

MANAGEMENT OF SUPPLY CHAIN COSTS ASSOCIATED WITH PART PROLIFERATION

By

DAVID GEREZ

Bachelor of Science in Electrical Engineering
University of Texas at Austin, 1993

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

Master of Science in Mechanical Engineering
and
Master of Science in Management

In Conjunction with the Leaders for Manufacturing Program at the
Massachusetts Institute of Technology
June 2003

© 2003 Massachusetts Institute of Technology. All rights reserved.

Signature of Author _____

Department of Mechanical Engineering
Sloan School of Management
May 2003

Certified by _____

Kamal Youcef-Toumi
Professor of Mechanical Engineering

Certified by _____

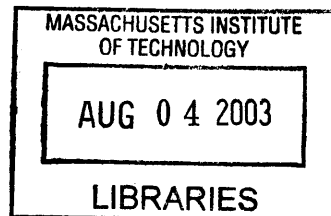
Roy E. Welsch
Professor of Statistics and Management, Director CCREMS

Accepted by _____

Ain A. Sonin
Chairman, Committee on Graduate Students
Department of Mechanical Engineering

Accepted by _____

Margaret Andrews
Executive Director of Masters Program
Sloan School of Management



BARKER

MANAGEMENT OF SUPPLY CHAIN COSTS ASSOCIATED WITH PART PROLIFERATION

By
DAVID GEREZ

Submitted to the Department of Mechanical Engineering and the Sloan School of Management on May 18, 2003 in partial fulfillment of the requirements for the degrees of Master of Science in Mechanical Engineering and Master of Science in Management

ABSTRACT

Dell's mass-customization manufacturing model, which allows the company to offer cost-effective, individually configured computer systems, has been a key source of competitive advantage. The expansion in the catalog of customer options, however, has led to increased supply chain costs. A three-part approach is presented for limiting these costs without adversely affecting the customer's experience.

The most important task is to quantify the total cost of the decision to add or retain a single part. The resulting model reflects the impact on the entire supply chain of maintaining a part, incorporating costs that are embedded in a variety of cost centers at the manufacturer and its suppliers. While it is impossible to precisely quantify this cost, the model provides sufficient accuracy to serve as a useful decision tool.

The second step is to apply the model to a sample of real parts. This helps to establish procedures for selecting parts that should be removed, overcoming problems such as the allocation of cost and revenue. The second objective of the case study is to establish the potential savings from addressing the part-proliferation problem.

Finally, the cost model and procedures of the first two stages are used as inputs to a set of management procedures that identify and discontinue unnecessary parts. These procedures take into account not only financial considerations, but also qualitative concerns such as strategic planning. This process ensures that the parts catalog is systematically managed to maintain the optimal offering of customer options.

Thesis Supervisors:

Kamal Youcef-Toumi, Professor of Mechanical Engineering

Roy E. Welsch, Professor of Statistics and Management, Director CCREMS

ACKNOWLEDGEMENTS

I would like to thank the Leaders for Manufacturing Program at MIT for its support of this work and for the learning opportunities it has provided.

I would also like to thank my thesis advisers, Roy Welsch and Kamal Youcef-Toumi, for their guidance and assistance throughout this project.

Finally, I would like to thank everyone at Dell who made this possible and helped to provide an invaluable learning experience.

TABLE OF CONTENTS

List of Figures	9
List of Tables	10
Definitions.....	11
1 Introduction.....	13
1.1 Background.....	13
1.1.1 Company History.....	13
1.1.2 Dell's Manufacturing Model.....	13
1.2 Motivation.....	20
1.3 Thesis Objectives	22
1.3.1 Alternative Approaches	22
1.3.2 Thesis Approach	24
1.4 Thesis Deliverables.....	26
1.5 Thesis Overview	26
2 Cost Model.....	28
2.1 Supply chain Process	28
2.1.1 Part Number Creation Process	31
2.1.2 Planning Process	31
2.1.3 Manufacturing Process.....	32
2.1.4 Service Process	33
2.2 Model Application	34
2.3 Assumptions.....	36
2.4 Cost Drivers	40
2.4.1 WWP.....	41
2.4.2 IT.....	41
2.4.3 CoC	41
2.4.4 Supply Chain Management.....	42
2.4.5 Hub.....	42
2.4.6 Logistics.....	50
2.4.7 Factory	51
2.4.8 Service.....	51
2.5 Model Results	53
2.6 Model Approval.....	59
2.7 Summary	59
3 Case Study	61
3.1 part-selection methodology.....	61
3.1.1 Allocation Problem	61
3.1.2 Proposed Allocation Methodology	65
3.2 Business Case.....	70
3.3 summary.....	72
4 Management Process	74
4.1 Existing Process.....	74
4.2 Proposed Process	75
4.2.1 Profile, Planning, and Implementation Phases	76
4.2.2 Deployment and Management Phases	77

4.2.3	Service Phase	81
4.3	summary.....	82
5	Summary and Recommendations	83
	References.....	86

LIST OF FIGURES

Figure 1: Continuum of manufacturing strategies	14
Figure 2: Part-proliferation problem.....	21
Figure 3: Part-proliferation trend.....	21
Figure 4: Generalized approach to optimizing market variety	24
Figure 5: Overview of supply chain processes	30
Figure 6: Inventory effect of offering 2 parts instead of 1	43
Figure 7: Pareto diagram of non-scaleable costs	54
Figure 8: Pareto diagram of scaleable costs.....	55
Figure 9: Simplified example of BOM explosion.....	62
Figure 10: Extended example of BOM explosion	64
Figure 11: Allocation of discounts.....	65
Figure 12: Databases for case study.....	68
Figure 13: Case study results (not to scale)	70
Figure 14: Case study results	71
Figure 15: Proposed management process.....	76

LIST OF TABLES

Table 1: Cost model scenarios	37
Table 2: Timing-related differences between scenarios	56
Table 3: Part-type-related differences between scenarios	57
Table 4: Cost model coefficients	58

DEFINITIONS

BOM	Bill of Materials
CAPM	Capital Asset Pricing Model
Core Team	Team incorporating all functions involved in a product's life cycle, e.g. engineering, marketing, and manufacturing. Responsible for all product decisions in the PRP. See Section 4.1.
CoC	Center of Competence. Organization with various roles, the most important in this context being the definition of pricing. See Section 2.1.
DAO	Dell Americas Operations. The scope of the project is limited to Dell's manufacturing operations in the US and Brazil.
DSI	Days' Sales in Inventory. A measure of inventory levels in terms of the average daily sales of the part. This inventory is stored in the hub.
Hub	A warehouse, operated by a third-party contractor, used to store safety stock. Inventory in the hub is kept on suppliers' books until it passes the factory loading dock a few hours before assembly. Also called "SLC." See section 2.1.2.
Part	See "Part number."
Part Number	Also called "part." Can be a subassembly consisting of individual, lower-level part numbers. Because these subassemblies are not stored separately, they do not incur significant cost. Part number in this thesis therefore denotes the lowest level in the BOM tree, i.e. a purchased part.

Peripheral	Any component or accessory of a base system for which a customer is offered a choice of options. Examples include hard drives, memory, and monitors.
PG	Product Group. Organization primarily responsible for product development and marketing. See Section 2.1
PRP	Phase Review Process. Set of procedures and guidelines developed by Dell to manage all products throughout their life cycles. The PRP is based on the concept of the core team, in which members from all important functions, e.g. engineering, marketing, and procurement, are jointly responsible for all decisions concerning their product. See Section 4.1.
SCM	Supply Chain Management (See Section 2.1)
SLC	Supplier Logistics Center. See “Hub.”
SKU	Stock Keeping Unit. Items sold to customers, which break down into part numbers in the BOM.
WACC	Working Average Cost of Capital. In this thesis, used to determine the holding cost of inventory.
WWP	Worldwide Procurement. Organization responsible for strategic supply chain planning and for the negotiation of supplier contracts. See section 2.1.

1 INTRODUCTION

1.1 BACKGROUND

The purpose of this background section is to provide a context in which to place the problems addressed by this thesis, rather than presenting a comprehensive survey or history of manufacturing. A brief history of Dell is followed by a description of how Dell's strategy of mass customization fits within a broader manufacturing and supply chain framework, and the issues that arise from this strategy.

1.1.1 Company History

Since its modest beginnings in 1984 in Michael Dell's dormitory room at the University of Texas, Dell has grown into a leading computer manufacturer, with a global presence and annual sales in excess of \$33 billion. This astonishing success was based on Dell's Direct Model, which enabled it to deliver cost-effective PCs manufactured to each customer's individual requirements. While other manufacturers also implemented similar models, they were unable to match Dell's efficiency and the resulting market dominance.

1.1.2 Dell's Manufacturing Model

For most of the 20th century, the mass production of standardized products was the dominant manufacturing model. The Ford Model T, which was available in any color as long as it was black, epitomizes the extreme of this type of manufacturing system. At the opposite end lies pure customization, in which every order is designed and built specifically to the individual customer's requirements. Examples of this approach include traditional crafts industries like jewelry, as well as large-scale construction

projects. Most industries lie somewhere along the continuum between these two extremes; Lampel and Mintzberg divide this continuum into the five stages shown on Figure 1:¹

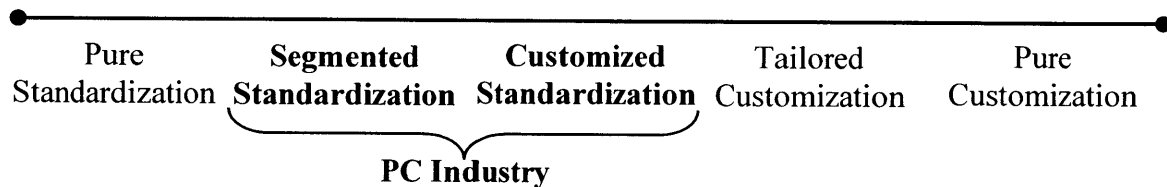


Figure 1: Continuum of manufacturing strategies

Traditionally, PC manufacturers have followed the segmented-standardization strategy, building systems to stock in large volumes for general customer segments, and selling the systems through wholesalers and distributors; any customization is accomplished through rework, either by distributors or at the point of sale. This results not only in the quality problems and added costs inherent in reworking products, but also in significant losses from a large, rapidly depreciating inventory (a general rule of thumb is that PC inventory depreciates by 1% per week).

Dell's model, on the other hand, is based on customized standardization, in which systems are individually tailored to each customer's requirements during final assembly. This type of strategy is more generally known as mass customization. While it is often promoted as a panacea for any manufacturing company, mass customization can only be

¹ Joseph Lampel and Henry Mintzberg, "Customizing Customization," *MIT Sloan Management Review*, Fall 1996, 21-30.

effective in those industries that satisfy the three primary requirements defined by Zipkin² – elicitation, logistics, and process flexibility – and can integrate these elements with a reliable communication flow. As described in the following sections, Dell has successfully implemented strategies that fulfill these requirements, leading to its current domination of the computer industry.

1.1.2.1 Elicitation

Dell's strategy of dealing directly with customers gives it an advantage in satisfying the first requirement for mass customization: *elicitation*, or a system of obtaining information from customers about their requirements for each order. The Direct Model provides Dell first-hand knowledge about customer needs, avoiding the misinterpretations and delays that can occur when this information has to travel through an indirect distribution channel. By controlling the customer's selection process, primarily on the website but also through telephone sales, Dell is able to determine exactly how the order should be configured, as well as coaxing the customer toward Dell's preferred alternative, avoiding the selection of scarce components. In addition to providing information about the customer's individual configuration, this direct relationship allows Dell to gauge the overall market, adjusting its product offering and pricing points in response to competitive conditions.

² Paul Zipkin, "The Limits of Mass Customization," *MIT Sloan Management Review*, Spring 2001, 82.

1.1.2.2 Logistics

The second necessary capability is *logistics*: the transportation and tracking of products through the entire value chain, from supplier to customer. In Dell's case, this process includes three main stages:

- The shipment of inbound material by third-party logistics providers, primarily through bulk airfreight, ocean cargo, and trucks.
- The movement and tracking of bar-coded components within the assembly line.
- Shipment directly to customers through package-shipment companies like FedEx and United Parcel Service (UPS).

Because these logistics solutions are available to any company, they do not constitute a competitive advantage for Dell.³ They do, however, constitute a necessary, though not sufficient, component of a mass-customization manufacturing model.

1.1.2.3 Process Flexibility

The third main requirement for mass customization is a manufacturing process with the flexibility to individually configure each system in a cost-effective manner. While improvements in design modularity, information technology, and fabrication methods have been important contributors to Dell's leadership in process flexibility, the real key has been the selection and implementation of its supply chain strategy, which is a hybrid

³ While logistics do not provide a conclusive competitive advantage, Dell's efforts to make its logistics processes more efficient have resulted in significant cost savings, providing an incremental contribution to Dell's ability to compete on price.

between the two types of inventory and material-flow systems: push and pull. By definition, this system only affects inbound inventory, as the build to-order model implies that there is no outbound inventory.

Push

In a push system, materials are produced and procured in anticipation of forecast demand, and are held in inventory until they are used in final assembly. These systems rely on two main inputs: the bill of materials (BOM) to explode anticipated demand for the end product into forecasts for individual components, and the forecast, also called the master production schedule. The most important of the five basic properties of forecasts defined by Nahmias is that they are always wrong.⁴ As a result of this unavoidable inaccuracy, the only way to ensure adequate availability of the manufacturer's end product is to maintain extensive safety stocks, tying up working capital and incurring storage charges.

Push systems have historically been more prevalent than pull systems, and define the strategies of some of Dell's competitors. While this ensures a high level of material availability, it also results in high costs; these costs are particularly pronounced in the case of rapidly depreciating components likethose used in computers . The upstream stages of Dell's supply chain are also push systems, as the long lead times inherent in the processing of many of these inputs, such as semiconductors, require the maintenance of extensive inventories. At a point further downstream, the supply chain becomes a pull system as Dell draws down the inventory in response to customer orders.

⁴ Steven Nahmias, *Production and Operations Analysis*, 4th ed. (New York: McGraw-Hill/Irwin, 2001), 57.

