Cost of Stockouts in the Microprocessor Business and its Impact in Determining the Optimal Service Level

by

Maria Claudia Sonnet

B.S., Mechanical Engineering, Universidad Tecnologica Nacional, Argentina (1997)

Submitted to the Sloan School of Management and the Department of Mechanical Engineering on May 6th, 2005 in partial fulfillment of the Requirements of the Degrees of

Master of Business Administration
and
Master of Science in Mechanical Engineering

In conjunction with the Leaders for Manufacturing Program at the
Massachusetts Institute of Technology
June 2005

© 2005 Massachusetts Institute of Technology. All rights reserved.

The author hereby grants to MIT permission to reproduce and to distribute publicly paper and electronic copies of this thesis document in whole or in part.

Signature of Author ______________________
Sloan School of Management
Department of Mechanical Engineering
May 6, 2005

Certified by ______________________
Stanley B. Gershwin, Thesis Supervisor
Senior Research Scientist, Department of Mechanical Engineering

Certified by ______________________
Donald B. Rosenfield, Thesis Supervisor
Senior Lecturer, Sloan School of Management

Accepted by ______________________
David Capodilupo, Executive Director of Masters Program
Sloan School of Management

Accepted by ______________________
Lallit Anand, Chairman, Committee on Graduate Studies
Department of Mechanical Engineering

MASSACHUSETTS INSTITUTE OF TECHNOLOGY
SEP 0 1 2005
LIBRARIES

BARKER

Quantification of the Cost of Stockouts in the Microprocessor Business
by
Maria Claudia Sonnet

Submitted to the Sloan School of Management and the Department of Mechanical Engineering on May 6^{th}, 2005 in partial fulfillment of the Requirements of the Degrees of

Master of Business Administration
and
Master of Science in Mechanical Engineering

Abstract

In order to develop optimal inventory policies, it is essential to know the consequences of stockouts and the costs related to each kind of stockout; at Intel, however, such costs have not yet been quantified. The primary goal of this internship is to quantify the cost of stockouts, focusing in the microprocessor business.

The first stage of this thesis consists of describing the different consequences of stockouts. In a stockout situation, customers may opt to buy alternate products at either higher or lower price (buy up or buy down), postpone the purchase until the product is available (postponed sales), buy the product from a non-authorized distributor (sales lost to the open market) or buy a non-Intel product (sales lost to the competition). Each of those consequences has a different financial impact, so we quantify stockout cost for different stockout situations and different service levels. This analysis is conducted at an aggregated level and also by product line.

The results of this project show that stockouts have a high financial impact in the microprocessor business. Due to the high margins, each sale lost to the competition means losing a significant dollar amount, which may easily out weigh the inventory related costs of that product. The result is service levels that are higher than Intel had believed would be appropriate. The quantification of the financial impact of the different stockout situations will be a valuable input into further supply chain optimization analysis including adjustment of finished good inventory levels.

Thesis Supervisor: Stanley Gershwin
Senior Research Scientist, Department of Mechanical Engineering

Thesis Supervisor: Donald Rosenfield
Senior Lecturer, Sloan School of Management

This page is intentionally left blank.

Acknowledgements

I would like to thank the Leaders for Manufacturing Program and Intel Corporation for providing the resources and support to make this project possible. It was a great learning experience that had a decisive impact in my professional growth.

The success of this internship was possible thanks to the endless support, collaboration and enthusiasm of my project advisor, Dennis Arnow, who allowed me to learn, grow and have fun through this amazing experience. Thank you also to my thesis advisors, Don Rosenfield and Stan Gershwin, for making this thesis possible through their advice and guidance throughout the process. Thanks also to Tony Newlin, Jen Rigoni, Rich Bridge, William Lo, Richard Post and Jim McGuire from Intel Corporation. All of them had a direct impact on this project by offering their valuable time, insights and knowledge.

I would also like to thank my husband, Bruno Tordini, for his unconditional support to my career; my friend Lou Chios, for convincing me to apply to MIT when it looked like an impossible dream; and the LFM Class of 2005, for making these last two years an unforgettable experience.

Finally, I would like to dedicate this thesis to my grandparents, Maria and Raul Sonnet and Maruca and Atilio Bersi. They complemented my education as the role models from whom I learnt integrity, respect for the others and the importance to follow my ideals.

This page is intentionally left blank.

Table of Contents

This page is intentionally left blank.

Chapter 1 – Introduction and Overview

1.1. Project Objective

This internship project was sponsored by Intel's Inventory and Demand Management Group and took place at their Santa Clara site, California, from June 2004 to January 2005. The Inventory and Demand Management Group, or IDM, is part of a highly matrixed organization that sets a common direction for multiple groups at Intel relating to supply network optimization.

In 2003, the IDM group sponsored LFM intern Joseph Levesque to study the impact of variability in Intel's supply chain. The outcome of that project was an estimate of the inventory targets needed to meet different service levels. Building on Levesque's work, Intel's IDM Group sponsored this new project with the primary goal of quantifying the cost of stockouts for different service levels. By contrasting such costs to inventory related costs, it is possible to determine the optimal relationship between inventory and service level and hence, to develop optimal inventory policies.

The project deliverables were to quantify the financial impact of stockouts aiming to develop optimal inventory strategies; to make recommendations for changes in the current service levels based on those findings; and to provide a framework for evaluating Intel's inventory policies in the future.

The scope of this project was limited to the microprocessor business, which accounts for over 80% of Intel's business. However, the ultimate goal was to provide a framework for reviewing Intel's inventory policies in other business units.

1.2. The Inventory Management Problem

Inventories play a key role in a company's operations, and have a decisive impact on its performance. Inventories allow companies to buffer fluctuations in production and demand, increase order fill rate and achieve higher customer satisfaction; but on the other hand, keeping inventory involves costs and risks. Balancing these advantages and disadvantages is a complex task that depends on a wide variety of factors, such as demand patterns, demand forecast accuracy, manufacturing process (including lead time and yield variability), ordering process, service requirements and costs factors (including cost of stockouts and inventory holding cost). These factors and their impacts on inventory management may vary considerable across industries; however, there is a trend to believe that the lower the inventory the better. This belief is due in part to the success of the Toyota Production System, which revolutionized the automotive industry in the 1970s. The Toyota Production System or TPS, based on the principles of Jidoka, Just-in-time (JIT) and Kaizen, is a major factor in the reduction of inventories and defects in Toyota plants, and is also extended to its suppliers. In fact, one of the goals of JIT is to achieve "zero inventories". However, this does not literally mean that plants have to eliminate all their inventories; it is just a goal to motivate workers and management to apply continuous improvement techniques to reduce inventory levels. The result was a system that eliminated the need for large amounts of inventories, and that lead Toyota to a sound success. In the 1980s, the Japanese automotive industry overtook that of the US, so American companies started to adopt some TPS techniques in their operations. Inventory reduction became a popular policy (enthusiastically supported by Finance managers); however, in order to maximize a company's performance, rather than just reduce inventory it is necessary to determine the optimal inventory level; that is, the right balance between inventory and service level.

There is extensive operations management literature about inventory modeling. Peterson and Silver (1985), Nahmias (1993) and Hopp and Spearman (1996) have proposed different models to address different demand and production scenarios. However, they share some common basic insights. The ones that are closely related to this work are the following (Hopp and Spearman, 1996)

- There is a tradeoff between customer service and inventory. Under conditions of random demand, higher customer service levels (i.e. fill rates) require higher levels of safety stock.
- There is a tradeoff between variability and inventory. If replenishment frequency and customer service remain fixed, then the higher the variability (i.e. standard deviation) of demand, the more inventory we must carry.

1.3. Industry Overview

The microprocessor industry is characterized by highly variable demand and complex manufacturing processes involving long lead times; causing the supply chain management to be particularly challenging. On top of that, this is a high clockspeed industry in which obsolescence occurs virtually everyday. The more inventory in the chain, the higher the obsolescence costs; and the faster the clockspeed, the higher the obsolesce costs (Fine, 1998). Given the variability of demand and production yield, and the threat of obsolescence, it is complex to find an optimal inventory level—that one that minimizes the probability of stockouts while keeping inventory holding costs and scrap costs as low as possible.

On the other hand, high margins make this industry profitable and attractive. Each additional unit sold brings a large contribution that may be lost totally or partially when a stockout occurs. As a result, in order to develop optimal inventory policies, it is essential to know the consequences and costs of stockouts.

1.4. Company Overview

Intel was founded in 1968 to build semiconductor memory products. Today, the company supplies the computing and communications industries with chipsets, microprocessors, boards, flash memory and other applications used in computers, servers and networking and communications products. Accordingly, Intel's mission is to keep a leading position as a supplier to the IT industry.

The company is divided into two main businesses, the Intel Architecture (IA) Group, and the Intel Communications Group (ICG).

The Intel Architecture (IA) Group focuses on microprocessors (CPU's) for three different kinds of platforms:

- Desktop (Desktop computers)
- Mobile (Notebook computers)
- Enterprise (Servers and workstations)

The CPU business accounts for over 80% of Intel's business. These products have a relatively slow growth in comparison with the ICG products, and it is more difficult for buyers to switch suppliers than that of ICG products. The main competitor for IA products is AMD, a relatively small (though growing) player.

The Intel Communications Group (ICG) focuses on wireless connectivity products for desktop, mobile and enterprise applications. Unlike IA products, ICG products are in a fast-growth stage, with strong competition – which makes it easier for customers to switch suppliers. Main competitors for the ICG products include strong players like Qualcomm, Samsung and Texas Instruments.

The different characteristics of the business units involve different implications for supply chain management. For the purposes of this project, we will focus in the CPU business, so the scope will be limited to the IA Group.

1.5. The Supply Chain

Intel's supply chain can be schematized as follows (Chow, 2003).

