's North Berwick Parts Center. Chapter 2 reviews relevant industry, company, and plant background information, and how these factors lead to the current state of capacity planning in North Berwick. Chapters 3 – 6 walk through the four phases of model development in-depth, and how they were applied through data collection and analysis to the company. Chapter 7 discusses the results obtained through application of the model. Finally, Chapter 8 reviews lessons learned and suggests extensions of this research both in scope within the same plant, and to other plants and industries.

Figue 1 shows sample output from the prototype model implemented as a decision aid. In these charts, two new machines (179705MT-8109) are added to combine operations and remove some of the constrained load on machine 179705B-8109. A layout of typical raw prototype model output is attached in Exhibit 1.

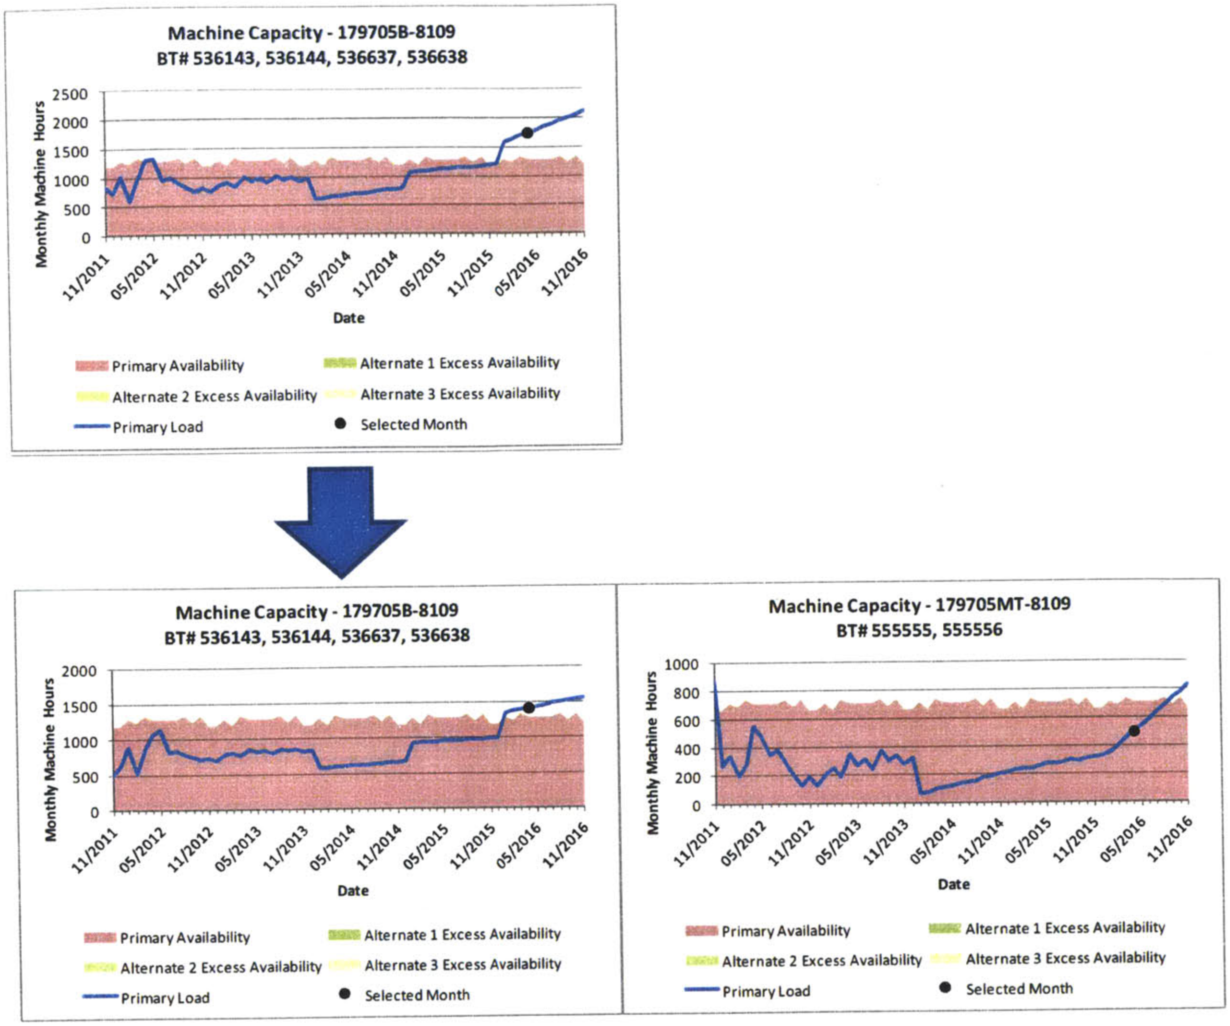

Figure 1 – Sample Model Decision Aid (179705B-8109)

2 Chapter 2 – State of Capacity Planning at the North Berwick Parts Center

Pratt and Whitney is a leading manufacturer of gas turbine engines, primarily for the propulsion of military and commercial aircraft. The North Berwick Parts Center (NBPC) manufactures over 1,200 different parts from many different sections of Pratt's various engine models. Most of these parts are static within the engine, allowing them to be made of metal rather than composite or ceramic materials. Because of this, most manufacturing processes at NBPC tend to be large-scale computer numeric controlled (CNC) removal machine tools such as mills, lathes, and grinders, with a few other types of processes as well. Almost all of the equipment required, however, is very expensive. As a result, planning for capital investment has always been a critical process at NBPC. The need for capital investment can be driven by several factors, including equipment end-of-life replacement, new technology investment, and capacity expansion to accommodate anticipated new work. The work of this thesis will focus on the third type of capital investment.

2.1 Relevant Industry Background

2.1.1 Market Forces

The United States domestic aircraft, engine, and parts manufacturing industry has three primary markets. The commercial and private aviation market, accounting for 45.9% of industry revenue in 2011, consists primarily of commercial passenger airlines and private cargo transport operators. The US government is estimated to have contributed 29.1% of industry revenue in 2011 for military and emergency services aircraft. The remaining 25% of 2011 industry revenue is estimated to have come from exports to global markets. (Samadi, 2011)

The commercial and private aviation market's demand driven by passenger airlines will respond in kind to demand for air travel. Likewise, cargo transport operators' demand for aircraft is driven by demand for cargo shipments. Because aircraft tend to have very long lead times, and require large amounts of capital investment, this sector's demand will respond primarily to long-term trends with regard to the aforementioned factors. In the short-term, fluctuations in passenger and shipment demand will result in aircraft purchase deferrals or delays. Furthermore, competition in the commercial passenger airline industry demands that players invest in newer aircraft as they become available (Samadi, 2011) Functionally, for aircraft, engine, and parts manufacturers with a diverse customer base, these dynamics mean that there is fairly good visibility to schedule ramps in the commercial sector as long-term trends and commercialization of new aircraft both tend to be fairly slow-moving and highly visible forces.

The military sector, on the other hand, tends to be subject to a lot more fluctuation. While the government funds development of new aircraft generally based upon a well-planned purchase contract, demand is subject to rapid and unexpected changes due to political action in response to world events. For example, Lockheed Martin's year-end Securities and Exchange Commission (SEC) filings in 2001 reflect the optimistic uncertainty surrounding the military aviation sector's outlook after the attacks on September 11th, and the government's Quadrennial Defense Review (QDR):

"The President's proposed budget for the U.S. Department of Defense (DoD) for fiscal year 2003 and beyond reflects the above-mentioned transformation of national defense policy and responds to increased needs for homeland security and defeating terrorism. Budget increases are projected for operational readiness and personnel needs, as well as for both the procurement and the research and development accounts. While there is no assurance that the proposed increased DoD budget levels will be approved by Congress, after over a decade of downward trends, the current defense budget outlook is one of growth. [Lockheed Martin]'s experience and capabilities are well aligned with U.S. defense priorities. Uncertainties remain, however, relative to the level of growth and the amount of the budget that will be allocated to the investment accounts." (Lockheed Martin Corporation, 2002)

Consequently, military demand is subject to substantially more long-term uncertainty than the commercial sector.

2.1.2 Industry Structure

The production of aircraft is a massive undertaking, utilizing resources spread to an increasing degree all around the globe. Air framers such as The Boeing Company and Airbus, a subsidiary of the European Aeronautic Defense and Space Company perform final assembly of the aircraft, with many of the individual systems, such as propulsion systems, wing, and landing gear manufacturing subcontracted out to various suppliers. These first-tier suppliers assemble the systems for the air framers. However, they often utilize a fabrication supply base for the actual production of their parts, creating a second tier of suppliers.

This succinctly describes the structure of the original equipment manufacturing (OEM) side of the aerospace industry. However, because of the extremely high capital cost of aircraft, there is also an active aftermarket for both individual parts and modular systems for maintenance, repair, and overhaul of the aircraft. This means that, while first and second-tier suppliers may act as such for the OEM market, many of these same companies interact with the end users or maintenance service providers in the aftermarket, often selling spare or replacement parts and systems directly.

The dynamics affecting demand in the aftermarket can be substantially different from those affecting OEM sales. Spare parts and modules (or "spares") sales are affected by the installed base of the equipment of which the spare is a part, the part's design and purpose, and the end user's utilization of the installed equipment. Per a conversations with engineering staff at Pratt & Whitney, blade outer air seals (BOAS) are designed to be a wear part that requires periodic refurbishment or replacement. They act as a blade seal in the engine's high-pressure turbine (HPT). The seal is created between the rotor blade and a static coating on the BOAS. This coating deteriorates with use, and as a result must be replaced. On the other hand, some parts may rarely if ever require replacement, such as static support structures.

For a parts manufacturer, the aftermarket can act as a buffer against OEM downturns, providing continuous part volumes based solely on the amount of previously sold equipment. However, aftermarket parts support will be required for the lifetime of that installed base. Consequently, acting in this capacity can also be a liability to a parts manufacturer, requiring production of legacy parts long after production of newer equipment designs has reshaped the factory's operations.

2.1.3 Implications for Capacity Planning

The variability intrinsic to the military sector subjects demand forecasts to large and sometimes unexpected swings when inflection points are reached. This makes capacity planning for manufacturers or equipment sets largely dedicated to the production of military parts difficult, and requires continuous updating of investment plans in order to keep investment in-line with expected need.

The spares market also requires a similar approach, necessitating that manufacturers or equipment sets dedicated primarily to spares production frequently update their investment plans. However, the extremely long duration that spares production is required to continue for a mature engine set also requires that smart decisions be made with regard to the extent to which an aging set of spare parts with trailing demand should continue to be manufactured in-house. Research has shown that planning under demand uncertainty can lead to underinvestment in appropriate capital. (Fuss & Vermeulen, 2008)

2.2 Manufacturing at Pratt & Whitney

Pratt & Whitney's (PW) gas turbine engine manufacturing operation acts as an OEM subcontractor to air framers as well as an aftermarket maintenance, repair, and overhaul (MRO) service provider to end users. PW's operations are spread across various facilities around the world. However, engine final assembly is located primarily in Middletown, CT and as a result, there is a high concentration of facilities located in the northeastern United States and parts of southeastern Canada. The Middletown location performs final assembly of the engine from subsystems, or "modules" that are produced at various other locations, dubbed "mod centers." These mod centers build modules from a combination of internally fabricated parts, and parts purchased from the vendor base.

PW produces engines for both the military and commercial sectors of the aerospace market, and also provides MRO services to both of these markets as well. PW's major gas turbine engine products are shown below in Table 1. PW also produces a number of small and medium-sized engines, as well as rocket engines and ground-based gas turbines. However, for the purposes of this research, the focus will be on commercial and military gas turbine aircraft engines, as these make up the overwhelming majority of the work at NBPC.

	Engine Model	Associated Aircraft
Commercial	PW2000	Boeing 757
	PW4000	Boeing 747, 767, 777; Airbus A300, A310, A330
	PW6000	Airbus A318
	GP7000	Airbus A380
	V2500	Airbus A319, A320, A321
	PurePower® PW1000G	Bombardier Cseries; Mitsubishi Regional Jet; Airbus A320neo; Irkut MC-21
Military	F100	F-15 Eagle, F-16 Fighting Falcon
	F117	C-17 Globemaster III
	F119	F-22 Raptor
	F135	F-35 Joint Strike Fighter

Table 1 – Pratt & Whitney's Major Gas Turbine Engine Products (Pratt and Whitney, 2011)

As shown above, PW has several products in each of the major markets, military and commercial. At the time of this research, the PurePower® PW1000G engine, or geared turbofan (GTF) engine, has recently gone into early-stage production for test engines, and the F135 engine program has just begin delivering production engines for sold aircraft. These two engine programs are expected to increase in volume of OEM production in the coming years. The other eight engine programs are fairly mature, and continue to sell both into the OEM market and the aftermarket.

2.2.1 Engine Manufacturing

The Middletown engine center's production schedule is built from an aggregate of projected demand from the company's various customers. Because of the extremely high price tag on the fully assembled engines, customers' demand projections are given several years in advance, and fluctuations in this schedule tend to be relatively small, subject to the aforementioned market conditions. From this baseline engine schedule, PW has an enterprise resource planning (ERP) system that generates module delivery schedules based upon when each module for each engine will be needed at the engine center, backing off the expected lead time for each module. Synchronizing the schedule in this manner serves to reduce in-process inventory, creating a pull system between the mod centers and the engine center. From this module schedule, each component part's schedule is generated by the same process. However, at the component level, there is a wide variety of different types of parts with different degrees of internal manufacturing content. Many small parts, such as nuts and bolts, rivets, and other standard hardware are purchased and stored on-site. Some of the larger manufactured parts are sourced to the vendor base, causing a manufacturing schedule to be released to each supplier based on the ERP schedule. In addition, some parts are manufactured internal to PW within the various mod centers. Each of these internally manufactured part schedules is then used to generate a raw material requirements schedule. The raw materials for each component part may be a combination of small parts held in inventory on-site, and make-to-order castings or other more expensive materials conducive to a pull system.

Each ERP-generated schedule is based on a standard lead time associated with each particular part, which incorporates historical order-to-delivery performance data coupled with a buffer against unforeseen local delays. Additionally, each part schedule is built based on an assumed scrap factor, often manually managed by the materials planning organization. While the engine schedule may provide fairly good visibility, these two manual levers can serve to obscure the visibility of the actual engine schedule from module centers such as NBPC.

2.2.2 Spares Manufacturing

In addition to the engine center's assembled engine requirements from the customer, PW also sells spare parts and modules to its customers for maintenance and testing. The sales model for spares is substantially different from that for engines. While engine schedules are often planned well in advance and based upon pre-determined contracts, spares are offered in an a-la-carte fashion. Each salable spare has a standard lead time and price that can be quoted to the customer. The spares schedule is then generated as the customer orders parts subject to this lead time. Depending on how the customer decides to place these orders, however, this system can create problems for the manufacturing floor. Many customers have contracts for stable spare parts deliveries based on a make-to-stock system. This, however, is not always the case.

Because lead times are based on an average of demonstrated order-to-delivery turnaround data, there is no case-by-case allowance for what is going through the shop floor at a given time. Standard lead times are updated periodically, and since business is cyclical, this eliminates some of the variability from period-to-period. However, it does not automatically account for changes in the product mix, or, more importantly, for the size of the order. This is a particularly relevant issue for military spares, as the military tends to order large amounts of parts, very infrequently, in line with its budgeting cycle. These bulk orders then create ERP requirements that may be unreasonable for the shop floor to meet. In order to accommodate these orders, the materials organization will then work to smooth them and pre-build the parts based upon tribal knowledge of what the shop's capability is to produce each part number.

If this manual smoothing process still leads to insufficient capacity, the organization can either work with the customer to extend the lead time, sequester parts previously allocated for engine deliveries (this is particularly common for small fluctuations or "hot" parts with requirements within the quoted lead time), or utilize second source manufacturers to make up any projected shortfall.

2.2.3 Management and Information Systems

United Technologies' operating system, called "Achieving Competitive Excellence" or ACE, governs the factory structure and operations at PW. ACE focuses on principles of lean manufacturing to establish processes and tools for managing operations on and off the shop floor. Additionally, ACE provides a set of metrics upon which departments', manufacturing cells', and plants' assessments and comparisons are based. ACE metrics are posted on boards located in every department throughout the plant, enabling at-a-glance assessment of what is going on in a particular area.

PW uses SAP software as its ERP platform. SAP brings together many of the different data sources scattered throughout the various departments in PW, enabling financial reporting and management to be linked directly with production scheduling and materials planning, as well as the other support organizations such as maintenance, quality, and purchasing. Because every business is different, many of these modules were designed specifically for PW when the ERP system was initially implemented. ERP software platform implementation has been extensively studied ever since it evolved from manufacturing resource planning (MRP) software platforms in the 1970s. However, it has been suggested by previous research that planning modules in these software packages tend to be used less than intended, due in part to their inability to deal with uncertainty. (Halsall, Muhlemann, & Price, 1994)

With PW, as with many companies, this implementation has been largely successful for day-to-day and quarter-to-quarter business. However, several of the support modules, including the one for capacity analysis, have fallen into disuse. In speaking with a member of the operations staff regarding the reason for this module's abandonment, the latent issues became apparent. "You may be able to build me the best 747, but that doesn't make me a pilot. [The model] isn't worth anything if it takes [an advanced degree] to run it." (Interview of CT Operations Employee, 2011) The capacity module built into SAP was designed to be extremely flexible within the constraints of the data systems and user interface upon which it is based. Unfortunately, this has made for a system that is very difficult and more importantly, time-intensive to use.

2.2.4 Implications for Capacity Planning

Because the spares ordering process is so dynamic, it is necessary for the purposes of long-term planning to maintain an "off-ERP" forecast of projected spares orders. This process is managed by each of the engine program offices. However, because these forecasts exist outside of the ERP system and are non-standardized, this information can be difficult to obtain and aggregate, and any ERP-based planning systems do not fully take it into account.

For long-term planning, PW's manufacturing structure allows for limited schedule flexibility for parts slated for engine contracts, with slightly more flexibility in the schedules for spare parts. However, for a parts center such as NBPC, the entire delivery schedule is somewhat inflexible as parts schedules are determined explicitly from assembly floor. The exceptions are when schedule can be pulled in to "pre-build" to projected demand, and the day-to-day activity of scheduling work on machines (such as process batching).

ACE is focused on characterizing and helping to control well-bounded processes, or processes whose procedures and impacts are confined to a well-defined organizational niche. For example, PW's capital equipment procurement office (CEPO) has well-characterized process maps and controls on the management of capital projects across the various module centers. However, this process is bounded; it begins after capital justifications have already been established, and ends with the delivery of equipment. The reason for this is that these boundaries approximate the extent of the responsibility of CEPO on capital projects. However, the processes leading up to and following this process fall under the responsibility of the myriad module center management teams. This can be attributed as much to the processes being outside of the ACE operating system, as to the highly variable needs of the module centers. In speaking with the CEPO Manager, the scale of this variability became apparent:

"Berwick has done a good job refurbishing machines. [Another mod center] likes to upgrade their best machines. . . [A different mod center] and Berwick were the top spenders, but [another mod center] used to be. . .[A mod center in] Georgia has the biggest difficulties. They have the most expensive equipment, and the lowest volume of capital projects. They purchased one machine for $3 million with a 14-month lead time. . . The [third mod center] gets a lot less capital, because they have a hard time spending the money they are allocated." (Interview with CEPO Manager, 2011)

Because of these variable needs, standardization of the capital justification and post-installation processes has been left to the discretion of the module centers, leading to varying degrees of rigor from center to center, and department to department.

2.3 The North Berwick Parts Center

The North Berwick Parts Center is one of PW's primary module centers. However, it is different from all of the other module centers in that it does not produce a module for a specific section of the engine.

Instead, NBPC acts as an OEM parts manufacturer for complex parts required by many of the other module centers for module assembly. In addition, NBPC performs manufacturing process development activities for new hardware, some of which are later transitioned out to the vendor base for full-scale manufacturing, while some are scaled up in-house and added to the shop's production load. Lastly, NBPC also houses a section of PW's aftermarket MRO business, called Global Service Partners (GSP).

2.3.1 Management Structure

NBPC's management is generally geared toward operations. An organizational chart is shown below in Figure 2. The shop floor is divided into six primary OEM business units, with a seventh just beginning to split from the unit of which it was previously a part. Each of these seven business units (BUs) has a business unit manager (BUM) responsible for all of the unit's operations. In addition to the seven OEM BUMs, there is also a manager for machine tool services (MTS), NBPC's in-house tool manufacturing and refurbishment shop, a manager for the entire GSP business on-site, and a manager from the "shared manufacturing services" section of the shop floor, which includes common batch processes such as heat treat furnaces, brazing and plating. Each of these 10 individuals reports to the plant's product manager, who reports to the general manager. In addition to the product manager, each of the shared services groups including human resources, finance, materials planning, manufacturing engineering, design engineering, maintenance and facilities has a manager that reports to the plant's general manager.