Pull

In a pull system, on the other hand, materials are manufactured and procured only in response to actual orders. One category of pull system is a Just in Time (JIT) scheme, in which inventory is kept to a minimum (or is nonexistent) and the only material present is that needed for systems currently in process. Although the term “lean manufacturing” is often used interchangeably with “JIT”, it should be properly viewed as a superset of JIT; Womack and Jones classify pull as the fourth of the five principles of lean thinking that lead to higher productivity through the elimination of waste:⁵

1. *Specify Value* in terms of actual customer needs.
2. *Identify the Value Stream*, from raw material to end customer.
3. Organize the process *flow* to match the value-creating activities for each product.
4. *Pull* product through the system only as customers need it.
5. Continuously improve processes to achieve *perfection*.

As mentioned in the preceding section, Dell’s supply chain is a hybrid between a push and a pull system. The point at which this transition from push to pull takes place is not entirely clear, as it depends on Dell’s level of commitment to buy the inventory (see Section 2.1.2), but, because the transition is upstream of Dell, the company generally considers its supply chain to be a pull system. It is important to keep in mind, though, that the requirements of Dell’s environment and business model actually call for a hybrid

⁵ James P. Womack and Daniel T. Jones, *Lean Thinking: Banish Waste and Create Wealth in Your Corporation* (New York: Simon & Schuster, 1996), 9-28.

system. The first factor precluding a pure pull system is the combination of unpredictability in Dell's demand with extremely long manufacturing lead times for some components. While Dell partially reconciles this discrepancy by influencing demand through pricing incentives, some safety stocks will always be needed somewhere in the system. Additionally, JIT requires very close relationships with a small number of suppliers, based on open communications and a high degree of trust;⁶ this type of approach is generally not conducive to multiple sourcing. Dell, on the other hand, uses multiple sourcing to obtain price reductions, and has somewhat adversarial relationships with many suppliers. Finally, JIT generally requires geographic proximity with suppliers in order to limit transit times. In Dell's case, though, the trend is to increasingly seek lower-priced Asian components, resulting in large in-transit inventories.

Given its market and operational constraints, Dell has selected a hybrid push-pull inventory system that allows it to fulfill the third requirement of mass customization: process flexibility.

1.1.2.4 Communications

Finally, the three main requirements of mass customization (elicitation, logistics, and process flexibility) must be integrated into an effective system. Because accurate delivery of each uniquely configured system involves multiple functional organizations, such as sales, engineering, manufacturing, and logistics, continuous communications across the organization are necessary. Dell's matrix organization, with its extensive

⁶ Nahmias, 392.

network of cross-functional links, ensures that these communications take place, effectively binding the three main components of a mass-customization model.

This mass-customization model has allowed Dell to tailor each system to the individual customer's requirements while remaining cost competitive, leading to its current dominance of the PC market. While generally highly successful, this approach has led to the problem described in the following section.

1.2 MOTIVATION

Figure 2 illustrates the problem addressed by this thesis. The key to Dell's continued growth in market share and profitability has been its build-to-order model, which allows it to offer cost-effective systems built to each customer's precise requirements. As the personal computer market has matured, customers have grown increasingly sophisticated, and have therefore demanded a growing array of options when customizing their systems; these options are in turn supported by a growing inventory of distinct part numbers. This proliferation of parts has led to higher part-management costs, such as purchasing headcount, inventory storage, and logistical expense.

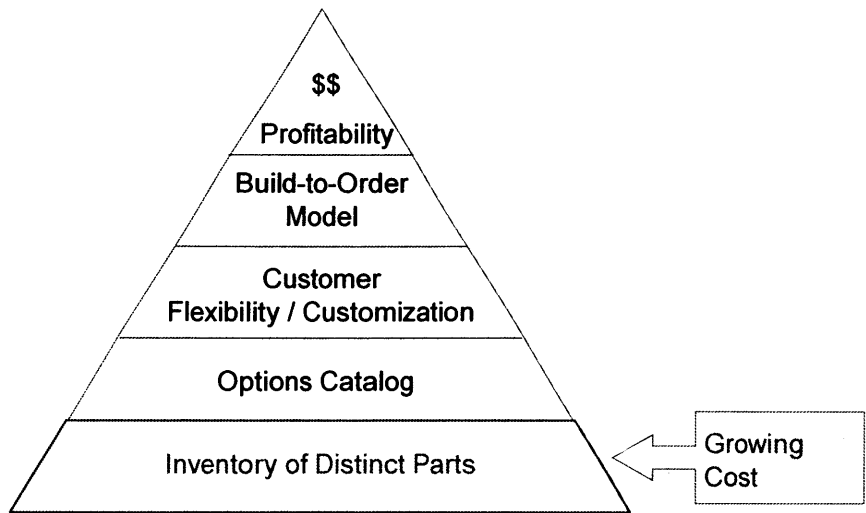


Figure 2: Part-proliferation problem

In addition to the long-term growth in the active-part count, there has been a tendency for parts to remain in the catalog even after demand had dropped below levels that would justify their existence. These parts slowly accumulated until somebody decided to initiate a part-cleanup effort, at which point the part count would suddenly drop, only to start slowly growing again. Figure 3 shows the resulting saw-toothed shape of the part count.

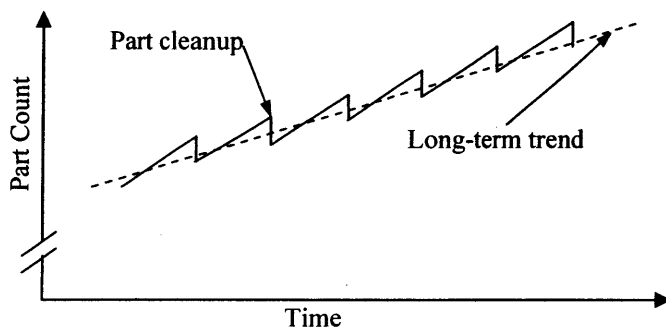


Figure 3: Part-proliferation trend

It is worth noting that no data on the extent of this problem were available; even determining the historical part count proved to be difficult, and in the end not worth the effort. It was therefore taken as a given, in developing this thesis, that part proliferation was a serious problem. The real severity of the problem is actually of secondary importance; more important is the fact that it was perceived to be a problem at Dell, and significant resources were being devoted to addressing various aspects of part proliferation. Therefore, simply having a systematic process in place to avoid needless proliferation would be of value, as it would prevent the waste of resources on inefficient ad hoc efforts.

1.3 THESIS OBJECTIVES

This thesis describes a method of striking a balance between controlling the cost of part management by limiting unnecessary and unprofitable options, and offering sufficient options to maintain a positive customer experience and ensure continued sales growth for Dell. While the costs of part management should certainly be limited where possible, it is essential not to erode a source of Dell's competitive advantage in the process.

1.3.1 Alternative Approaches

The open-ended nature of the initial project mandate necessitated a lengthy discovery phase before arriving at the most appropriate approach for solving the problem. Of the numerous other ideas that were considered and not selected, the two most important are presented in this section, with two objectives: to demonstrate the constraints that led to the selection of the solution proposed in this thesis, and to provide information that might be of use in addressing the part-proliferation problem under different conditions.

A generalized solution might have followed the recommendations given by Child et al⁷ to optimize market variety. This approach, shown on Figure 4, calls for a comprehensive review of the product offering. While this framework, if implemented, would certainly lead to improvements in the product offering, it was not selected because it reflects, to a certain degree, decisions that were already being made in other contexts at Dell. Moreover, its broad scope made its implementation unrealistic, both because of the limited duration of this project and the political issues associated with completely overhauling the product offering.

⁷ Peter Child, Raimund Diederichs, Falk-Hayo Sanders, and Stefan Wisniowski, "SMR Forum: The Management of Complexity," *MIT Sloan Management Review*, Fall 1991, 74-75.

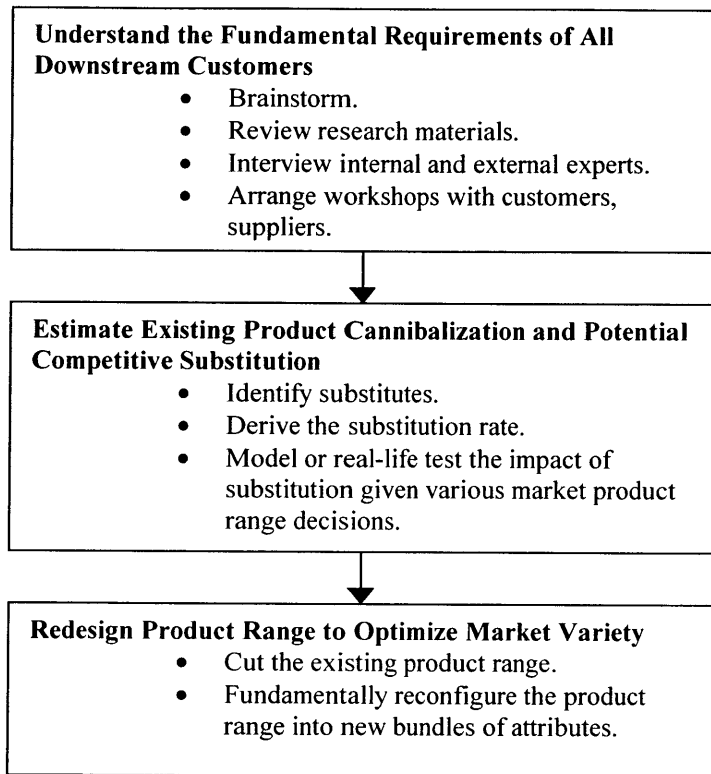


Figure 4: Generalized approach to optimizing market variety

Another solution would have been to make the management of every single part more efficient, rather than changing the part count. While this would be desirable, the high stakes involved – the Dell Americas procurement organization buys over \$10 billion a year in parts – would make far-reaching changes to Dell’s supply chain strategy highly risky, and unlikely to happen during the short duration of this project.

1.3.2 Thesis Approach

After considering the preceding alternatives, the selected approach was to develop a set of management procedures to systematically review all customer options (also called peripherals) both before their introduction and periodically during their production life, and eliminate those that are deemed unnecessary. The decision would take place at a

meeting, in which both quantitative and qualitative concerns are taken into account; these concerns include customer experience, long-term strategic market placement, and the financial benefit of introducing or keeping the part. Because the financial benefit is the only one of these inputs that was not already available in some form, the focus of this thesis is the development of a methodology for performing a cost/benefit analysis of each part. The cost is found by quantifying all of the indirect, hidden costs that are incurred in introducing a new part, incorporating the results into a comprehensive model. The benefit is found by assigning a share of the revenue from the high-level assemblies that customers buy to each individual part number.

It would be valuable to validate the proposed approach after its implementation in order to ensure that it really does solve the problem. Unfortunately, the nature of the part-proliferation problem precludes any conclusive validation. It would be possible to track the part count after the implementation of the part-management process, but this would rely on some arbitrary reduction target that did not represent any real measure of success. Even if such a target were selected with confidence, the other forces driving the part count, such as consumer demands or supplier-selection considerations, would confound any gains from this management process. It would also be unrealistic to directly measure financial savings, as the cost of managing each part varies over time, for example as purchasing procedures become more efficient or inventory is trimmed, again confounding any gains from the management process presented here.

Given these limitations to stipulating a verifiable set of criteria for success, it is more relevant to use more subjective, open-ended criteria. This process provides a forum that

forces the review of part numbers, ensuring that no parts are forgotten and continue incurring costs long after their useful lives. By forcing the organization to actively consider each part, the procedures ensure that at least the most inappropriate parts will be eliminated. Moreover, the cost model provides visibility into the cost of managing a part, ensuring better-informed decisions in the future. Even in the unlikely event that it did not result in the elimination of parts, it would ensure that the indirect costs of the decision to add or retain a part were considered in a rational, systematic manner. Finally, the comprehensive approach taken in this thesis will ensure that fewer resources are wasted on ad hoc efforts to eliminate part numbers. The benefits, though unquantifiable, will therefore be twofold: a reduction in part numbers and the associated costs of managing them, and the preemption of wasteful localized efforts to address the problem.

1.4 THESIS DELIVERABLES

Due to the limited duration of this project, it was not feasible to fully implement all of the recommendations of this thesis. It was agreed ahead of time that the deliverables would consist of a set of recommendations to Dell management, along with a reasonable business case for implementing those recommendations. Additionally, a model of the costs of managing a part number was to be completed and delivered.

1.5 THESIS OVERVIEW

This document will begin with the presentation, in Section 2, of a model that comprehensively quantifies the incremental costs that are incurred upon the addition, or retention, of a single part number. The cost model is the focus of this thesis for two reasons: it took the bulk of project time, and it is the only phase that was expected to be

completed during the project timeframe. Because an understanding of the process of managing a part is a prerequisite to understanding the associated costs, Section 2 begins with an overview of this process, and continues with a description of the cost model. Section 3 presents the results of a case study that served two purposes: to establish a suitable method of allocating revenue and cost to each part, and to establish the business case for implementing a management process to limit the proliferation of parts. Section 4 describes this management process, which uses the cost model of Section 2 and the revenue-allocation methods of Section 3 as inputs. Finally, this thesis concludes with a brief summary and the recommendations of Section 5.

2 COST MODEL

The first and most time-consuming phase of the project was to find the cost to Dell of managing a single part number, i.e. the incremental lifetime cost to the company of adding one part number (or, conversely, the savings resulting from the removal of one part number). This phase was divided into six tasks, each of which is described in its own subsection:

1. Map the supply chain process to understand where costs are incurred
2. Evaluate the model's ultimate application, to judge the required accuracy
3. Make the necessary simplifying assumptions
4. Quantify all the cost drivers
5. Consolidate the most important cost drivers into a single model
6. Certify the model as Dell's official cost of managing a part

Because the results are based on a series of assumptions, it is unrealistic to provide a single value for the cost. Therefore, the model outputs a range of values that incorporates the propagation of all best-case and worst-case assumptions, providing a more realistic estimate of the cost.

2.1 SUPPLY CHAIN PROCESS

The first step in developing a cost model is to understand the work performed in creating and managing a part, and the organizations that are affected by that work. A new part incurs cost through the four processes shown on Figure 5: part-number creation, planning, manufacturing, and service. Note that the figure is a simplification: given

Dell's highly-networked structure, there are actually countless communication and feedback paths between the organizations.