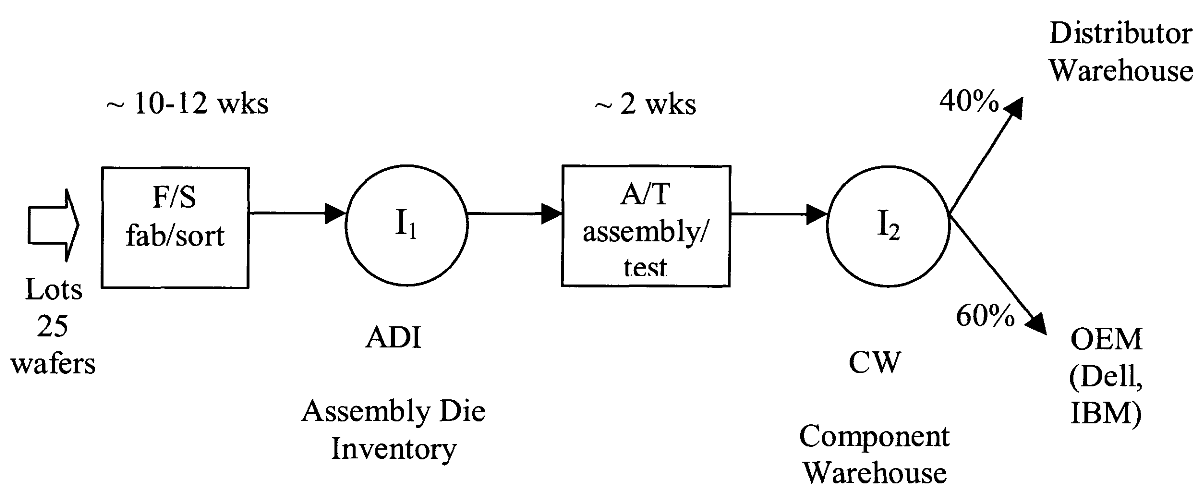

The supply chain can be divided into two main steps, the fabrication (F/S) and the assembly and test (A/T). The fabrication process consists on transforming the silicon wafers into finished die. Such die are later assembled and tested, and finally packaged into either boxes or trays, according to the distribution channel to be served. There are two main channels: original equipment manufacturer (OEM's) and Intel official distributors. Each channel has its own characteristics, resulting in different implications from the customer service point of view.

OEMs are large companies such as IBM, Dell and Hewlett Packard that purchase microprocessors directly from Intel and use them to assemble their own products. Each OEM purchases large quantities of products, hence the ordering process is complex, involving long-term planning, price bargaining and negotiations. OEMs are Intel's largest market, accounting for approximately 60% of CPU sales. Products serving this channel are desktop, mobile and enterprise and are packaged in trays.

Official Intel distributors purchase products from Intel and resell them to their customers. Such customers are small shops that assemble their own computers. Since scale of each of these businesses is very small compared to that of the OEMs, they buy relatively small quantities and they have virtually no bargaining power, making their purchasing process relatively simpler. Sales to Distributors account for the remaining 40% of CPU sales, but is growing fast. Products serving this channel are mostly desktop and are packaged in boxes.

Besides the official distributors, there are also unofficial ones (gray market) that sell both Intel and non-Intel products. They do not supply directly from Intel but from OEMs that eventually sell excess inventories.

Chapter 2 – The Cost of Stockouts Problem

2.1. Definition of Stockout

Before starting to investigate the causes of stockouts, it was necessary to adopt a clear definition of what would be considered a stockout for the purposes of this project.

After interviewing with people from different areas within the company, the conclusion was that there was not a formal, company-wide definition of stockout. So, as a first approach, we defined stockouts as

Stockout Rate = 100% – Service Level

where

Service Level = % of demand requested by customers with requested date and quantity satisfied (Chow, 2003).

According to this, a stockout would be "the inability to satisfy a customer's requirement either in specification or time". But after investigating the supply process further, it was necessary to adjust this definition by adding the concept of forced stockouts.

A forced stockout happens when the company launches a new product –upgrading an old one– thus trying to compel customers to switch from the old product to the new one. To do so, Intel intentionally causes stockouts of the old product.

On the other hand, the situation opposite to a stockout happens when a determined product is available in the moment the customer requires it, which for the purposes of this project will be defined as "in stock".

According to the definition, all the following situations classify as stockouts:

- Buy up
- Buy down
- Postponed sale

- Sale lost to the open market
- Lost sale (mostly to the competition)

2.2. Description of Different Stockout Cases

2.2.1. Buy Up

A buy up happens when a customer agrees to buy a higher priced product because the one required is not available. In this situation, Intel actually wins in terms of revenue, because higher price products yield higher margins. The downside is that Intel loses in terms of goodwill, because the customer is forced to buy a more expensive product.

Even though such product also offers a higher speed, the customer's priority tends to be price and not performance. This tends to happen especially with end of life products, since the company intentionally forces stockouts of such products to compel customers to buy the new, upgraded ones.

2.2.2. Buy Down

A buy down situation happens when a customer agrees to buy a lower price product because the one required is not available. In this situation, Intel loses revenues and margins because lower price products yield lower margins.

2.2.3. Postponed Sale

This happens when the customer agrees to wait for the required product to be delivered at a later date.

2.2.4. Sale Lost to the Open Market

This situation happens when a customer gets an Intel product from the open market because it is not available from Intel or its distributors. We define as open market all non-official distributors selling both Intel and other brand products. Such non-official distributors supply from OEMs that sell their overstocks. Hence, in this stockout situation the customer still buys an Intel product, but not through the official channels.

2.2.5. Lost Sale

Lost sales happen when customers facing a stockout decide to buy a product from the competition. This stockout situation involves not only the revenue lost with that sale but also two important secondary consequences. The first one is the risk of losing a whole stream of sales (six months, one year, etc.) because the customer may adapt its platform to a competitor's product. The second one is to lose the customer forever, which is less likely to happen, but at the same time, very difficult to quantify.

Chapter 3 – Quantification of Stockouts

3.1. Introduction

As mentioned in Chapter 1, Distributors and OEMs are channels of different nature and accordingly, their characteristics, policies and procedures vary considerably.

When looking at the stockout issue, the main problem is to get data from the customer point of view. As a first approach we modeled the decision process in case of a stockout applied to a generic customer, to be later extended to Distributors and OEMs.

3.2. Decision Tree in Case of a Stockout for a Generic Customer

Classical inventory theory considers that stockouts generate a one-time revenue loss due to the lost sale. However, it does not consider the loss of goodwill caused by the disappointment that may lead customers to change their purchasing habits. Schwartz (1966) made a contribution to this area by investigating how the demand changes after a stockout ("perturbed demand"). According to Schwartz's model, since the customer request is not completely satisfied (in terms of either product or time) there is a loss of goodwill that might cause to lose the customer in the long term.

Oral, Salvador, Reisman and Dean (1972) built up on Schwartz's model and proposed a customer decision tree following a stockout to evaluate the costs of stockouts. This model can be adapted to analyze the decision alternatives and consequences following a stockout for an Intel distributor customer (Figure 1). According to this model, when a stockout situation occurs, the customer has the following options.

(a) Still buys from this or an alternate Intel distributor, accepting either a different product or a different date.

(b) Buys an Intel product from the gray market

(c) Buys a competitor's product

In this model, it is assumed that in case (a), since the customer has accepted a different product or date, he is satisfied and will come back for future business

(RETURN), and that in cases (b) and (c), since the customer did not negotiate, he is probably unsatisfied and might either come back or not (TERMINATE). Each event has an associated unitary cost C_n^k, where n is the event index and k is the product index, and a conditional probability P_n^k