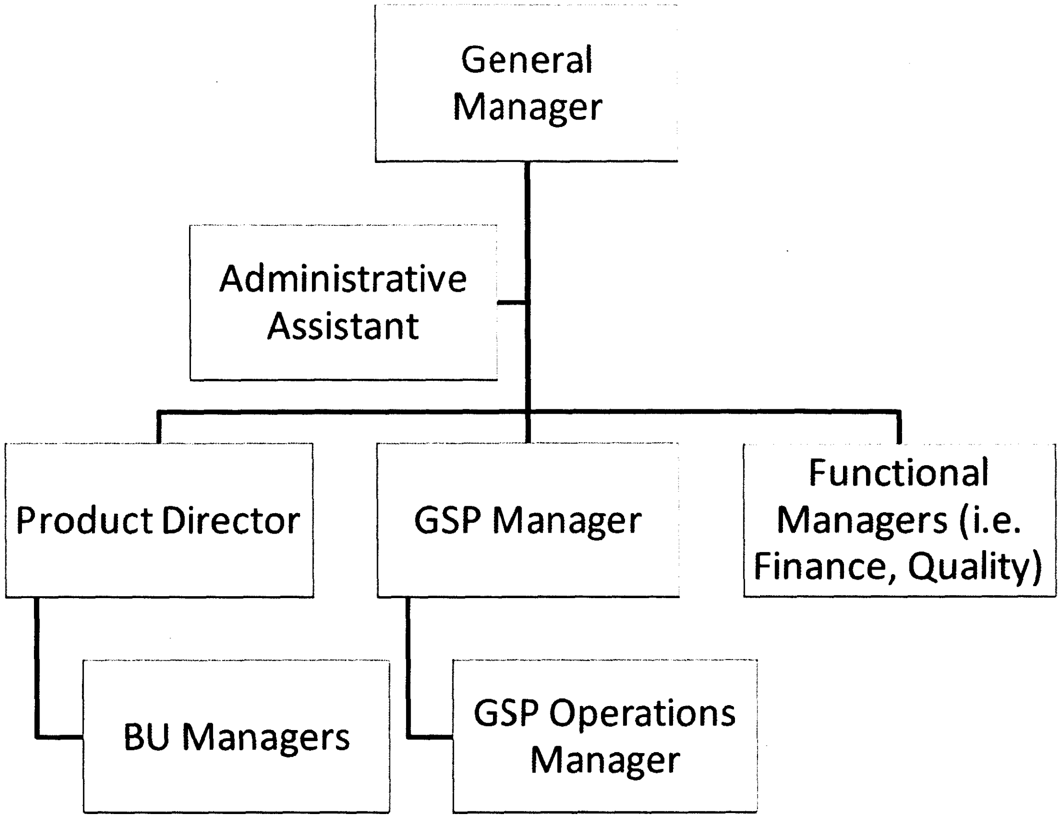

Figure 2 – NBPC Organizational Chart

However, in spite of the organizational chart, the general manager is heavily involved with the BUMs and operations managers, with each individual present both at the product managers' and the general managers' staff meetings, while this is often not the case for other functional mid-level managers. This keeps all levels of management well-informed about what is happening on the shop floor, but also incentivizes unit managers to act autonomously, as they are held individually and independently accountable for their own metrics up to the top level of local management.

BUMs typically have a staff of manufacturing cell leaders or shift supervisors, as well as a materials controller and an ACE leader that often act as information shuttles between the shop floor and the BUM. Each cell leader or shift supervisor manages the shop technicians and engineering technicians that run the equipment on the floor.

In addition to the BUM, each BU has a technical supervisor who manages a group of manufacturing engineers (MEs). While the technical supervisors officially report through the functional manufacturing engineering organization, they work very closely with the BUMs. The technical supervisor often acts as a stand-in for the BUM in meetings when the BUM is not available. The MEs provide engineering support for shop floor operations, including new process development, process troubleshooting, and change management.

2.3.2 OEM Parts Manufacturing at NBPC

As previously mentioned, OEM parts manufacturing is divided into seven business units. Each unit manufactures a set of parts, typically differentiated by the parts' function, shape, and location in the engine. Each business unit is divided into several (3-5) manufacturing cells, each of which has a defined footprint with specific manufacturing equipment, geared toward the subset of parts that are allocated to it. Because of the highly diverse product mix, however, each cell must produce a substantial number of different parts, many of which follow similar, but not identical, processes. This results in some parts having very straightforward flow paths through the cell, and others having convoluted paths. Additionally, some parts must travel between cells, or even between business units for specific operations or sets of operations. Cell layout decisions are made at the business unit level, and are typically designed to create simple flow paths for the part or parts that represent the majority of the volume in the cell. However, equipment is very expensive to move, and as product mix changes and engines enter different life cycle phases, cell layouts become outdated and inefficient.

As a result of this hybrid design of the job shop and the focused factory, and because each part number often requires a different jig or work holder, work is typically run in small batches, with a cell setting up to run a particular part number for a fixed amount of time or number of pieces, and then setting up to run a different part number. Batch sizes are typically determined by the BU's shop floor controller. In some cells, this introduces inefficiency into the system, where setup times are long and process times are short. The trade-offs between this batch size and the utilization of equipment and its effects on capacity is a topic that has been well-studied. (Li & Meissner, 2011) However, because of the extensive amount of turning and milling operations performed on the parts at NBPC, and the degree of automation present in the equipment, process times for these types of operations tend to be far greater than their corresponding setup times, resulting in efficient machine usage.

Before a new part is allowed to be released to the shop floor, it must have a fully detailed process plan outlined in PW's data system, including specific data on every individual operation required to manufacture the part. Each operation must be assigned to a primary piece of equipment, designated by a six-digit commodity code (CC) with an optional alphanumeric suffix. These CCs are used to differentiate machines from one another based upon their capability. The six-digit number serves to identify the type of machine, while the suffix is usually used to call out subtle differences, such as manufacturer, or ancillary machine capability not explicated by the CC. Each operation is also assigned a specific cell location for the piece of equipment on which it is to be performed. This designation of CC, suffix, and location is not necessarily unique, and so individual machines are marked with a brass tag (BT) number for machine-to-machine differentiation, but this number is not specified in a part's operation profile. It is therefore assumed that machines with identical CC, suffix, and location information are considered interchangeable with respect to a particular operation. In addition to equipment assignment, each operation is also required to have an estimated total process time, batch setup time, and process lot size.

Lot size is different from batch size in that a lot is processed in a single operation, while a batch indicates a quantity of parts to run through a cell before changing it over to another part number, regardless of how many operations this requires. Batch size is typically consistent by part from operation to operation, while lot size can change drastically depending upon the type of operation. Process time estimates are split into machine automatic time, characterized by the time the machine is working on the part, manual time, characterized by the total time the operator is occupied on the particular operation (including lot setup, on- and off-machine gauging), and man-machine time, which is the time that the operator is working on the machine, but the machine cannot be running (such as on-machine gauging and lot setup). In a sense, man-machine time is time when both the operator and the machine are occupied, manual time is time when the operator is busy, and automatic time is when the machine is in use and the operator is free. Splitting process time into these three categories enables differentiation between machine utilization and manpower utilization, and determination of which, on a particular operation, is the rate determinant.

Before and after an operation is performed, the operator must scan a barcode on the part and on the machine, indicating the beginning and end of the operation, and each workstation contains all the standard procedures for every operation assigned to it. These procedures call out the preceding and proceeding operations to ensure sequence consistency. However, some operations also have alternate operations or alternate machines on which they can be run. It is therefore left to the discretion of the business unit staff which of these operations to utilize; though it is generally accepted that whichever operation is listed as the "primary" is to be used whenever possible.

If an operator notices a defect on a part before or after the operation is performed, it is pulled out of the cell into a separate area for a process called "dispositioning". This process assesses the degree of severity of the defect, and whether or not the material can be salvaged, must be reworked, or must be scrapped. Often, salvaging a particular part requires that the business obtain approval from the customer, or a modification to the customer's accepted standard, which can take a substantial amount of time to process, depending upon the magnitude of the change.

Every part, after all of its necessary manufacturing operations have been performed, goes through a final inspection operation before it is boxed and shipped to the customer (mod center or end user). Some final inspection operations are more detailed than others, based upon the specification on the part, or how established the manufacturing process is.

NBPC's OEM business is treated within PW as a cost center, meaning that the plant does not report revenues. Instead, parts are assigned a "delivered hours" value, a measure of the standard amount of value-added work put into the part at the plant. Delivered hours (H_{Dit}) is the metric that determines how much work a particular business unit (i) did in a given time period (t). Because of the clocking system, wherein employees must scan a barcode to indicate the beginning and end of each operation on a particular part, there is real-time data available for how long each part took. Periodically, this data is aggregated to reevaluate how much each part is "worth" in delivered hours. In order to assign a dollar value to a part, each business unit reports its cost structure, and is assessed an "unburdened" shop rate (R_{Uit}) by taking the total unit costs (C_{it}) and dividing it by total delivered hours within the same time period. Plant overhead costs (C_{OHt}) are then distributed by volume to each business unit, creating a "fully burdened" shop rate (R_{Bit}) in dollars per delivered hour. These relationships are explained mathematically below.

$$h_j = Delivered\ hours\ for\ part\ \mathrm{j}$$

$$Q_{jt} = Quantity\ of\ part\ \mathrm{j}\ delivered\ in\ time\ \mathrm{t}$$

$$H_{Dit} = \sum_{all\ j} h_j \times Q_{jt}$$

$$R_{Uit} = \frac{C_{it}}{H_{Dit}}$$

$$R_{Bit} = R_{Uit} + \frac{C_{OHt}}{\sum_i H_{Dit}}$$

This relationship is important because, when project valuation is assessed for capacity expansion, it is determined by the dollar value of the additional delivered hours enabled by the equipment purchase, evaluated at the fully burdened shop rate for the given business unit.

2.3.3 Development at NBPC

North Berwick designs many of the parts that it runs on its shop floor, through its on-site design engineering organization. These individuals work through the engine program-specific integrated product teams to design parts that fit the required envelope. Being collocated with the manufacturing engineers and shop floor personnel enables rapid feedback for incremental design changes to accommodate more efficient manufacturing.

Development work is run on the same equipment as production parts, meaning that development learning curves can affect production part deliveries if they are not properly anticipated. Business units doing development work also tend to dedicate to it their most skilled machinists, which can in turn affect the quality of production work in situations where a skilled operator is required for a particular production operation. Because of the large range of part types produced at NBPC, when a new engine program is in development, a lot of new process development is pushed onto NBPC at once.

2.3.4 Maintenance, Repair, and Overhaul Business at NBPC

NBPC also has on-site operations for PW's maintenance, repair, and overhaul business called Global Service Partners (GSP). Due to independent aftermarket quality standards, GSP work must be kept separate from OEM work. The GSP business operates quite differently from the OEM business. GSP customers will send an engine or module to PW for maintenance, repair, or overhaul. The modules are then disassembled, and assessed regarding the extent and nature of the repair required. Work is then sent to the GSP business unit that specializes in the required type of repair.

While in the OEM business, parts with the same part number are considered interchangeable, and can be shifted to fill different orders for the same part number, this is not the case for GSP. Traceability is required on the specific part, identified by serial number, rather than part number. Because of this, and because every repair is unique, GSP operations do not use the same data systems as OEM operations. Because of the differences in the nature of the manufacturing process, quality standards, data systems, and business model, GSP and OEM work, although collocated within NBPC, are kept in different areas of the shop floor.

However, OEM and GSP do require many of the same manufacturing processes, which means that they also require much of the same capital equipment. While most GSP cells can perform their operations without the use of OEM equipment, there are some processes, such as coating, that are shared between the two businesses, and there are some times when, due to GSP loading, parts must utilize OEM fabrication equipment in order to meet the schedule. When this happens, the time is allocated to the GSP business in a process called "charge out".

2.3.5 Implications for Capacity Planning

The fact that BUMs are held accountable up to the plant manager level can incentivize autonomous decision-making from BU to BU. This can lead to inefficient allocation of capital resources, as units act on their own behalf in order to maximize local benefit, rather than optimizing allocations on a plant-wide scale. For example, if a business unit requires additional 4-axis milling capacity based upon its approximation of demand forecasts' impact on the shop floor, the BU management may elect to purchase a new machine. However, because of the lack of communication between BUs, there may be sufficient available capacity on 4-axis mills elsewhere in the plant. It is often left to the BUs themselves to seek out these cross-utilization opportunities, and since this is neither easy to do, nor are the BUs incentivized to do it, it rarely happens. Exceptional cases include those where there is a shortage on a scarce asset, such as a very specialized piece of equipment, in which case management at both the plant and BU levels are typically aware of the status of each piece of equipment in the shop. Other common exceptions are cases where a shortage is so severe that it is already affecting deliveries, in which case management is forced to search for a short-term fix; in these times, buying a piece of equipment is typically not an option.

The coexistence in the shop of development work with production work has extreme implications for capacity planning as it relates to the engine development cycle. When new engines are in development, development work can significantly affect production capacity. However, as this development work progresses, constraints can quickly change or be eliminated. The uncertainty in the development process, and the progressive reduction of this uncertainty requires that forecasted capacity be continuously monitored as experience is gained in the process, because future demand may be heavily driven by expected production loads for the parts under development, and a small change in the expected process time, or the process itself can result in huge downstream shifts in expected capacity constraints.

Finally, the charge-out process requires that the GSP side of the business be taken into consideration, as this can amount to a substantial piece of the load on a particular machine, depending upon the machine type. While this expected load cannot be calculated with the same degree of rigor as the OEM load, it must be accounted for in some respect.

2.4 Capital Procurement and Strategic Capacitation

Pratt and Whitney's scope of operations requires a substantial investment from year to year in manufacturing equipment. Per United Technologies' 2011 Annual Report, Pratt and Whitney's capital expenditures for fiscal year 2011 were $290 million, or approximately 2.2% of net sales. (United Technologies Corporation, 2012) For mod centers such as NBPC, capital procurement is handled through CEPO. As mentioned before, CEPO's internal process is well-documented, but the parts of the process outside of CEPO's responsibility are not. Through communication with CEPO and participating parties in the capital process at NBPC, a process map was laid out for the full life cycle of a capital project, from inception to release to production. This process can be divided into four phases, as shown below in Figure 3. Estimates for the duration of each phase were determined by dividing the capital procurement process into 76 discrete steps, and conducting a survey of each of the primary stakeholders' estimates of the duration of each step for which he or she is responsible. Minimum and maximum estimates were determined by aggregating the minimum and maximum estimated durations for each process step, respectively.

	Budget Approval	Specification	Procurement	Installation
Minimum Estimate	16 days	20 days	152 days*	9 days
Cumulative Minimum	16 days	36 days	188 days	197 days
Maximum Estimate	83 days	93 days	637 days*	36 days
Cumulative Maximum	83 days	176 days	813 days	849 days

* Procurement phase variability is driven by variability in equipment build times, which can be anywhere from 84 to 420 days, depending on the type of equipment purchased.

Figure 3 – Module Center Capital Procurement Process

The budget approval phase involves the identification of needed equipment, budgetary quoting, and capital plan approval. After a project has been added to the capital plan, a formal capital justification is constructed by the business unit, while CEPO works with both the business unit and the vendor to develop a detailed specification, and request a detailed quote. When the specification phase is finished, a purchase order is issued and approved, and the piece of equipment is built by the vendor. At the end of the procurement phase, CEPO and business unit staff will travel to the vendor's manufacturing site and conduct a test of the machine before it is shipped to the mod center. After this is finished, the facility must execute its installation plan (developed in the "specification" phase) and make connections to the piece of equipment. When this is complete, and basic on-site testing is done, the piece of equipment can be released to production. This entire process can take anywhere from 6 to 28 months, depending upon the type of equipment and the priority of the project.

2.4.1 Capital Project Planning and Budget Approval

Budget approval is the phase during which a new capacity planning tool has the most value for NBPC, since this is the phase where all of the long-term planning is done. Typically, technical supervisors will, throughout the course of the year, develop a "wish list" of needed equipment. Ideas may come directly from the shop floor, from home-grown planning tools, or from their own intuition. When an idea is developed, budgetary quotes must then be obtained. A budgetary quote enables rough capital allocation and justification without the need for developing a detailed specification, a process that can take weeks. The budgetary quotes also enable business units to vet their wish lists, in order to remove projects which may not be economically justified, or may not be needed right away. When this is complete, the technical supervisors are responsible for putting together, for each project they wish to add to the capital plan, a "quad chart" business case. This quad chart describes, at a very high level, what equipment is to be purchased, what the project schedule is, how the money will be repaid, and where responsibility lies. These quad charts are then presented to plant management prior to the annual capital budget meeting. Plant management aggregates all of the business units' projects into a single list, and prioritizes them. This list is then partitioned by the year the capital is required (business units are to plan capital needs three years in advance). In the late summer or early fall, the annual capital planning meeting involves CEPO and representatives from each of the module centers. Mod center representatives bring their respective capital lists for the upcoming year, and the previously allocated capital budget for all of the module centers is then divided between them. On each module center's list of capital projects, a "water line" is set at a certain priority level, and only projects "above the water line" are added to the capital plan. Mod center representatives then return and inform their business units which projects have been approved.

This begins the specification phase of the capital projects, which starts with a formal project valuation and capital justification. In some sense, due to the fact that the projects have already been approved by the time a formal justification is completed, this justification becomes more of a formality than an actual decision tool. Projects that do not appear justified are often cited as having "strategic" reasons for execution. In practice, few approved projects are reconsidered because of an unfavorable valuation model.

2.4.2 Implications for Capacity Planning

Because of the long cycle time of the capital procurement process, equipment purchases on the capital budget for a given year are often not usable for production until more than a year later. For this reason, development of a three-year capital plan requires a demand forecast of four to five years in order to anticipate future capacity needs. However, since an itemized capital plan is only approved one year in advance, very little rigor is put into mathematically assessing this long-term need. Long-term capital planning is often based on intuition or, in some cases, a transformative vision of a particular area. For example, a technical supervisor may have immediate needs for replacement of equipment that is end of life and for current capacity constraints, and this may drive the current year's capital plan. In the future, however, he or she may know that work currently under development will scale to production volumes, and may base a capital plan on equipment dedicated to this work, rather than assessing the needs of the unit as a whole based on aggregate expected demand.

2.5 Capacity Planning Process

Capacity planning at NBPC is carried out differently by different business units, with varying results. Each business unit's basic properties are described below in Table 2. (Note that because business unit 851 had not yet matured, it was not included in most of this research.) A unit is defined as having "significant new work" if it is the defined source of more than 5 part numbers on any of the development engines. A unit is defined as being "high-piecerate" if its mean batch size is greater than ten units. Planned capital investment is normalized by the maximum value of investment, shown as a percentage of this value.

Business Unit	Primary Products	# Part #s	% of Total Delivered Hours, June '10 to June '11	Significant New Work	High-Piecerate	2012-2014 Planned Investment for Capacity Expansion (Normalized)
810	HPC Stators	418	20%	X		27%
820	Outer Air Seals	106	20%	X	X	100%
840	Seals and Shrouds	360	17%	X		22%
850	Blades and Vanes	182	15%		X	0%
851	Module Assembly	--	2%	X		17%
860	LPC Stators and Cases	245	9%	X		45%
870	Mechanical Components	283	14%	X		81%

Table 2 – NBPC Business Unit Properties

A survey of the business units' practices in capacity planning reveals that the degree of sophistication in capacity planning methodologies varies from completely qualitative and very infrequent to regular quantitative analysis. All business units used tools based in offline spreadsheets to conduct capacity analysis. When asked why the preexisting SAP module was not used, technical supervisors cited the aforementioned operational shortcomings of the ERP-based system. Namely, the fact that the schedule is unreliable due to the low visibility on spares orders and the variability of military demand, as well as the lack of good operation time study data. Oddly, although many business units had their own time studies for parts, these numbers were never updated in the existing system. The added complexities of charge out, development and early-stage production work, and maintenance downtime are also absent from the model. Finally, as previously mentioned, the lack of a user-friendly interface and output rendered the tool inadequate for the needs of shop floor management.

In short, in spite of the relatively well-executed ERP and management system implementation, processes such as capital investment for capacity planning that span multiple departments have remained a challenge. Because the process spans multiple systems, and must take into account risks that are not characterized in the existing systems, individuals responsible for the planning process must choose between single-system built in tools that require trivializing approximations of externalities, or must engage in the onerous task of performing independent model development. Factoring in variability in unit structures and needs, and the resultant system becomes fairly unreliable.

Given the current state, there is significant room for improvement in terms of providing a means for management to prioritize capital projects from business unit to business unit, providing a standard language for plant staff to use when discussing capacity, and creating a model that provides a more accurate representation of what is happening on the shop floor. In light of these potential benefits, a capacity planning tool was developed for NBPC, following the framework proposed by Ozturk et. al. After localized validation, the tool was then used as an aid in identifying alternatives to BU 840's capital plan for capacity expansion.

3 Chapter 3 – Phase I: Selection of Modeling Solution

Given that the existing SAP module for capacity planning has been deemed insufficient to meet the capacity planning needs of the business units, off-the-shelf software solutions were explored. There are many different options for capacity planning software on the market. These software packages tend to fall into one of the following categories:

- Stand-alone ERP systems
- ERP add-ons
- Shop floor schedulers
- Custom solutions (not off-the-shelf)

The first category would not be a viable option, because SAP is deployed company-wide, and ERP changeover costs can be extremely expensive. Moreover, such a system would require several years to implement, and there is no guarantee of a better result than the existing ERP system.

The second type of software, an ERP add-in that could essentially "bolt on" to the existing system, would unfortunately suffer from many of the same problems of the existing ERP module. Because of the inflexibility of the ERP system post-implementation, planners must keep "off-ERP" databases of the information they require in order to aid in planning decisions. Because the capital planning cycle only happens annually, however, it is not viewed as a priority for updating the ERP system. For these reasons, the second category of software tools is not a feasible option.