Two organizations are involved throughout the four processes: IT⁸ and WWP. IT provides the information-management infrastructure to maintain product files and part databases. WWP (Worldwide Procurement) is responsible for the high-level design and strategic management of the entire supply chain. Its functions include establishing the supplier base, negotiating supplier contracts and pricing, and allocating a percentage of Dell's total purchases to each supplier. WWP has a strategic role, and is not involved in the day-to-day purchasing and expediting of materials.

⁸ Note that IT costs are in reality distributed among many organizations. Conceptually, though, it is useful to aggregate them into the single entity shown on Figure 5.

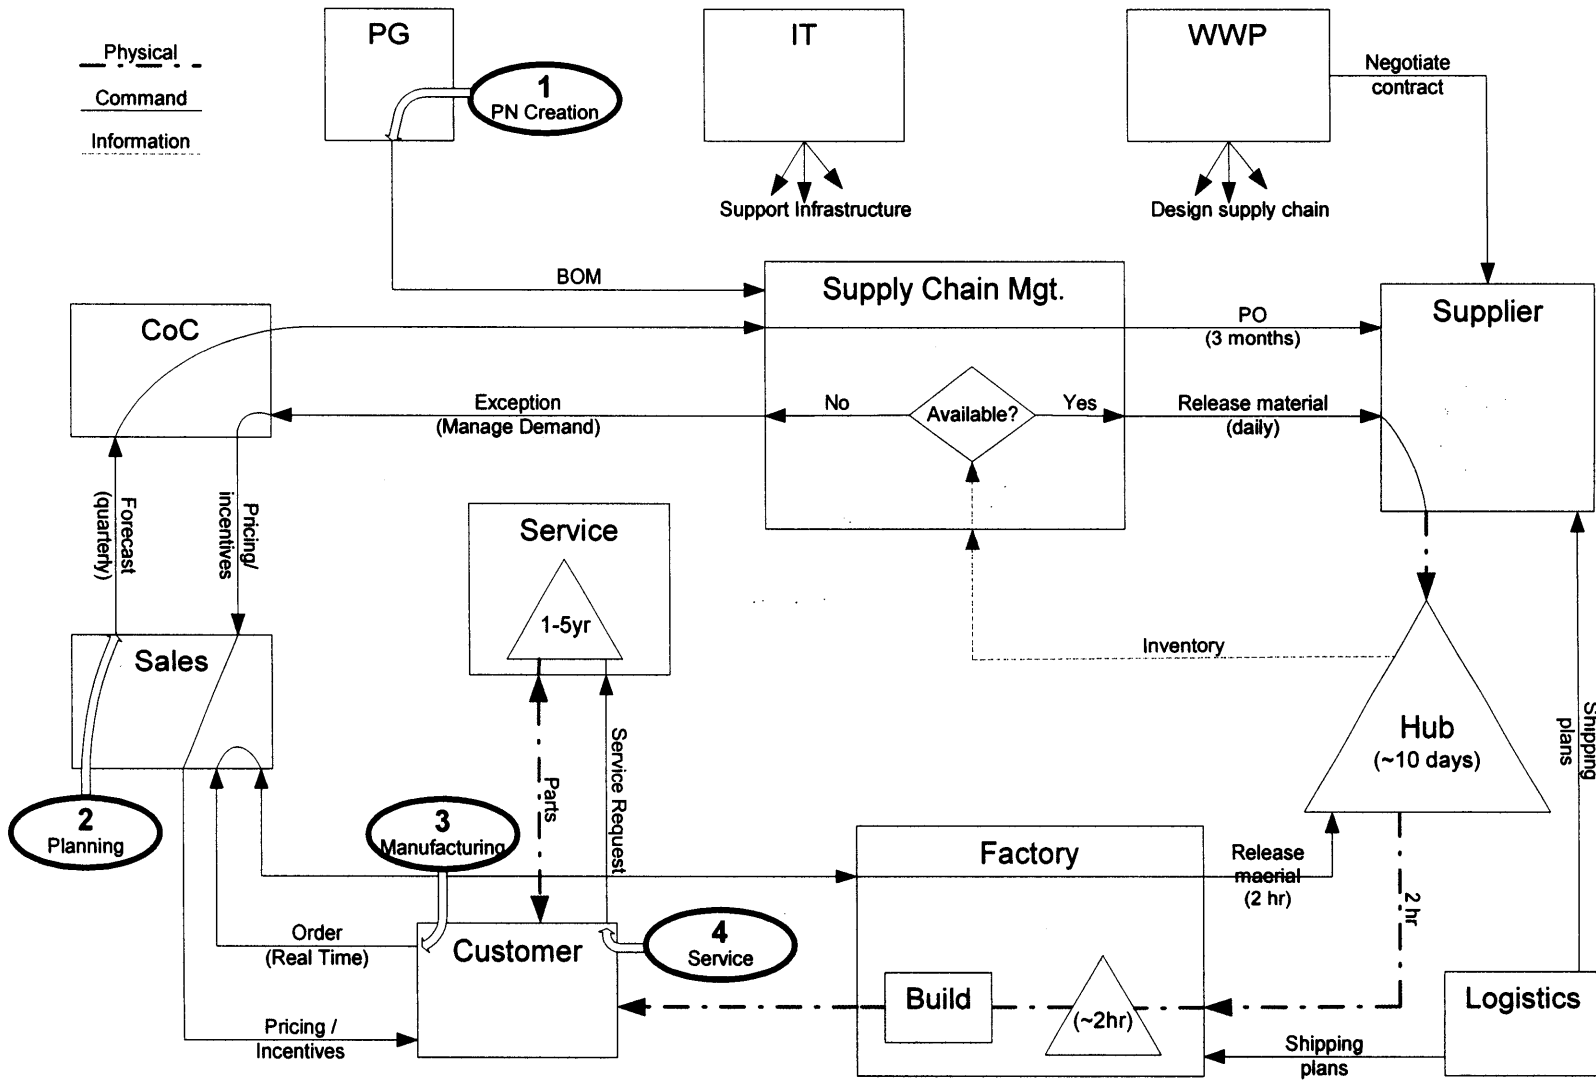


Figure 5: Overview of supply chain processes

2.1.1 Part Number Creation Process

A new-part introduction results in significant work by engineering, marketing, and other organizations. Because these upstream costs vary greatly from part to part, a model would not have accurately represented the true cost of any given part. More importantly, these costs are generally visible and already accounted for, so there is no need to model them. These upstream costs were therefore left outside the scope of the model. After these processes have been completed, the product file, including its Bill of Materials (BOM), is distributed throughout the organization, including Supply Chain Management (SCM).

2.1.2 Planning Process

After a part is created, it becomes part of the planning process. The long lead times inherent in the computer supply chain force Dell to forecast demand and order parts with a far longer horizon than the two hours used for stocking the factory. This quarterly forecasting process begins with Sales, which passes the forecast to the Center of Competence (CoC), where it is adjusted and passed on to SCM. After further hedging, SCM issues purchase orders (POs), which are used by suppliers to schedule their production up to three months ahead of time. Although Dell does not take ownership of incoming material until it physically enters the factory, these purchase orders represent a contractual commitment to purchase a given amount. While Dell generally applies pressure on suppliers to absorb excess material when the forecast proves to be too high and efforts to stimulate consumer demand are ineffective, suppliers have on occasion used this contractual obligation to force Dell to absorb the loss.

2.1.3 Manufacturing Process

While Dell orders material to forecast during the planning process, it pulls material to order as part of the manufacturing process. This process starts with a customer order, which is processed by Sales before being passed to the factory. The factory then updates its manufacturing plan, and schedules production of the system within the five-day order-to-ship target. A few hours before it is ready to assemble the system, the factory releases material from the hub, which sends trucks to the factory every two hours. The hub is a warehouse, operated by a third party under contract to Dell, in which suppliers hold inventory until it is pulled by the factory; all inventory in the hub is supplier owned. Dell takes ownership as the material passes the factory door, and places the boxes on the pick-to-light racks of each assembly line. After assembly, the system is shipped to the customer through the carrier selected by Logistics. The Logistics group also negotiates contracts with carriers for shipments from suppliers to the hubs.

At the same time, SCM continuously monitors inventory in the hub, and issues daily orders to suppliers to release part of the material that had been ordered when the PO was issued. While SCM has traditionally managed inventory to maintain 10 days' sales in inventory (DSI) in the hubs, suppliers, which must cover the carrying costs, frequently save money by maintaining less inventory while still ensuring adequate service levels, based on their experience with the uncertainty of Dell's demand. Apparently, suppliers are managing inventory to maintain a desired service level, while Dell generally tries to maintain a fixed level of inventory. While there was an effort underway to explicitly change the focus of SCM's inventory-management efforts, this had not been implemented at the time of writing.

If the parts necessary to fulfill customer orders are not available from the supplier, and expediting efforts fail, SCM responds by managing customer demand to match the available supply. In this case, a signal is sent to the CoC, which generates pricing incentives or promotions to try to steer customers away from the scarce components. This ability to quickly and precisely manage the demand and supply of every part in the computer is seen as one of the keys to Dell's success.

2.1.4 Service Process

After selling a computer, Dell must maintain a spare-parts inventory through the duration of the longest warranty, normally three to five years after the product's introduction. Dell can continue to buy some of these components during the warranty period, but it must buy a lifetime inventory of those that are discontinued. This inventory is also replenished by parts that have been refurbished after being removed from repaired systems. The service inventory thus consists of a small safety stock of parts that are available on the market, a large long-term stock of discontinued parts, and a number of used parts.

2.2 MODEL APPLICATION

Before examining the cost incurred by each of the organizations shown in Figure 5, it is important to consider the model's ultimate application, which dictates how precisely these costs have to be calculated. As described in section 2.3, accurately accounting for all of the differences between parts would have required a highly complex model or collection of models. Because time constraints made this approach impractical, it was necessary to strike a balance in the model between accuracy and simplicity.

A lack of time is not, however, a valid reason to avoid developing a complex model, if it is in fact necessary. A more important issue is the level of complexity that is necessary or appropriate for the model's ultimate application. A number of considerations led to the decision to keep the model simple, at the expense of some accuracy.

First, the cost model will only be used to identify parts for further evaluation, not as an automatic filter to eliminate unprofitable parts. Subjective concerns, such as perceived customer needs or long-term strategic considerations, will ultimately carry more weight in the decision. There is no point quantifying the cost with high precision when the resulting number will be just one input – albeit an important one – to a fundamentally subjective decision. The case of 64MB memory modules provides a good example of such a situation. These modules, which had already been superseded by several generations of technology, commanded very low margins, and appeared to be losing money. However, these low-cost parts allowed Dell to advertise low-end systems for only \$599, which was important not only because it helped reinforce Dell's reputation for value, but also because many customers who were attracted by the low price were

subsequently convinced to buy to more expensive systems. In general, such subjective considerations turn out to be more important than the quantifiable costs and benefits of each individual peripheral.

Additionally, the model has to be easily understood if it is to be credible. Most people would not take the time to understand an intricate model, but would not be inclined to accept its output on faith if they did not know how it was derived. Given the alternative of a highly accurate model that is never used, it is preferable to use an imperfect model whose assumptions, though broad, are understood by the user.

Finally, the model has to be maintainable. Given the rapid pace of change at Dell, the inputs to the model, such as costs and headcounts, are likely to change over time, and should be periodically updated. Because significant resources are unlikely to be available for this task, the model should be simple and easy to update. The goal was for one person to spend approximately one day per year updating it.

For these reasons, the emphasis throughout the project was on keeping the model usable and simple, while reflecting – on average – the true cost of part management. When considering the simplifying assumptions that were incorporated into the model, it is important to remember that the choice was not between a highly accurate model and an approximate one. The real choice was between approximating the cost or ignoring it entirely. Thus, the general philosophy in developing the model was that it is better to account for costs approximately than ignore them exactly.

In other applications, the relative importance of these factors may be different, making a more complex (or simpler) model more appropriate. The thought process behind this tradeoff, though, will be similar.

2.3 ASSUMPTIONS

Once the general level of accuracy required from the model was understood, a number of simplifying assumptions could be made.

Most importantly, the model provides the cost for an average part, ignoring the unique circumstances that drive the costs of individual parts; these differences are illustrated by several examples:

1. A single buyer for mechanicals (e.g. cables, bezels, labels) handles ten times as many parts as a buyer for processors.
2. The standard deviation of the forecast error varies by as much as a factor of five for different memory parts.
3. Expediting charges for a single chassis part exceeded the average per-part amount by five orders of magnitude.
4. Dell is liable for Excess & Obsolete (E&O) charges on Dell-specific chassis parts, while commodities such as memory can always be sold on the spot market when they are no longer needed.

Accounting for all of these variables would have required a multi-dimensional matrix of independent models. By averaging out the costs over different part types, the problem was simplified to the four scenarios of the two-by-two matrix shown on Table 1. The columns define the timing of the decision: the first column shows the lifetime cost of

introducing a new part, while the second column shows the lifetime cost of retaining an existing part (at this time, some costs have been sunk and are therefore irrelevant). The rows define two types of parts in terms of the decision being made. In the first case, the decision is whether to offer a single part or nothing at all, so the decision to add the new part does not affect any existing parts. An example of this type of part is a Bluetooth card: if it is introduced, a certain number will be sold, but if it is not introduced the demand for Bluetooth cards will not affect demand for other peripherals. In the second case, the decision is whether to offer a selection between two parts or just a single part, satisfying the same total demand in either case. For example, every computer needs a graphics card, but Dell could either force every customer to buy the same model, or it could offer a selection of graphics cards. For simplicity, the top-left scenario is used as the baseline in quantifying the cost; this value is then adjusted to describe the other three scenarios.

	Before Part Introduction	After Part Introduction
Unique (1 part instead of none)	Model 1 (Baseline)	Model 1 Minus sunk costs
Replaceable (2 parts instead of 1)	Model 2	Model 2 Minus sunk costs

Table 1: Cost model scenarios

Another simplification is that all headcount costs are allocated to part numbers as linear functions, rather than the step functions that would more accurately represent the fact that the company does not pay for a fractional worker as the workload goes up, but rather allows the workload to accumulate until an employee is hired. This type of cost allocation is a standard practice, and yields a reasonable approximation.