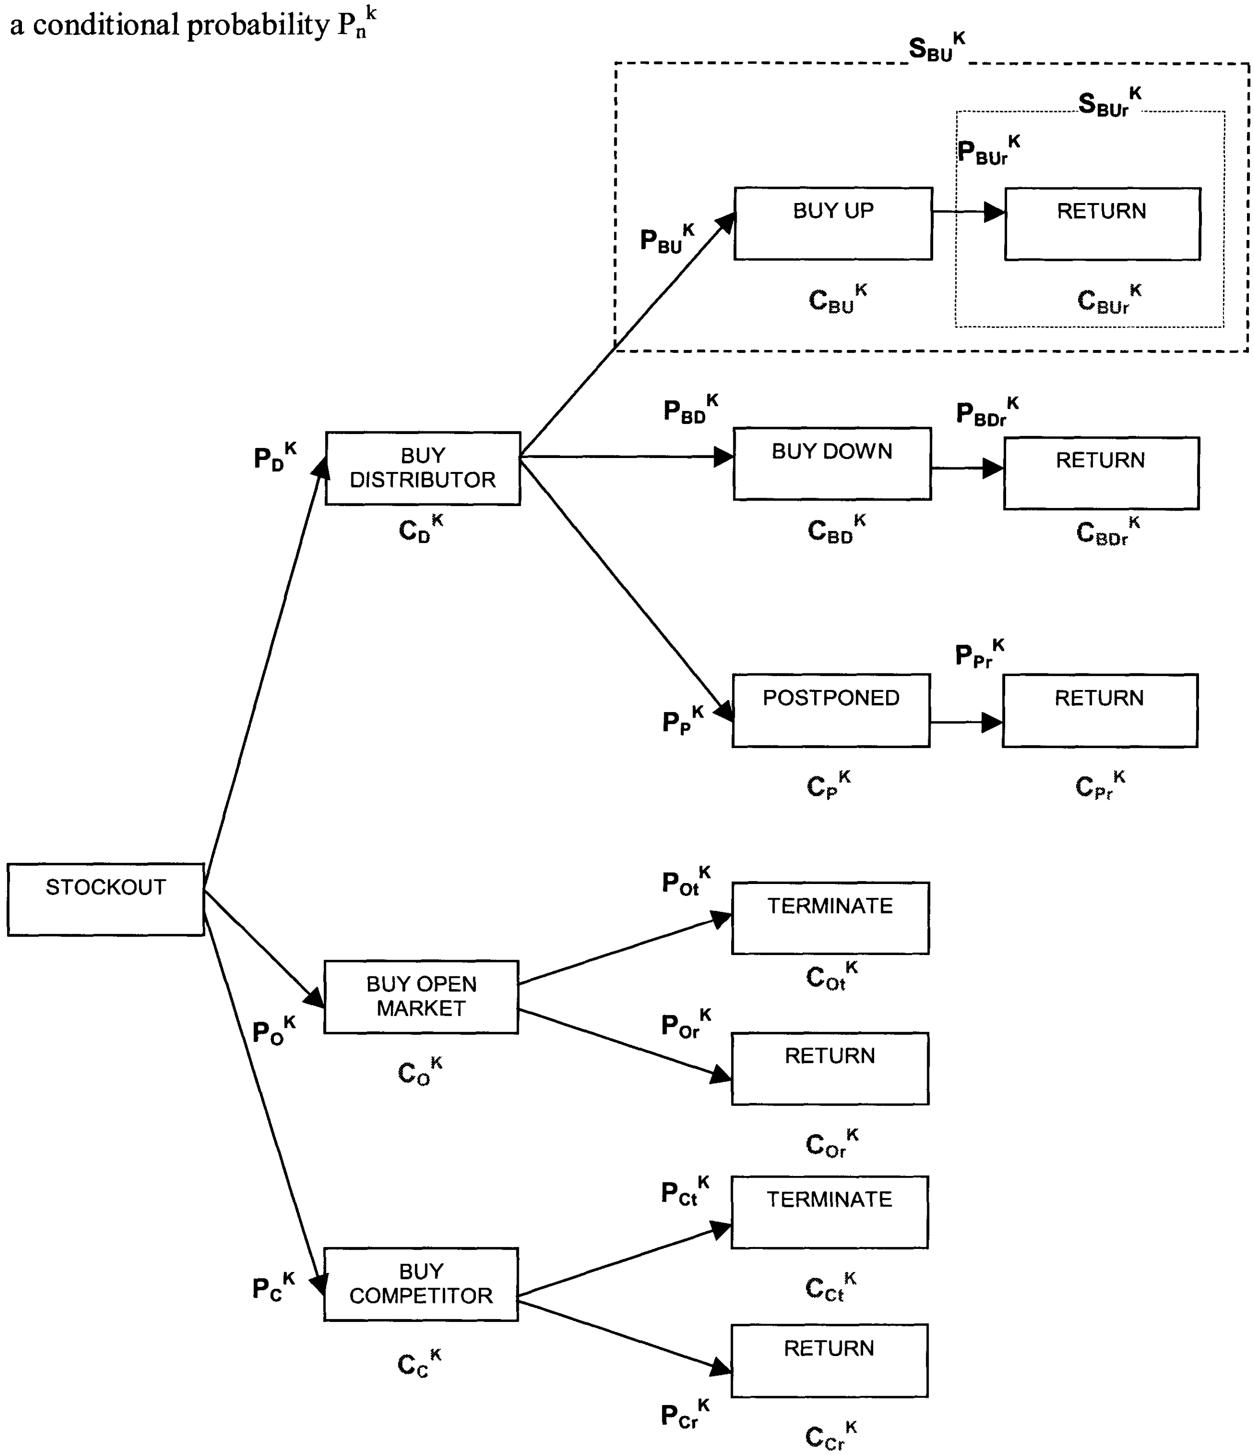

Figure 1

where

P_n^k = Probability for event n to happen for product k

C_n^k = Cost associated with event n for product k

let us define

S_n^k = Partial cost associated with event n for product k

so analyzing each node we get the following equations:

$$S_{BUr}^k = P_{BUr}^k C_{BUr}^k$$

$$S_{BU}^k = P_{BU}^k (C_{BU}^k + S_{BUr}^k) = P_{BU}^k (C_{BU}^k + P_{BUr}^k C_{BUr}^k) \rightarrow \underline{\text{Cost of buy up}} \quad [1]$$

$$S_{BDr}^k = P_{BDr}^k C_{BDr}^k$$

$$S_{BD}^k = P_{BD}^k (C_{BD}^k + S_{BDr}^k) = P_{BD}^k (C_{BD}^k + P_{BDr}^k C_{BDr}^k) \rightarrow \underline{\text{Cost of buy down}} \quad [2]$$

$$S_{Pr}^k = P_{Pr}^k C_{Pr}^k$$

$$S_P^k = P_P^k (C_P^k + S_{Pr}^k) = P_P^k (C_P^k + P_{Pr}^k C_{Pr}^k) \rightarrow \underline{\text{Cost of a postponed sale}} \quad [3]$$

$$S_O^k = P_O^k (C_O^k + P_{Ot}^k C_{Ot}^k + P_{Or}^k C_{Or}^k) \rightarrow \underline{\text{Cost of losing a sale to the open market}} \quad [4]$$

$$S_C^k = P_C^k (C_C^k + P_{Ct}^k C_{Ct}^k + P_{Cr}^k C_{Cr}^k) \rightarrow \underline{\text{Cost of losing a sale to the competition}} \quad [5]$$

so the expected unit stockout cost for item k will be

$$E^k = P_D^k (C_D^k + S_{BU}^k + S_{BD}^k + S_P^k) + S_O^k + S_C^k$$

In order to continue with this analysis, we assume that no costs are incurred by the company if a customer returns after experiencing a stockout, so

$$C_{BUr}^{k} = C_{BDr}^{k} = C_{Pr}^{k} = C_{Or}^{k} = C_{Cr}^{k} = 0$$

If a customer terminates business for an item k, it can be assumed that such customer will resume business after a time Tk [months]. The cost of a terminated business relationship on a customer basis can be approximated as

$$C_{Ot}^{k} = C_{Ct}^{k} = Y = \frac{g_k \lambda_k D_k}{N_k}$$

where

D_k = annual dollar sales of product k

N_k = number of customers

g_k = gross profit % of item k

λ_k = discount factor for item k (at compound interest rate i)
$= [(1+i)^{Tk} - 1] / i\,(1+i)^{Tk}$

However, for the purposes of this study we will focus on the immediate financial impact of stockouts, disregarding the long-term effects. The assumption that the loss of goodwill has virtually no impact on the optimal service level was validated by the sensitivity analysis and will be discussed in the Recommendations section. In consequence, we will disregard C_{Ot}^{k} and C_{Ct}^{k}.

We also assume that if the customer accepts to postpone a purchase, buy up or buy down, he is satisfied and will come back again for new businesses. Hence,

$$P_{BUr}^{k} = P_{BDr}^{k} = P_{Pr}^{k} = 1$$

and applying the beforehand mentioned assumptions to [1], [2], [3], [4] and [5], we get that

$S_{BU}^{k} = P_{BU}^{k} C_{BU}^{k} \rightarrow$ Cost of buy up

$S_{BD}^{k} = P_{BD}^{k} C_{BD}^{k} \rightarrow$ Cost of buy down

$S_{P}^{k} = P_{P}^{k} C_{P}^{k} \rightarrow$ Cost of a postponed sale

$S_{O}^{k} = P_{O}^{k} C_{O}^{k} \rightarrow$ Cost of losing a sale to the open market

$S_{C}^{k} = P_{C}^{k} C_{C}^{k} \rightarrow$ Cost of losing a sale to the competition

so the total cost of stockouts for product line k will be:

$$S^{k} = S_{BU}^{k} + S_{BD}^{k} + S_{P}^{k} + S_{O}^{k} + S_{C}^{k}$$

$$S^{k} = P_{BU}^{k} C_{BU}^{k} + P_{BD}^{k} C_{BD}^{k} + P_{P}^{k} C_{P}^{k} + P_{O}^{k} C_{O}^{k} + P_{C}^{k} C_{C}^{k} \qquad [6]$$

3.3. Estimation of Probabilities for each Stockout Situation

The analysis aiming to estimate the probability for each stockout situation to happen (P_{n}^{k}) is based on the following already existing frameworks.

1) MMBP (Microprocessor Marketing & Business Planning) Customer Survey (Q3 2003)

2) MMBP Geographic Representatives Survey (Q3, 2004)

3.3.1. Distributor Customer Survey

In the third quarter of 2003, the MMBP (Microprocessor Marketing & Business Planning) conducted a survey among Intel distributor's final customers (small shops that assemble and sell their own computers) in order to analyze their behavior when the products they required were not available in either specification or time.

The survey was sent to customers in both mature and emerging markets in all three Intel geographic areas: EMEA (Europe, Middle East and Africa), APAC (Asia Pacific) and NAMO (North America) and was answered by more than 300 final customers. They were asked what percentage of times they found the product they requested in both specification and time (in stock situation), and what were their most typical reactions when the product was not available in specification, time of both.

The results are displayed in the table below. For confidentiality reasons, the actual figures are not shown. It is important to bear in mind that since this information comes from the Distributors channel, it can be only applied to desktop (DT) products.

	P_n for DT Dist
$P_{Buy\ Up}$	Moderate
$P_{Buy\ Down}$	Low
$P_{Postponed}$	High
$P_{Open\ Mkt}$	High
P_{Lost}	Low
Total	100.0%

Table 1

3.3.2. Geographic Representatives Survey

In order to complement the above data, the MMBP group conducted a new survey aiming to obtain information on the OEM's behavior in case of stockouts. This survey was answered by the Geographic Representatives, who manage the ordering process between OEM's and Intel in the different geographic markets served. The information is broken down by product line: desktop (DT), mobile (MOB) and server (SVR).

	P_n		
	DT OEM	MOB	SVR
$P_{Buy\ Up}$	Null	Null	Null
$P_{Buy\ Down}$	Low	High	Low
$P_{Postponed}$	High	High	Extremely high
$P_{Open\ Mkt}$	High	Moderate	Null
P_{Lost}	Moderate	Low	Low
Total	100.0%	100.0%	100.0%

Table 2

3.3.3. Estimation of Probabilities for each Stockout Situation at Aggregated Level

From Tables 1 and 2, we know P_n for each distribution channel and product line. Since the volumes sold through each channel are known, we can estimate P_n at the aggregate level by doing a weighted average (Table 3). For example, $P_{Buy\ Down}$ on Table 3 is the following weighted average.

$$P_{Buy\ Down\ Aggregated} = P_{Buy\ Down\ DT\ Dist} * Vol_{DT\ Dist} + P_{Buy\ Down\ DT\ OEM} * Vol_{DT\ OEM} + $$
$$+ P_{Buy\ Down\ MOB} * Vol_{MOB} + P_{Buy\ Down\ SVR} * Vol_{SVR}$$

where

$P_{Buy\ Down\ Aggregated}$ = Probability of buy down at the aggregated level

$P_{Buy\ Down\ DT\ Dist}$ = Probability of buy down for DT line in the distributors channel

$Vol_{DT\ Dist}$ = Sales volume for DT line through the distributors channel

$P_{Buy\ Down\ DT\ OEM}$ = Probability of buy down for DT line in the OEM channel

$Vol_{DT\ OEM}$ = Sales volume for DT line through the OEM channel

$P_{Buy\ Down\ MOB}$ = Probability of buy down for MOB line

Vol_{MOB} = Sales volume for MOB line

$P_{Buy\ Down\ SVR}$ = Probability of buy down for SVR line

Vol_{SVR} = Sales volume for SVR line

	P_n at aggregated level
$P_{Buy\ Up}$	Very low
$P_{Buy\ Down}$	Low
$P_{Postponed}$	High
$P_{Open\ Mkt}$	High
P_{Lost}	Moderate
Total	100.0%

Table 3

3.4. Estimation of Financial Impact for each Stockout Situation

The next step is to determine the cost associated with each stockout situation (C_n^k). For the purposes of this study we will focus on the short-term financial impact of stockouts, disregarding the long-term effects due to loss of goodwill.

3.4.1. Financial Impact of Buy Ups

Buy ups happen when due to the stockout of a given product, customers opt for buying an alternate product that is higher in price than the originally requested one. Intel has estimated the average increase on price in case of buy ups, but for confidentiality reasons the actual figure cannot be shown so we will call it U. Since the cost of producing any chip is roughly the same, it can be assumed that buy ups result in a revenue increase to Intel. In other words, buy ups can be considered "negative costs" and can be estimated by

$$C_{\text{Buy Up}} = -\ \text{Buy Up Volume} * \text{ASP} * \text{U\%}$$

where ASP is the average selling price and U% the average ASP uplift when customers buy up.