Most stand-alone capacity planning tools are actually shop-floor schedulers with capacity planning capability as an ancillary feature. These types of software solutions can be extremely useful in certain circumstances. NBPC has endeavored to implement one of these tools in the past to aid in scheduling. However, it was found that the level of granularity at which the scheduler operated rendered any plan it created almost instantaneously irrelevant because of variability on the shop floor, in line with the assessment of Tenhiälä et. al. The amount of data required to characterize the system and keep the plan relevant, because of the complexity of the shop, required so many additional resources that the package was ultimately abandoned. For this reason, shop floor schedulers have been avoided at NBPC.

Finally, the option of a custom solution was ignored, as hiring a third party to develop the solution would only extend the schedule and increase the cost over internal development. Having explored the off-the-shelf software space, it was determined that a custom, internally-developed solution would provide the proper functionality in the proper timeframe, with sufficient flexibility to aid in decision-making, but without requiring excessive data maintenance.

Ultimately, the tool would need to migrate to a stand-alone application with ERP connectivity. However, in order to satisfy the immediate need of an updated capacity plan, a prototype was to be developed for the purposes of model validation and initial analysis. This prototype was developed as a spreadsheet-based user interface driven by custom-coded modules. Because of the flexibility that spreadsheet-based applications provide, and since all previous capacity planning tools were spreadsheet-based, this was viewed as the medium with the fewest barriers to adoption. As a stand-in for database connectivity in the initial prototype, it was decided that data in the form of raw reports from the myriad data systems would drive the tool, enabling easy updating of the prototype without having to repurpose or reformat database reports.

In situations where immediate impact is a concern, but a custom software application must be developed, prototyping can be useful, as it provides a proof-of-concept for functionality and some aspects of user experience, without the long development time of a standalone application. ERP connectivity, per a conversation with the IT department, would require at least 6 to 12 months to develop, with a correspondingly large budget. The use of a prototype avoids these extended timeframes whilst still providing the functionality and giving the results required in the near term.

4 Chapter 4 – Phase II: User Requirements Specification

Specification of user requirements was approached on two fronts. First, it was determined what the user experience had to be in order for the tool to provide the necessary value to the target users. Next, it was determined what critical parameters affect machine tool capacity, and how those should be taken into account.

4.1 User Experience

The output of the desired capacity planning tool would be consumed by both BU management and plant management. However, the tool was to act as a decision aid in a different capacity for each. In speaking with plant management, several critical issues came to light. Per a conversation with the plant manager, management's priority is being able to ensure that the plant can keep the commitments it makes to the product teams. Because NBPC has to bid against external competition for a lot of the work, to over-commit and under-deliver could hurt the reputation of the plant, and cause more work to be sent elsewhere. However, management does not have the time to sift through data tables and spreadsheets to look for the results they require. In this way, management needs an at-a-glance way to determine whether or not the plant can meet its commitments. Additionally, since commitments are made to product teams, management must understand easily which products are impacted when a constraint is identified. Finally, management does not have time to use a model that must run for hours, so keeping the calculation time short is critical if management is to adopt the tool.

Through this conversation, several critical aspects of the user experience from management's perspective became clear:

- Simplicity of report outputs (at-a-glance results)
- Ease of use
- Quick runtimes (< 15 minutes)

BU management, on the other hand, has different priorities. While it is viewed as a requirement in all cases that the software package be easy to use, reports be simple, and runtimes be short, decisions required at the business unit level require more at-a-glance information to be available. The critical decisions of the two different user groups are outlined below in Table 3.

User Group	Decision	Information Required
Plant Management	What work to commit to product teams	Plant-wide machine constraints
	How to prioritize projects	Impact of projects on capacity
	How to delegate work by business unit	Impact of schedule change on machine capacity
BU Management	Which machines to purchase and when to make the purchases	BU-wide machine constraints, time-series
	Which parts to outsource and which to keep	Capacity impact by part #
	Whether to send parts to other business units	Capacity machine-by-machine, considering location

Table 3 – Critical Management Machine Capacity Decisions

In light of the requirements assessed herein, it further suggests that a model approach in line with Tenhiälä's "capacity requirements planning" is appropriate. This requires a certain degree of aggregation, optimally suggested at the monthly level, as this is the typical time unit used for long-term schedule planning. In order for the model to help in building the three-year capital plan, it must look forward 60 months (as mentioned in Section 2.4.2).

4.2 Parameter Requirements

In addition to the user interaction with the model, it is also necessary to understand the technical requirements of the model. To discover this, various stakeholders across the plant were asked to give their opinions as to which factors affect machine tool capacity, beyond the standard considerations of process times and forecasted schedule. These recommendations, coupled with the author's assessment of potential additional factors helped to establish the framework for how the model calculates capacity.

Per user experience requirements of at-a-glance capacity assessment, it was determined through mock-ups that an easy way to express whether or not capacity is sufficient is to split capacity into two factors, machine availability, and machine load. This way, in a time-series chart, one can easily look and identify where load exceeds availability, indicating a capacity constraint.

Baseline determinants and modifying parameters for both machine load and machine availability are shown below in Table 4.

	Baseline Determinants	Modifying Parameters
Machine Availability	Quantity of machines Staffing policy	Maintenance downtime Charge out Utilization efficiency Operator time away from the machine Absenteeism
Machine Loading	Operation times Schedule forecast	Development work Part quality "Drop-in" orders short-to-lead time Reverse vendor assist/insourcing Alternate operations

Table 4 – Factors Affecting Capacity

Each of these factors has a potential impact on capacity, depending upon the business unit, part type, or process type. Each of the modifying parameters was thought to have a potentially significant impact on machine capacity.

Maintenance Downtime

Some maintenance downtime is scheduled, and therefore controllable, while some is not. Machines require periodic maintenance, and must be cleaned and calibrated in order to keep them running to the specifications required of them. Typically, a particular machine will have a recommended preventive maintenance procedure involving a periodic inspection of the machine. This requires a shut-down, which causes some availability on that machine to be lost. However, according to the maintenance department, inspections and preventive maintenance on any major piece of equipment typically takes the machine down for one or two days, and normally is scheduled at most a couple of times per year. In addition to scheduled downtime caused by periodic preventive maintenance, there are also situations where there can be scheduled but unanticipated downtime. These can occur if, for example, a machine filtration system fails, but a backup is available. In this case, the machine can continue to run, but there must be a scheduled time period wherein the filtration system can be fixed, requiring that the machine be down. While preventive maintenance has a relatively small impact on total machine availability, this second type of scheduled downtime can become a fairly substantial hindrance.

Unscheduled downtime is a well-known issue with respect to machine availability. The issue of how often and for how long a machine is down is one that has been studied extensively in both the professional and academic arenas. The difficult part of planning for unscheduled machine downtime is that, when a machine does go down, there can be extremely variable amounts of time until that machine is available again. Some "machine down" situations can be fixed on the spot, and some can take a machine out of service for weeks or even months at a time, depending upon the nature of the breakdown. Both scheduled and unscheduled types of downtime have a potentially significant impact on capacity, and therefore both should be considered a modifier, taking away machine availability.

Charge Out

Charge out, as mentioned in Chapter 2, is time when OEM equipment is being used for aftermarket operations. Because of the vastly different nature of the aftermarket business, an appropriate tool for capacity analysis of the GSP shop would require a different type of model. Instead of looking at a schedule forecast by component, aftermarket capacity would require looking at the relative frequency of the various types of problems diagnosed when an engine or module comes in for overhaul or repair, combined with usage models of the installed base. For this reason, charge out is viewed as an externality to the model, and therefore takes away from availability for load that is considered internal to the model.

It is believed that the occurrence of OEM equipment being used for aftermarket work tends to be fairly consistent within each BU and over time, but varies from BU to BU. Because charge out can hurt a BU's metrics, as it can cause the unit to lose potential delivered hours, the BUMs are keenly aware of how often, for how long, and when their equipment is being used for GSP work.

Utilization Efficiency

This modifier acts as a catch-all for systematic process efficiency issues, of which there can be many, depending upon the type of work done. Common types of systematic process inefficiency include those caused by cell configuration for a particular part, those caused by imbalance in work centers wherein one operator is running two machines, and many other specific situations where machines may be left idle owing to the nature of a process.

As an example of the first type, in BU 850, which is a high-volume unit, there exists a cell that runs four creep-feed grinders. The cell runs large batches, because the required part quantities are so high. This enables the set up of a fairly balanced cell with good single-piece flow, keeping in-process inventory low, while simultaneously keeping throughput high. However, this typically entails having each machine in the cell set up to run a particular operation, and passing parts from machine-to-machine. Parts in the cell may require two, three, or four creep-feed grinding operations. Basic cell configurations are shown in Figure 4. Configuration A represents a setup with four creep-feed grinding operations, B has three, and C has two. For the operations with two and four operations, all four machines can be utilized, either in series or in parallel. However, for Configuration B, because of line balance, the fourth machine cannot be utilized in parallel with any of the others, leaving it idle. Therefore, depending on the relative demand of the parts in Configuration B versus those in Configurations A and C, equipment utilization will vary.

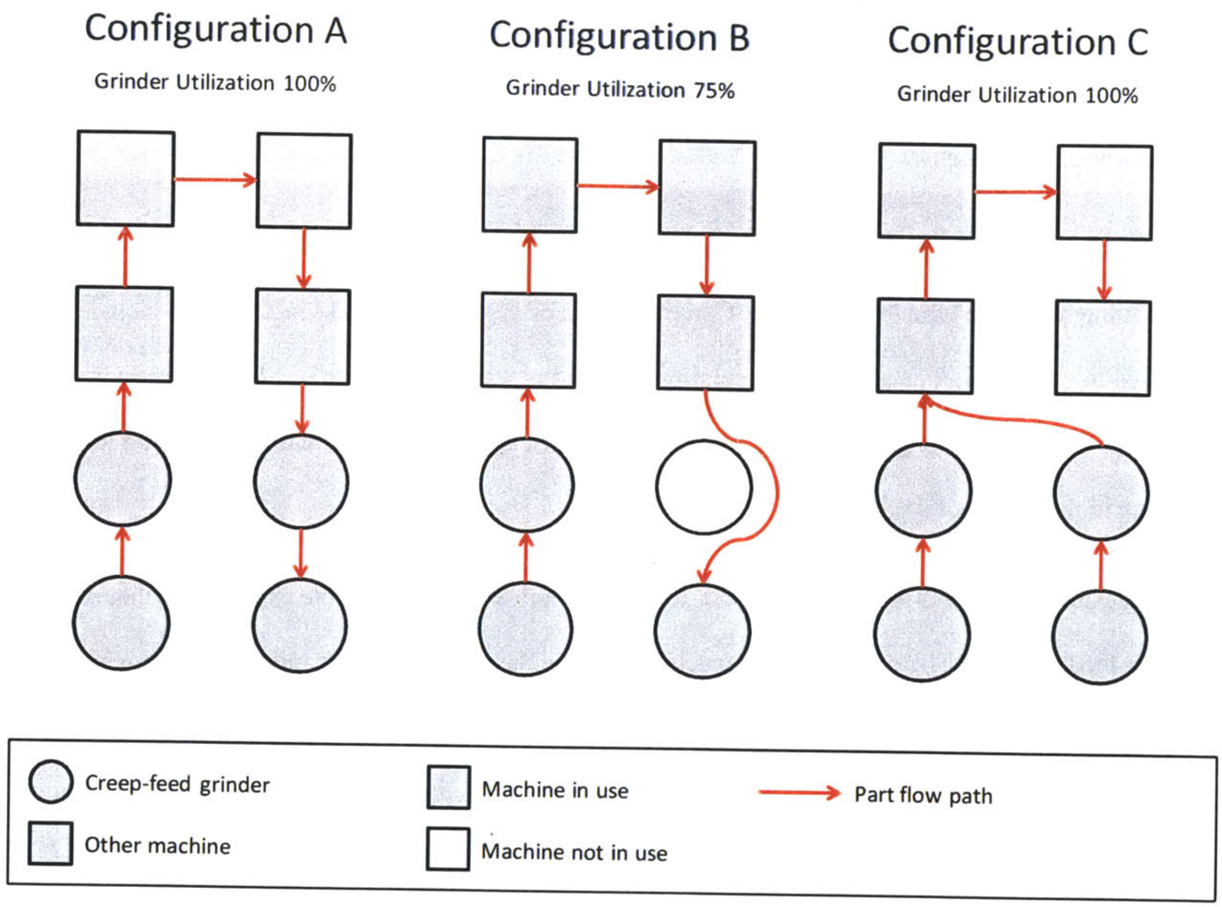

Figure 4 – Basic Cell Layouts, BU 850 Creep-Feed Grinders

The second common type of systematic process inefficiency is created by an imbalance in the work allocated to a particular operator, responsible for more than one machine. A single operator's station is called a work center. Work centers may include one or more than one piece of equipment. When an operator is responsible for more than one machine, as is often the case, there must be enough automatic time on each machine for the operator to do all the manual tasks required on the other machines. If this is not the case, machines may regularly sit idle, waiting for operators to certify dimensions or change the part. For example, suppose an operator is assigned two machines, Machine A and Machine B. Assume that neither machine is starved for work. Machine A runs Operation 25 on Part X, and Machine B runs Operation 240 on Part Y, with the following properties:

Part #	X	Y
Operation #	25	240
Machine	A	B
Manual Time	20	25
Automatic Time	60	15

Table 5 – Hypothetical Work Center Configuration

Assume, for simplicity's sake, that all manual time is done at the beginning of the operation, and all automatic time is done afterwards (though this is rarely the case). Looking only at cycle time, it appears that every 80 minutes, the operator should be able to run one part X and two part Ys. However, if the operator starts a part Y, and as soon as the automatic time has started, begins a part X, by the time part X is set up, the first part Y will have been idle for 5 minutes. This is a systemic issue which contributes to a certain amount of lost capacity on machine B. Additionally, a similar effect occurs on machine A, because while the two sets of operations appear balanced, because of the 5-minute idle time, the system actually operates in an unbalanced fashion. This type of dissonance between the cycle of operation on one machine and that on the other machine can have significant effects on capacity.

In addition to these two types of design-related operational inefficiency, there are many other specific cases that apply based upon the particular configuration of the operation, machine, part, or work center.

Operator Time Away from the Machine

There are times during the day when, although a work center is staffed, the operator is required to perform activities other than running the machines. These can include staff meetings, plant-wide presentations, or any other non-production-related activities. At NBPC, this time is referred to as direct charging indirect time, or DCI. This name refers to the fact that shop floor operators are referred to as "direct" employees, and time spent producing parts is referred to as "direct" time, while time spent on non-production activities is referred to as "indirect" time. Because the machine is not staffed during these periods of time, this modifier removes availability.

Absenteeism

In spite of staffing policies, there are inevitably times when people simply do not show up to work. While on a plant level this is relatively small, losing a shift becomes much more significant at the individual machine level.

Development Work

When a part is in development, theory suggests, and research supports that it often takes significantly longer to manufacture than when it is mature. (Argote & Epple, 1990) NBPC is no exception to this research. This symptom can be caused by a number of problems uniquely caused by new process development.

When a new part's operation profile and production plan are released to the shop floor for the first time, there are many unknowns. While almost all critical operations are done on computer numeric controlled (CNC) machines run by programs, referred to anachronously as "tapes", these tapes often are released to the shop floor with errors. Because of this fact, prior to running the first part of a particular tape, the operator is required to do a "dry run" of the program. This normally entails running the full tape without actually loading a work piece, and watching to make sure there are no odd or incorrect moves. Additionally, tool changes and other interruptions in the program must be executed properly. This process occurs at the same speed as the actual operation, meaning that the first part will take at least twice as long as the standard operation time. The dry-run process also normally uncovers errors, which must then be fixed before a part can actually be run. Additionally, operators must ensure that the proper tools and gauges are available. Any time a nonstandard cutting tool is found to be out-of-stock, or a set of tooling has not been received, additional machine time is lost.

In addition to issues related to the setup of the operation, the process of actually manufacturing the first few parts is also prone to significant variability, which can lead to large amounts of rework and scrap. For example, based upon the programmed settings for a particular mill cut, the dry run may appear to run flawlessly. However, when a part is inserted and material is physically being removed, unexpected tool deflection or tool chatter could cause the resultant part to fall outside of the specified tolerances. When this happens, not only must a part be reworked or scrapped, but the process must also be redesigned in order to fix the problem.

Finally, there are often design changes that occur during development. Depending upon how significant the design change is, this can set the process back to the beginning. In speaking with people on the shop floor, impressions suggest that development parts can take as much as 10 times longer than their expected standard times. For operations with long process times, this can have a very significant impact on capacity, sometimes tying up machines for weeks or months at a time.

Part Quality

The quality of parts affects capacity because, quite simply, if a bad part is scrapped, another must be produced to fill the demand for it. Depending upon how far along its manufacturing process the bad part was when it was caught, this means that certain operations will need to be done twice in order to produce only one good part, effectively loading machines more than should be expected based upon the baseline schedule. For some parts, part quality has such a small impact it can almost be discounted. However, for some parts, especially low-volume parts that are complex or new, quality can have a significant impact on the machines on which the part is manufactured.

"Drop-In" Orders Short to Lead Time

As mentioned in Chapter 2, the commercial spares market operates in an a-la-carte fashion, quoting standard lead times for all salable spares. However, there are times when a customer may have an emergency situation, and requires parts more quickly than the quoted lead time. This type of unexpected demand can have a significant impact on machine capacity. Because this load is not visible on the schedule, standard planning methods will not account for it, and it becomes a machine load modifier.

Reverse Vendor Assist/Insourcing

Like "drop-in" orders, these types of situations lead to demand in excess of what is expected based upon forecasts and existing schedules. However, unlike "drop-in" orders, these situations are more difficult to predict, and can have much more substantial impacts to shop load. For example, at one point a fairly substantial parts supplier went out of business. Many of this supplier's part requirements were shifted to NBPC, causing a significant swell in demand on the plant. While this type of occurrence is difficult to plan for, it nonetheless has a notable impact on machine load and machine capacity.

Alternate Operations

While every part released to the shop floor has a production plan, complete with each operation assigned to a specific commodity code in a specific location, there are often alternate operations developed on similar, but not identical machinery, which can be utilized in the event that the primary machine is not available. While alternate operations are typically similar in process time to their primary counterparts, utilization of an alternate operation is something which is not typically taken into account in a baseline machine load calculation, but when utilized, will take away from the alternate machine's availability.

5 Chapter 5 – Phase III: Design of Data Structure and Information Flow

Having established the desired user experience and parameter dependencies, the next task is to find sources for data to determine the parameter values. Historical data is used as a method of first-order control for decision making with regard to factors dependent upon future system performance. In other words, parameters determined by historical data are assumed to be stable over the planning period. The degree of data aggregation, however, should be rationalized by first looking at the validity of the stability assumption over the suggested period. Unfortunately, due to a lack of available data enabling testing of the 60-month stability of each parameter, the model is built on this assumption based upon the qualitative assessment of the plant's technical staff, to be validated as data is gathered after implementation.

Data sources must be identified first to validate the model. Then, in order to ensure the model continues to be used, data maintenance must be automated. The degree of automation, however, must also be rationalized. Because there is substantial switching cost in both dollars and labor in augmenting or changing existing data systems, this should only be done where it is necessary. Each parameter is therefore first assessed for its historical average impact on monthly capacity, and then for the variability thereof. Parameters with high variability and high average impact are determined to be the highest-priority cases for data automation. Those with high variability and low average impact are determined to be the second priority for data automation, and finally those with low variability are the lowest priority, regardless of average impact. Where historical data to assess this was available, quantitative analysis was performed. Where it was not, the opinions of the technical staff were used instead. The results are shown below in Table 6.

	Parameter	Historical Data Available?	Scale of Average Impact	Scale of Impact Variability
Machine Availability	Quantity of machines	No	High	High
	Staffing policy*	No	High	--
	Maintenance downtime	Yes	Low	High
	Charge out	Yes	High	High
	Utilization efficiency	No	High	High
	Operator time away from the machine	Yes	High	Low
	Absenteeism	Yes	Low	Low
Machine Loading	Operation times	No	High	High
	Schedule forecast	Yes	High	High
	Development work	Yes	High	High
	Part quality	Yes	High	High
	"Drop-in" orders short to lead time**	No	Low	High
	Reverse vendor assist/insourcing**	No	Low	High
	Alternate operations*	Yes	Low	High

* Staffing policy and alternate ops are viewed as controllable variables and therefore measurement is unnecessary

** Drop-in orders was grouped into forecast variability and reverse vendor assist into charge out, as that is where they are manifest

Table 6 – Summary of Parameter Impact Analysis

5.1 Parameter Impact Analysis

All parameter impact analyses using data are based on the assumption that individual data points used for empirical distributions are independent and identically distributed. For each parameter for which adequate data is available, an analysis is conducted to characterize the type and degree of variability. For example, if it is suggested that a given parameter has significant variability from business unit to business unit, data are separated by business unit. Then, normal quantile plots are constructed to determine whether each particular business unit's data set is normally distributed. If the data appear to be normal (with an approximately $y=x$ straight-line normal quantile plot), an appropriate statistical test is conducted to determine the probability that each data set comes from the same distribution as each other data set. This can be done by conducting pairwise t-tests comparing each business unit to each other business unit, or by using standard analysis of variance to conduct an f-test, which determines the likelihood that the explanatory variable's (in this case, "business unit") impact on unexplained variability is by chance. A low f-test probability implies that it is likely that the business units' respective distributions of values are significantly different from one another. Pairwise t-tests only provide valuable information, however, if the analysis of variance f-test yields a significant statistic.