The third assumption, described in detail in section 2.4.5.1, is that overall inventory levels are left unchanged when a new part is introduced in the replaceable scenarios of Table 1. Starting from the baseline that a given total demand is fulfilled with a single part, an additional part is introduced so that the same total demand is now divided between the two part numbers, each of which has its own independent demand function. The resulting decline in forecast accuracy can be dealt with in one of two ways: increasing inventory to maintain the previous service level, or keeping total inventory constant and accepting a decline in service level. The model assumes that the second approach is taken, but both methods are described in detail in section 2.4.5.

The fourth assumption is that the Dell corporate working average cost of capital (WACC) can be used to approximate the cost of the working capital tied up in inventory, called inventory holding cost in this thesis⁹. The WACC represents the opportunity cost of tying up capital in inventory instead of investing it in a project with a risk profile typical for Dell. In reality, some types of parts will represent a lower risk than a typical Dell investment because they can be resold on the secondary market, while others will carry a higher risk because they have highly unpredictable demand and limited resale value. Therefore, theoretically, a different risk-adjusted opportunity cost could be assigned to each class of inventory. This would clearly have been unrealistic, and was in any event unnecessary, as it would have only resulted in a very marginal improvement to a few line items of the total part-management cost. Given the facts that Dell at the time held large

⁹ In some contexts, holding cost would include expenses such as storage charges and obsolescence that are accounted for separately in this model.

cash reserves that earned a far lower return than the WACC, and that any additional cash for inventory would come from these reserves, one could also make the argument that the true opportunity cost of inventory is this (lower) rate of return on cash. This argument, however, ignores the fact that, because inventory carries a higher risk than cash, any investment in inventory should also yield a higher return. For these reasons, a single opportunity cost, the WACC, was used as a reasonable approximation of the opportunity cost of holding inventory.

Counting the number of parts, which might appear straightforward, actually required a number of assumptions. The basic problem is that the Dell part database contains many parts that do not add to the company's costs, for example obsolete parts that are no longer in the inventory, supplier-managed parts (e.g. screws and labels) that are not tracked by Dell, and customer bundles. The parts list was pared down to approximately 2,600 parts by applying several filter criteria to the database:

1. Only parts with an open purchase order or forecast over the following 20 weeks
2. No supplier-managed parts
3. No bundles or subassemblies
4. No DellPlus (customer-specific) parts

Six more assumptions were used in developing the model:

1. The average part is in production for one year.
2. The average part is then in the service process for an additional 2-3 years.
3. Only Dell Americas Operations are accounted for.
4. Sales expenses are excluded.

5. Only costs downstream of engineering and part qualification are modeled.
6. Tooling expenses are excluded.

The last two assumptions relate to costs that are generally known for each specific case; these were left out of the model because there is no need to estimate costs that can be counted directly.

2.4 COST DRIVERS

Only after mapping the work involved in managing parts and determining the level of accuracy required in this application was it possible to quantify each of the contributions to the total part-management cost. The problem was divided in terms of the organizations shown on Figure 5; each of the blocks was examined separately, and is described in a subsection below.

The general approach was to first interview people in each organization, focusing on how to quantify their own part-management costs. These interviews often provided new ideas for other possible cost drivers, each of which was carefully considered, even when it was apparent that it was not significant or even relevant. This approach helped ensure that the model was exhaustive (within the constraints described in section 2.3) and, more importantly, would be viewed within the organization as being exhaustive. After all of the potential costs had been gathered, they were ranked, and only the most important ones were included in the final model; this aggregation process is described in section 2.5.

Some costs do not vary with the annual consumption of a part; for example, the workload in the Worldwide Procurement organization is the same regardless of how many graphics

cards of a given model are shipped every year. Other costs vary with production volume, and are scaled by the annual material spending on the part; for example, the likelihood of losing a sale because a particular graphics card is out of stock increases with the percentage of computers ordered with that card.

2.4.1 WWP

The Worldwide Procurement (WWP) cost is a straightforward allocation of headcount spending. Because the department has global responsibilities, only the portion of the budget allocated to Dell Americas Operations was considered. This was then multiplied by 40 to 80%, which covers the range of reasonable assumptions about the proportion of employees' time that would be affected by the addition or removal of parts. Finally, this part-management budget was divided by the number of parts to arrive at the per-part cost.

2.4.2 IT

The cost of maintaining the IT infrastructure was found by combining the operating and capital budgets of the departments responsible for all of the parts databases, and dividing by the total number of parts stored in those databases. The cost per part ended up being trivial.

2.4.3 CoC

Center of Competence (CoC) cost is another headcount allocation. The total departmental budget was divided by the number of parts. Because the resulting cost was trivial, there was no need to make assumptions about the proportion of workload affected by part numbers.

2.4.4 Supply Chain Management

The cost of Supply Chain Management (SCM) is also headcount allocation; the total departmental budget was simply divided by the number of parts. Because SCM is responsible for the day-to-day management of suppliers and inventories, it was reasonable to assume that its entire workload was driven by the number of parts.

2.4.5 Hub

The hub was the most important cost driver, not only because of its impact on the total cost per part, but also because of the work required to quantify the cost. As described in section 2.1.3, the hub is a warehouse that is run by a third party under contract to Dell. All suppliers are expected to maintain a Dell-mandated level of safety stock in the hub, traditionally ten days' sales in inventory (DSI). Because inventory remains on suppliers' books until it physically enters the factory several hours before system assembly, suppliers have a strong incentive to limit hub inventory, and frequently maintain significantly less than ten DSI.

In the context of the cost model, the important issue is not the absolute level of inventory, but the change in inventory when a new part number is added. Figure 6 shows the incremental effect on inventory of offering a choice between two parts instead of offering a common part; this corresponds to the replaceable scenario of Table 1.

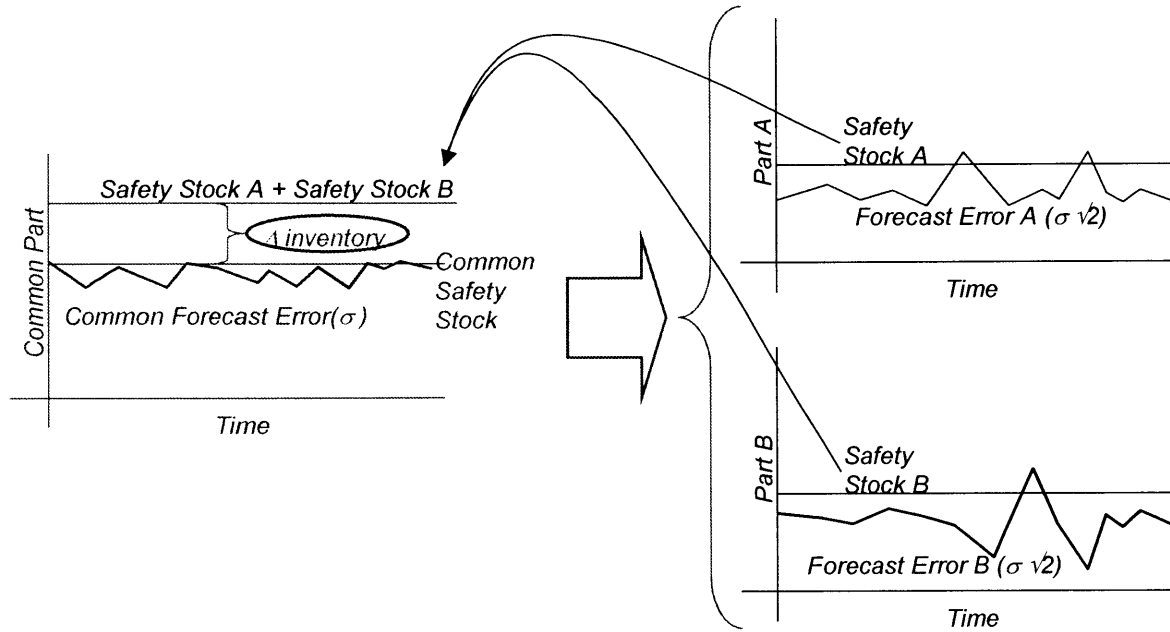


Figure 6: Inventory effect of offering 2 parts instead of 1

The demands for each of the separate parts (d_1 and d_2), as well as the common part (d_c), are unpredictable, and are assumed to follow a normal distribution. This demand consists of a forecast (F), which is assumed to be non-random because it is based on long-term trends instead of short-term demand fluctuations, and a normally distributed forecast error (e). Assuming that the introduction of the separate parts does not lead to any change in overall sales, the relationship between the common part's forecast error (e_c) and that of the separate parts (e_1 and e_2) can be found as follows:

$$d_c = d_1 + d_2 \quad \text{because } d \equiv F + e,$$

$$F_c + e_c = (F_1 + e_1) + (F_2 + e_2) \quad \text{because unchanging overall sales imply that } F_c = F_1 + F_2,$$

$$F_1 + F_2 + e_c = F_1 + F_2 + e_1 + e_2$$

$$e_c = e_1 + e_2$$

Because the variance of the sum of the individual forecast errors ($e_c = e_1 + e_2$) is equal to the sum of the variances of the individual errors, the standard deviation of the overall forecast error (σ_c) equals $\sqrt{\sigma_1^2 + \sigma_2^2}$. If demand between the two parts is evenly split, $\sigma_1 = \sigma_2 = \sigma_c / \sqrt{2}$.

The standard deviation of forecast error drives the safety stock required to maintain a given service level, as described in Section 2.4.5.1. Figure 6 shows that, in order to maintain a fixed service level, inventory levels must be increased when a common part (whose safety stock is driven by σ_c) is replaced by two parts (whose safety stock is driven by $2\sigma_1 = 2\sigma_2 = 2\sigma_c / \sqrt{2} = \sigma_c \sqrt{2}$).

There are two ways to deal with the increased forecast uncertainty arising from the addition of a new part: maintaining a constant service level by increasing inventory levels (as shown on Figure 6), and accepting a decline in service level while keeping inventory fixed. For the sake of simplicity, these choices are assumed to be mutually exclusive, so there can be no partial increase in inventory levels combined with a partial decline in service level. Additionally, it is assumed that a stockout of the part in question always results in lost sales, as customers will not accept a substitute part. Section 2.4.5.1 addresses the fixed-inventory assumption that was used in the final model, while Section 2.4.5.2 describes the constant service-level assumption that was not used but could serve as an alternative under other circumstances.

2.4.5.1 Fixed Inventory Assumption

Under the fixed-inventory assumption, Dell responds to the added uncertainty associated with an additional part by maintaining overall inventory levels fixed, and accepting a decline in service level. This decline in service level results in two types of cost: factory overtime and lost sales, both of which are scaled by the annual material expenditure on the part. These costs are different for the unique and replaceable scenarios of Table 1, and are described separately below.

Unique Scenarios

After a stockout incident, the factory often runs overtime shifts to catch up with its backlog. Because overtime is also driven by other causes, which are not tracked, the only way to find the annual overtime expense related to stockouts was to ask a factory scheduler to provide a subjective estimate of the proportion of overtime attributable to stockouts. This value was multiplied by total annual overtime spending, and the result was divided by the number of parts to arrive at the average per-part overtime cost under the unique scenarios.

While it is impossible to precisely quantify how many sales are lost as a result of quoting customers extended lead times in response to a part shortage, the CoC was able to provide a reasonable estimate based on their assessment of several specific incidents in which lead times exceeded two weeks – well above Dell’s standard five days. For each incident, the sales forecast was treated as the trend (after controlling for forecast inaccuracy by comparing to actual sales before and after the lead time incident), and the shortfall in actual sales during the incident was attributed to the stockouts. As with the

factory overtime, this total cost was divided by the number of parts to arrive at a per-part cost for the unique scenarios, relying on the implicit assumption that each part has an equal probability of causing such an incident.

Replaceable Scenarios

The stockout cost under the replaceable scenarios was found by first determining the increase in stockouts as the new part is added, and then quantifying the additional costs associated with these stockouts.

The increase in stockouts was found using a model that had been used at Dell to evaluate the impact on service levels of changes to inventory policy. This model was based on the standard safety-stock equation for a continuous-review inventory policy with stochastic lead times, which is described in detail by Simchi-Levi¹⁰ et al.:

¹⁰ David Simchi-Levi, Philip Kaminsky, and Edith Simchi-Levi, *Designing & Managing the Supply Chain: Concepts, Strategies & Case Studies*, 2nd ed. (New York: McGraw-Hill, 2003), 58-62.

Note that Simchi-Levi et al. use the standard deviation of demand instead of the standard deviation of forecast error, as was done here. The use of forecast error is an extension of the equation for situations, like this, in which the forecast varies with time.

$$SafetyStock = z_0 \times \sqrt{\mu_{leadtime} \times \sigma_{error}^2 + \mu_{demand}^2 \times \sigma_{leadtime}^2}$$

z_0 = z-value of the standard normal distribution such that $P(z < z_0) = \text{service level}$

$\mu_{leadtime}$ = mean supplier lead time

$\sigma_{leadtime}$ = standard deviation of supplier lead time

μ_{demand} = mean demand

σ_{error} = standard deviation of forecast error

Because gathering information on the last four inputs was a very time-consuming process, a single commodity, memory, was selected as a proxy for all other parts. It is impossible to directly measure $\mu_{leadtime}$ and $\sigma_{leadtime}$, so the master scheduler for memory provided a range of reasonable values based on his judgment. For example, lead time is nominally three months, but the scheduler said that he can always get any part he needs within three to six weeks, so the best- and worst-case values for $\mu_{leadtime}$ were three and six weeks. Similarly, the scheduler estimated $\sigma_{leadtime}$ to be between two and five days; this value represents the variation about the promised lead times ($\mu_{leadtime}$) of three to six weeks. The standard deviation of the forecast error (σ_{error}) was found by reviewing available historical data for 20 memory parts that the master scheduler felt were representative of the spectrum of memory parts, and calculating the pooled variance based on the 10-20 weeks of available data. These model inputs closely matched those for a collection of five hard-drive parts, verifying that the memory parts were reasonably representative of all part types.