3.4.2. Financial Impact of Buy Downs

When customers buy down, they buy an alternate product that is lower in price than what they originally requested. Intel has estimated the average hit on price in case of buy downs (i.e. how much less customers pay on average), but for confidentiality reasons the actual figure cannot not be shown so we will call it D. Since the cost of producing any chip is roughly the same, it can be assumed that when a buy down happens, it will result in a revenue decrease to Intel. Buy downs can be estimated by

$$C_{\text{Buy Down}} = \text{Buy Down Volume} * \text{ASP} * \text{D\%}$$

where ASP is the average selling price and D% the average ASP hit when customers buy down.

3.4.3. Financial Impact of Postponed Sales

Sometimes customers facing a stockout situation are willing to wait until the requested product is available. In such a situation, Intel does not lose the sale and there is no impact on the revenue, but there is an opportunity cost associated with the time value of money. OEMs are more likely to wait for an order to be delivered since they primarily demand high end products that have no potential replacement, while distributor's customers will wait a few days only. As a result, the average wait can be approximated as 30 days, and the cost of postponed sales can be estimated from

$$C_{postponed} = \text{Postponed Volume} * \text{Average Contribution} * \text{Time Value of Money}$$

where

Average Contribution = ASP – Average Variable Cost

Time Value of Money (TVM) = $(1+\text{WACC})^{\text{Average Days of Delay}/365} - 1$

WACC = weighted average cost of capital = 15%

Average Days of Delay = 30

3.4.4. Financial Impact of Sales Lost to the Open Market

As mentioned in 2.2.4., this situation happens when customers get Intel products from the open market, or non-official distributors selling both Intel and other brand products. For the purposes of this study, we consider that when customers supply from the open market, they buy Intel products only (Otherwise, such sales would fall under the "lost sales" classification). Since the open market supplies from OEMs that sell their overstocks, and OEMs supply directly from Intel, it can be assumed that Intel is not really missing such sales so the financial impact is zero.

$$C_{OpenMkt} = 0$$

3.4.5. Financial Impact of Lost Sales

Lost sales happen when customers facing a stockout either buy a competitor's product; hence the financial impact is the loss of the whole margin, so the cost of lost sales can be estimated by the following formula.

$$C_{Lost} = \text{Lost Volume} * \text{Average Margin}$$

There is also another type of lost sale, known as sale evaporation. Sale evaporations happen when for some reason customers decide to cancel an order and not to buy any product at all. Since sale evaporations are relatively rare, and their financial impact is the same than that of sales lost to the competition, so for the purposes of this study we will refer to them as lost sales.

3.5. Cost of Stockouts Model

Now we know the different situations associated with stockouts and the probability for each of them to happen for the current service level; and we can estimate the financial impact associated with each of them. Applying formula [6] of Chapter 3.2., we can finally get the cost of stockouts for the current service level by adding up the financial impacts of each stockout situation.

$$S^k = P_{BuyUp}{}^k\, C_{BuyUp}{}^k + P_{BuyDown}{}^k\, C_{BuyDown}{}^k + P_{Postponed}{}^k\, C_{Postponed}{}^k + P_{OpenMkt}{}^k\, C_{OpenMkt}{}^k + P_{Lost}{}^k\, C_{Lost}{}^k$$

where k corresponds to the aggregated level.

In order to estimate S^k for the current service level, we need to estimate how the stockout rate breaks down. We can do so by multiplying the figures in Table 3 times the stockout rate. The figures cannot be shown for confidentiality, so we will use the following nomenclature:

L%: low rate

M%: moderate rate

H%: high rate

so we can express the results as

$P_{Buy\ Up} = L\%$

$P_{Buy\ Down} = M\%$

$P_{Postponed} = H\%$

$P_{Open\ Mkt} = H\%$

$P_{Lost} = M\%$

Now we need to estimate each stockout situation cost for the current service level. From 3.4., we know that:

$C_{Buy\ Up}$ = - Buy Up Volume * ASP * U%

$C_{Buy\ Down}$ = Buy Down Volume * ASP * D%

$C_{Postponed}$ = Postponed Volume * Average Contribution * TVM

$C_{OpenMkt} = 0$

C_{Lost} = Lost Volume * Average Margin

where

Buy Up Volume = L% * Average Demand

Buy Down Volume = M% * Average Demand

Postponed Volume = H% * Average Demand

Buy Up Volume = M% * Average Demand

Data such as ASP's, average contribution, average margin and average demand cannot be disclosed due to confidentiality issues, but is readily available at Intel.

Moreover, if we assume that the customer response will maintain a linear relationship (in terms of buy up, down, postponed sales, etc.) as service level varies, we can model the cost of stockouts for different service levels. The outcome of this model is shown in Graph 1.

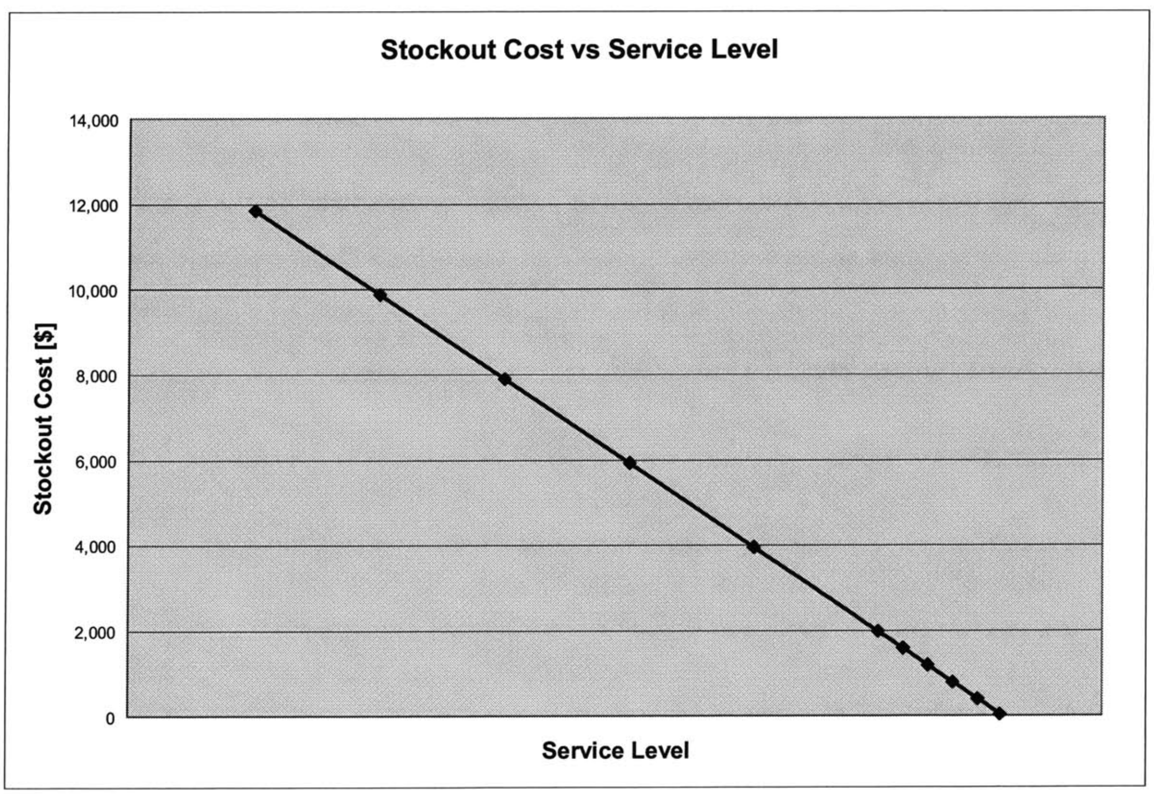

Note: scale normalized to protect confidentiality

Graph 1

By definition, stockouts decrease with service increases, and so does stockout cost. But each stockout case (Buy up, buy down, postponed, lost) will have a different weight in the total stockout cost, so we need to compare each of them in order to determine which are the main stockout cost drivers.

Graph 2 shows the financial impact for each stockout case as a function of service level (Sales lost to the open market, since their cost is assumed to be zero, are not shown for clarity) and Graph 3 shows the cost by stockout case as a percentage of total stockout cost. We can see that the main cost driver is lost sales and the next in importance is buy downs. Postponed sales cost is relatively small and roughly offset by buy ups, which can be considered a "negative cost".

Graph 4 shows that in terms of stockout volume, the largest is due to postponed sales, closely followed by sales lost to the open market, which account for most of the stockout volume, while the remaining is shared by buy downs, lost sales and to a minor extent, buy ups.

The main conclusion here is that lost sales, which account for a relatively small part of the total stockout volume, drive most of the stockout cost. This is due to the fact that when a sale is lost, the whole margin is lost.

As mentioned in Chapter 1, the microprocessor business is characterized by high margins and low variable costs. Hence, each lost sale results in a big financial loss, while in the case of buy ups or buy downs, the result is just a slight change in revenue, and in the case of postponed sales, the loss is just the time value of money, because the sale is going to be carried on later.