Quantity of Machines

Plant records do not include the historical evolution of the plant equipment list, meaning that variability over time cannot be assessed, though it is suggested that this is low, as machines are infrequently acquired, retired or moved. However, it is widely known that different areas have different amounts of the various machine types based upon the different types of manufacturing processes contained therein, and the types of parts typically made. Additionally, the plant as a whole tends to have substantially more of certain types of machines than others. For example, over all manufacturing cells, the mean number of vertical lathes is approximately 1.54 per cell, with cell-to-cell standard deviation of 2.5 machines. However, for 5-axis mills, there are 0.30 machines per cell, with standard deviation of 0.78 machines. It stands to reason that, based upon different cells' manufacturing needs, they would require different equipment. Because the number of machines available in a particular area defines baseline availability for that asset type, and because the variability is so substantial from machine-to-machine and area-to-area, machine quantities and locations are a high priority for data automation.

Staffing Policy

How BUMs decide to staff their machines and cells is considered, for the purposes of understanding machine capacity, to be a controllable variable. Logically, it would be unwise to deduct availability from a machine simply because it has not been staffed 100% of the time, based upon historical data. This could lead to misleading conclusions regarding recommended equipment purchases or cell capability. For this reason, staffing policy is left to the discretion of the user. This way, if a BUM elects to run his or her unit for 3 shifts, 7 days per week, the tool will give him or her a good idea of what the unit's machine capability is given this strategy, and likewise for a strategy of 3 shifts, 5 days per week. Historically, it has been the practice of the plant to keep the operation staffed for 3 shifts, 5 days per week. Keeping this variable as a lever controllable by the user enables the tool to reflect any future policy changes, or aid in assessing the capacity impact of potential changes.

Maintenance Downtime

The maintenance organization keeps very detailed records for unexpected machine downtime. The reason for this is that, when a machine goes down, the operator initiates a ticket in the ERP system indicating the machine is down, and the maintenance technician then closes the ticket when it is running again. The duration that the ticket is open is considered a good approximation of the amount of time that the machine was down. However, when downtime is planned, whether due to a scheduled preventive maintenance, or due to a scheduled replacement, there is no data available. Total monthly downtime data for each of the 32 prior months was analyzed by machine brass tag number. Mean monthly downtime over the entire data set ($n = 23{,}200$) is approximately 1%, which is fairly small. However, variability within the data set is nearly 4%, which is substantial. Additionally, an extreme skewness estimate of 9.8 suggests a long upper tail meaning that, in some cases, downtime can amount to a very significant percentage of monthly machine availability. A kurtosis estimate of 151 indicates a large amount of peakedness; specifically, there are a substantial amount of data points (17,524 or 75.5%) with a value of zero, indicating there was no accumulated machine downtime for that particular machine in that month. A cumulative distribution of this data is shown below in Figure 5.

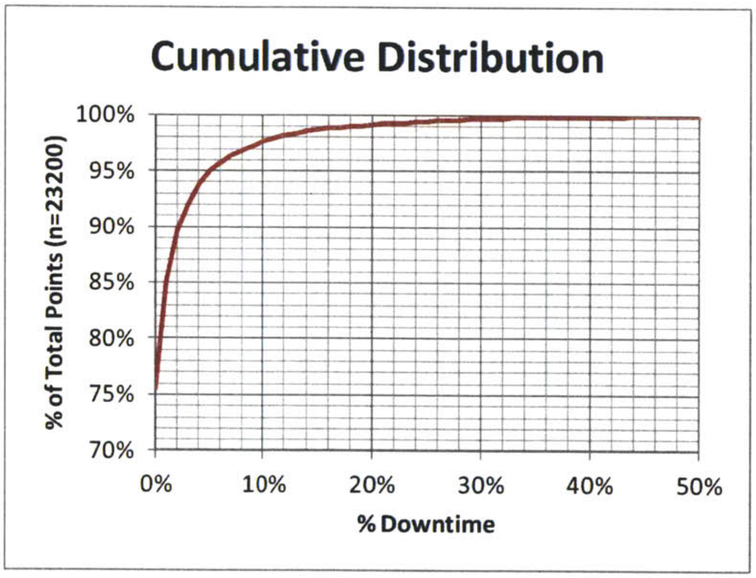

Figure 5 – Cumulative Distribution of Machine Downtime as % of Monthly Availability

Because the unplanned downtime data is available in automated form, embedding the full analysis of expected downtime, by machine, by area, with an allowance for month-to-month variability is possible to utilize without additional investment in resources. However, the system must be augmented with the planned downtime data in order for it to be comprehensive. As such, a project was initiated to change the system in this manner.

Charge Out

Charge out data is aggregated only at the business unit level by the finance department. This monthly aggregate data was analyzed for the 30 months prior to the study. For the entire data set (n = 180), mean charge out as a percent of total business unit delivered hours was 15%, with an estimate of standard deviation of 14%. Skewness and kurtosis estimates were 1.20 and 0.16, respectively. Splitting out by business unit, it also became apparent that variability was significant at this level. First, tests for normality were conducted by creating quantile plots against the normal distribution, shown below in Figure 6.

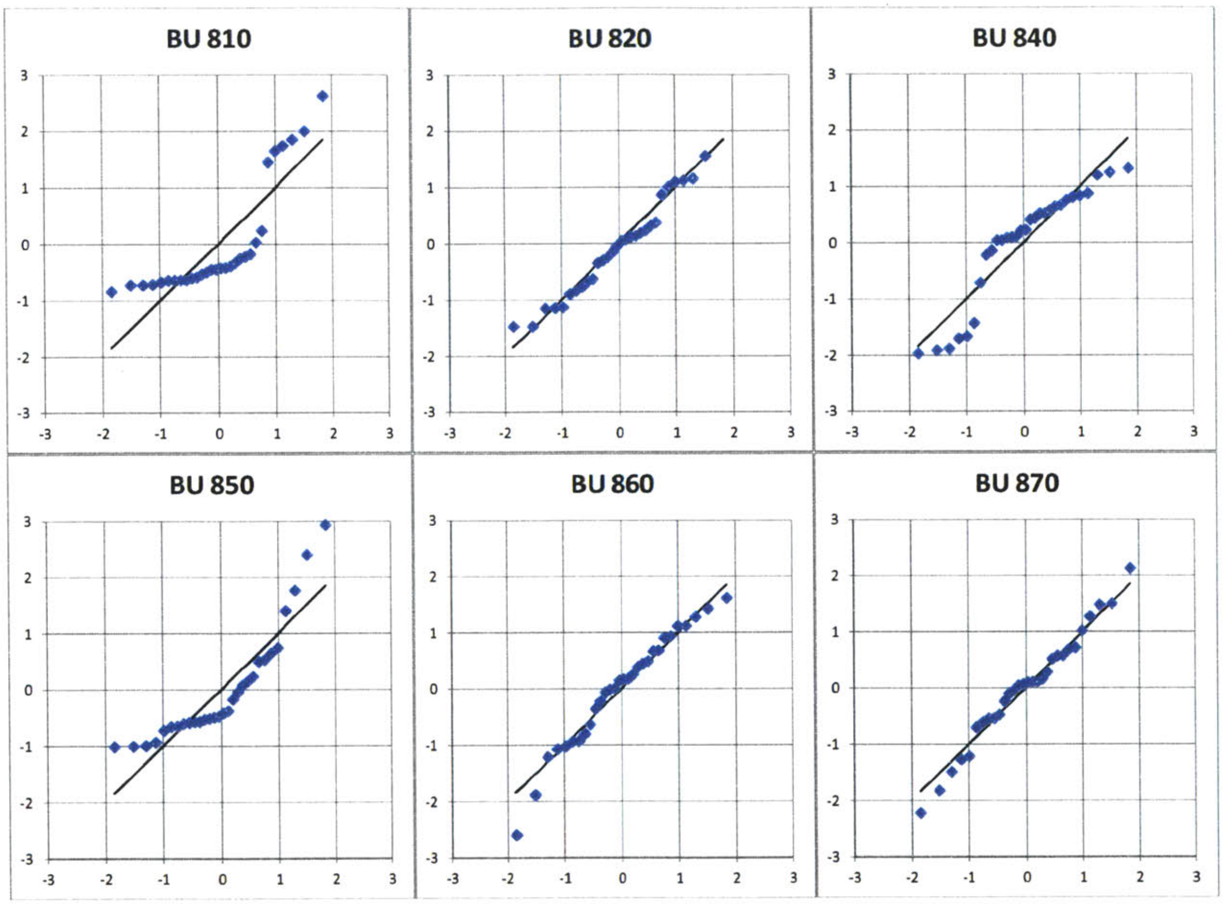

Figure 6 – Quantile Plots by Business Unit of Monthly Charge-Out vs. Normal Distribution

Units 820, 860 and 870 appear to be normal, while 810, 840 and 850 are not. A one-way analysis of variance was conducted, using the F-test to determine the probability that the variance between business units is different from that due to pure error. Results yielded an estimated F-statistic of 21.2, with a p-value of 1.3×10^{-16}, indicating statistical significance.

Heteroscedastic two-tailed t-tests were also conducted, comparing each business unit's data set to each other business unit. The only pair that failed to reject the null hypothesis that the units' distributions are alike (p-value > 0.05) was that comparing BU 850 to BU 870; however, these results are unreliable due to BU 850's apparently non-normal data. The results are shown below in Table 7.

BUs	810	820	840	850	860
820	3.21E-41				
840	5.27E-06	1.06E-33			
850	7.04E-03	2.91E-47	1.12E-10		
860	1.24E-18	1.64E-36	3.20E-08	2.51E-27	
870	4.72E-03	6.69E-38	1.90E-10	9.09E-01	7.22E-24

Table 7 – P-Values from Student's T-Test Comparing Monthly BU Charge-Out

No data was available down to the machine level. However, in speaking with operations staff, it became apparent that instances of charge out were fairly consistent from machine to machine, though it is recommended that this be validated when and if the data is available. Because this data is not available in an automated form, it is recommended that this be changed in order to accommodate accurate capacity calculation, available by area and over time. However, for the proof of concept, this remains a manually-controlled parameter.

Utilization Efficiency

No historical data is available regarding machine utilization efficiency, because no data of this sort has ever been collected. However, it is suggested that the impact may be substantial in some areas. It has also been suggested that the degree to which utilization impacts capacity is dependent upon the nature of the business unit, as well as the operations approach of each cell therein. Typically, designed-in utilization inefficiency is at the part level, so the temporal dependence of machine-level utilization will therefore be linked with the types and relative quantities of parts being run. If a high fraction of parts being run have substantial designed-in utilization inefficiency for a particular machine, this machine's capacity will be reduced more than if that mix changes in future demand periods. For the purposes of prototype development, this will remain a manually entered parameter. However, it is recommended that a study be initiated to understand the degree to which utilization impacts plant-wide machine capacity.

Operator Time Away from the Machine (DCI)

Like charge out, DCI is only available in aggregate form, by business unit. The same analysis was run for DCI as for charge out (30 months, over 6 business units; n = 180), with substantially different results. Mean impact, like charge out, was fairly significant at 17% of monthly hours charged. However, standard deviation was substantially lower, at only 6%. Skewness and kurtosis estimates were 0.49 and -0.11, respectively. Monthly data by business unit was plotted on normal quantile plots, shown below in Figure 7.

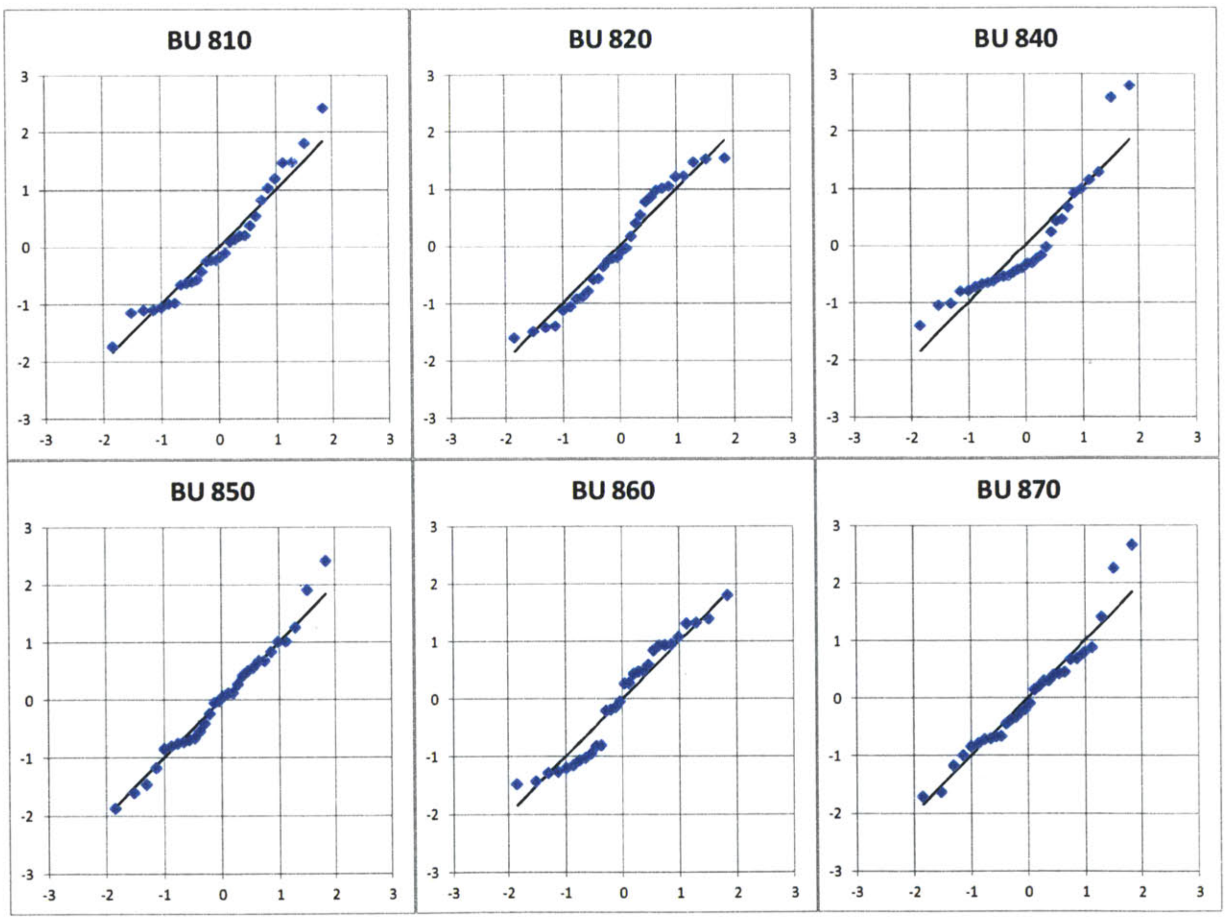

Figure 7 – Quantile Plots by Business Unit of Monthly DCI vs. Normal Distribution

Most of the data appears normal. One-way analysis of variance comparing the BUs to one another led to an F-statistic estimate of 0.422, with a corresponding p-value of 0.83. This analysis failed to reject the null hypothesis that the units' distributions are alike, suggesting little difference in DCI from business unit to business unit. Due to this, unit-to-unit t-tests were not performed. Coupled with the relative stability of the process over time, it is recommended that DCI be treated as a manually-controlled variable, as implementation of automated data systems to track it by area, time, or machine would likely have little effect on machine-to-machine capacity. However, it is recommended that a study be conducted to understand the impact of DCI variability at the machine level.

Absenteeism

Data on absenteeism is available. However, because absenteeism can theoretically be compensated for by unit management, it was determined that this factor would not be considered for the model. If, for example, the plant manager knows that a particular month has, on average, 3% absenteeism, this can simply be compensated for by overstaffing by 3.1%. Thus, it is not deemed relevant to the concept of understanding machine capacity for capital planning.

Operation Times

Operation times and schedule are for machine load what quantity and type of machine are to availability. Operation times are a critical component that make up the baseline load, and therefore must be data-linked in order to understand capacity. As was previously mentioned, the data system for operation times acts as a gateway to production release. Prior to the implementation of the current system, wherein part cost is based upon actual average labor charged by part number and the corresponding shop rate, part cost was assessed based on standard operation times. When this was the case, it was important that data quality in this system be high, and that it be maintained. However, when the new system was implemented, usage of the operation time database dropped, and was limited to production readiness reviews for new work and for capacity analysis. When this occurred, maintenance of the system dropped off, and the effort put into ensuring the accuracy of the data disappeared. Because of the criticality of this information for capacity analysis, however, it is recommended that maintenance of this system be restored.

With regard to operation time variability, very little is known. Because of the large number of operations in the building (> 35,000 individual operations), continuous monitoring of actual on-machine operation time cannot be economically executed with manual methods, and automatic measurement capability does not currently exist in the plant. Consequently, application of this capacity tool is best suited to machines on which operation time is primarily automatic (and subsequently more repeatable), allowing for a higher degree of consistency. For operations with lower ratios of automatic to manual time, it is recommended that an analysis be done as to the extent of effort required to assess this impact on machine capacity. For the purposes of the current prototype, however, it is assumed that variability of operation times does not significantly impact machine capacity.

Schedule Forecast

Schedule forecast, for all of the reasons previously identified in Chapter 2, is often inaccurate. Fortunately, data is available regarding previous forecasts, though not very much of it, and not in an automated form. While baseline schedule is one of the critically important building blocks of capacity analysis, forecast inaccuracy or variability is very difficult to characterize. Because of the myriad dynamics associated with the different engine and part markets, schedules can be unexpectedly augmented or reduced. An attempt was made to determine whether there is any bias inherent in the forecast, which cause over- or under-estimation of capacity requirements. Fourteen months worth of schedule forecasts were analyzed, and the difference between forecasted and realized demand for approximately 900 parts in five individual months were considered, looking at forecast horizons ranging from one to ten months. In other words, for a 10-month horizon, the predicted demand by part 10 months prior to June (last August) was compared to June's realized demand. The same was done for May, April, March, and February. The data set for each horizon therefore was comprised of approximately 4500 individual data points. Points were then grouped into three bins; overestimates, underestimates, and accurate estimates. Results are shown below in Figure 8.

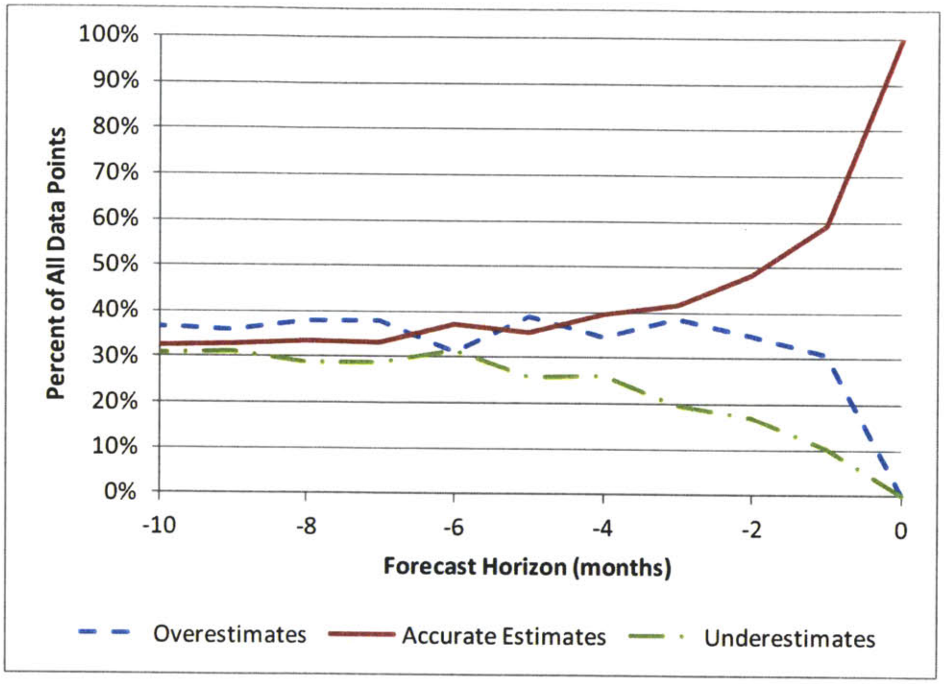

Figure 8 – Aggregate Realized Forecast Inaccuracies by Horizon

It is notable that for almost all horizons, the forecast appears to be skewed toward overestimates of demand; this is especially the case for shorter horizons. This may suggest that, while schedule augmentations (which manifest as underestimates) are continually and quickly added to the forecast, leading to a slowly declining curve as horizon decreases, schedule reductions (manifest as overestimates) tend not to be updated until they are just about to be realized. While this data gives no clue as to the scale of overestimation, it is recommended that further analysis be done into what the impact of forecast bias and variability may be on plant capacity. However, without data to accommodate this analysis, or a framework for quantifying the net effects, schedule is assumed to be accurate.

Development Work

Development work data is available in a very limited capacity, and not in a form that lends itself to data automation. However, in talking with operations and manufacturing engineering staff, it is held as a general rule that the first run of a new development part can take up to 10 times as long as its standard expected time. Maintaining traceability on machine time dedicated to development work can be difficult because the ERP system is not set up to track actual operation times when, for example, a part is left on an idle machine because an error was found in the program, or a possible defect is identified. Typically, when estimating the expected time for development parts, a learning curve is used. The traditional Crawford experience model follows the following mathematical relationship between the amount of pieces produced to date (X) and the direct labor content of the last unit (Y):

$$a = Labor\ content\ of\ the\ first\ unit$$

$$b = \log_2 L$$

$$Y = aX^b$$

The value "L" is referred to as the learning rate. Subject to the constraint $50\% \leq L < 100\%$, the learning rate represents the amount of labor content remaining every time the total number of units produced has doubled. In other words, processes with steeper or "faster" learning curves have *lower* learning rates, and processes with shallower or "slower" learning curves have *higher* learning rates. (Argote & Epple, 1990)

Using this model, understanding the expected future impact of development work requires knowledge of not only development part schedules, not available in the ERP system, and expected process times, not available in the operation time database, but also (often empirically-derived) learning rates and first-unit inflation factors by area and machine, as well as cumulative number of parts produced. None of this data is available. However, conservative estimates can be made of the Crawford model parameters using historical development part aggregate process time data, and these estimates can be (conservatively) applied to each individual operation. Process time estimates are available, as they are determined for the purposes of providing quotes to the end customer for development engines, and part schedules are maintained off-ERP by the engine program office. Access to the latter enables tracking of cumulative part deliveries as well.