The best- and worst-case values of these four inputs were applied to the inventory model to calculate the safety stock needed to maintain a reasonable range of service levels (95-99%) for a typical part. The demand for that hypothetical average part was then split into two independent and equal demand functions, so that the mean demand of each equals half the mean demand of the combined part. As explained in Section 2.4.5, each of the parts would then have a forecast error with a standard deviation of $\sigma_1 = \sigma_2 = \sigma_c / \sqrt{2}$, where σ_c represents the forecast error of the combined part. The updated standard deviation of the forecast error (σ_{error}), the updated mean demand (μ_{demand}), and the original values of lead-time mean and standard deviation ($\mu_{leadtime}$ and $\sigma_{leadtime}$) were then applied to the model to find the service level with two parts. This value was compared to the service level with a single part to find the decline in service level after a single part is replaced by two.

Finally, the factory overtime and lost sales amounts from the unique scenarios were multiplied by the decline in service level resulting from the addition of the second part to arrive at the cost under the replaceable scenarios.

2.4.5.2 Constant Service Level Assumption

While the assumption of a constant service level was not used in the final model, it is presented here because it represents an alternate way of calculating cost, and provides some useful lessons about the problem. Under this assumption, a cost is incurred only under the replaceable scenarios of Table 1, in which the baseline condition is that all demand is met by a single part. Each of the two parts' suppliers will charge for the cost of managing its total inventory, which includes both the inventory necessary in the

baseline condition and the additional inventory necessary to support the decision to offer two separate parts, so the incremental cost of offering two parts instead of one is not visible and must be included in the cost model. Under the unique scenarios, the baseline is not having any part at all, so the cost of maintaining inventory for a single part is embedded in the single supplier's material cost, which would be considered separately from the cost model.

The inventory model described in Section 2.4.5.1 could first be used to find the additional safety stock needed to maintain a fixed service level after the new part is added. This increase in inventory then results in two costs: working capital and storage charges. The cost of working capital, i.e. inventory holding cost, could be found by multiplying the value of the additional inventory by the company's working average cost of capital (WACC), which represents the opportunity cost of tying up capital in inventory instead of investing in a different project with a risk profile equivalent to that of the company as a whole. The reasons for using the WACC as the opportunity cost were described in section 2.3.

Storage charges assessed by Dell's warehouse contractor are the second cost associated with the additional inventory. While Dell's current contract with the warehouse operator stipulates a charge on each pallet that is shipped through the hub, regardless of the storage time, it is clear that increasing the average storage time per pallet by adding to the days' sales in inventory represents a cost to the supply chain that will eventually be passed on to Dell. In order to avoid the need to revisit the constantly renegotiated storage-pricing scheme every time the model is updated, Dell's blanket cost of factory

space of \$3.00 per square foot per month could be used. This pricing structure presents a problem, as the model's user must now know not only how much is spent on a part, but also the floor space occupied by the part.

These items represent the costs to Dell's entire supply chain of holding the inventory necessary to maintain service levels after the addition of a single part number. Because Dell does not take ownership of the inventory until it passes the factory loading dock a few hours before assembly, though, it is not entirely clear how much of this cost increase is absorbed by suppliers and how much is passed on to Dell through higher material pricing. One could assume that some reasonable fraction is passed on, for example 75-100%.

The constant-service-level assumption was challenged by the belief that any costs that are passed on to suppliers do not affect Dell, and should not be considered. It proved to be very difficult to get those few individuals who believed this to change their minds. In order to avoid the issue, as well as the need for awareness of the floor space occupied by a part, the fixed-inventory assumption of Section 2.4.5.1 was used in the final model.

2.4.6 Logistics

The addition of a new part results in added forecast uncertainty, as described in the preceding section, resulting in more frequent expediting of material. Although in reality only a limited set of part categories (such as chasses and speakers) require expediting, all parts are assumed to be equally likely to result in expediting expense. This is a result of the general decision, discussed in Section 2.3, to develop a single model for all part categories. The cost for a unique part (as defined in Table 1) was found by dividing total

annual expediting expenditure by the number of parts, and scaling this value by each part's annual material spending. The cost for a replaceable part was found by multiplying this value by the increased probability of stockouts when two parts are used instead of one.

2.4.7 Factory

In addition to the overtime described in section 2.4.5.1, the factory incurs costs associated with the kitting line, in which the components for each computer are placed on a tote for subsequent assembly, and the boxing line, in which items such as mice and keyboards are put into the box after system assembly. These parts are transferred by workers standing in front of pick-to-light (PTL) racks, which indicate the part to be picked by activating a small light. Because each worker is assigned two racks, which contain a fixed number of lights (and part numbers), the cost of this labor is driven by the number of part numbers, not by the volume of material being loaded onto the rack by stockers. It was therefore included in the model by dividing the total annual labor cost by the number of parts.

2.4.8 Service

While most parts are used in production for a short time, assumed to be one year, spare parts must be available to service the longest customer warranty for a total of up to five years (assumed to be three to four years on average). This range was necessary because some parts can be bought at any time, while others must be stocked for years after they are discontinued. Three types of cost are incurred in supporting this spare-parts inventory: headcount, capital, and excess and obsolete.

The lifetime headcount cost was found by subjectively estimating the fraction of the Service organization's annual labor expenditure attributable to part numbers, dividing by the number of parts in the service inventory, and multiplying by the three to four years that the part is supported.

Inventory holding cost is the cost of capital (WACC) applied to the average inventory per part, over three to four years. This inventory consists of two types of parts, each of which has a unique cost of capital associated with it: brand-new purchased components, and remanufactured customer returns. The cost of capital for the new components would be the same as for production parts, and the WACC is a close approximation, as discussed in Section 2.3. It is somewhat less clear, though, how much capital is tied up in remanufactured parts, as these were already paid for by the customer who purchased the original system; the real cost would depend on the proportion of remanufactured parts in the inventory, the valuation method for those parts, and the scrap rate on returned parts, among other factors. These issues are addressed in detail in the literature on reverse-logistics policy, for example by Teunter and van der Laan¹¹, but are not critical in this context, as the objective is not to define inventory policy, but to account for the costs involved. Given this objective, the WACC can be used as a reasonable approximation for the cost of holding all of the service inventory, regardless of its origin.

¹¹ Ruud Teunter and Erwin van der Laan, "On the Non-Optimality of the Average Cost Approach for Inventory Models with Remanufacturing," *International Journal of Production Economics*, September 1, 2002, 67-73.

To calculate the cost of excess and obsolete material (i.e. inventory write offs), the total annual write-off amount was divided by an estimate of the number of parts that are removed from service in any year.

2.5 MODEL RESULTS

After all of the individual cost drivers were found and quantified, they were integrated into a comprehensive model of the cost of managing a part number. Because of the need for simplicity in the model (see Section 2.2), the Pareto Principle, which states that 80% of the results in a system are produced by 20% of the effort, was used to isolate the line items that accounted for 80% of the total cost, and exclude the remainder from the model. While it took significantly more than 20% of the total effort to arrive at 80% of the cost, this exercise was consistent with the general spirit of the Pareto Principle. Because this approach calls for ignoring approximately 20% of the total cost, it might appear to make the model unnecessarily inaccurate. The high degree of uncertainty in the cost drivers, and the resulting wide bands in the estimate, however, mean that this omission does not significantly degrade the model's output. Figure 7 shows the cost components that are not scaleable by the part's annual material expenditure, while Figure 8 shows those that are scaleable; both figures also show which of the drivers were included in the final model. One of the scenarios of Table 1, a unique part being considered before its introduction, was used in selecting the most relevant cost drivers. All of the cost drivers that were considered are included in Figure 7 and Figure 8; even those that are not affected by the addition of a new part but were mentioned in the course of an interview are shown, to demonstrate that they were considered. For the sake of brevity, most of the costs that were excluded from the final model were not described in detail Section 2.4.

Annual non-scaleable cost for one part number

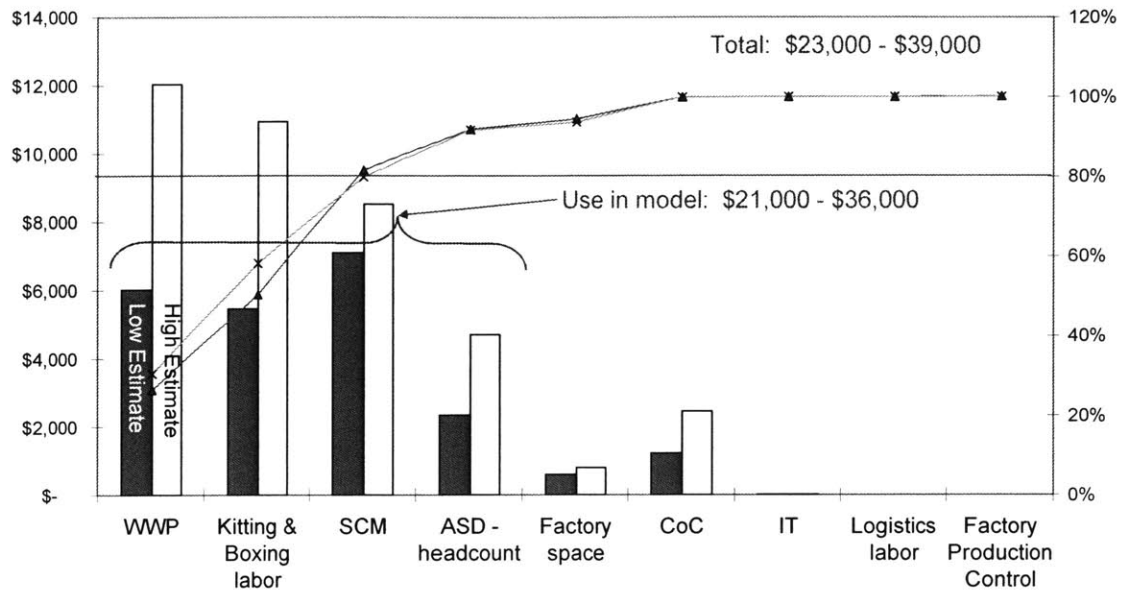


Figure 7: Pareto diagram of non-scaleable costs

Annual scaleable cost for one part number (% of material spend)

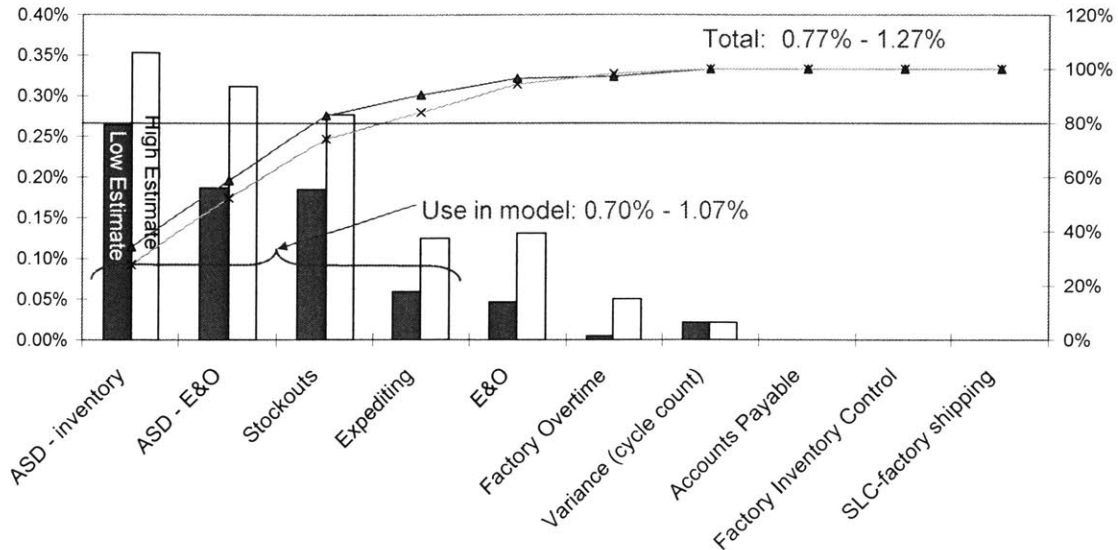


Figure 8: Pareto diagram of scaleable costs

Figure 7 shows that four non-scaleable costs were used in the final model:

1. Worldwide Procurement (WWP) headcount
2. Factory Kitting and Boxing Labor
3. Supply chain Management (SCM) headcount
4. Service (ASD) headcount

As shown on Figure 8, four costs that scale with the part's annual material expenditure were included in the final model:

1. Service (ASD) inventory
2. Service (ASD) excess and obsolete (E&O)
3. Stockouts
4. Expediting

Each of these costs was then adjusted to derive the costs for the remaining three scenarios (see Table 1). The differences between the costs associated with the introduction of a new part and the savings from removing an existing part are shown on Table 2, while the differences between a replaceable and a unique part are shown on Table 3.

Before Part Introduction			After Part Introduction		
	Low	High		Low	High
Stockouts	0.18%	0.28%	Stockouts	0.18%	0.28%
ASD - inventory	0.26%	0.35%	ASD - inventory	0.00%	0.18%
Expediting	0.06%	0.12%	Expediting	0.06%	0.12%
ASD - E&O	0.19%	0.31%	ASD - E&O	0.00%	0.16%
WWP	\$6,021	\$12,043	WWP	\$6,021	\$12,043
Kitting & Boxing	\$5,468	\$10,973	Kitting & Boxing	\$5,468	\$10,973
ASD - Headcount	\$2,350	\$4,700	ASD - Headcount	\$2,350	\$4,700
SCM	\$7,103	\$8,524	SCM	\$5,273	\$6,328

Table 2: Timing-related differences between scenarios

Unique	Stockouts	0.18%	0.28%
	ASD - inventory	0.26%	0.35%
	Expediting	0.06%	0.12%
	ASD - E&O	0.19%	0.31%
	WWP	\$6,021	\$12,043
	Kitting & Boxing	\$5,468	\$10,937
	ASD - Headcount	\$2,350	\$4,700
	SCM	\$7,103	\$8,524
Replaceable	Stockouts	0.11%	0.55%
	ASD - inventory	0.05%	0.07%
	Expediting	0.06%	0.12%
	ASD - E&O	0.03%	0.06%
	WWP	\$6,021	\$12,043
	Kitting & Boxing	\$5,468	\$10,937
	ASD - Headcount	\$2,350	\$4,700
	SCM	\$7,103	\$8,524

Table 3: Part-type-related differences between scenarios

The individual cost components were finally combined into two sets of coefficients – each consisting of a low-to-high range – for each scenario: a constant dollar amount and a percentage of the part’s annual material spending. These coefficients are shown on Table 4:

	Before Part Introduction			After Part Introduction		
Unique		Low	High		Low	High
	Scaleable	0.70%	1.07%	Scaleable	0.24%	0.73%
	Not scaleable	\$21,000	\$36,000	Not scaleable	\$19,000	\$34,000
Replaceable		Low	High		Low	High
	Scaleable	0.25%	0.81%	Scaleable	0.17%	0.75%
	Not scaleable	\$21,000	\$36,000	Not scaleable	\$19,000	\$34,000

Table 4: Cost model coefficients

The cost of introducing a unique part (the upper-left scenario) with a forecast annual consumption of \$5,000,000 provides an example of the application of the model:

$$\begin{aligned}
 \text{Lifetime supply chain cost (\$)} &= 0.70\% \times (\text{annual spend}) + \$21,000 \\
 &= 0.70\% \times (\$5,000,000) + \$21,000 \\
 &= \$56,000
 \end{aligned}$$

The model was implemented as a spreadsheet with one input: the annual material spending on the part in question. The cost output by the model can be used as an estimate for the majority of parts, particularly those whose low volumes do not justify the time required for more detailed analysis. If this cost is questioned or considered unrealistic, it can easily be exploded into its individual components for further analysis. Finally, some parts, such as motherboards and chasses, are important enough to justify a significant investment of time in the analysis. In this case, the model provides a framework for evaluating the sources of cost, as well as a methodology for quantifying

the individual cost drivers. The model is therefore useful not only for the cost estimate that it generates, but also for the visibility it provides into the supply chain cost structure.