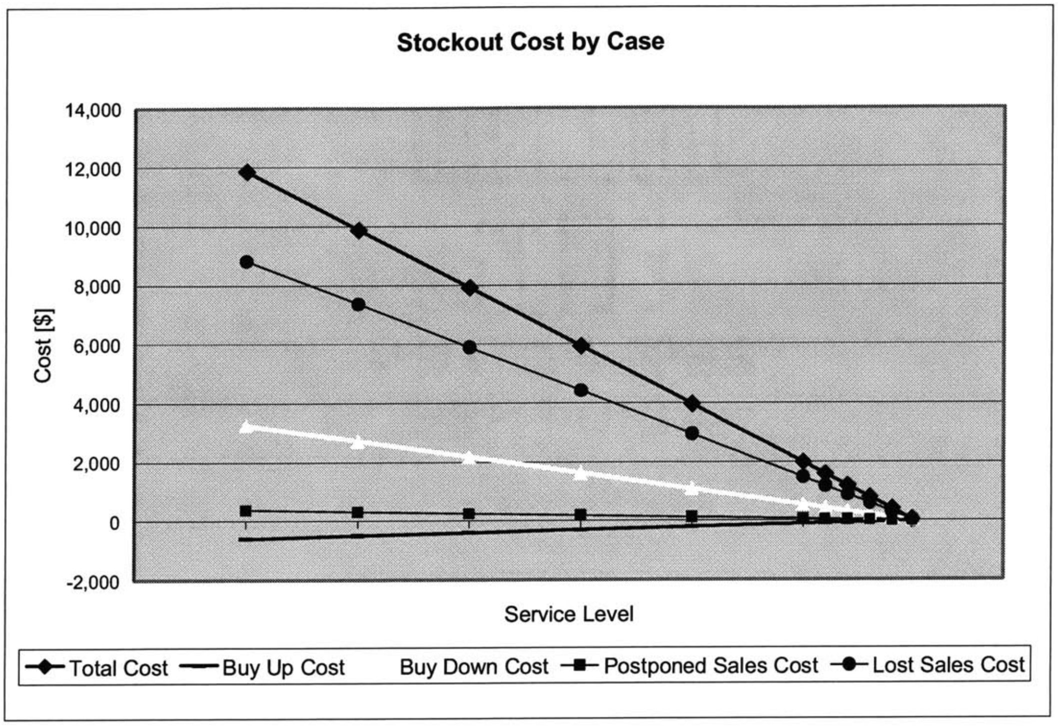

Note: scale normalized to protect confidentiality

Graph 2

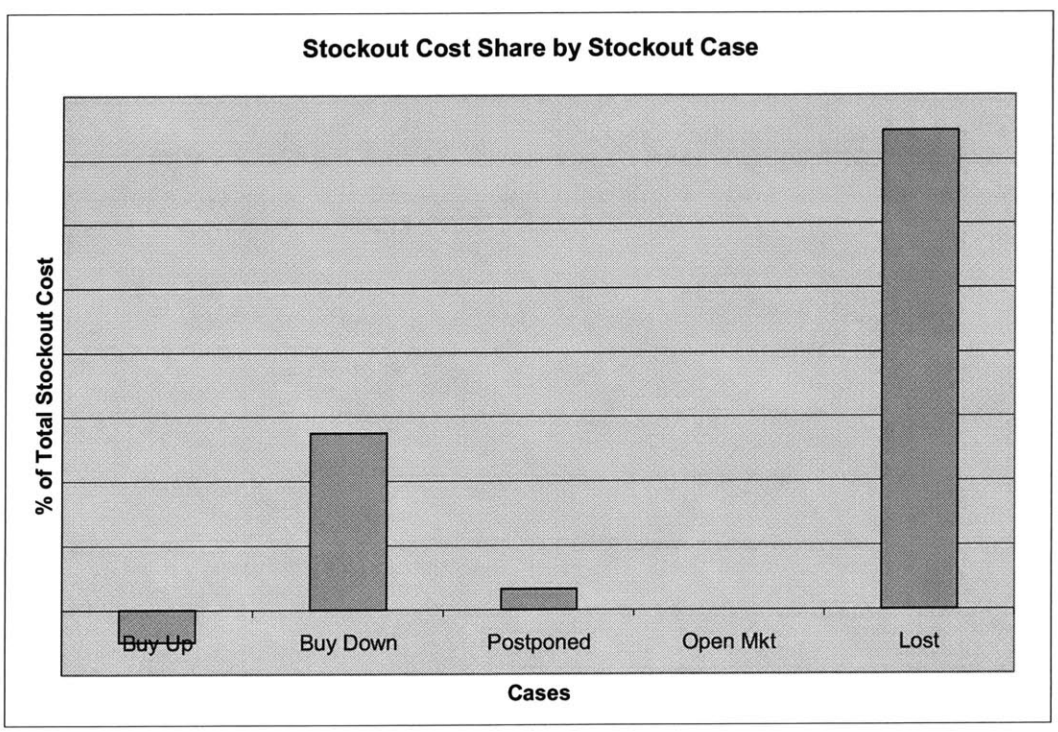

Graph 3

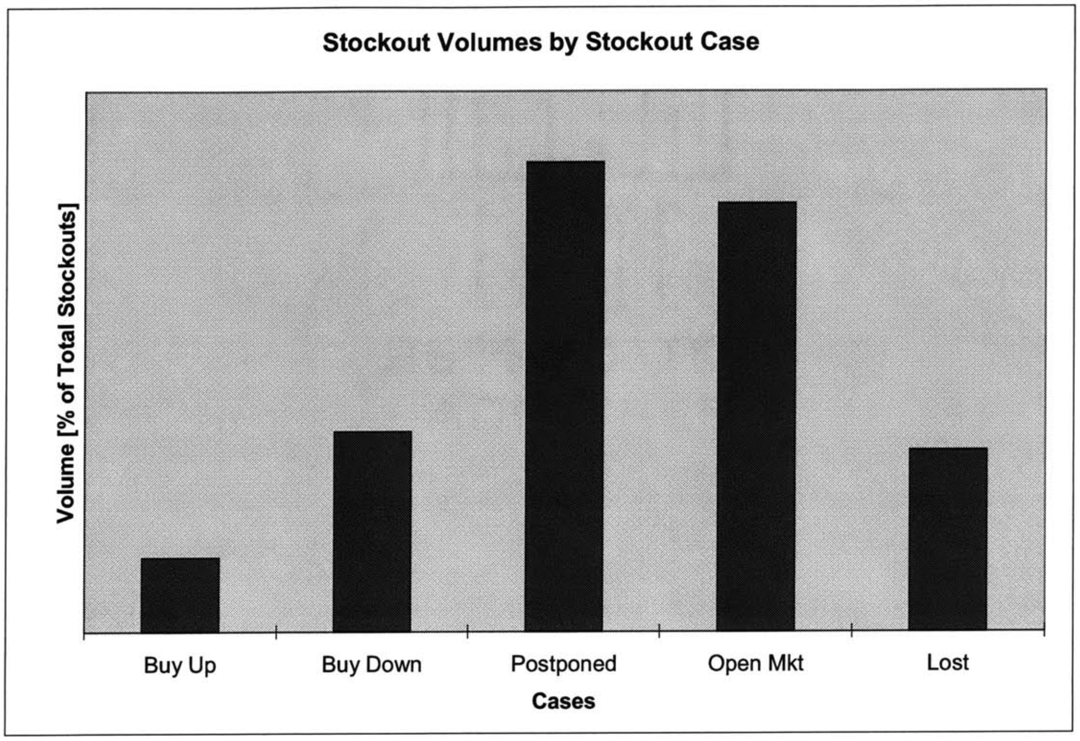
Stockout Volumes by Stockout Case
Volume [% of Total Stockouts]
Buy Up
Buy Down
Postponed
Open Mkt
Lost
Cases

Graph 4

Chapter 4 – Determining the Optimal Service Level

4.1. Introduction

Graph 1 shows that stockout cost decreases as service level increases. However, in order to increase service level, it is necessary to increase inventory, which involves increasing other costs such as inventory holding cost and inventory scrap cost. Hence, the optimal service level will be that at which the addition of all costs (stockout cost, inventory holding cost and scrap cost) reaches a minimum.

4.2. The Inventory Holding Cost

Nahmias (1989) defines inventory holding cost as "the sum of all costs that are proportional to the amount of inventory physically on hand at any point in time". Hence, in order to estimate the inventory holding cost for different service levels, it is necessary to know the safety stock required to achieve each service level. At Intel, this issue was researched by Levesque (2003). His work explores the variability in Intel's supply chain, and based on that, calculates the inventory targets required to meet different service levels. On the production side, Levesque analyzed the variability of in throughput time, yield and other related factors; and on the demand side, he evaluated the forecast error and used it as a proxy for demand variability.

Levesque based his model on a version of the base-stock model described by Zimmerman et al (1974). Under this system, the safety stock can be calculated as

$$\text{Base Stock} = \mu_d (r + \mu_{LT}) + z\,[\sigma_d^2 (r + \mu_{LT})]^{1/2} \qquad [7]$$

where:

μ_d = average demand rate over lead time

r = review period

μ_{LT} = average lead time or throughput time

z = safety factor calculated from service level

σ_d = variability in demand

The average demand times the review period in [7] represents the cycle stock, which is relevant for warehousing systems and batch production; however, since Intel's production system is continuous, Levesque removes so the cycle stock (μ_d r) from his analysis and focuses on safety stock. The safety stock portion of the base-stock model is then extended to account for variability in lead time as shown below. The additional term in the equation represents the demand variability caused by variability in lead time.

$$\text{Safety Stock} = z\,[\sigma_d^2\,(r + \mu_{LT}) + \mu_d^2\,\sigma_{LT}^2\,]^{1/2} \qquad [8]$$

where:

μ_d = average demand over lead time

σ_{LT} = variability in lead time

Similarly, the model can be extended to account for variability in production yield (Black, 1998).

$$\text{Safety Stock} = z\,[\sigma_d^2\,(r + \mu_{LT}) + \mu_d^2\,\sigma_{LT}^2 + (\sigma_S^2 \mu_d\, \mu_{LT}) / \mu_S\,]^{1/2} \qquad [9]$$

where:

μ_S = average yield

σ_S = variability in yield

Levesque then he applied these data to the two-node base-stock model shown below to calculate the inventory targets required to meet different service levels.

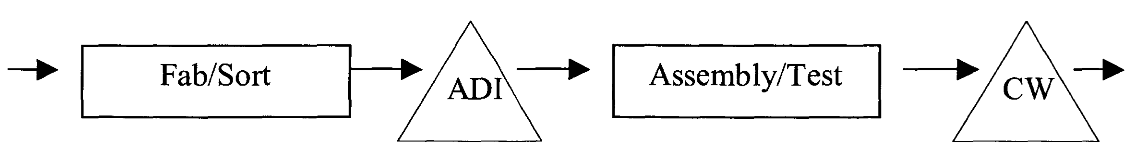

ADI: assembled die inventory

CW: component warehouse

Aggregating all Intel's CPU products, for both distribution channels, the relationship between service level and security is shown in Table 4 (actual values are not shown for confidentiality).

Overall Fill Rate	Overall Stockout Rate	Weeks of safety stock
70.0%	30.0%	W_{70}
75.0%	25.0%	W_{75}
80.0%	20.0%	W_{80}
85.0%	15.0%	W_{85}
90.0%	10.0%	W_{90}
95.0%	5.0%	W_{95}
96.0%	4.0%	W_{96}
97.0%	3.0%	W_{97}
98.0%	2.0%	W_{98}
99.0%	1.0%	W_{99}
99.9%	0.1%	$W_{99.9}$

Table 4

Knowing the safety stock in terms of weeks of inventory corresponding to each service level and the aggregated demand, it is possible to calculate the safety stock in terms of units.

$$\text{Safety Stock [units]} = \text{Weekly Demand} * \text{Safety Stock [weeks]}$$

The next step is to determine the inventory costs for each service level. According to Nahmias (1989) the components of the holding cost include "a variety of seemingly unrelated items" such as the cost of providing the physical space to store the items; taxes and insurance; breakage, deterioration and obsolescence; and the opportunity cost of alternative investments.