Because development work acts as such a unique entity and schedules, designs and processes can change so quickly when a part is under development, it is unreasonable to think that systems for maintaining this information can be automated. Development work also comes along infrequently, meaning that the amount of effort required to manually manage this information is substantially less than that required for production parts. For this reason, the implementation of development work is left as a manual task. However, because operation profiles must be developed in order to quote development parts, it is recommended that these profiles be maintained in the standard operation time database, enabling the use of a single source of operation time data, with little incremental effort.

Part Quality

Since part quality relates to scrap factors, there is very little data available in any usable form. While scrap factors are used to augment production schedules in order to meet expected demand, the maintenance of these scrap factors is managed by materials planning, often based upon tribal knowledge of the specific set of parts for which each materials planner is responsible. However, data is available with regard to the number of times a particular operation has been completed.

Using this data, analysis was performed assessing the ratio of the number of times each intermediate operation was performed over the prior 32 months to the quantity of that part that went through final inspection in the same timeframe. End effects from part lead time should be minimized, since lead times are small (on the order of 6-8 weeks) relative to the aggregation window of 32 months. As a rough filter, only operations tagged as primary machining operations were considered, and any values less than one, or with zero final inspection pieces were considered to be erroneous, and were omitted. The resultant data set of 8,014 points had a sample mean and standard deviation of 112% and 42%, respectively. Skewness and kurtosis estimates were extremely high at 37 and 2180, respectively, indicating a highly peaked distribution with a very long upper tail. The cumulative distribution is shown below in Figure 9. Scrap factor here is given as the multiple applied to a particular operation in order to yield 100% of target deliveries.

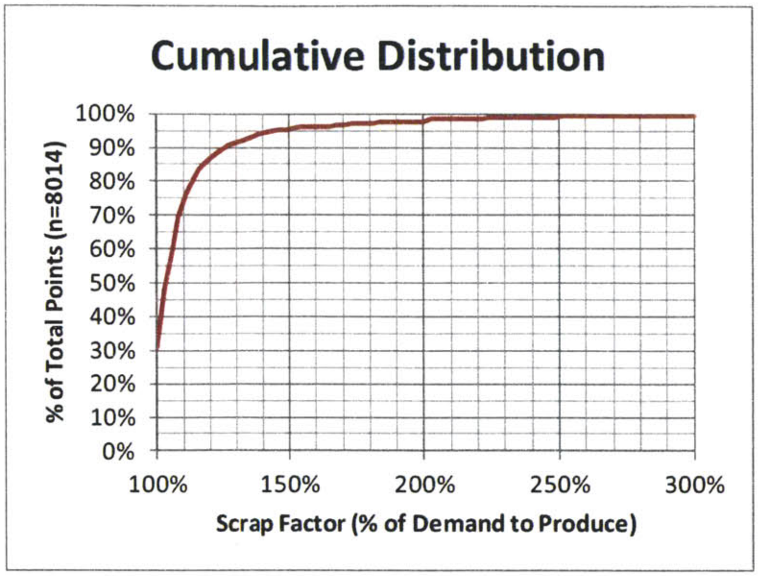

Figure 9 – Cumulative Distribution of Operation Scrap Factor as a % of Delivered Parts

Mean impact and variability of part quality are both very substantial, and because automated data is available, this data link can be automated without significant incremental effort.

Alternate Operations

Although data is available detailing historical usage of alternate operations, it is not deemed necessary as an automated parameter because the utilization of alternate operations is controllable. Because of this, it must be available as an option for scenario testing, as it does have systemic capacity impacts, but it will be assumed in every base case that only primary operations are utilized.

A summary of all parameters' priority level, automation availability and corrective actions for alignment between the two is shown in Table 8.

	Parameter	Automation Priority	Automated Data Available?	Corrective Action
Machine Availability	Quantity of machines	1	Yes	
	Staffing policy	N/A	N/A	
	Maintenance downtime	2	Yes*	Augment automated report to include scheduled downtime
	Charge out	1	No	Automate aggregate charge-out report
	Utilization efficiency	1	No	Conduct study on machine staffed efficiency to assess scale of impact
	Operator time away from the machine	3	No	
	Absenteeism	3	No	
Machine Loading	Operation times	1	Yes	
	Schedule forecast	1	Yes	
	Development work	1	No	Automate development part profiles and op time estimates
	Part quality	1	Yes	
	"Drop-in" orders short to lead time	2	N/A	
	Reverse vendor assist/insourcing	2	N/A	
	Alternate operations	N/A	N/A	

* Automated data is avaliable for unplanned downtime, but not for planned downtime

Table 8 – Parameter Data Automation Requirements

Maintenance downtime requires a corrective action to include scheduled or planned downtime into the same report for unplanned downtime. This requires a small change on the part of the maintenance staff, in that they must open and close their own requisitions for planned downtime. No major changes are required to the ERP system to implement this, simply administrative changes to procedures.

Charge out reports are aggregated from raw clocking data in the ERP system. The report needs to be automated and a study should be conducted as to the extent of the impact of charge out on individual machines. If this impact is found to be substantial, further corrective actions should be performed.

Utilization efficiency is something that is not fully understood as it relates to machine capacity. A study should be conducted to assess the impact of designed-in inefficiency by machine, work center and part number on machine capacity. Following this study, the need for further corrective action should be reassessed.

Development part profiles should be loaded into the current operation time database. This would enable automation to piggyback on the system already in place. Because development parts require a profile to be created for quoting, this will simply require that this profile be entered into the system, again needing very few incremental resources.

5.2 Data Sources and Flows

Data for the model must come from several disparate sources including ERP databases, other databases, and manual entry. Data sources for the prototype model are shown in Figure 10.

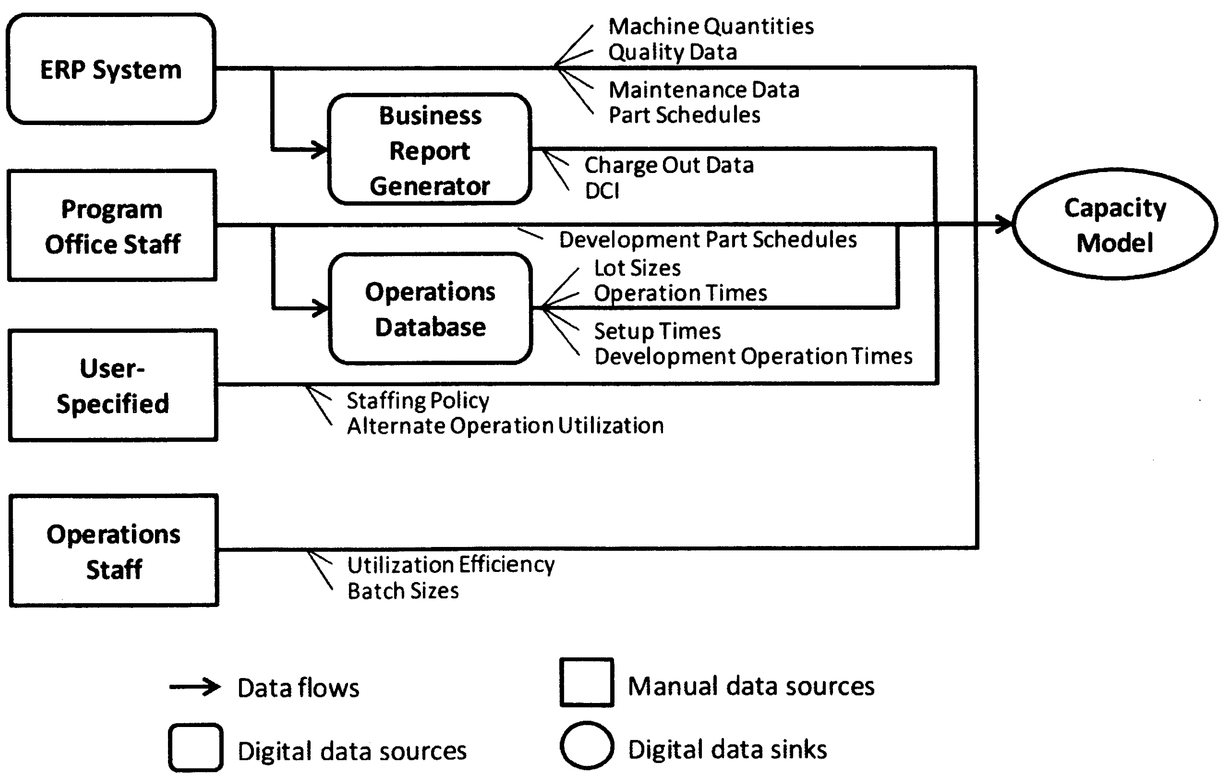

Figure 10 – Data Sources for Prototype Model

Data that is entered manually, but informed by particular digital sources is labeled as being fed from the source, although the connection is not automated. All digital data sources are in some way interconnected, though these connections are not shown where they are not relevant.

Having identified data sources for all parameters, and having initiated corrective actions to bring data sources in line with what is required for capacity planning, the project can move into phase IV for implementation and validation.

6 Chapter 6 – Phase IV: Implementation and Validation

With the user experience and functionality requirements specified and data sources for all parameters identified, the final step is to build the functionality that connects the required outputs to the specified inputs. Implementation is divided into two tasks: building a mathematical model, and designing a user interface, which are highly interconnected. The way data is to be reported will determine how calculations must be performed. While, in practice, this process is iterative, requiring multiple attempts of building a user interface, testing its functionality, and garnering feedback from the user base, the final prototype model will herein be explained starting with the user experience, which then informs the mathematical backbone.

Following implementation, validation must be performed to verify that the model actually yields accurate results. This process is also iterative, as validation reveals faults in analytical methods, as well as issues of data quality, which can be a large factor in an organization with numerous legacy data systems like NBPC.

6.1 User Interface Implementation

The design of the user interface is guided by revisiting the critical information that must be provided by the model to the user (see Table 3). To aid plant management in each of its required decisions, the model must provide plant-wide analysis of machine constraints, depending upon various user-defined scenarios (project configurations or schedule alterations). In order to aid BU management in its decisions, the interface must be capable both of BU-wide constraint identification (which is implicit in the aforementioned plant-wide analysis), as well as detailed analysis of particular area-specific machine constraints, specifically by part number. In addition to these functionality profiles, one more critical feature was identified through iteration: detailed analysis must include a way to see the rough contribution levels of each critical parameter. This aids both as a tool to inform management of the reason for a particular constraint while also serving as a method of error-checking. When issues of data fidelity come into play if an invalid constraint is identified, being able to quickly identify if a particular parameter is contributing disproportionately to the constraint aids in understanding where the issue may reside. This set of requirements suggests the creation of two separate reports: one that looks at plant-wide constraints, and one that looks in-depth at a particular constraint.

Plant-Wide Constraint "Hot Sheet" Report

For plant management and for BU management's analysis of its entire area, a report must be generated that shows every constraint within the plant, identifying the area in which the constraint exists. This is best conveyed as a table. While charts can show aggregated "capacity metrics" against one another, they cannot give the specificity required to fully understand what constraints exist, and where they exist. Capacity is understood on the shop floor, as well as in the management offices, in terms of percent utilization. For this table, percent utilization should therefore be the critical measure of capacity. To aid in management's decisions regarding the efficacy of various project or schedule scenarios in removing constraints, management should be able to set a threshold for a piece of equipment to be flagged as "constrained". A sample table is shown in Figure 11.

Rank	Work Department	Brass Tag	Description	Commodity Code	Max Loading %	Current Loading %
37	8317	789773	FLUORESCENT PENETRANT INSPECT	576200	113.64%	94.96%
37	8317	790659	PRODUCTION FPI CHECK	576200	113.64%	94.96%
38	8109	536143	MITSUI SEIKI MACH CTR 5A FANUC 11M	179705B	104.24%	41.83%
38	8109	536144	MITSUI SEIKI MACH CTR 5A FANUC 11M	179705B	104.24%	41.83%
38	8109	536637	MITSUI HR6A MCH CTR 4AX FANUC 11MF	179705B	104.24%	41.83%
38	8109	536638	MITSUI HR6A MCH CTR 4AX FANUC 11MF	179705B	104.24%	41.83%
40	8422	541483	G&L 36" VERT TURNING CENTER	112425A	103.27%	63.59%
40	8422	538629	36 G&L VTL 2A NUMERIPATH 8002-L CNC	112425A	103.27%	63.59%
40	8422	538632	42 G&L VTL 2A NUMERIPATH 8002-L CNC	112425A	103.27%	63.59%
40	8422	538631	42 G&L VTL 2A NUMERIPATH 8002-L CNC	112425A	103.27%	63.59%
40	8422	538630	36 G&L VTL 2A NUMERIPATH 8002-L CNC	112425A	103.27%	63.59%
40	8422	541492	G&L 36" VERT TURNING CENTER	112425A	103.27%	63.59%
41	8211	529611	#I/1 ELB CAM MASTER GR 4AX AB 8650	153159	101.89%	44.48%
41	8211	529615	#I/1 ELB CAM MASTER GR 4AX AB 8650	153159	101.89%	44.48%
42	8214	539969	PROCECO ACQUEOUS CLEANING SYS SI300	992000	99.38%	49.28%
43	8418	536875	RAYCON PINNACLE EDM 60AMP	198206	98.53%	13.89%
44	8211	529620	#I/2 ELB CAM MASTER GR 5AX AB8650	153159A	97.81%	72.20%

Figure 11 – Sample Plant-Wide Constraint Report

The "Rank" column codifies the relative severity of the constraints, with lower ranks indicating more severe constraints. "Work Department" indicates the location of the particular constraint within the plant. BU's are split into specific work departments, so a BUM or technical supervisor can filter this list looking only at his or her delegated work departments. "Brass Tag" indicates the specific piece of equipment that is constrained, as well as "Description". "Commodity Code" gives the six-digit indicator with the optional alphanumeric suffix that identifies the machine type and functionality. Per the description in Chapter 2, capacity is analyzed by commodity code, suffix and location. It is assumed that multiple brass tags with the same commodity code, suffix and location have work evenly distributed across them, resulting in multiple entries with identical loads but dissimilar brass tags. Entries make it on the list based upon their maximum load over the full forecast range of the model. If this maximum load exceeds a user-specified value, the entry makes the list. "Current Loading %" is included as an error-checking feature. The user can look at this value and verify if it is accurate with the shop floor.

Machine-Specific "Deep Dive" Report

In order to aid BU management in understanding and planning for future capacity constraints, a second report must look in-depth at individual commodity code, suffix, and location entries' capacity, and the factors determining it. The report should contain a central time-series chart that compares monthly machine availability with monthly machine load over the full forecast period of the model, per the initial user specification. This report should also give a detailed breakdown of the user-selected month in three charts. One chart identifies the top ten part numbers by machine load for the selected month, as well as whether the parts are for spares or engines. Separating this out gives visibility to the risk inherent in the schedule. If a particular part's schedule consists entirely of engine-order parts, it may be more reliable than one consisting entirely of spare parts. The other two charts identify the critical parameters' impact in the selected month on availability and loading, respectively. A sample report is attached in Exhibit 1.

The time-series chart shows monthly availability in machine-hours in red, with the blue load line overlaid. Any time the load line exceeds the availability, the machine is constrained. The chart also identifies the machines by brass tag, indicating the quantity and specific machines being analyzed. The user has the ability to stack excess availability of alternate machines on top of primary machine availability as well, giving an allowance for alternate operations. The two pie charts break down the impact of the various critical parameters on machine availability and load. On the availability chart, the "Available" pie section represents what is seen in red on the time-series chart. Each additional slice represents a piece of availability that was removed based upon the critical parameters. Noticeably absent is an indicator of charge out. This was incorporated to the manually specified variable "Max Loading Buffer". This is meant to be a lever that can allow users to account for any variables which are not present in the other data feeds. Because of the presumed significant variability from machine to machine of charge out, but lack of data, it was presumed that this be left to the user to specify, as the likely user, unit management, will have a better idea of the machine-specific charge out than any current data feed. The loading pie chart breaks out baseline load (defined by the schedule forecast and standard operation times), development load, and quality allowance. Additionally, two aspects of the operation profile are split out. Setup time, which is a measure of the inter-batch setup (not intra-batch, inter-part setup, which is incorporated into standard operation time) gives an indicator of the impact of batch size on load, as well as an additional error trigger. If-necessary operations are operations that are only performed as needed for a particular part. For example, if a part has a coating that can be damaged during part processing, an if-necessary operation may be developed for re-coating the part in the event that it is damaged at a certain point in the process. These if-necessary operations are assigned a typical percent occurrence rate, defined from the same data set as the quality allowance factors. This was split out because, in some cases, these if-necessary operations are performed much more often than management believes. This may cause management to overlook what could be a valuable indicator of problems with the process. Finally, the fourth chart identifies the top ten part numbers (disguised) that make up the machine load, and what types of orders these parts are for.

User Interface

The prototype user interface is not designed for the required simplicity, but rather for full functionality. Tables like the one shown in Figure 12 are used for user input. Data feeds, which are not automated in the prototype, come from additional worksheets, each housing the output from a standard database query. Scenario analysis is performed by manually editing the source data. However, this mirrors the desired UI functionality, which should be based on scenarios defined by replicating data feeds with the desired user-specified changes.

Machines	Machine ID	Shifts/Week	DCI %	Efficiency	Max Loading
Primary	153336-8727	15	13.1%	100.0%	80.0%
Alternate 1		15	13.1%	100.0%	80.0%
Alternate 2		15	13.1%	100.0%	80.0%
Alternate 3		15	13.1%	100.0%	80.0%
Hot Sheet		15	13.1%	100.0%	75.0%

Figure 12 – Sample User Input Table

In this table, the user selects a "Machine ID" for deep dive analysis, which is simply a combination of commodity code, suffix, and location from a drop-down, populated with every combination in the plant. The user then selects a staffing policy, and specifies DCI, efficiency, and max loading buffers. The user can then do the same thing for alternate machines whose excess availability the user wishes to add to the primary machine's availability. For hot sheet analysis, the user specifies a uniform staffing policy, DCI and plant-wide efficiency. The max loading buffer, in this case, sets the loading threshold for machines to make it onto the hot sheet.

6.2 Mathematical Model Implementation

Once the desired outputs are specified, it is possible to understand how the input data sets must be manipulated to generate the output information. Mathematical models will be described for both machine availability and machine load, as these two are calculated separately.

Machine Availability

Variables used in calculating machine availability are specified in Table 9. Machine availability is calculated the same way regardless of whether it is being done for a machine deep dive or hot sheet analysis.

Parameter	Units	Variable	Source
Machine ID	none	m	User-defined
Month	none	t	Forecast data
# Days in Month t	days	$D(t)$	Deterministic function of t
# Machine m	machines	$N(m)$	Function of m using equipment list
# Shifts per Week	shifts/week	S	User-defined
Utilization Efficiency	%	η	User-defined
Max Loading Buffer	%	β	User-defined
DCI	%	δ	User-defined
Maintenance Downtime	hours/month	$\mu(m)$	Calculated
Baseline Availability	hours/month	$A_0(m,t)$	Calculated
Loss to Inefficiency	hours/month	$A_\eta(m,t)$	Calculated
Loss to Max Loading Buffer	hours/month	$A_\beta(m,t)$	Calculated
Loss to DCI	hours/month	$A_\delta(m,t)$	Calculated
Availability	hours/month	$A(m,t)$	Calculated

Table 9 – Machine Availability Variable Definitions

The equations for machine availability are as follows:

$$A_0(m,t) = N(m) \times \frac{D(t)}{7\frac{days}{week}} \times S \times 8\frac{hours}{shift}$$

$$A(m,t) = \eta \times \beta \times (100\% - \delta) \times (A_0(m,t) - \mu(m))$$

$$A_\beta(m,t) = (100\% - \beta) \times \eta \times (100\% - \delta) \times (A_0(m,t) - \mu(m))$$

$$A_\eta(m,t) = (100\% - \eta) \times (100\% - \delta) \times (A_0(m,t) - \mu(m))$$

$$A_\beta(m,t) = \delta \times (A_0(m,t) - \mu(m))$$

Based on these equations, maintenance downtime loss is taken directly off of baseline availability. This means that it is assumed that all percentage parameters (η, δ, β) apply to time when the machines are "up", or not undergoing repair or maintenance. DCI is removed next. This assumes that the operator is performing direct tasks at a particular machine as a percent of the time that machine is "up". Next, efficiency is removed, which assumes that, of the time the operator is performing direct work on the machine, a certain amount is lost to inefficient design. Finally, maximum load buffers are taken off of the remaining machine availability.