2.6 MODEL APPROVAL

The final step in completing the model was to have it approved by high-level management as Dell's official cost of managing a part. Although it did not significantly improve the model's output, this process was politically critical for the long-term credibility of the model – without this legitimacy, the model would be questioned any time it calls for the removal of a specific part that some constituency feels should be retained for other reasons. Operational and financial managers from each organization whose costs were included in the model were asked for their approval in an extended series of individual meetings. This process caused problems when members of one of the organizations asked to approve a single line item found that the overall model results did not support their agenda. Because it was impractical to change the model in such a way that their input was not needed, it became necessary to first appeal to them to compromise, and finally to request help from management in resolving the disagreement. This experience provides a useful illustration of the political obstacles that can obstruct even a seemingly objective, neutral cost-accounting effort.

2.7 SUMMARY

This section described the most important output of the project: a model of the incremental costs incurred when a single part number is added to the inventory or, conversely, the savings associated with the removal of a single part. The model is valid for most types of parts, including those that are engineered into a basic system and are

never considered by customers, but its development was motivated by the need to provide an input to a management process, described in Section 4, to eliminate unprofitable customer options. The model was derived using a six-step process:

1. Map the supply chain process to understand where costs are incurred
2. Evaluate the model's ultimate application, to judge the accuracy required
3. Make the necessary simplifying assumptions
4. Quantify all of the cost drivers
5. Consolidate the most important cost drivers into a single model
6. Certify the model as Dell's official cost of managing a part

As it was unrealistic to precisely quantify these hidden costs, step two was critical in determining the appropriate level of simplification. As with many cost-accounting problems, the solution may not be perfectly accurate in every single case, but it leads to the correct decisions, on average. This result, though not ideal, is far more desirable than the alternative of entirely ignoring the cost of part management.

3 CASE STUDY

The second phase of the project was to apply the cost model to historical sales data for a sample of parts, with two objectives: to refine methodologies for selecting parts for elimination (Section 3.1), and to determine the potential savings from the part-review process described in Section 4, justifying its implementation (Section 3.2). September 2002 sales of Optiplex – a desktop-computer line for business customers – were evaluated, and the results were then scaled by revenue to estimate the potential savings across Dell’s Americas operations. This relatively small sample size, which was selected to keep the task manageable, provides sufficient information to refine the part-selection methods and to establish a reasonably reliable business case.

3.1 PART-SELECTION METHODOLOGY

The first objective of the case study is to determine how revenue and cost should be allocated to individual parts in order to establish which should be candidates for elimination. The problem is presented in Section 3.1.1 and the proposed solution, which resulted from the trial of different methods, is presented in Section 3.1.2.

3.1.1 Allocation Problem

The calculation of a part’s profitability is complicated by the two problems that are presented in the following subsections: the allocation of revenue and cost as the bill of materials (BOM) tree is exploded, and the application of system-level discounts to individual SKUs.

3.1.1.1 BOM Explosion

Because revenue is earned on the stock keeping units (SKUs) at the top of the BOM tree, while costs are incurred on the individual part numbers at the base of the tree, a methodology had to be developed to appropriately allocate revenue and cost as the BOM is exploded. This could be done in one of two ways: allocating SKU revenue to parts, or allocating part costs to SKUs.

Figure 9 shows a hypothetical example of the allocation problem. Three components of part profitability are known:

1. The revenue associated with each SKU
2. The material cost of each part number
3. The supply chain cost associated with each part number (\$2 in this hypothetical example)

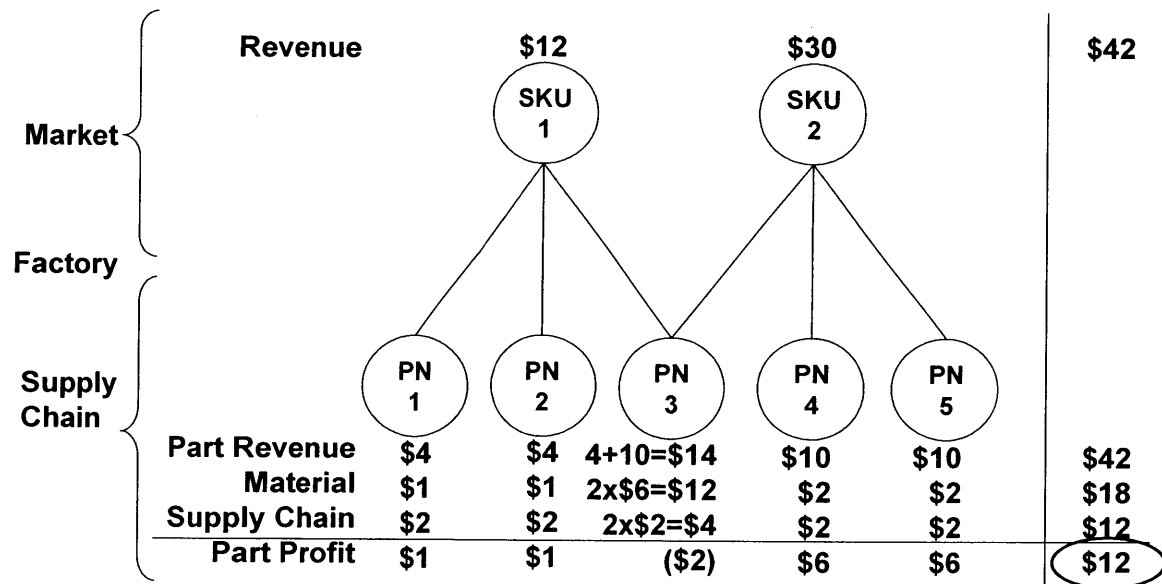


Figure 9: Simplified example of BOM explosion

Without making any assumptions about the revenue-allocation method, the total profit from selling these two SKUs is \$12. The simple approach of allocating the revenue from each of the two SKUs evenly to each of its constituent parts results in the per-part profit shown on the bottom row. Under this methodology, part 3 appears to be unprofitable, and would therefore be eliminated. As a consequence of this decision, however, neither SKU 1 nor SKU 2 could be manufactured, resulting in an overall profit of \$0 instead of \$12; the entire exercise was therefore counterproductive.

A possible solution to this problem would be to allocate revenue to each part number in proportion to its cost. Examples could easily be shown, though, in which this method also results in unprofitable decisions.

As an alternative to directly identifying part numbers for elimination, as in the preceding example, unprofitable SKUs could be identified. Any part numbers that were not used by the remaining – profitable – SKUs would then be removed from the lineup. This indirect method of identifying candidates has the advantage of entirely avoiding the cost-revenue-allocation problem, as the cost of every part can be clearly and unambiguously associated with the SKU with which it was sold. The only significant drawback is that the elimination of a part depends on the elimination of every single SKU that uses that part in its BOM; as long as a part is needed by a single SKU, it has to be retained.

The example of Figure 9 showed a simplified scenario with only two SKUs. In reality, as illustrated in Figure 10, there will be many more relationships between SKUs and part numbers, further complicating the problem.

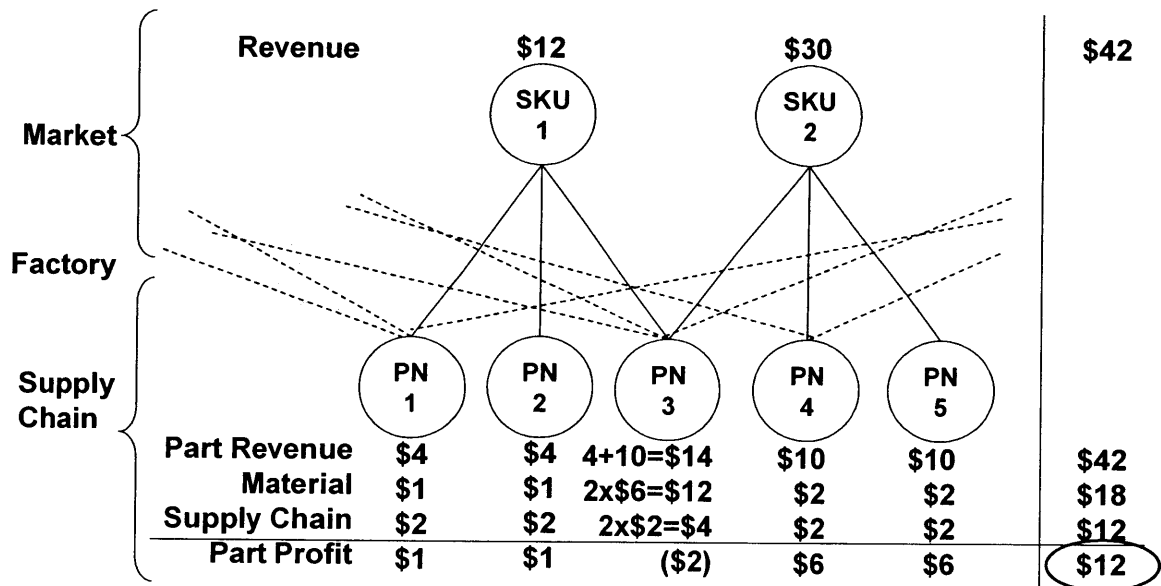


Figure 10: Extended example of BOM explosion

3.1.1.2 Allocation of System Discounts

Figure 11 shows an example of the second problem associated with determining part profitability: the allocation of system discounts. Due to limitations in Dell's order-management software, as well as simple convenience, sales people often assign discounts, which they offer as an inducement to purchase an entire system, to a single SKU within the system (for example a monitor or DVD drive). In this example, they assign the entire \$30 system discount to SKU 3, making it appear unprofitable. Clearly, it would be ideal to distribute this discount among all three SKUs in order to present a more realistic picture. Because the order-management software is outside the scope of this project, though, it is necessary to accept the fact that some SKUs will have unrealistically high discounts assigned to them, and find another way of dealing with the problem.

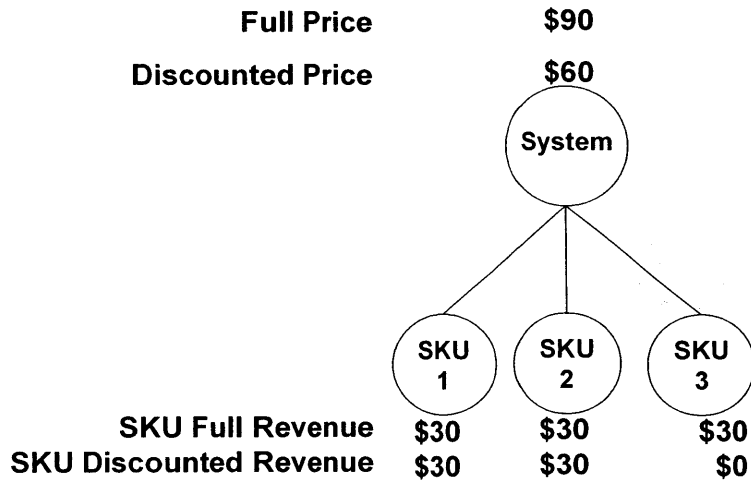


Figure 11: Allocation of discounts

3.1.2 Proposed Allocation Methodology

Given the scale of the problems presented in the previous sections, with thousands of relationships between SKUs and part numbers and millions of individual transaction discounts, it is clearly unrealistic to precisely determine the profitability of each part. Instead, it is more appropriate to find a generalized allocation methodology that on average provides the correct guidance, even though it may be incorrect in specific instances. This philosophy, which was also applied to the cost model of Section 2, is a common way of addressing the problems inherent in all cost-accounting exercises. After accepting that the allocation method has to provide, at best, an approximation, it is important to apply the method to real data in order to determine whether the level of inaccuracy is acceptable.

As mentioned in Section 3.1.1.1, the preferred method for identifying parts for elimination would be based on eliminating unprofitable SKUs, retaining only the parts

that are used in the remaining SKUs, and removing the unused parts. Because no part can be removed if it is used by any remaining SKUs, it is important to apply this technique to real sales data in order to determine whether it results in the elimination of a sufficient number of parts.

Sales figures for each SKU, which included shipped volume, revenue, and material cost, were first gathered from the appropriate database (called the Data Warehouse). A separate database was then used to find every SKU to part number relationship, allowing the SKU sales figures to be exploded into individual part numbers. The third step was to compare these figures, which were derived from the sales database, with the version from the factory's own database. Because the factory list turned out to be incomplete due to the exclusion of several classes of parts, the sales figures from the Data Warehouse were used for the analysis.

This conversion of SKU sales data to part shipments might at first sight appear to be the most straightforward part of the case study, but it was actually the most difficult.