At Intel, the inventory holding cost was modeled by Bridge (2004) as a function of the variable cost of building inventory. The inventory holding cost modeled by Bridge includes other related costs such as opportunity cost, inventory obsolescence, facilities and staff, etc, but no scrap costs associated to excess inventories. The actual result of his work cannot be disclosed for confidentiality reasons, so we will refer to it as I%, where I expresses the inventory holding cost as a percentage of the variable cost of building inventory. Now we can quantify the inventory holding cost for each service level as follows.

Inventory Holding Cost [$] = I%*100 *Average Variable Cost * Safety Stock [units]

Data such as variable cost and demand cannot be disclosed due to confidentiality issues, but is readily available at Intel, so the inventory holding cost for each service level can be easily calculated.

4.3. The Scrap Cost

In order to complement the inventory holding cost aforementioned, it necessary to calculate the scrap cost due to excess inventory. According to Intel's historical data, such cost is 1.4% of the inventory volume.

Scrap Cost [$] = Average Variable Cost * Scrap Rate * Safety Stock [units]

Data such as variable cost and demand cannot be disclosed due to confidentiality issues, but is readily available at Intel, so the scrap cost for each service level can be easily calculated.

4.4. Total Inventory Cost

Adding together the inventory holding cost and the scrap cost, we get the total cost related to inventory, which unlike stockout cost, tends to increase with service level. Results are shown in Graph 5.

4.5. The Optimal Service Level

From Graphs 1 and 2, we know that stockout cost decreases with service level, while inventory cost increases. Hence, for each service level, we will have a cost related to stockouts and a cost related to inventory. The addition of both, or total cost, is shown in Graph 6. Hence, the optimal service level is that where the total cost, or addition of stockout and inventory costs, reaches a minimum.

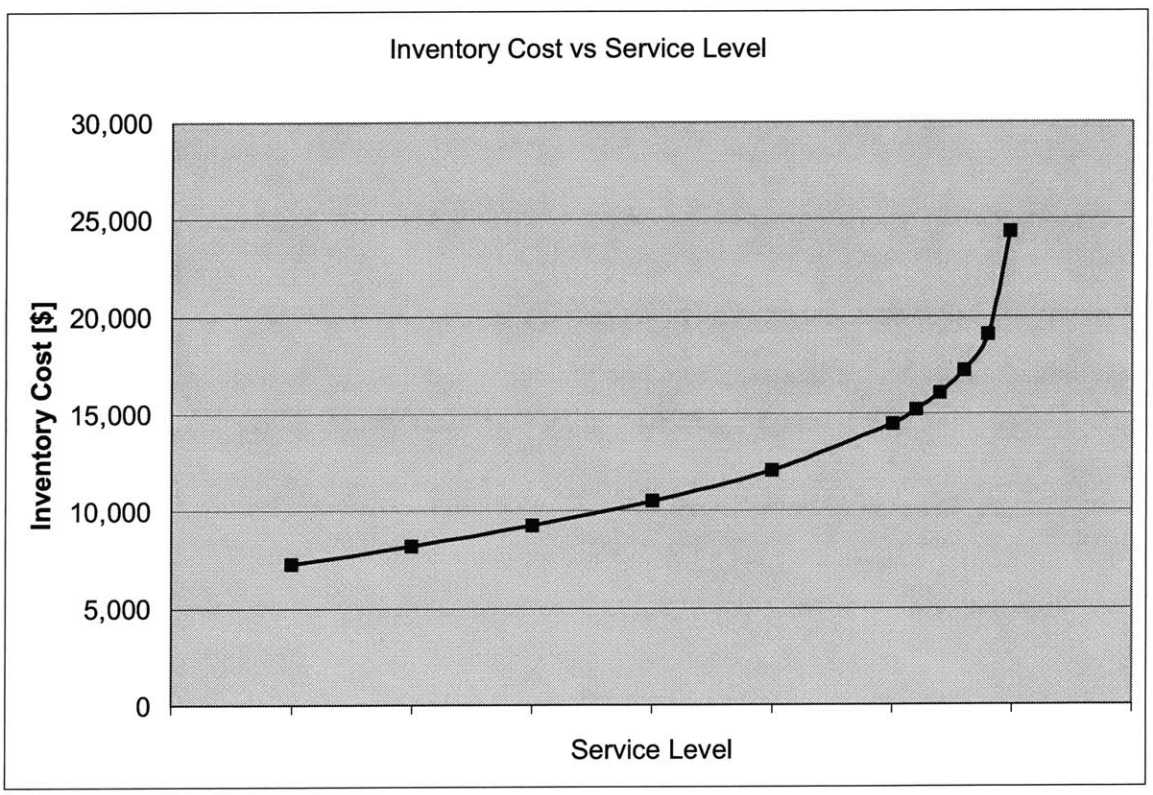

Note: scale normalized to protect confidentiality

Graph 5

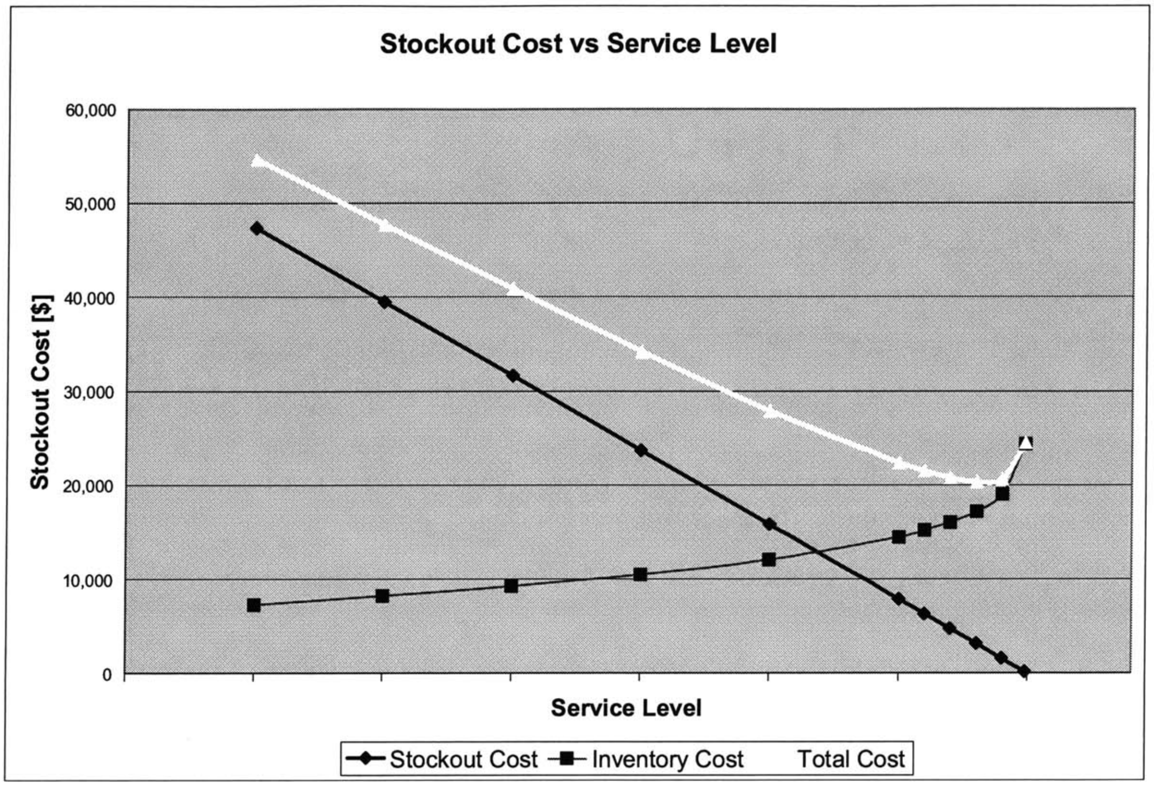

Note: scale normalized to protect confidentiality

Graph 6

Even though stockout costs decrease while inventory costs increase, we can see that the total cost keeps decreasing and only picks up slightly for extremely high service levels close to 100%. In other words, the impact of inventory costs is almost negligible in comparison to stockout costs, causing the total cost to reach a minimum for a service level close to 100%.

Chapter 5 – Cost of Stockouts for Different Product Lines

5.1. Introduction

On Chapter 3, we estimated the stockout cost at the aggregated level, and we concluded that despite its relatively small volume, lost sales are the main stockout cost drivers. However, at a more granular level some generalizations no longer apply. At the aggregated level, we used an average selling price, but the price range varies greatly across product lines. This causes stockouts in high price product lines to have a larger impact that in low price ones. Moreover, customer behavior may be completely different across product lines. One of the main differences is that some product lines are marketed mostly to OEMs, and OEMs never buy up; if they get a higher price product due to a stockout, they have the market power to negotiate the same price as the stocked out item. The other important difference is that some product lines encounter virtually no competition, causing lost sales to be almost inexistent.

In sum, customer responses in case of a stockout are very different for each segment, and so are prices and margins, even though the variable cost is roughly the same. As a result, in order to quantify the cost of stockouts by product line, we need to analyze the customer behavior and find out what the stockout consequences are in each case.

As mentioned in Chapter 1, the main product lines are Desktop (DT), Mobile (MOB) and Enterprise or Server (SVR). Furthermore, the Desktop and Mobile lines can be segmented as Performance, Mainstream, Seam and Value, depending on their performance and price, and the Server line can be segmented in multi-processor (MP) and dual-processor (DP) servers.

The different stockout consequences, prices and sales volumes by product line are summarized in Table 5.

	Price	Volume	Stockout Consequences				
			Buy Up	Buy Down	Postponed	Open Mkt	Lost
Desktop							
Performance	High	Very Low	Null	Moderate	Extremely high	Very low	Very low
Mainstream	Low	High	Low	Moderate	High	High	Low
Seam & Value	Low	High	Low	Low	High	High	Mod
Mobile							
Performance	High	Very Low	Null	High	Very high	Null	Low
Mainstream	Moderate	Moderate	Null	High	Very high	Null	Low
Seam & Value	Low	Low	Null	Moderate	Very high	Null	Mod
Server							
MP	Very High	Very Low	Null	Moderate	80%	Null	Very low
DP	Moderate	Low	Null	Very low	80%	Null	Mod

Table 5

5.2. Desktop (DT) Product Line

Desktop chips are marketed through distributors and OEMs. Performance chips are sold to OEMs, while Mainstream, Seam and Value chips are sold to both OEMs and distributors.