Calculating maintenance downtime requires the following additional variables:

Parameter	Units	Variable	Source
Brass Tag	none	b	Equipment list
Machine ID	none	M(b)	Function of b using equipment list
Historical Month	none	t	Maintenance data
Machine-Specific Historical Maintenance Downtime	hours/month	$\mu_s(b,t)$	Function of b, t using maintenance data
Quantile	%	q	User-defined

Table 10 – Maintenance Downtime Variable Definitions

$$[\mu_s(t)] \ni \sum_{\{b|M(b)=m\}} \mu_s(b,t) \ \forall \, t$$

$$\tilde{\mu} \in [\mu_s(t)]$$

$$\{\mu(m) | \Pr(\tilde{\mu} \leq \mu(m)) = q\}$$

An empirical distribution of maintenance downtime by machine ID [$\mu_s(t)$] is constructed using 24 months of historical data. This requires calculating the sum of the downtime of all machines whose IDs match the user-specified ID, by month, as shown in the first equation above. An empirical cumulative distribution is then built under the assumption that the probability downtime will be less than the minimum or greater than the maximum of the data set is zero, and using simple linear interpolation between defined intermediate quantiles, as shown in the third equation above. This assumes pooled machine time, and historical monthly data points are independent and identically distributed. Projecting this distribution into the future assumes that the prior two years worth of data is demonstrative of the future 5 years' performance. While equipment is known to degrade over time and become less reliable, this was viewed within the plant as a reasonable assumption, though it is recommended that it be tested when sufficient data is available.

Machine Load

Variables required for the calculation of machine load are defined in Table 11.

Parameter	Units	Variable	Source
Part Number	none	p	Forecast data
Operation Number	none	o	Operation database
Business Unit	none	u(m)	Derived from machine ID
If Necessary Operations	none	[i]	Operation database
Production Parts	none	[P]	Forecast data
Development Parts	none	[D]	Development forecast
Total Operation Machine Time	hours/part	τ(o,p)	Operation database
Machine ID	none	M_0(o,p)	Operation database
Inter-Batch Setup Time	hours/batch	σ(o,p)	Operation database
Batch Size	parts/batch	B(p)	Operations staff
Lot Size	parts/lot	λ(p)	Operation database
Part Forecast	parts/month	F(t,p)	Forecast data
Quality Factor	parts/part	ξ(p,o)	Calculated
If Necessary Factor	parts/part	ι(p,o)	Calculated
Development Batch Inflation Factor	hours/hour	Δ(u,t,p)	Calculated
Total Load	hours/month	L(t,m)	Calculated
Baseline Load	hours/month	L_0(t,m)	Calculated
Development Load Contribution	hours/month	L_ψ(t,m)	Calculated
Quality Load Contribution	hours/month	L_ξ(t,m)	Calculated
If Necessary Load Contribution	hours/month	L_i(t,m)	Calculated
Setup Load Contribution	hours/month	L_σ(t,m)	Calculated
Part Load Contribution	hours/month	L_p(t,m,p)	Calculated

Table 11 – Machine Load Variable Definitions

The total load L(t,m) comprises the sum of all the individual load components shown in Table 11:

$$L(t,m) = L_0(t,m) + L_\xi(t,m) + L_\sigma(t,m) + L_i(t,m) + L_\psi(t,m)$$

Where the individual terms are given by the sum of all forecasts multiplied by the sum of all operation times τ(o,p) assigned to the machine divided by the part's lot size λ(p) multiplied by a factor characterizing the particular load component.

$$L_0(t,m) = \sum_{p\in[P]} F(t,p) \times \sum_{M_0(o,p)=m;\ (o,p)\notin[i]} \frac{\tau(o,p)}{\lambda(p)}$$

$$L_\xi(t,m) = \sum_{p\in[P]} F(t,p) \times \sum_{M_0(o,p)=m;\ (o,p)\notin[i]} (\xi(p,o) - 1) \times \frac{\tau(o,p)}{\lambda(p)}$$

$$L_\sigma(t,m) = \sum_{p\in[P]} F(t,p) \times \sum_{M_0(o,p)=m;\ (o,p)\notin[i]} \xi(p,o) \times \frac{\sigma(o,p)}{B(p)}$$

$$L_i(t,m) = \sum_{p\in[P]} F(t,p) \times \sum_{M_0(o,p)=m;\ (o,p)\in[i]} \iota(p,o) \times \left(\frac{\sigma(o,p)}{B(p)} + \frac{\tau(o,p)}{\lambda(p)}\right)$$

$$L_\psi(t,m) = \sum_{p\in[D]} F(t,p) \times \sum_{M_0(o,p)=m} \Delta(u,t,p) \times \left(\frac{\sigma(o,p)}{B(p)} + \frac{\tau(o,p)}{\lambda(p)}\right)$$

$$L_p(t,m,p) = \begin{cases} \displaystyle\sum_{M_0(o,p)=m} \Delta(u,t,p) \times \left(\frac{\sigma(o,p)}{B(p)} + \frac{\tau(o,p)}{\lambda(p)}\right) & \forall\, p \in [D] \\ \displaystyle\sum_{M_0(o,p)=m;\ (o,p)\notin[i]} \xi(p,o) \times \left(\frac{\sigma(o,p)}{B(p)} + \frac{\tau(o,p)}{\lambda(p)}\right) & \forall\, p \in [P]; (o,p) \notin [i] \\ \displaystyle\sum_{M_0(o,p)=m;\ (o,p)\in[i]} \iota(p,o) \times \left(\frac{\sigma(o,p)}{B(p)} + \frac{\tau(o,p)}{\lambda(p)}\right) & \forall\, p \in [P]; (o,p) \in [i] \end{cases}$$

The following equations all assume zero variability in setup time, operation time, batch size and lot size. The equation essentially subdivides machine load first into two groups: development and non-development work. It then divides non-development work into two more groups: if-necessary and primary operations. Following this division, there is one more: setup and primary operation contributions.

Additional variables required for the calculation of if-necessary and quality factors are shown in Table 12.

Parameter	Units	Variable	Source
Historical Month	none	t	Quality data
Date of Last Operation Profile Update	none	$t_u(p)$	Operation database
# Parts Through Operation	parts	$\pi(o,p,t)$	Quality data
# Parts Through Final Inspection	parts	$\pi_f(p,t)$	Quality data

Table 12 – Quality/If-Necessary Factor Variable Definitions

$$\iota(p,o) = \frac{\sum_{t \geq t_u(p)} \pi(o,p,t)}{\sum_{t \geq t_u(p)} \pi_f(p,t)} \ \forall\, p \in [i]$$

$$\xi(p,o) = \frac{\sum_{t \geq t_u(p)} \pi(o,p,t)}{\sum_{t \geq t_u(p)} \pi_f(p,t)} \ \forall\, p \notin [i]$$

Calculation of each of these factors assumes that the average value is representative of the data set over a given period of time. Using a historical two-year, or 24-month window, the total number of parts put through each operation is summed up, as is the total number of parts put through the corresponding final inspection operation. If an update has happened to the part profile within the two-year timeframe (which could indicate a discontinuity in the data), these values are calculated only for the duration of the updated profile. The equations for if-necessary and quality factors are calculated in an identical fashion, but have different properties.

If-necessary operations generally should have values less than 1. By definition, an if-necessary operation is one which is not required for every part. When establishing the part profile, an estimate of the if-necessary factor is developed for each occurrence of an if-necessary operation. However, historical data reveal the actual realized if-necessary factors. When there is insufficient data based upon recent profile updates, or simply low-volume parts, the operation profile estimate is used instead. However, as more data continues to be collected on these operations, estimates should become more accurate due to automated data links.

Quality factors on primary operations should have values greater than or equal to 1. The inherent assumption here is that, for any particular operation, at least the number of parts that made it through final inspection will have to have come through the previous operations, as well as the potential for additional parts that did not make it to final inspection due to quality defects. Where insufficient data is available, the global average quality factor is used instead. This factor again will become more accurate as part profiles become more mature.

Additional variables required for the calculation of development batch inflation factors are shown in Table 13. The following calculations are only for part numbers contained in the [D] development part number set.

Parameters	Units	Variable	Source
Learning Rate	%	Λ	Manual curve-fitting
Realization Rate	hours/hour	$\rho(u)$	Historical clocking data
Total # Parts Delivered to Date	parts	$\pi_D(t,p)$	Program office staff
First Piece Inflation Factor	hours/hour	$\phi(u)$	Calculated
Batch Algebraic Midpoint Unit	parts	$K(t,p)$	Calculated
First Batch Unit - ½	parts	$\Pi_1(t,p)$	Calculated
Last Batch Unit + ½	parts	$\Pi_2(t,p)$	Calculated
Crawford Inflation Factor	hours/hour	$Y(u,t,p)$	Calculated

Table 13 – Development Batch Inflation Factor Variable Definitions

$$\Delta(u,t,p) = \begin{cases} Y(u,t,p) & \forall\, Y(u,t,p) \geq 1 \\ 1 & \forall\, Y(u,t,p) < 1 \end{cases}$$

$$Y(u,t,p) = \varphi(u) \times K(t,p)^{\log_2 \Lambda}$$

$$\varphi(u) = \frac{\rho(u)}{25^{\log_2 \Lambda}}$$

$$K(t,p) = \left[\frac{F(t,p) \times (1 + \log_2 \Lambda)}{\Pi_2(t,p)^{1+\log_2 \Lambda} - \Pi_1(t,p)^{1+\log_2 \Lambda}}\right]^{-\frac{1}{\log_2 \Lambda}}$$

$$\Pi_1(t,p) = \pi_D(t,p) - \frac{1}{2}$$

$$\Pi_2(t,p) = \pi_D(t,p) + F(t,p) + \frac{1}{2}$$

Calculation of the development batch inflation factor uses a Crawford learning curve, (Liao, 1988) calculated off the 25th piece realization rate by business unit, based upon an empirically fit learning rate. Realization rates are calculated manually using clocking data. For any given business unit, the realization rate is equal to the ratio of total direct hours charged in the previous month to total standard delivered hours for the month. While this ratio applies well to labor rates, it breaks down in typical production manufacturing for machine hours because, in many cases, one operator runs more than one machine. When this is the case, it cannot be assumed that, though the operator will take 1.3 or 1.4 times as long to do his work, each machine will have the same inflation factor. However, for development manufacturing, work centers are configured such that one operator runs only one machine. Because of this, labor realization rates can more reasonably be applied to the development manufacturing machine time model, because total machine time will always be less than or equal to total clocked manual time.

In speaking with program office staff, it is used as a heuristic that the 25th piece of any development project tends to converge to the particular unit's realization rate. This is the point at which processes are considered "production-ready". Consulting with manufacturing engineering staff led to the approach of continuing the learning curve decline until the inflation factors reached 1. At this point, it is assumed that no further cycle time can be taken out of the process. The reason for this assumption is that many of these parts have substantial amounts of time spent on the machine. While manual labor practices can be streamlined to reduce total cycle time, as well as scrap rates to reduce rework, this intrinsic cycle time is difficult to reduce without process redesign, which is not assumed for the purposes of capacity planning.

Fitting an appropriate learning rate involved consultation with the development program staff, and a manual process. Because of the irregularity of development work, there is not a substantial amount of data showing trends. However, one particular business unit had data available for a set of parts. Unfortunately, there was no analysis available for individual operation inflation rates, but the analysis was available at the part level. Anchoring the curve at the unit's realization rate for the 25th piece, the curve's learning rate was adjusted until approximately 90% of the historical data fell below the fit curve. Assuming the data is a representative set, it can then be predicted that development 90% of future development work should also realize favorable to the fit curve. These criteria corresponded to a learning rate of approximately

75%, which was then assumed for all development work in the plant. The final curve and historical data are shown in Figure 13.

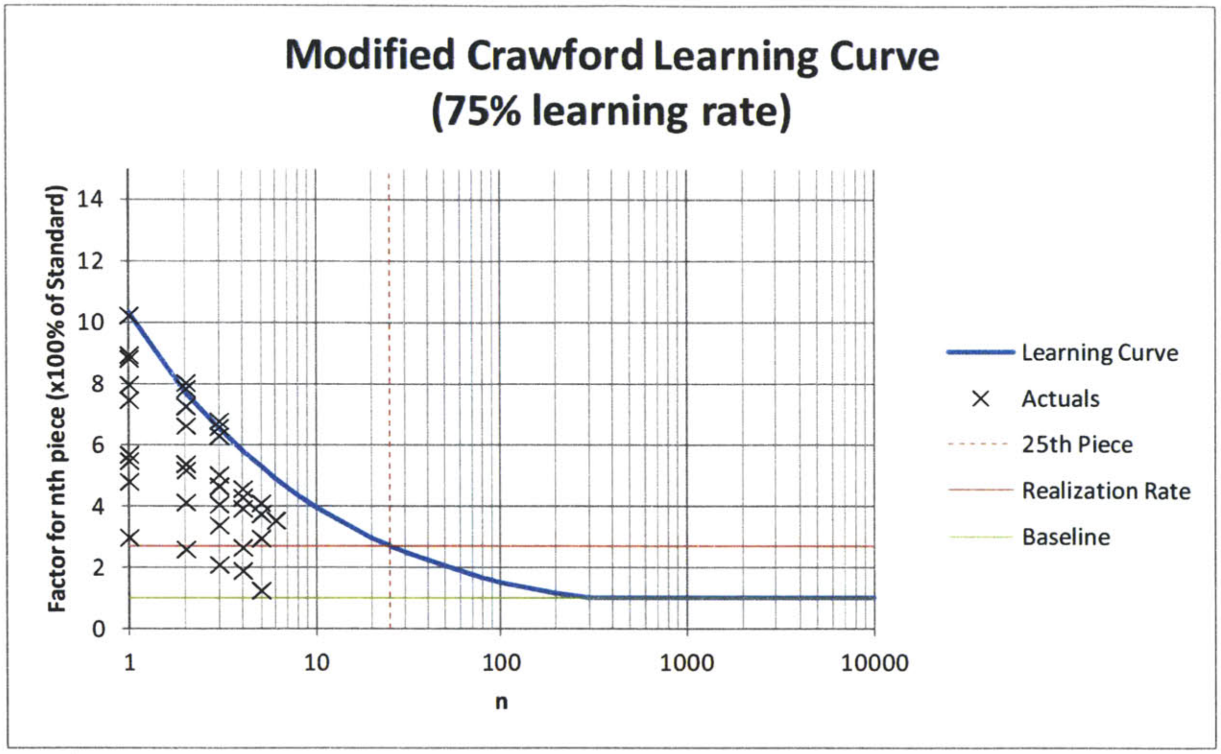

Figure 13 – Learning Curve Parameter Fit

The set of historical data points had n=37, and of those points, 33 fell below the 75% curve. Because this data is so dispersed, and many parts will actually have realized inflation rates likely significantly lower than the model, it is important that machines with loading heavily dependent on development work look more closely at the extent to which this model holds. However, because of the relatively rapid learning rate, it is also advisable that equipment purchases not be made on the basis of short-term early-stage development constraints. Furthermore, as more data is available, the validity of this methodology should be tested across business units.

6.3 Model Validation

With a fully-functioning prototype, the next step is to validate its results. Ideally, full historical data sets could be validated against historical realized machine loads. However, none of this data is available, because tracking of real-time machine utilization is something that was not historically done in the plant. As a rough-cut validation process, hot sheet entries were verified with the shop floor based upon current loading. Feedback data from the plant was designed to be in "yes/no" form. If the current loading is found to be inaccurate, it was then asked whether the actual machine utilization was higher or lower than suggested. Standardized queries were as follows:

Current Loading	Question	Response	Follow-Up Question	Response	Result
90% +	Is this machine constrained?	Yes			Valid
		No	Is this machine utilized more or less than suggested?	More	Valid
				Less	Overestimate
60 - 90%	Does this machine run three shifts?	Yes			Valid
		No	Is this machine utilized more or less than suggested?	More	Underestimate
				Less	Overestimate
30 - 60%	Does this machine run two shifts?	Yes			Valid
		No	Is this machine utilized more or less than suggested?	More	Underestimate
				Less	Overestimate
0 - 30%	Does this machine run one shift or less?	Yes			Valid
		No	Is this machine utilized more or less than suggested?	More	Underestimate
				Less	Valid

Table 14 – Validation Data Collection Method

Initial validation data suggested that the model was much more accurate when applied to traditional removal fabrication equipment (such as lathes, mills, EDM, laser cutting, and grinders), due in large part to data fidelity issues regarding nonstandard equipment (such as coating booths, deburr benches, welders, and heat treat furnaces). This data set is shown graphically in Figure 14 and numerically in Table 15.

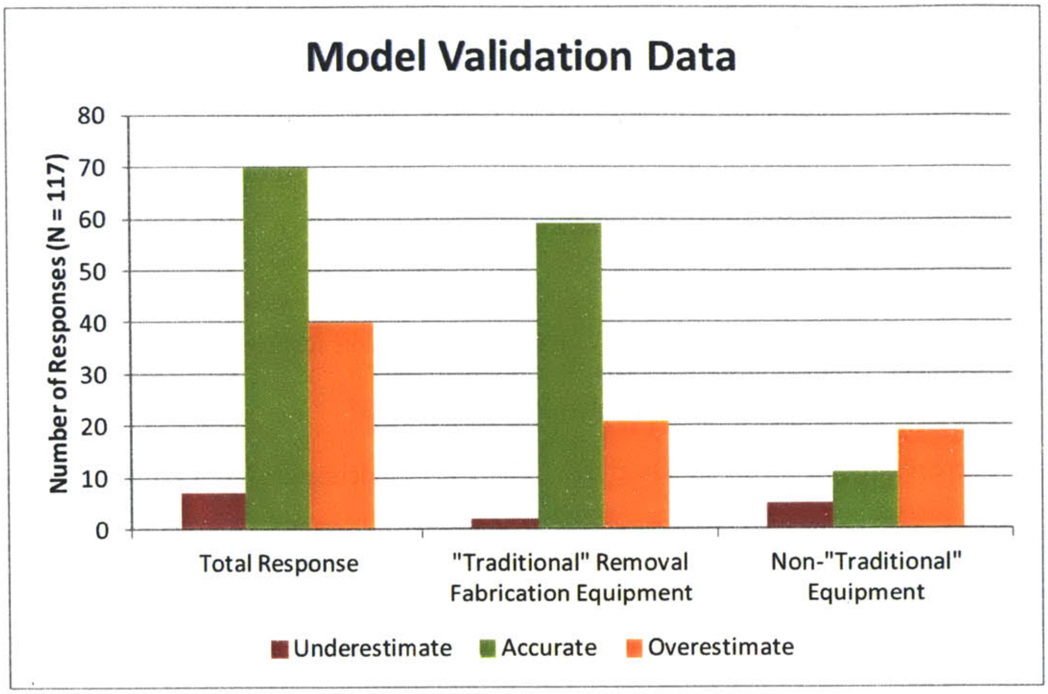

Figure 14 – Current Loading Model Validation Data

The validation data indicate that, of 117 "hot sheet" entries tested, 60% of the estimates of machine loading were approximately accurate, while in 34% of cases, loading was overestimated. The remaining 6% of cases indicated an underestimation of loading.

Accuracy of Current Loading	Total Response	"Traditional" Equipment	Non-"Traditional" Equipment
Underestimate	6%	2%	14%
Accurate	60%	72%	31%
Overestimate	34%	26%	54%

Table 15 – Validation Data Parsed by Equipment Type

As shown above, the model accuracy improves to 72%, with 24% overestimation and only 2% underestimation for traditional manufacturing equipment, with non-traditional equipment showing only 31% accuracy. While projects were incited to rectify the data issues behind this discrepancy, for the purposes of application of the prototype, the scope was limited to the aforementioned types of traditional manufacturing equipment.

7 Chapter 7 – Application in Business Unit 840

Having validated the model for use on traditional manufacturing equipment, it was applied in a case study on business unit 840's capacity requirements. The application of the tool followed the following basic process:

1. Run "Hot Sheet" analysis to identify all flagged constraints within the unit
2. Verify those constraints' current state with the shop floor
3. Where results are valid, run "Deep Dive" analysis to understand what is causing the constraint
4. Develop potential solutions for addressing the constraint

In addition to these four steps, the analysis extends into the capital justification process with the following additional steps:

5. Perform economic analysis on alternative solutions
6. Select the most favorable solution, and prepare a business case

These final steps are considered outside the scope of this research, which has focused specifically on developing tools to aid in scenario analysis and understanding capacity constraints and requirements.

7.1 Hot Sheet Analysis and Validation

Business unit 840 contains five departments, 8315, 8422, 8616, 8617, and 8618. Running the hot sheet analysis with a threshold of 75% of maximum load, a plant-wide constraint list was obtained. Filtering the list of predicted constraints by location, and looking only at those under unit 840's jurisdiction, the results in Figure 15 were considered.