Because Dell does not use a centralized Enterprise Resource Planning (ERP) system, the data had to be gathered from five independent databases that often yielded contradictory data. This lack of readily available, reliable data was a hindrance not only in this project, but in several others that the author heard about at Dell. While this could be interpreted as a necessary – if unfortunate – side effect of Dell's highly flexible, fast-paced business model, the problem is actually more widespread; at Kodak, a mature business,

Rockwell's¹² implementation of a centralized database provided visibility to the product portfolio for the first time, leading to the realization of previously unseen part-consolidation opportunities. Dell had attempted an implementation of SAP, a leading ERP product, in the early 1990s, and had concluded that the software could not be adopted without sacrificing some of the flexibility that provided the foundation of Dell's success. The suitability of ERP systems in this type of environment is beyond the scope of this discussion; given the constraints of the existing infrastructure, it was necessary to manually gather the data and make assumptions to overcome the contradictions between the various databases.

The relationships among the various data sets are shown on Figure 12 to illustrate the complexity of the data-gathering problem.

¹² Clinton J. Rockwell, "A Framework For Optimizing the Supply Chain: A Case Study at Kodak" (LFM Thesis, Massachusetts Institute of Technology, 2002).

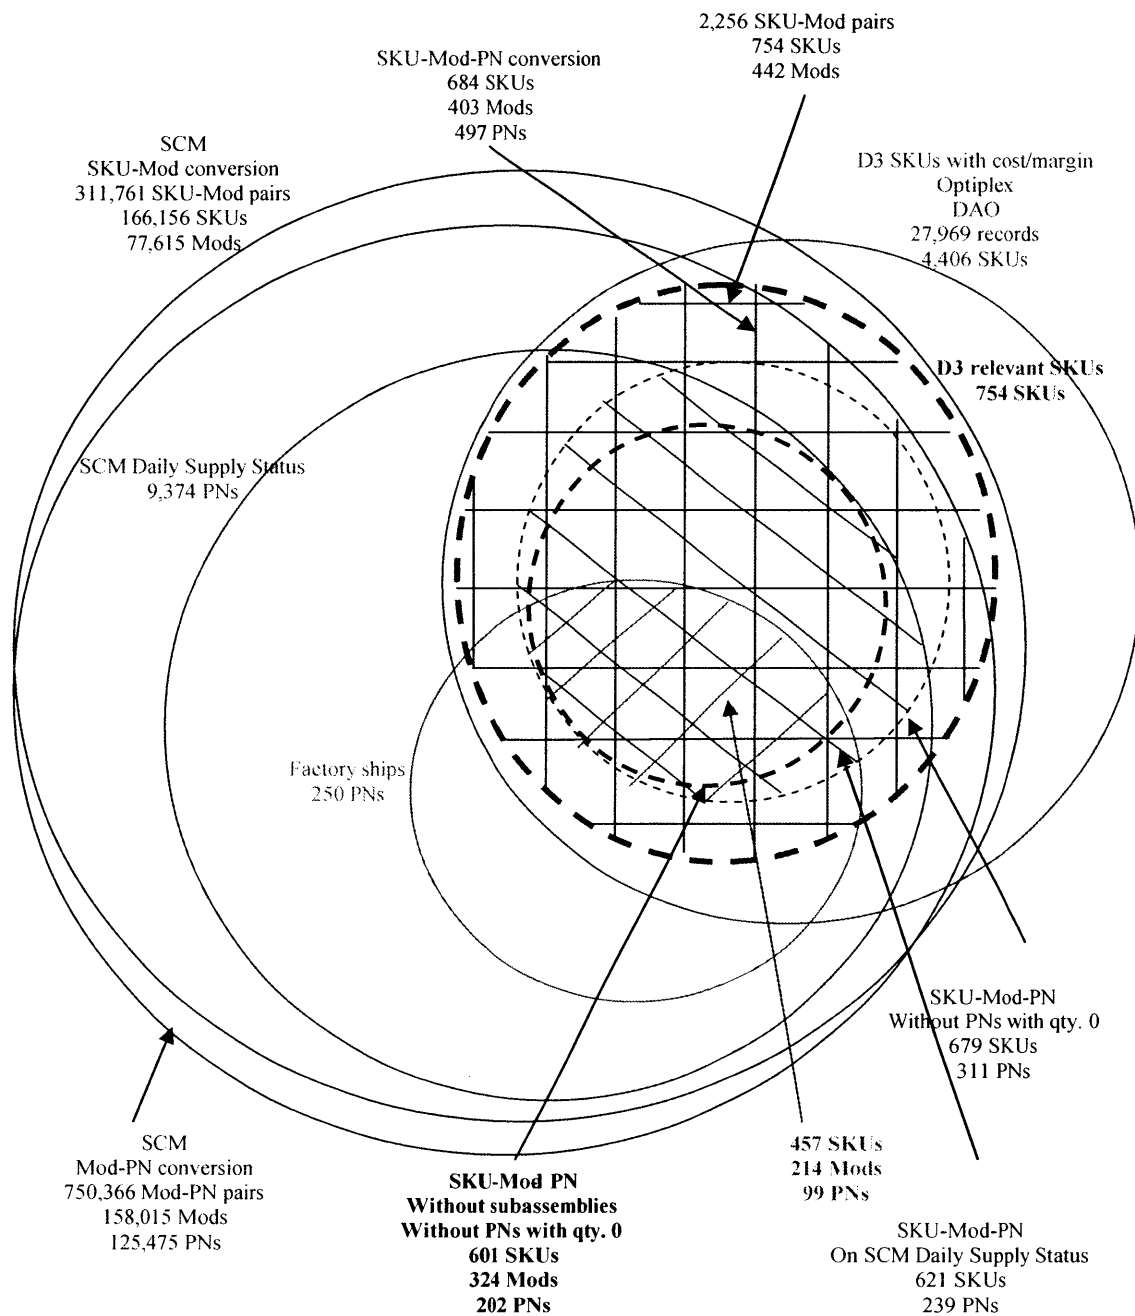


Figure 12: Databases for case study

The next step was to assign cost to each SKU, using the resulting SKU profit to identify candidates for elimination. Given the uncertainty in determining SKU profitability, best- and worst-case scenarios were used to provide a more realistic estimate of the potential savings. The discount-allocation problem described in section 3.1.1.2 was solved by using the retail price of each SKU as the high estimate of revenue, and the discounted price as the low estimate; the real revenue would lie somewhere in between if system discounts were properly allocated to the constituent SKUs. Like the revenue, the material cost of the parts constituting each SKU was then found directly in the Data Warehouse. The supply chain cost of each part constituting a SKU was found using the cost model of Section 2, with the non-scaleable portion allocated evenly to all SKUs using the part, and the scaleable portion applied to the part's material spending for the SKU in question. Unprofitable SKUs were then found by subtracting these material and supply chain costs from SKU revenue. Once the unprofitable SKUs were removed from the list, the parts needed to support the remaining SKUs were found. The parts present in the original list, but not this reduced list, could then be removed, as shown on Figure 13. These candidates for elimination were used as the input for the financial analysis of the following section.

	601	202	
327-442 SKUs eliminated	Existing Total		104-120 parts eliminated
	274 Remaining (High est.)	98	
	159 Remaining (Low est.)	82	
	SKUs	Parts	

Figure 13: Case study results (not to scale)

Because this analysis demonstrated that evaluating SKUs for profitability resulted in a significant reduction in the part count, it proved to be unnecessary to develop a methodology for assigning revenue to part numbers in order to evaluate individual parts for profitability. While screening SKUs was adequate in this application, it is possible that this approach would not work in other situations, for example when each part is used by many SKUs, some of which are profitable. The suitability of each approach depends on the structure of the bill of materials.

3.2 BUSINESS CASE

The second objective of the case study was to quantify the potential savings from eliminating unprofitable peripherals, justifying the implementation of a management process – described in Section 4 – to proactively evaluate all peripherals for suitability.

Figure 14 shows the results of the case study, including the identification of parts shown on Figure 13.

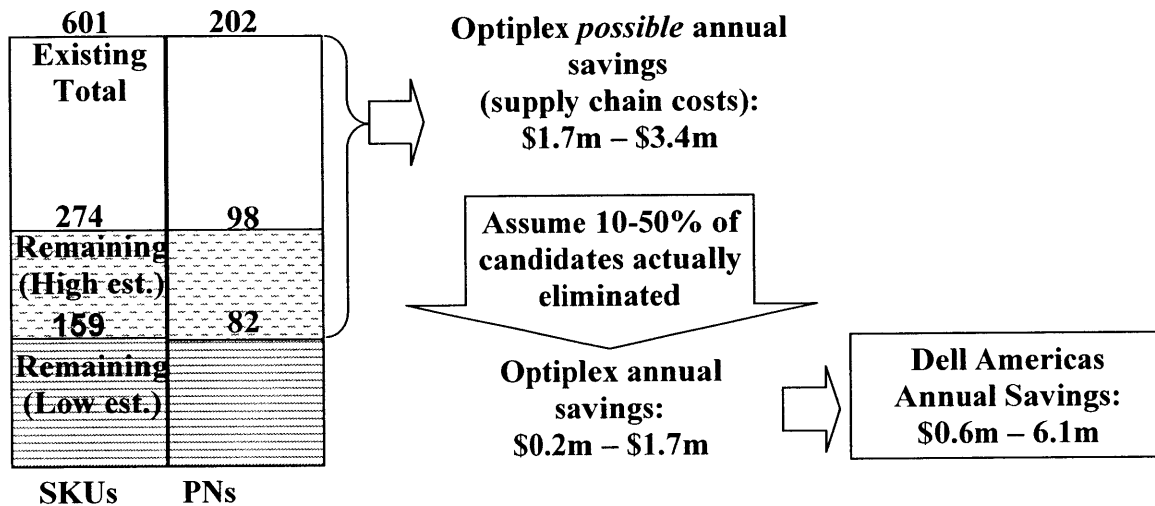


Figure 14: Case study results

The first step was to calculate the savings from removing the 104 to 120 unprofitable parts identified in the preceding section (see Figure 13). The assumption that all demand for a discontinued part would be displaced to another part helped to provide a conservative estimate, as it meant that only the non-scaleable component of supply chain cost was considered; the scaleable component would simply shift to the substitute part (see Section 2.5 for a full description of the cost components).

The preceding analysis revealed the potential savings from removing all parts that appear to be unprofitable. In reality, many of these parts would be retained after closer analysis, for example because they are needed for a broader strategic objective or because the revenue recorded in the database does not reflect their true importance. Because resource

constraints ruled out a thorough part-by part evaluation, it was conservatively assumed that, after closer evaluation, only 10-50% of these parts would actually be removed.

Finally, the potential savings from the single product line used in the case study were extrapolated across all products, based on the revenue of each product. This resulted in potential annual savings of \$600,000 to \$1,600,000 in Dell Americas Operations; the potential global savings would be higher, but were outside the scope of this project.

Given Dell's billions of dollars in annual revenues in the Americas, \$600,000 to \$1,600,000 does not represent a dramatic change to the economics of the business. This level of savings, though, is consistent with other efforts under way at the time. The consensus at Dell was that there were few dramatic savings opportunities left in this maturing business, and that the key to Dell's continued cost competitiveness was the aggregate effect of hundreds of relatively modest cost-reduction efforts. Viewed in this context, the potential savings from proactively limiting part proliferation are significant, and justify the implementation of a systematic process to manage the options catalog. Because this analysis was based on conservative estimates, the true savings are in fact likely to be higher.

3.3 SUMMARY

The first objective of the case study was to use the supply chain cost model of Section 2 to refine methodologies for identifying unprofitable parts that should be removed from the lineup, so that these methodologies could be integrated in the management process described in Section 4. The case study provided solutions to the two problems associated with calculating part profitability: the explosion of the BOM between revenue-earning

SKUs and the part numbers that incur costs, and the allocation of transaction discounts. The evaluation of a limited sample of historical sales data demonstrated that SKUs should be evaluated for profitability, and that only the parts needed to support the remaining profitable SKUs should be retained. The second objective was to establish a business case for the implementation of the management process described in the following section. While the potential \$600,000 to \$1,600,000 savings would not fundamentally change the nature of a multi-billion-dollar business, their scale is consistent with other cost-management efforts under way, which in combination will further improve Dell's cost competitiveness.

4 MANAGEMENT PROCESS

The final phase of the project was to develop a set of management procedures to proactively review all peripherals and eliminate those that are unnecessary or inappropriate. The management process relies on the cost model developed in Section 2 to determine the cost of maintaining a part, as well as the methodology of Section 3.1.2 to calculate part profitability. This quantitative information is supplemented by strategic and marketing analyses, providing a comprehensive overview of each part. The implementation of these procedures is justified by the case study of Section 3.2.

While the first two phases of the project (Cost Model and Case Study) were completed, the objective of this phase was to provide general recommendations that would be subsequently implemented by a Dell team.

4.1 EXISTING PROCESS

Because a stand-alone process would run the risk of being forgotten – or ignored – as those responsible for performing the work focused on their everyday priorities, it is important to integrate the part-review process into the company's existing procedures. The Phase Review Process (PRP) is a set of guidelines and procedures that govern the life cycles of all Dell products, from the early conceptual stage through obsolescence. The PRP is based on the concept on the Core Team, on which all organizations with a major stake in the product are represented, and which is ultimately responsible for the product. Organizations with an indirect role have places on Extended Core Teams, which in turn have a representative on the Core Team. There are three types of Core Team:

Peripheral, Platform, and Sustaining. Peripheral Core Teams are responsible for an entire class of peripherals, e.g. network cards, graphics cards, or monitors, and control each peripheral throughout its life cycle. Platform and Sustaining Core Teams, on the other hand, are responsible for a specific PC model, and are disbanded when they are finished with that platform. The Platform Core Team is responsible for the commercialization of a product, and hands it over to the Sustaining Core Team upon its production release. As this project is focused on peripherals, the Peripheral Core Team is the most important in this context.

4.2 PROPOSED PROCESS

The proposed part-review process, shown on Figure 15, is based on a review of every peripheral, both before its production release and periodically during its production life. Profile, Planning, and Implementation are the three phases of the PRP that take place before a product is introduced to the market, and are described in Section 4.2.1. These are followed by the Deployment and Management phases described in Section 4.2.2, which lead to the end of production life in the Service phase of Section 4.2.3.