Performance chip stockouts tend to result in either postponed sales or buy downs, and to a minor extent, in sales lost to the open market. Lost sales are extremely low since there is no competition in this segment; the few lost sales that may happen are due to evaporation. Buy ups are also nonexistent; not only because the main market is OEMs but also because there is no chance to buy up since they are the top priced items in the line.

Due to the high selling prices, stockouts are highly undesirable; however, such stockouts rarely result in lost sales, the main stockout cost driver, which tends to balance the financial impact.

In the case of Mainstream chips, the main difference is that the open market for them is more active, hence stockouts result in postponed sales, sales lost to the open market and buy downs, and to a lesser extent, in buy ups and lost sales. Mainstream chip prices are relatively low in comparison to Performance ones, which makes stockouts less undesirable; however, stockout consequences are worse due to the increase of buy downs and lost sales. Besides, sale volumes are much higher, so when service level goes down, losses due to stockouts scale.

The case of Seam and Value chips is completely different, since both the open market and the competition are aggressive here. Most stockouts result in sales lost to the open market, and lost sales increase dramatically. Postponed sales are also important, and there are also buy ups and buy downs, that tend to balance each other. Seam and Value margins are slightly lower than Mainstream ones, but stockout consequences are tougher due to the increase of lost sales. Sale volumes are also very high, so when the service level goes down, losses due to stockouts scale.

The variation of stockout cost with service level for each Desktop product line can be observed in Graph 7. It is clearly shown that Performance chip stockouts have the lower financial impact; this is because despite their high price, the competition is not significant so lost sales are virtually zero. It can also be observed that Mainstream and Seam & Value chip stockouts have the largest financial impact; this is because despite their relatively low prices, the competition in this segment is very aggressive and most stockouts result in lost sales and hence, in the loss of the full margin.

However, if we set sales volumes aside and look at unit stockout costs, it is clear that they are directly proportional to the selling prices; hence they are higher for Performance chips and lower for Mainstream, Seam and Value (Graph 8).

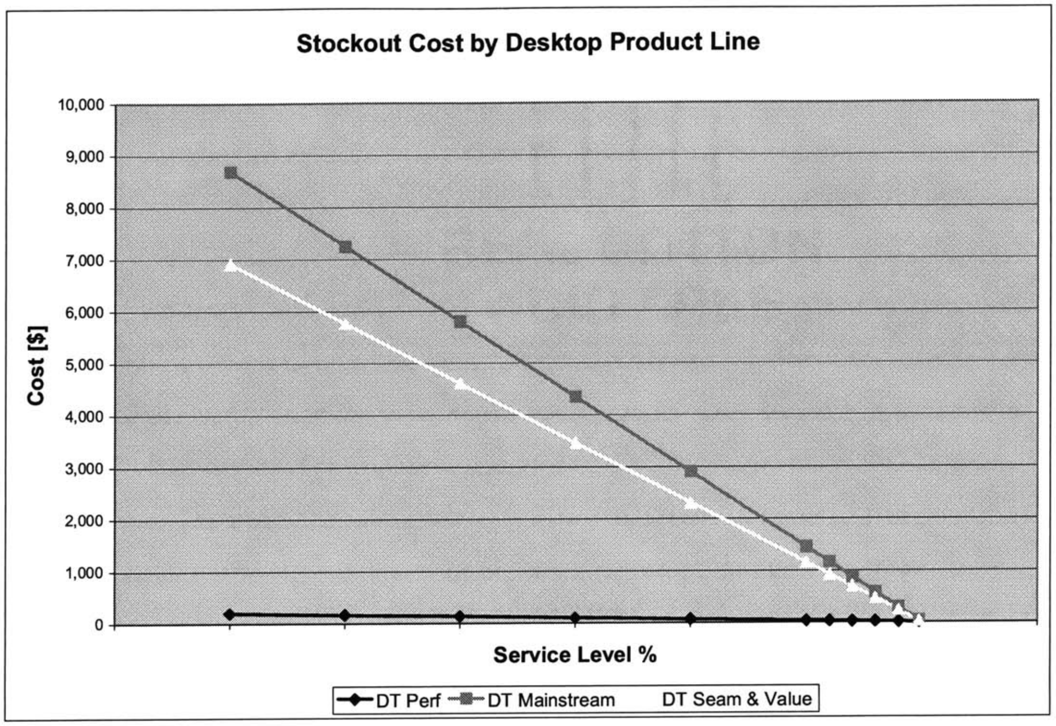

Note: scale normalized to protect confidentiality

Graph 7

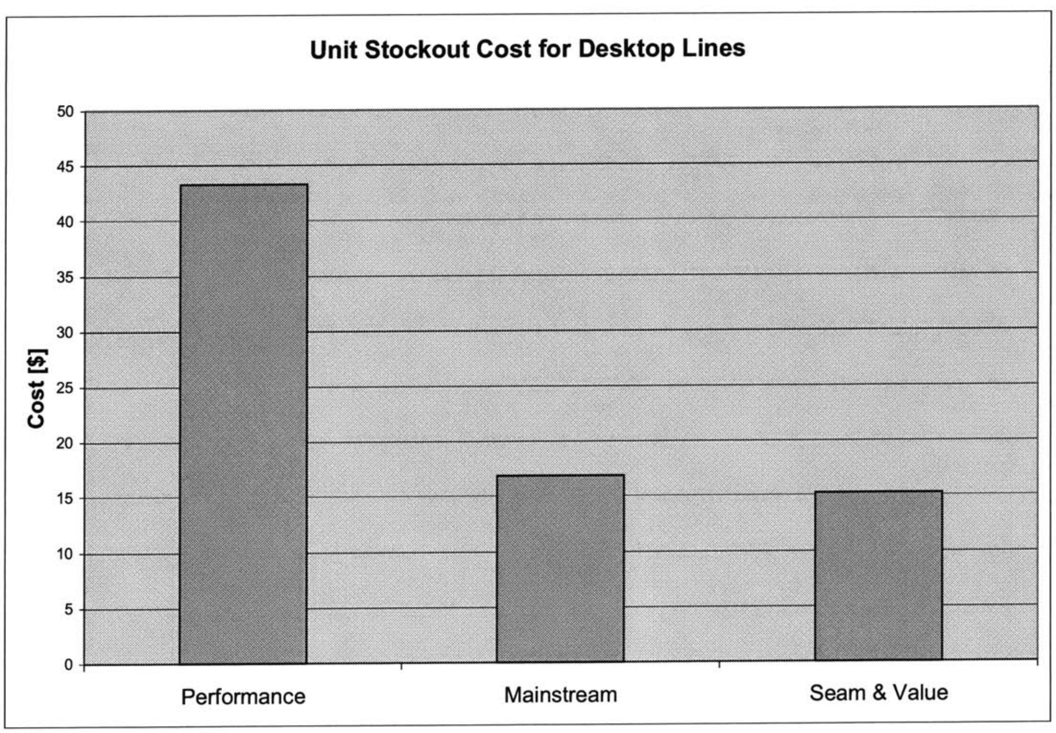

Note: scale normalized to protect confidentiality

Graph 8

5.3. Mobile (MOB) Product Line

Like the Desktop product line, the Mobile line comprises Performance, Mainstream, Seam and Value chips. About 95% of Mobile customers are OEMs, so for the purposes of this analysis we will consider OEMs to be the only market. There is no open market for Mobile chips, and since they are marketed to OEMs only, there are no buy up cases.

Performance chip stockouts tend to result in either postponed sales or buy downs, and to a minor extent, in lost sales. Like in the case of the Desktop line, Mobile Performance chip price is high, which makes stockouts highly undesirable; however, such stockouts rarely result in lost sales -the main stockout cost driver- which tends to balance the financial impact.

Customer response is the same for Mainstream chips. Prices are still important, but not as high as for the Performance chips. However, sale volumes are relatively high, so when service level goes down, losses due to stockouts scale.

For Seam and Value chips, customer response changes due to the presence of strong competition. Stockouts still result mostly in postponed sales; but now lost sales are almost as frequent as buy downs. However, margins and volumes are relatively low, which tends to offset the increase of lost sales.

The variation of stockout cost with service level for each Mobile product line is shown in Graph 9. It can be observed that Performance and Mainstream chips have very different financial impact despite having the same customer behavior patterns. Performance chips stockout costs are much lower than Mainstream. This is because Performance chips have a higher price while sales volume is relatively low. On the other hand, Mainstream chips have a relatively low price, while their sales volume is large. Seam and Value chips face the tougher customer response pattern, with high lost sales; however, their low price and sales volume offset the financial impact of lost sales.

If we set sales volumes aside, it is clear that unit stockout costs are directly proportional to selling prices; hence they are higher for Performance chips, intermediate for Mainstream and lower for Seam and Value (Graph 10).