Rank	Work Dept	Brass Tag	Description	Comm Code	Max Load %	Current Load %
19	8618	521529	48 G&L VTL 4AXCNC 804L NC	112485	158.52%	70.80%
33	8422	019400	MITSUI SEIKI MACHINE CENTER	179705	120.19%	120.19%
35	8422	540724	NEW ENGLAND GAS OVEN 850F	723000A	116.81%	112.06%
40	8422	541483	G&L 36" VERT TURNING CENTER	112425A	103.27%	63.59%
40	8422	538629	36 G&L VTL 2A NUMERIPATH 8002-L CNC	112425A	103.27%	63.59%
40	8422	538632	42 G&L VTL 2A NUMERIPATH 8002-L CNC	112425A	103.27%	63.59%
40	8422	538631	42 G&L VTL 2A NUMERIPATH 8002-L CNC	112425A	103.27%	63.59%
40	8422	538630	36 G&L VTL 2A NUMERIPATH 8002-L CNC	112425A	103.27%	63.59%
40	8422	541492	G&L 36" VERT TURNING CENTER	112425A	103.27%	63.59%
45	8315	523117	4515B J&L 2AX NC LATHE AB7360 CNC	165224	97.00%	62.85%
55	8422	540922	MONARCH VMC-RTB 50"	133219	87.91%	59.21%
55	8422	540921	MONARCH VMC-RTB 50" (NEW)	133219	87.91%	59.21%
71	8422	522323	48 G&L VTL 2AX CNC 802-L	112485	75.51%	48.85%
71	8422	522326	48 G&L VTL 2AX CNC 802-L	112485	75.51%	48.85%
71	8422	522328	48 G&L VTL 2AX CNC 802-L	112485	75.51%	48.85%
71	8422	522327	48 G&L VTL 2AX CNC 802-L	112485	75.51%	48.85%
71	8422	520411	48 G&L VTL 2AX CNC #802L	112485	75.51%	48.85%

Figure 15 – Initial Hot Sheet Results, Unit 840

Of these, several were found to be erroneous including the oven because of improper operation profile parameters for lot size, and the lathe in department 8618 because of schedule inaccuracy. Additionally, the lathe in department 8315 was also found to be a non-issue, as similar lathes are shared between 8315, 8616, and 8617, the latter two of which have sufficient open capacity. Likewise, the five 48" vertical lathes in department 8422 also are not problematic, as vertical lathe resources are pooled throughout the business unit, and sufficient under-utilized capacity is available to cover any borderline constraints that may arise therein. Filtering out these results, Figure 16 contains the remaining constraints for deep dive analysis.

Rank	Work Dept	Brass Tag	Description	Comm Code	Max Load %	Current Load %
33	8422	019400	MITSUI SEIKI MACHINE CENTER	179705	120.19%	120.19%
40	8422	541483	G&L 36" VERT TURNING CENTER	112425A	103.27%	63.59%
40	8422	538629	36 G&L VTL 2A NUMERIPATH 8002-L CNC	112425A	103.27%	63.59%
40	8422	538632	42 G&L VTL 2A NUMERIPATH 8002-L CNC	112425A	103.27%	63.59%
40	8422	538631	42 G&L VTL 2A NUMERIPATH 8002-L CNC	112425A	103.27%	63.59%
40	8422	538630	36 G&L VTL 2A NUMERIPATH 8002-L CNC	112425A	103.27%	63.59%
40	8422	541492	G&L 36" VERT TURNING CENTER	112425A	103.27%	63.59%
55	8422	540922	MONARCH VMC-RTB 50"	133219	87.91%	59.21%
55	8422	540921	MONARCH VMC-RTB 50" (NEW)	133219	87.91%	59.21%

Figure 16 – Constraints for Deep Dive Analysis, Unit 840

7.2 Deep Dive Analysis and Solution Development

Deep dive analyses for each of these three unique machine ID's (179705-8422, 112425A-8422, and 133219-8422) were performed. Upon deep dive analysis, it became apparent that machine ID 112425A-8422 did not pose a long-term capacity concern, because its maximum load is short-term. The remaining two machine IDs' deep dive analyses are attached in Exhibit 2.

Especially with the 4-axis vertical milling center 179705-8422, the perception on the shop floor was that a new machine was needed immediately. This is reflected in the deep dive analysis showing substantial spikes of machine overloading in the near term. This period, however, is followed by a multi-year demand lull in 2014 and 2015. Because of this, an immediate equipment purchase may be unnecessary, if a short-term fix can be found to handle the production load being displaced by so much development work. Situations like this represent a blind spot in the current process for capacity planning, and demonstrate the need for a tool such as the one being proposed.

A set of two two-phased solution alternatives were developed, each containing a short-term solution to address the immediate capacity constraints, and a long-term solution to address the future capacity constraint as demand picks up after the 2014 – 2015 demand lull. Short-term solutions included redistribution of all production work to a set of similar machines in business unit 810 (179704-8418), with the result shown in Figure 17. Alternatively, it was suggested in the short term that some development work be sent out to the vendor base, in case scaling of the development operation cannot happen quickly enough.

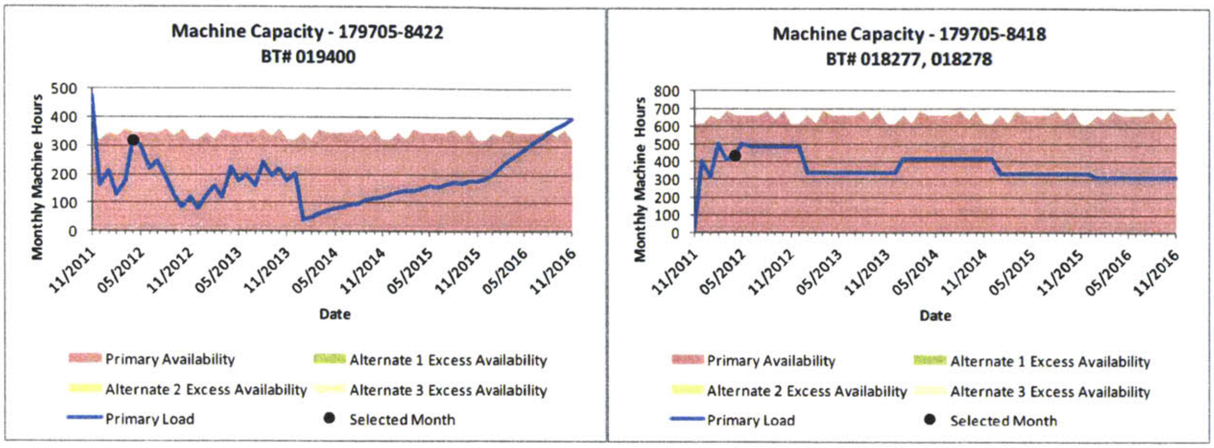

Figure 17 – Short-Term Work Redistribution

Long-term, however, equipment purchases will need to be made to accommodate the demand incumbent upon both 179705-8422 and 133219-8422 mills. Current capital plans suggested a move to 5-axis milling technology, which enables the combination of operations that are currently split between the two different types of milling machines. The result of the current capital plan is shown in Figure 18.

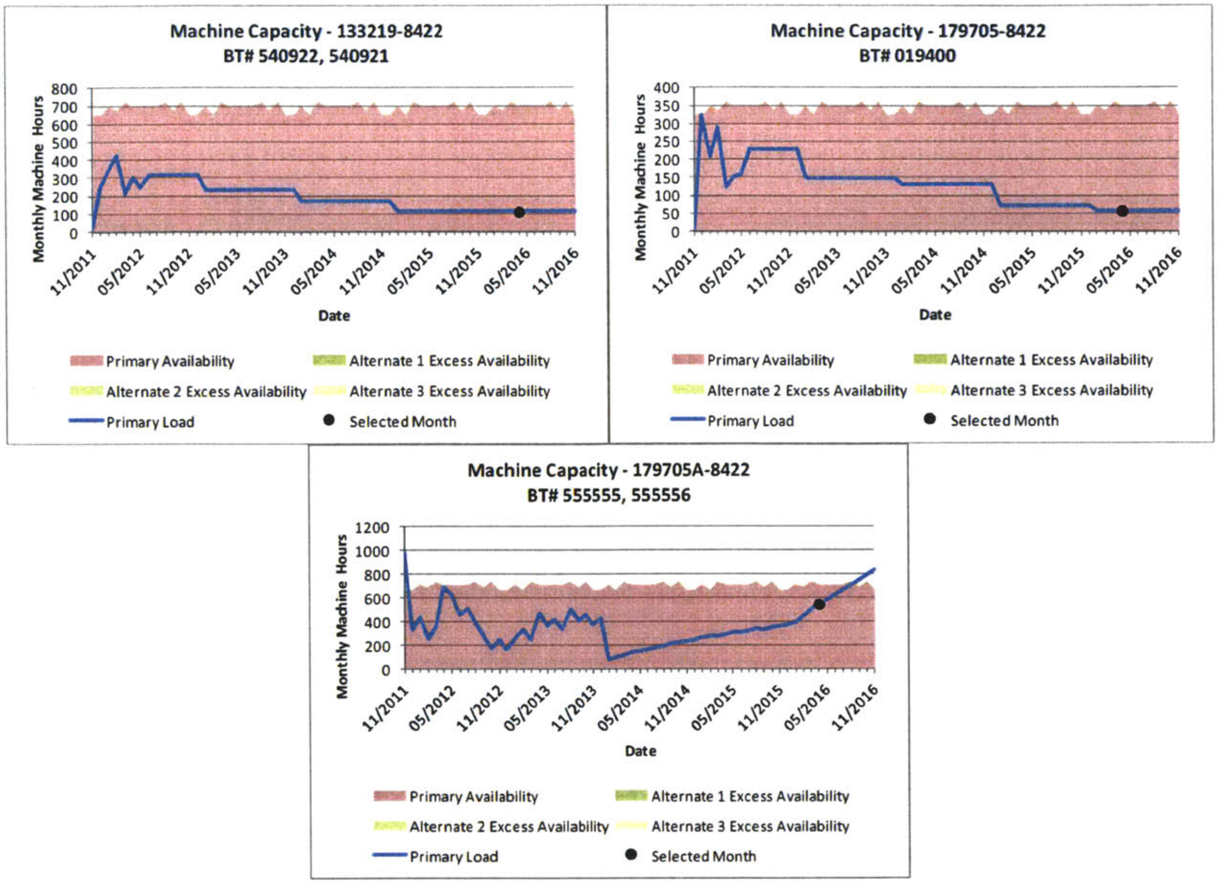

Figure 18 – Projected Impact of Current Capital Plan

As apparent in Figure 18, the current capital plan will leave the existing machines with almost no load by 2016. While this will enable one of the 133219-8422 machines to be removed, it may not be the best option. An alternative long-term solution is instead posed, wherein the purchase of one five-axis machine is deferred, with the more efficient result shown in Figure 19.

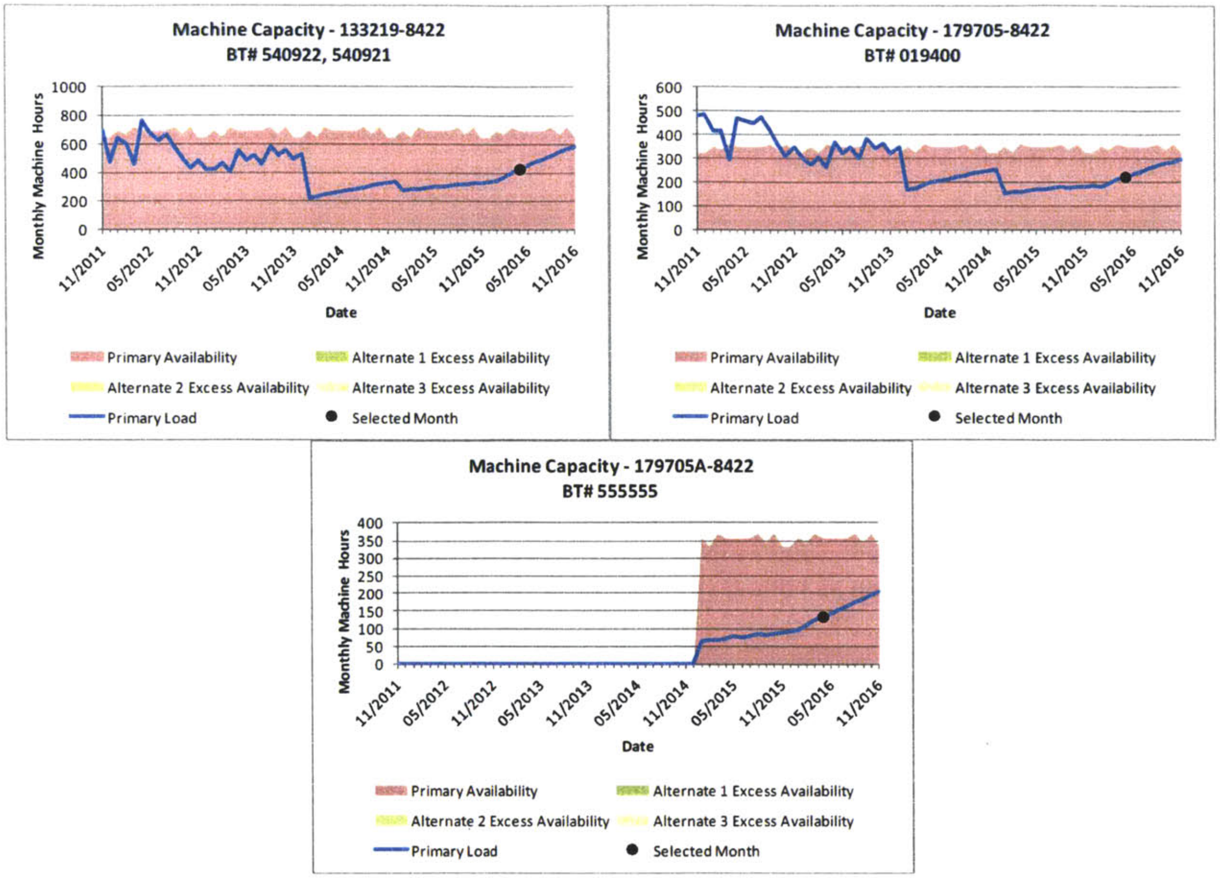

Figure 19 – Alternative Scenario Recommendation

In this scenario, all machines are well-loaded without being constrained, and the purchase of one five-axis mill may be deferred, creating a potentially significant savings in opportunity cost of capital. From this set of recommendations, the unit staff can then assess the economic opportunity of each alternative, and select the best one.

In practice, implementation of these types of scenarios involves substantially more work than is apparent on the surface, as moving parts from one department or machine to another requires development of new computer numeric control (CNC) programs, and development of alternate operations in the operation profile. However, having access to a simple way of understanding aggregate impact is of critical importance when missteps can lead to potential misallocation of millions of dollars of capital.

This process can be replicated in any other parts of the plant where constraints are identified, and as data system projects identified in Phase III of model development are completed, the scope of the model may encompass the entirety of the plant.

8 Chapter 8 – Lessons Learned and Extensions

8.1 Lessons Learned

The preceding chapters outlined in detail the process used for development and application of a static capacity model for capacity requirements planning at the North Berwick Parts Center. While many of the specific aspects of the model and this analysis are not extensible to other plants, industries, or areas of study, there a few central, extensible lessons. At an application level, mixing development work with production work can lead to misperceptions regarding capacity constraints, as demonstrated by the analysis on unit 840's milling capacity. At the manufacturing plant level, capacity is understood differently by different parties, and developing a uniform language by any method (here with a centralized tool) enables ease of communication. Finally, at the model development and implementation level, while a linear phased process helps in outlining the necessary steps toward implementation of a model, the process is in fact far from linear.

Application-Level Lessons

The process of capital planning happens infrequently in large manufacturing organizations such as Pratt and Whitney. Because of this, there is a tendency not to devote substantial resources at the operations level to codifying the process. While operations staff may have a fairly good handle on their shop floor needs, however, as processes get more and more complex, this becomes more difficult. Issues such as those demonstrated in the BU 840 analysis are inevitable in an environment as broad in scope and complexity as NBPC. In this kind of situation, developing tools to visualize data is of critical importance to enable intelligent decision-making. As complexity increases, so too does this need increase.

Manufacturing Plant-Level Lessons

From a managerial perspective, machine capacity is understood differently by different groups. Plant management sees machine capacity as answering the questions, "Can we make more?" and "How much more?". The tendency is to look at aggregate metrics in the long term, often without acknowledgement of the assumptions inherent in that aggregation. Operations-level management sees machine capacity as one of many constraints on a highly-constrained system. Typically taking a more localized, myopic view, operations-level management understands that capacity decisions cannot be made using aggregated metrics. However, the tendency is to focus on this complexity at a local level, as doing so systemically would be intractable. At this level, tools help to aid in communication between the two groups. Codification of the localized knowledge residing in operations staff can inform more accurate metrics for plant management. Likewise, this same codification can enable operations staff to understand the systematic implications of their localized decision-making. In model development, care must be taken to ensure that all potential consumers of model output are in agreement as to the type of data and user experience required.

Model Development Lessons

The four-phased approach taken ensures that all needs are met, and a model is developed that meets the needs of its users, and is fully functional. However, the process of validation does not occur linearly. In order to develop a model which is not only functional, but adoptable, and accurate, requires an iterative approach, and the construction and implementation of not only a model, but also a management system around the model. In Phase III, it is critical to acknowledge where sufficient data was not available and incite a project to fix this problem, thus improving upon the model being built. In Phase IV, noting that certain types of equipment tended not to fit the model well led to the acknowledgement of a significant issue with data quality. This must be accompanied with the inception of a project to fix this problem.

8.2 Recommendations for Future Work

The work performed to generate this thesis led to the creation of prototype model which is reasonably applicable to a limited subset of NBPC's capital equipment. The recommendations contained herein serve to make this model's approximation of the shop floor behavior closer to the truth, and for a greater subset of equipment, until it is eventually universally applicable to the entire shop. Within NBPC, it is recommended that all recommendations generated in the course of this analysis be implemented, and the model be used to take a closer look at the entire plant's long-term capital plan.

The successful implementation of this full-scale model warrants an analysis of its impact on the organization's decisions regarding capital investment. While the specific data analyzed was for NBPC only, it is likely that the problem of the lack of a comprehensive capacity planning tool is one that exists not only at NBPC, but elsewhere within Pratt and Whitney, and likely in other large manufacturing organizations dealing with large plants with significant amounts of complexity in process and product types.

9 Exhibits and Appendices

9.1 Exhibit 1

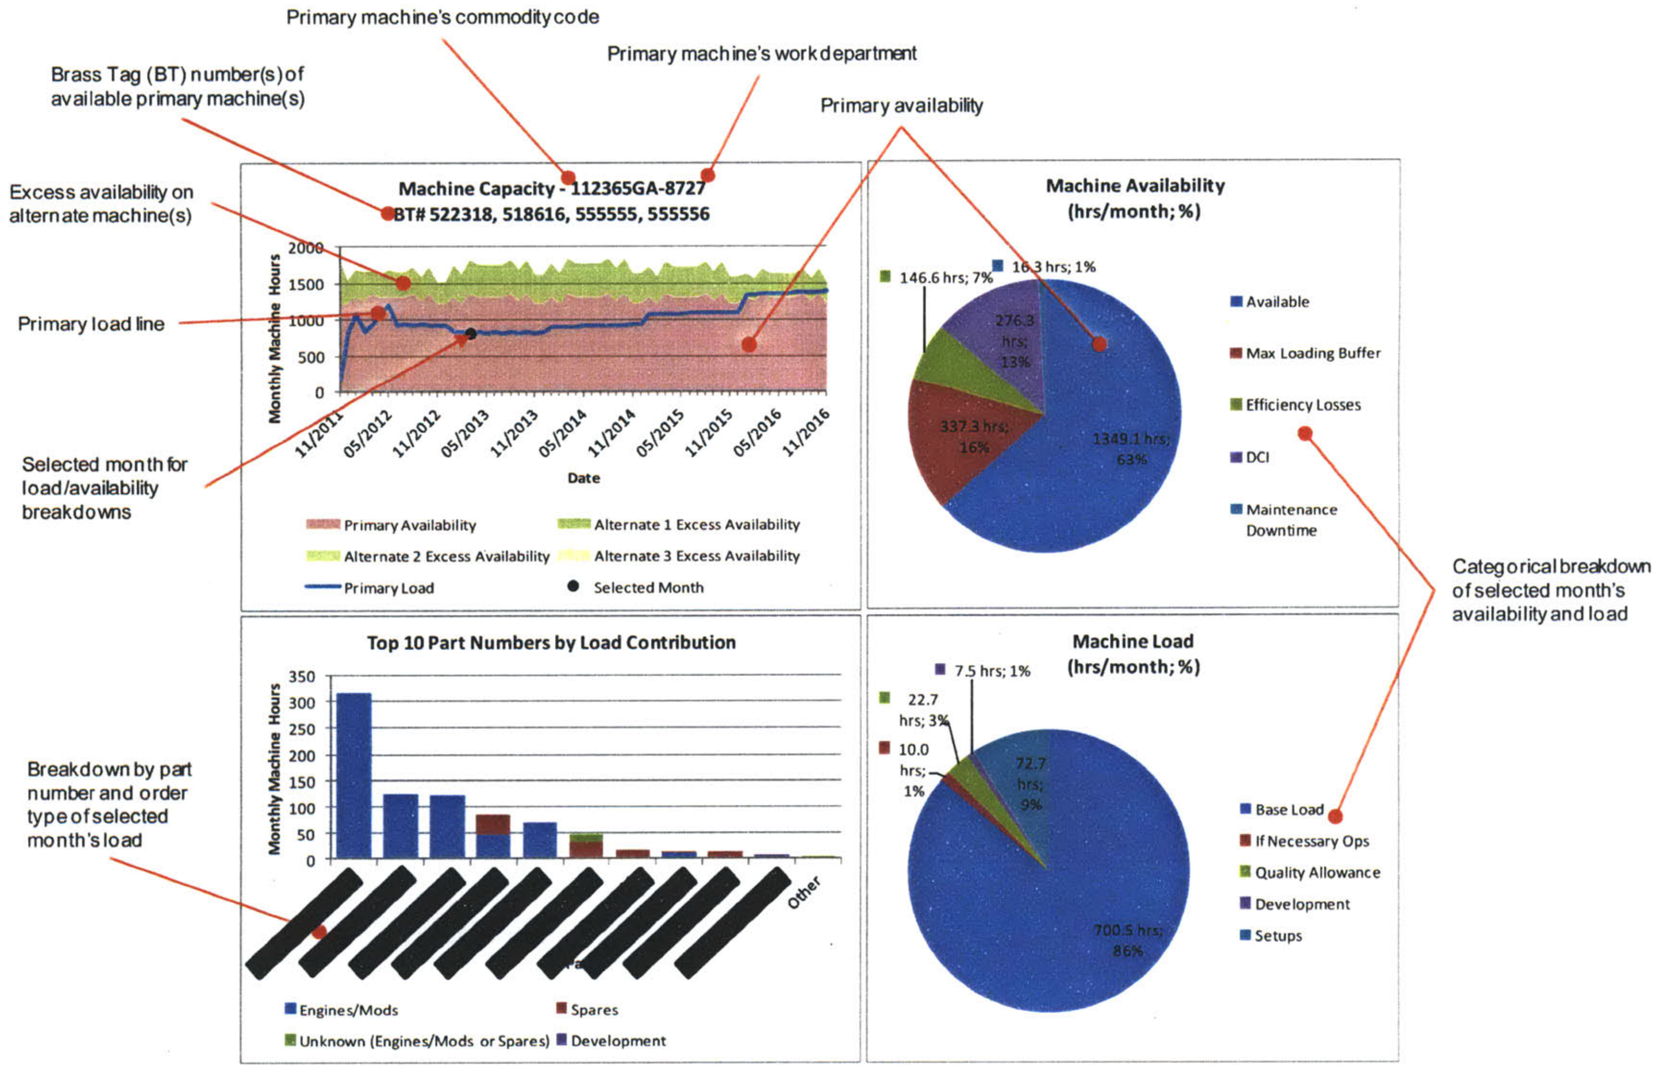

Primary machine's commodity code
Primary machine's work department
Brass Tag (BT) number(s) of available primary machine(s)
Primary availability
Excess availability on alternate machine(s)
Primary load line
Selected month for load/availability breakdowns
Categorical breakdown of selected month's availability and load
Breakdown by part number and order type of selected month's load
Machine Capacity - 112365GA-8727
BT# 522318, 518616, 555555, 555556
Monthly Machine Hours
2000
1500
1000
500
0
11/2011
05/2012
11/2012
05/2013
11/2013
05/2014
11/2014
05/2015
11/2015
05/2016
11/2016
Date
Primary Availability
Alternate 1 Excess Availability
Alternate 2 Excess Availability
Alternate 3 Excess Availability
Primary Load
Selected Month
Machine Availability
(hrs/month; %)
146.6 hrs; 7%
16.3 hrs; 1%
276.3 hrs; 13%
337.3 hrs; 16%
1349.1 hrs; 63%
Available
Max Loading Buffer
Efficiency Losses
DCI
Maintenance Downtime
Top 10 Part Numbers by Load Contribution
Monthly Machine Hours
350
300
250
200
150
100
50
0
[redacted]
Other
Engines/Mods
Spares
Unknown (Engines/Mods or Spares)
Development
Machine Load
(hrs/month; %)
7.5 hrs; 1%
22.7 hrs; 3%
10.0 hrs; 1%
72.7 hrs; 9%
700.5 hrs; 86%
Base Load
If Necessary Ops
Quality Allowance
Development
Setups