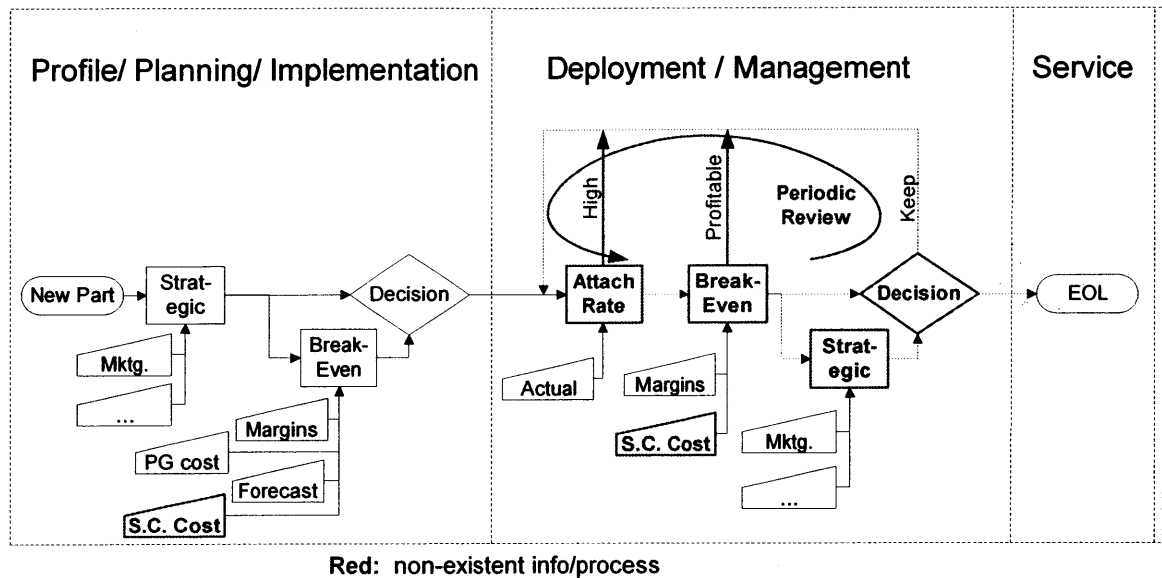


Figure 15: Proposed management process

4.2.1 Profile, Planning, and Implementation Phases

As shown on Figure 15, the decision to introduce a peripheral starts with a strategic analysis of the target market space. This is followed by a financial analysis, denoted Breakeven Analysis in the diagram, involving forecasts of volumes, margins, and development costs. Finally, the decision on the peripheral’s introduction is made by the appropriate Peripheral Core Team before transfer to the Deployment phase. This is the point at which a peripheral starts to incur costs; it is the transition between the left and right columns of Table 1.

Because the PRP already incorporates reasonably robust processes for introducing new peripherals, the only necessary change would be the inclusion of the indirect costs incurred by the supply chain as a result of the decision to add the new peripheral (S.C. Cost on Figure 15). This input, which is currently not considered, would be provided by the cost model of Section 2.

4.2.2 Deployment and Management Phases

Production starts and the first customers receive their systems at the beginning of the Deployment phase, meaning that some costs, such as production inventories and customer warranties, have been sunk. The savings from eliminating a part would therefore be lower than they would have been if the decision had been taken sooner.

The next opportunity to remove a part would be during the Management phase, in which a part is sustained for ongoing production and sale. This phase represents the bulk of a product's life (while the cost model assumes that it lasts one year, it could potentially go on for longer). As mentioned in Section 1.2, the default condition in the Management phase is for a part to remain in the catalog indefinitely, until somebody initiates a blanket part-cleanup effort. The proposed procedures would replace this ad hoc process with a systematic, planned review of every individual peripheral. This process could initially take place at monthly intervals, which could later be extended to quarterly periods as the most wasteful parts made their way through the system and there were fewer parts needing review.

The process shown on Figure 15 would be based on four analyses, which would be similar in purpose to those performed during the Profile, Planning, and Implementation phases, but would be implemented from scratch: Attach-Rate Analysis, Breakeven Analysis, Strategic Analysis, and the Decision. Each of these analyses, described in the following sections, represents an increasing amount of per-part work, so the parts list will be filtered in four stages as it moves through the sequence. If any of the analyses

indicates that a part is worth keeping, it will remain in the system until it is reviewed at the next periodic interval.

4.2.2.1 Attach-Rate Analysis

The Attach-Rate Analysis begins with a list of every single SKI and generates a list of peripherals with low sales volumes. All irrelevant SKUs, for example subassemblies and software, are first removed from the list to generate the subset of peripherals. These peripherals are then screened by attach rate, i.e. the percentage of shipped computers using the peripheral, to arrive at a reduced list of peripherals for the Breakeven Analysis of the next section. Any items with a sufficiently high attach rate would be automatically retained until the next review period. The attach rate threshold would be adjusted over time based on experience; 1% would be a useful starting point.

This analysis would be performed by the Supply Chain Management (SCM) organization, which is responsible for day-to-day materials procurement (see Section 2.1). There were three primary reasons for this selection, which was not contested by anybody:

1. SCM purchases parts across all product lines, and would therefore have access to the total consumption of a part, regardless of the products with which it is sold.
2. SCM already performs a more limited version of this analysis, so this would be a natural extension of its existing work.
3. Early on, SCM management expressed a willingness to take on this work.

4.2.2.2 Breakeven Analysis

The Attach-Rate Analysis of the previous section generates a list of peripherals for the Breakeven Analysis, which outputs a list of parts that appear to be unprofitable and should therefore be considered for elimination. Unprofitable parts are those whose sales volume appears to lie below the breakeven point, at which revenue (net of material cost and the scaleable component of supply chain cost) does not exceed the non-scaleable (i.e. fixed) component of supply chain cost. Because Dell policy dictates a 16% return on all peripherals, breakeven in this case means a minimum return of 16%, not zero. As described in Section 3.1, there is no perfect method to assign revenue and cost to part numbers, primarily because revenue is earned on the SKUs at the top of the BOM tree and cost is incurred at the part number level at the bottom of the BOM. The methodology developed in Section 3.1, however, provides a reasonable approximation, and is applied to all peripherals in the list to find those below the breakeven point. The limitations of the Breakeven Analysis are acceptable because it would not result in the automatic elimination of parts. Instead, the results provide guidance to the decision, in which strategic issues are considered and ultimately take precedence. Profitable parts are retained until the next periodic review cycle, while unprofitable parts are transferred to the Strategic Analysis.

There was some disagreement about which of two organizations should perform the Breakeven Analysis: the CoC or Product Group (PG) Finance. While there were valid reasons for assigning the task to either, with a reasonable chance of a successful implementation, PG Finance is, on balance, the appropriate organization. The primary reason is that the finance organization's approval affords the financial decision a degree

of legitimacy, enhanced by the perception of impartiality, that the CoC would not be able to match. Although the financial results are only meant to provide guidance, this credibility is essential in order to avoid endless arguments about the suitability of the assumptions in the cost model and the part-profitability calculations.

4.2.2.3 Strategic Analysis

The list of parts that were found to be unprofitable under the Breakeven Analysis then sent to the Peripheral Core Team for the Strategic Analysis. Any relevant issues that were not considered in the financial analysis are evaluated at this stage. This is an opportunity to consider any broader strategic or marketing concerns, which are by their nature subjective and can therefore not be easily described by a monetary figure. The following examples illustrate the types of issues that might be considered in this analysis:

1. Commitments to existing customers to provide a specific part as a general condition for wider, more profitable business.
2. The ability to advertise low-priced systems, which attract potential customers who can then be convinced to upgrade to more lucrative products.
3. Defensive presence in a market space in order to deprive competitors of a potential profit pool.

While it would not perform all of the detailed analysis, the Peripheral Core Team would be responsible for assigning responsibility to the correct representative. This is necessary because further data gathering or analysis may have to be performed by different organizations. For example, Marketing may evaluate current offerings in the peripheral's market space, or the CoC may review pricing strategies. After performing the assigned

analysis, the responsible team member would present the results to the Peripheral Core Team for the Decision, which is addressed in the following section.

4.2.2.4 Decision

The Peripheral Core Team would then use the results from the Breakeven and Strategic analyses for guidance in deciding whether a peripheral should be retained or sent to the Service phase. Because the results of the Strategic Analysis will in many cases be critical, but cannot be quantified, this decision will be subjective, and will take place at the team's regularly scheduled meetings. The decision must be taken by a Core Team, as it is the only forum in which all organizations with an interest in the decision are represented. The Peripheral Core Team was selected instead of the Sustaining Core Team because it has visibility across all platforms that use a peripheral. It may in fact consult with the Sustaining Core Teams for products that would be affected by the decision to eliminate a part, but it would ultimately be responsible for the decision.

4.2.3 Service Phase

Finally, any parts that the Peripheral Core Team decides should be eliminated would be sent into the Service Phase. The main activity during this phase, which takes place after the product's end of production life, is to support the warranties on any systems left in the field. Operations within the Service Phase would not be affected by this proposal; the only difference is that a larger set of parts would be passed on from the Management Phase.

4.3 SUMMARY

The third and final phase of the project was the development of a set of management procedures to systematically review all peripherals and remove those that are unprofitable or unnecessary. These procedures are to be integrated into Dell's existing Phase Review Process (PRP), which relies on Core Teams that include representatives from all organizations that have a stake in a product. The PRP already includes reasonably robust procedures for determining whether a new peripheral should be introduced; these would be supplemented by the supply chain model described in Section 2. A new set of procedures would also be put in place to periodically review all parts, progressively passing them through an Attach-Rate Analysis, a Breakeven Analysis that uses the methodologies developed in Section 3, and finally a Strategic Analysis. The appropriate Peripheral Core Team would then make the decision to retain the part or remove it from the product lineup.

5 SUMMARY AND RECOMMENDATIONS

The ability to customize products to each customer's unique requirements can be a source of competitive advantage, as it has been for Dell. It is important, however, to effectively manage the catalog of customer options in order to minimize the costs of maintaining the underlying parts. When addressing the part-proliferation problem in other settings, one should first consider the two alternative approaches described in Section 1.3.1: the comprehensive, top-down review of the product offering recommended by Child et al., and the optimization of the management of all parts, rather than the elimination of a subset of parts. Given the constraints of Dell's business environment and the current state of its supply chain, these approaches were found to be suboptimal, leading to the approach proposed here.

The first – and most important – step is to quantify the cost to the entire supply chain of managing a single part. Without this information, there is no way to determine whether the benefits, both tangible and intangible, of adding or retaining a part outweigh the costs, and the only argument against offering the part becomes a subjective, indefensible position that “complexity is bad.” Because it is impossible to precisely determine that cost at any instant, especially in a rapidly changing environment like Dell's, any model has to be an approximation that relies on a series of assumptions. The balance between model accuracy and simplicity is not an easy one to strike; here, the Pareto Principle is used to find the subset of cost drivers that account for 80% of the total cost. Given the broad bands of uncertainty in the cost estimates, and the fact that the supply chain cost is only one input to an ultimately qualitative decision, this provided an adequate level of

accuracy. It is important to keep in mind, though, that it may be necessary in some settings to establish a more accurate model in order to overcome political obstacles, as it is all too easy for opponents to use objections to the model's assumptions as a pretext for general opposition to a part-reduction effort. To overcome these problems once the appropriate level of accuracy has been selected, it is vital to afford the model legitimacy by obtaining high-level management approval. This cost model can be worth developing even if the objective is not a reduction in the number of parts, as the visibility into the cost structure that it provides would be valuable in other types of decisions. An important enhancement to the model would be the establishment of several broad categories of parts, doing away with the assumption that all types of parts incur the same costs (see Section 2.3).

Once the cost of part management has been determined, a sample of real parts should be evaluated. This helps to establish methodologies for selecting parts for elimination, solving problems like the allocation of cost and revenue. This process also provides an idea of the amount of work that will be required to collect data about individual parts. In the case of Dell, the integration of information systems is likely to be the greatest obstacle to the implementation of these recommendations (see Section 3.1.2). This case study is also important because it quantifies the potential savings, justifying the possibly disruptive implementation of a process to remove unnecessary parts.

Finally, the results of the cost model and case study are integrated into a set of management procedures. Because the decision about any part will always be partly subjective, as it requires consideration of unquantifiable inputs such as marketing or

strategic issues, this process should establish a setting in which all concerned parties can discuss the issues and reach a consensus. These meetings should be integrated into existing processes, such as Dell's Phase Review Process. Otherwise, the review of parts is likely to be neglected as more immediate problems are addressed, and the significant time required for the process is never set aside.

While some constituencies will always argue for more parts, and others for fewer, this process will ensure that the product lineup is defined not by chance or by the constituencies' political influence, but by an objective, systematic set of criteria. This results not only in direct savings through the elimination of parts, but also in the preemption of wasteful, ad hoc efforts to reduce the costs of part management.

REFERENCES

Child, Peter, Raimund Diederichs, Falk-Hayo Sanders, and Stefan Wisniowski. "SMR Forum: The Management of Complexity." *MIT Sloan Management Review*, Fall 1991, 73-80.

Lampel, Joseph and Henry Mintzberg. "Customizing Customization." *MIT Sloan Management Review*, Fall 1996, 21-30.

Nahmias, Steven. *Production and Operations Analysis*. 4th ed. New York: McGraw-Hill/Irwin, 2001.

Rockwell, Clinton J. "A Framework for Optimizing the Supply Chain: A Case Study at Kodak." LFM Thesis, Massachusetts Institute of Technology, 2002.

Simchi-Levi, David, Philip Kaminsky, and Edith Simchi-Levi. *Designing & Managing the Supply Chain: Concepts, Strategies & Case Studies*. 2nd ed. New York: McGraw-Hill, 2003.

Teunter, Ruud and Erwin van der Laan. "On the Non-Optimality of the Average Cost Approach for Inventory Models with Remanufacturing." *International Journal of Production Economics*, September 1, 2002, 67-73.

Womack, James P. and Daniel T. Jones. *Lean Thinking: Banish Waste and Create Wealth in Your Corporation*. New York: Simon & Schuster, 1996.

Zipkin, Paul. "The Limits of Mass Customization." *MIT Sloan Management Review*, Spring 2001, 81-87.