9.2 Exhibit 2

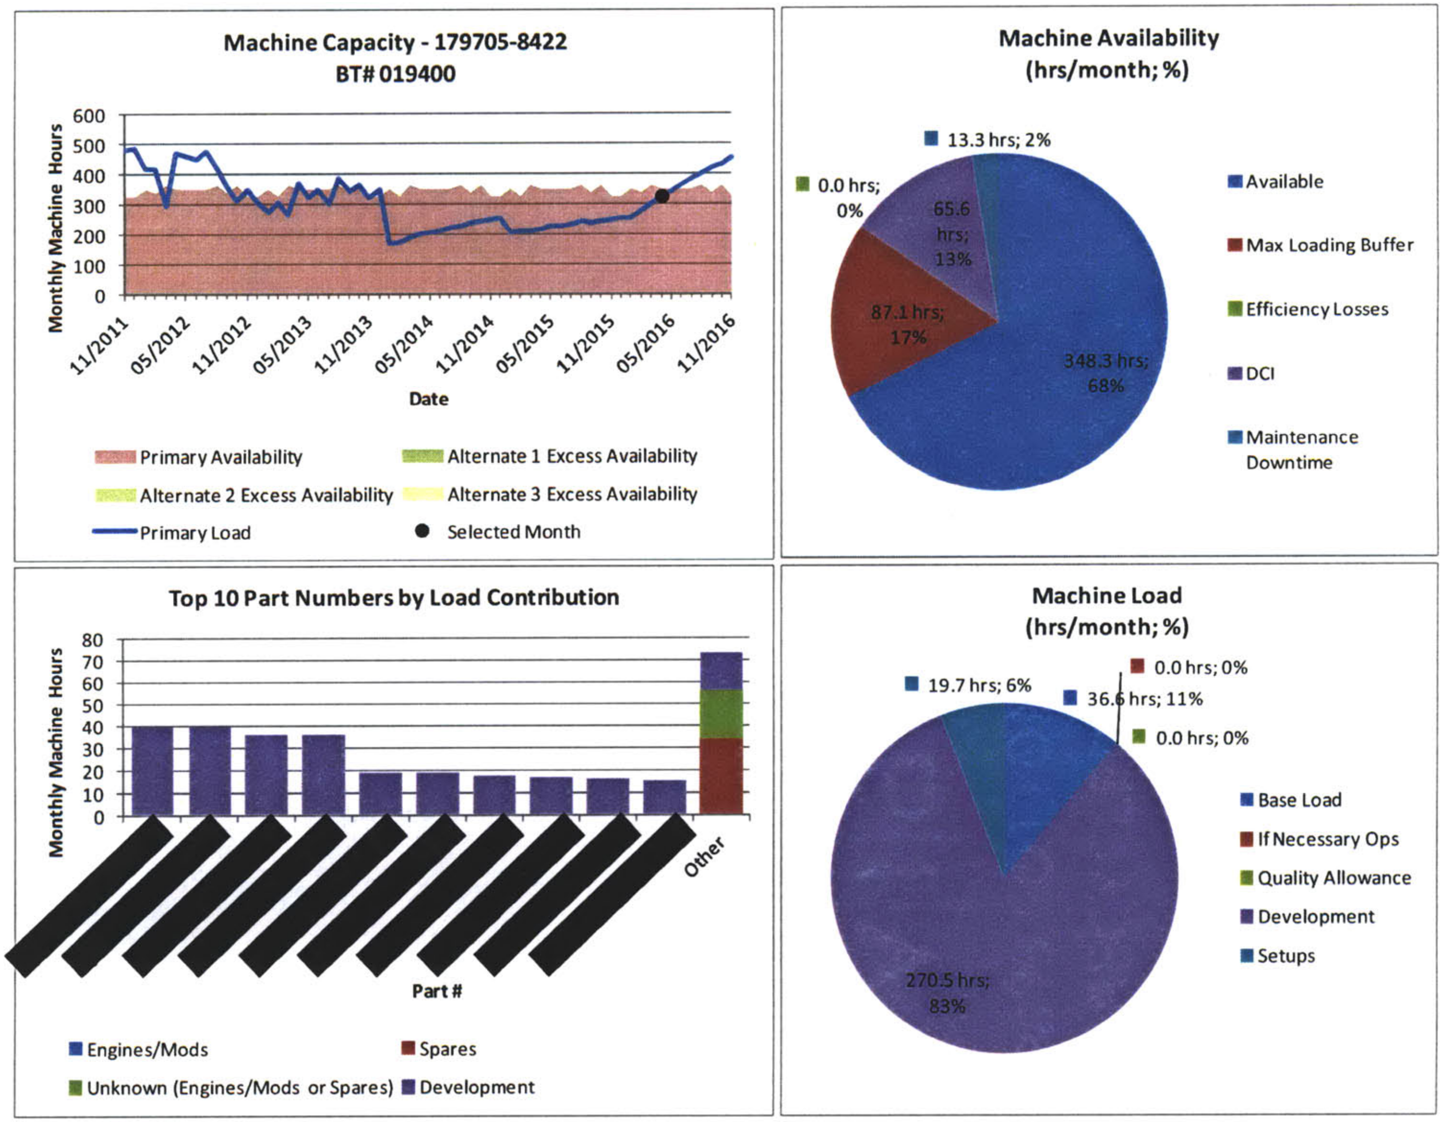
Machine Capacity - 179705-8422
BT# 019400
Monthly Machine Hours
600
500
400
300
200
100
0
11/2011
05/2012
11/2012
05/2013
11/2013
05/2014
11/2014
05/2015
11/2015
05/2016
11/2016
Date
Primary Availability
Alternate 1 Excess Availability
Alternate 2 Excess Availability
Alternate 3 Excess Availability
Primary Load
Selected Month
Machine Availability
(hrs/month; %)
13.3 hrs; 2%
0.0 hrs; 0%
65.6 hrs; 13%
87.1 hrs; 17%
348.3 hrs; 68%
Available
Max Loading Buffer
Efficiency Losses
DCI
Maintenance Downtime
Top 10 Part Numbers by Load Contribution
Monthly Machine Hours
80
70
60
50
40
30
20
10
0
[redacted]
Other
Part #
Engines/Mods
Spares
Unknown (Engines/Mods or Spares)
Development
Machine Load
(hrs/month; %)
0.0 hrs; 0%
19.7 hrs; 6%
36.6 hrs; 11%
0.0 hrs; 0%
270.5 hrs; 83%
Base Load
If Necessary Ops
Quality Allowance
Development
Setups

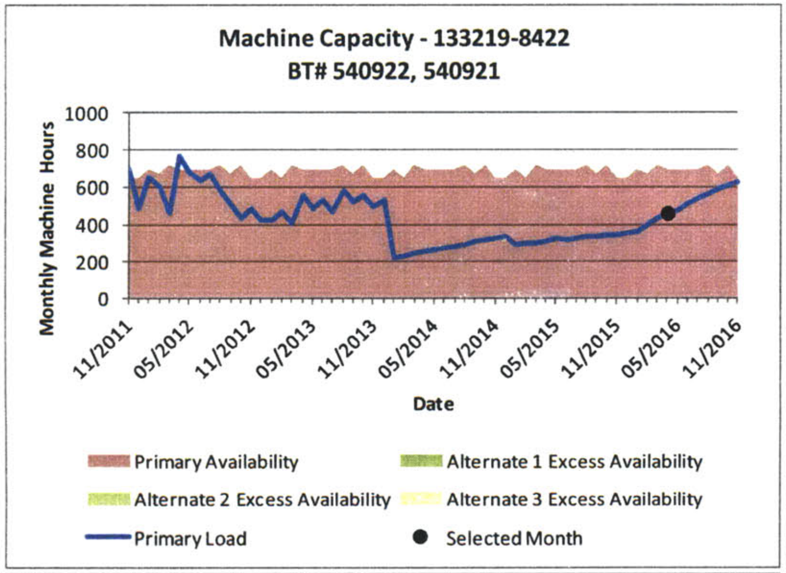
Machine Capacity - 133219-8422
BT# 540922, 540921
Monthly Machine Hours
1000
800
600
400
200
0
11/2011
05/2012
11/2012
05/2013
11/2013
05/2014
11/2014
05/2015
11/2015
05/2016
11/2016
Date
Primary Availability
Alternate 1 Excess Availability
Alternate 2 Excess Availability
Alternate 3 Excess Availability
Primary Load
Selected Month

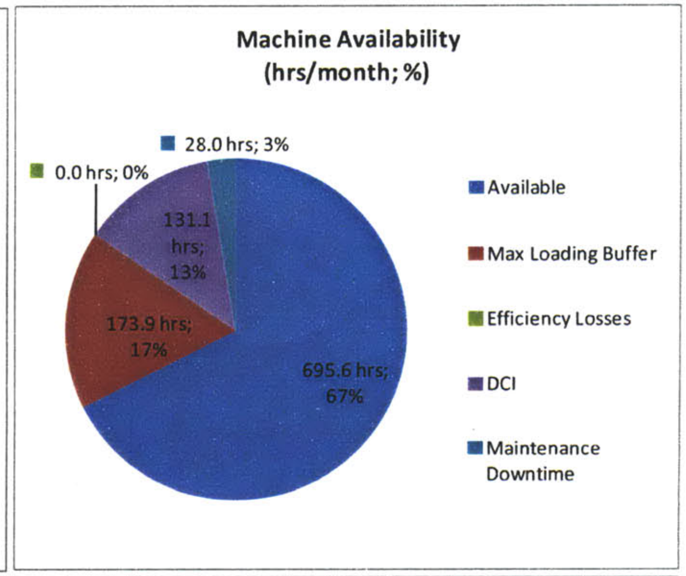
Machine Availability
(hrs/month; %)
28.0 hrs; 3%
0.0 hrs; 0%
131.1 hrs; 13%
173.9 hrs; 17%
695.6 hrs; 67%
Available
Max Loading Buffer
Efficiency Losses
DCI
Maintenance Downtime

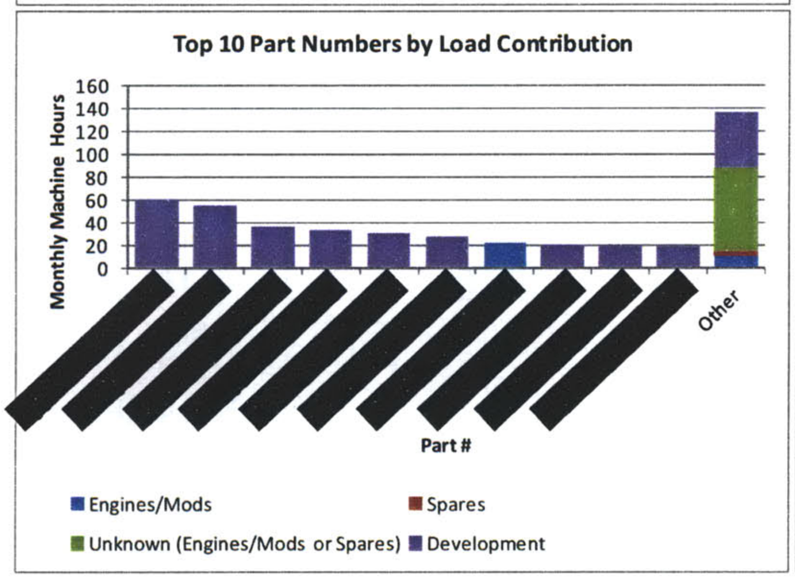
Top 10 Part Numbers by Load Contribution
Monthly Machine Hours
160
140
120
100
80
60
40
20
0
[redacted]
Other
Part #
Engines/Mods
Spares
Unknown (Engines/Mods or Spares)
Development

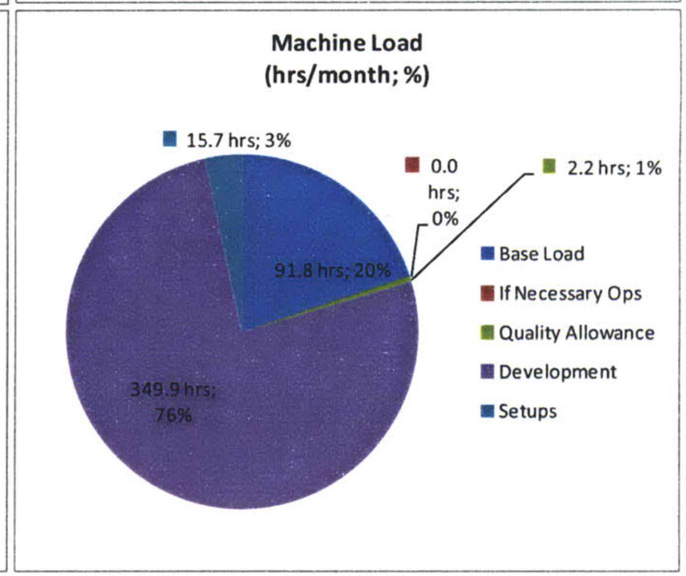
Machine Load
(hrs/month; %)
15.7 hrs; 3%
0.0 hrs; 0%
2.2 hrs; 1%
91.8 hrs; 20%
349.9 hrs; 76%
Base Load
If Necessary Ops
Quality Allowance
Development
Setups

9.3 Appendix A – Acronym Definitions

Acronym	Stands For	Definition
ACE	Achieving Competitive Excellence	Pratt & Whitney's production system
BOAS	Blade Outer Air Seal	A part of the gas-turbine engine that forms a friction seal with rotary parts
BT	Brass Tag	The unique number given to each machine tool in the shop
BU	Business Unit	A subdivision within one of Pratt & Whitney's module centers
BUM	Business Unit Manager	The manager of a business unit within one of Pratt & Whitney's module centers
CC	Commodity Code	A six-digit number indicating a machine tool's capability and function
CEPO	Capital Equipment Procurement Office	The organization within Pratt & Whitney responsible for the procurement of capital equipment
CNC	Computer Numeric Control	The process of using automated computer programs to run fabrication equipment
DCI	Direct Charging Indirect	Term for time hourly manufacturing employees spend on things other than manufacturing
DoD	Department of Defense	Branch of the US government that funds military aircraft development and production
ERP	Enterprise Resource Planning	A type of software used for managing resource allocations throughout large enterprises
GSP	Global Service Partners	Pratt & Whitney's aftermarket division
GTF	Geared Turbofan	Pratt & Whitney's newest engine redesign
HPT	High-Pressure Turbine	A section of the traditional gas-turbine engine downstream of the combustion chamber
IT	Information Technology	Business function responsible for information systems and technology
ME	Manufacturing Engineer	Individuals that provide engineering support to production
MRO	Maintenance, Repair and Overhaul	A term for the aftermarket operations at Pratt & Whitney
MRP	Manufacturing Resource Planning	A type of software used for managing resources for the manufacturing of products
MTS	Machine Tool Services	The organization within NBPC responsible for the design of tooling and re-manufacture of cutting tools
NBPC	North Berwick Parts Center	The subject of the research study; a Pratt & Whitney module center in Maine
OEM	Original Equipment Manufacture	In aerospace, the business of manufacturing new aircraft and components
PW	Pratt & Whitney	A leading manufacturer of aircraft engines
QDR	Quadrennial Defense Review	A process that occurs once every four years wherein the US government reviews the defense budget
SEC	Securities and Exchange Commission	Branch of the US government that oversees the financial reporting of corporations

10 Bibliography

Argote, L., & Epple, D. (1990). Learning Curves in Manufacturing. *Science*, *247*, 920-924.

Elbertsen, L., Benders, J., & Nijssen, E. (2006). ERP use: exclusive or complemented? *Industrial Management and Data Systems*, *106* (6), 811-824.

Fine, C. H., & Freund, R. M. (1990). Optimal investment in product-flexible manufacturing capacity. *Management Science*, *36* (4), 449-466.

Fuss, C., & Vermeulen, P. (2008). Firms' investment decisions in response to demand and price uncertainty. *Applied Economics*, *40*, 2337-2351.

Geng, N., & Jiang, Z. (2009). A review on strategic capacity planning for the semiconductor manufacturing industry. *International Journal of Production Research*, *47* (13), 3639-3655.

Halsall, D. N., Muhlemann, A. P., & Price, D. H. (1994). A review of production planning and scheduling in smaller manufacturing companies in the UK. *Production Planning & Control: The Management of Operations*, *5* (5), 485-493.

Ho, C.-F., Wu, W.-H., & Tai, Y.-M. (2004). Strategies for the adaptation of ERP systems. *Industrial Management and Data Systems*, *104* (3/4), 234-251.

(2011, August). Interview of CT Operations Employee. (M. Reveley, Interviewer)

(2011, July). Interview with CEPO Manager. (M. Reveley, Interviewer)

Kefeli, A., Uzsoy, R., Fathi, Y., & Kay, M. (2011). Using a mathematical programming model to examine the marginal price of capacitated resources. *International Journal of Production Economics*, *131*, 383-391.

Li, H., & Meissner, J. (2011). Competition under capacitated dynamic lot-sizing with capacity acquisition. *International Journal of Production Economics , 131*, 535-544.

Liao, S. S. (1988). The Learning Curve: Wright's Model vs. Crawford's Model. *Issues in Accounting Education , 3*, 302-315.

Lockheed Martin Corporation. (2002). *10-K Annual Report 2001.* Retrieved March 13, 2012, from http://www.lockheedmartin.com

Occhino, T. J. (2000). Capacity Planning Model: The Important Inputs, Formulas, And Benefits. *IEEE/SEMI Advanced Semiconductor Manufacturing Conference*, (pp. 455-458).

Ozturk, O., Coburn, M. B., & Kitterman, S. (2003). Conceptualization, design and implementation of a static capacity model. *Proceedings of the 2003 Winter Simulation Conference*, (pp. 1373-1376).

Pratt and Whitney. (2011). *Pratt & Whitney: An Overview.* Retrieved March 13, 2012, from http://www.pw.utc.com

Samadi, N. (2011, October). *IBISWorld Industry Report 33641A - Aircraft, Engine & Parts Manufacturing in the US.* Retrieved January 4, 2012, from www.ibisworld.com.

Tenhiälä, A. (2011). Contingency theory of capacity planning: The link between process types and planning methods. *Journal of Operations Management , 29*, 65-77.

Terwiesch, C., & Bohn, R. E. (2001). Learning and process improvement during production ramp-up. *International Journal of Production Economics , 70*, 1-19.

United Technologies Corporation. (2012, February 9). *10-K Annual Report 2011.* Retrieved March 13, 2012, from http://www.utc.com

Vranakis, S. K., & Chatzoglou, P. D. (2012). A Conceptual Model for Machinery & Equipment Investment Decisions. *International Journal of Business and Management , 7* (1), 36-58.