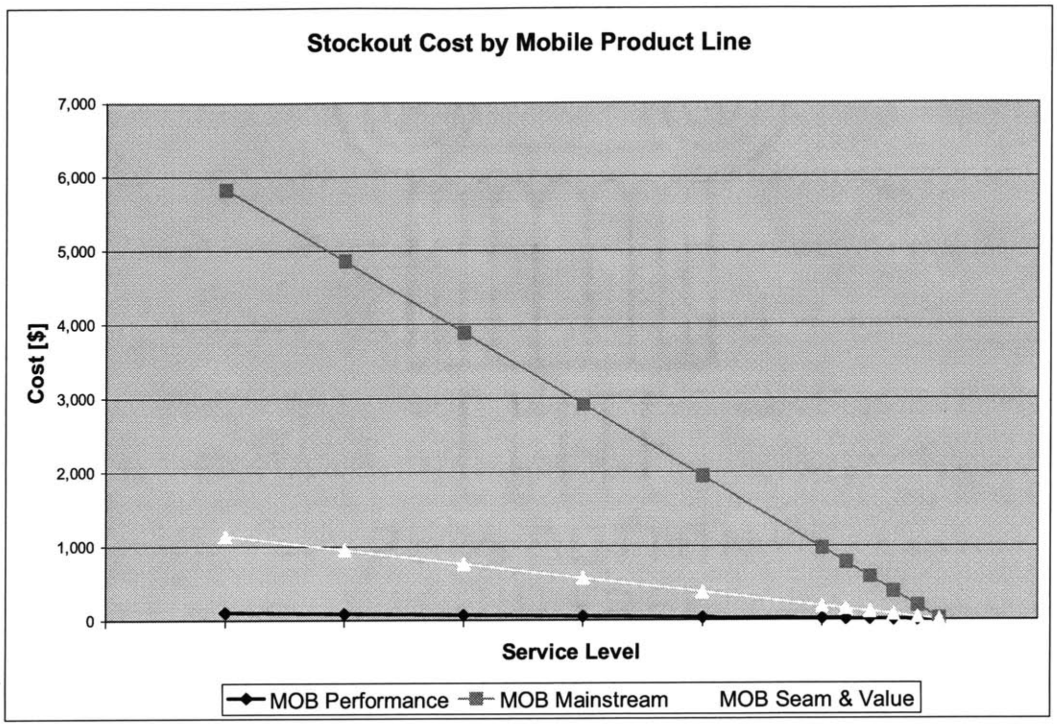

Note: scale normalized to protect confidentiality

Graph 9

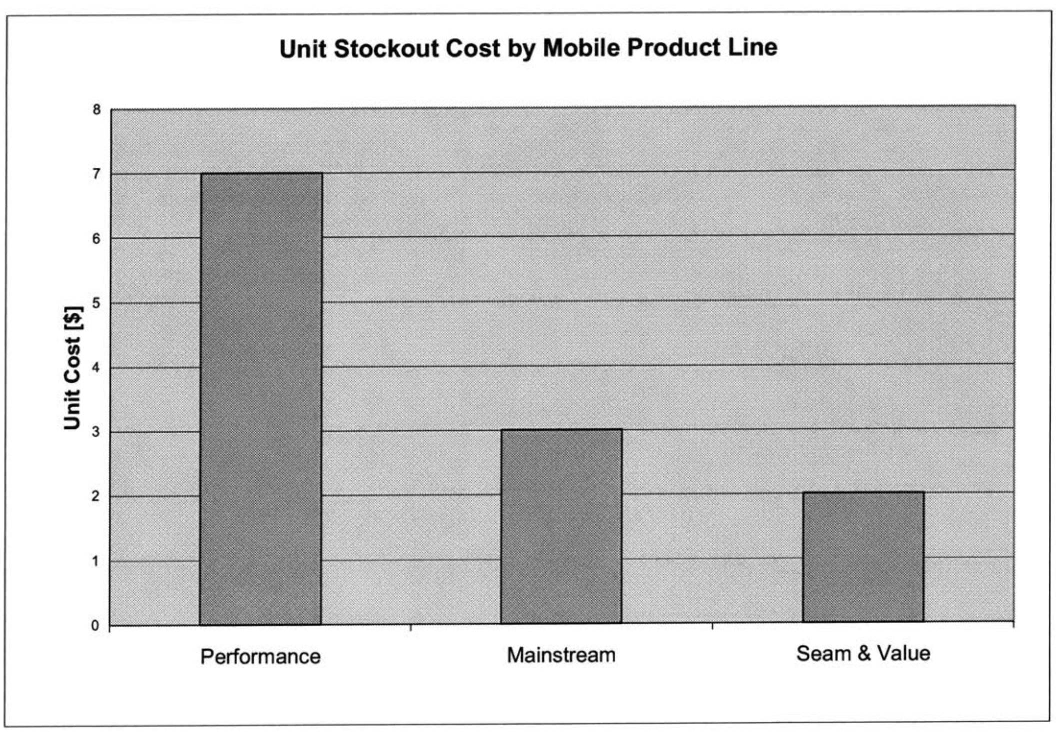

Note: scale normalized to protect confidentiality

Graph 10

5.4. Server (SVR) Product Line

The server line comprises multi-processor (MP) and dual-processor (DP) servers, where MP is the top of the line. About 90% of Server customers are OEMs, so for the purposes of this analysis we will consider OEMs to be the only market. Like in the case of the Mobile line, there is no open market for Server products, and since they are marketed to OEMs only, there are no buy up cases.

There is no competition for the MP line, so stockouts result mostly in postponed sales, and to a minor extent, buy downs. There is also a very small percentage of lost sales due to evaporation. MP margins are very high, which makes stockouts highly undesirable; however, such stockouts result mostly in postponed sales, which is not an important cost driver. Moreover, sales volume is very low, which helps balance the financial impact.

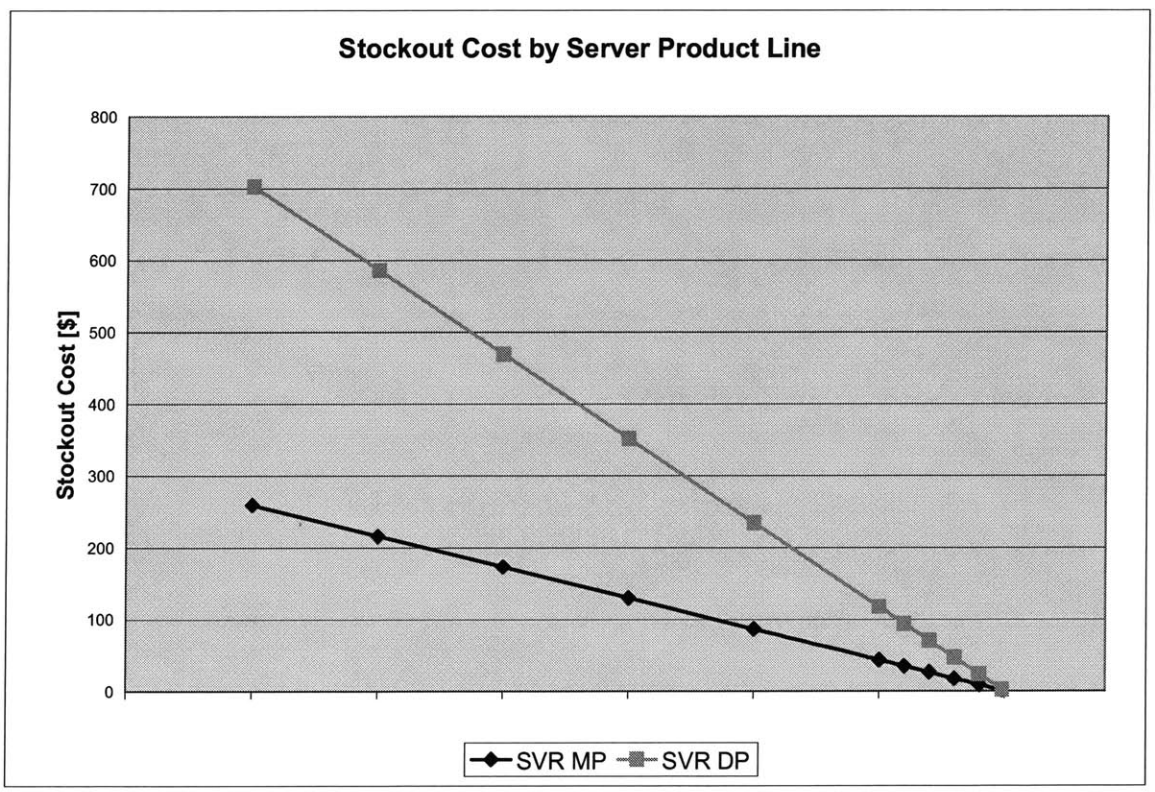

Note: scale normalized to protect confidentiality

Graph 11

Customer response is somewhat different for the DP line. Stockouts still result mostly in postponed sales, but there is an increase in lost sales. Prices are moderate and volumes are low, but still higher than for MP chips. As a result, DP chip stockout costs are higher, as can be observed in Graph 11. On the other hand, unit stockout costs are proportional to selling prices; hence they are much higher for MP chips (Graph 12).

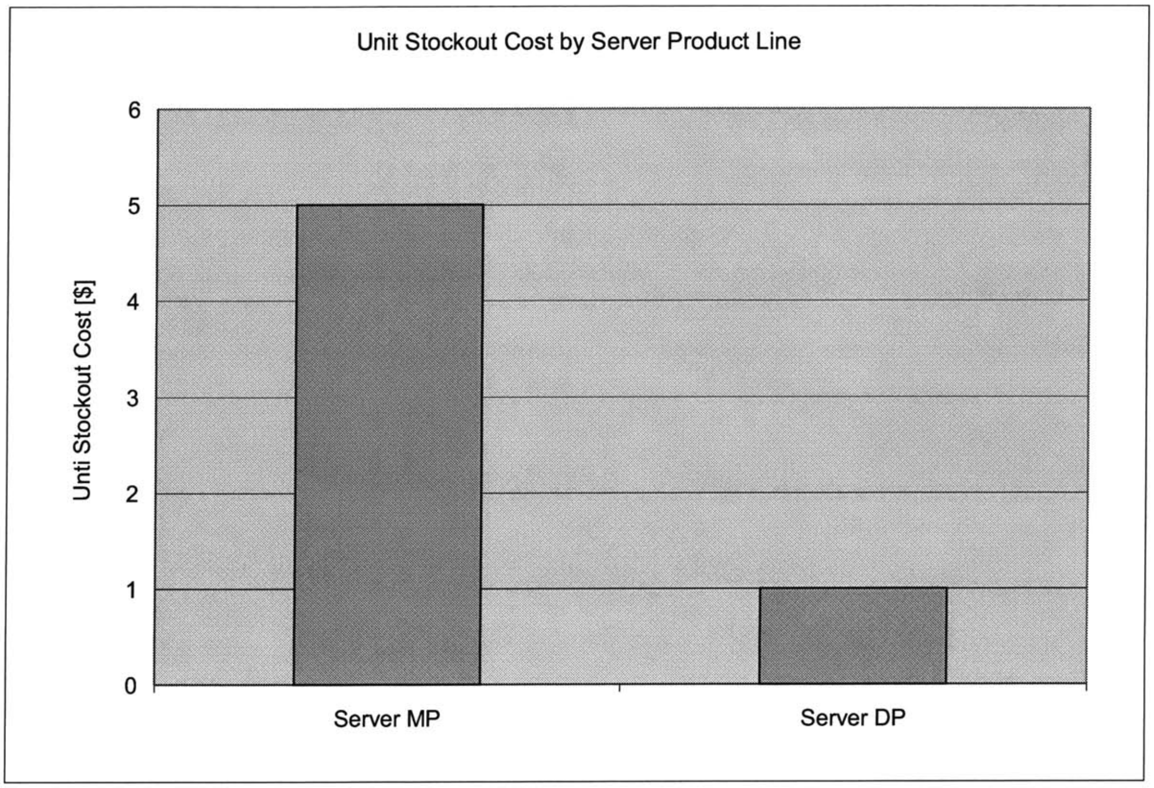

Note: scale normalized to protect confidentiality

Graph 12

Chapter 6 – Determining the Optimal Service Level by Product Line

6.1. Introduction

On Chapter 4, we estimated the optimal service level at the aggregated level. However, in order to maximize inventory efficiency, it is necessary to operate at optimal service level for each product line. Otherwise, it may happen that aggregated service level is optimal, but at a more granular level, some product lines are underserved.

In order to estimate the optimal service level for the different product lines, we need to estimate the inventory related costs (holding and scrap costs) and them to the stockout costs calculated on Chapter 5.

6.2. The Inventory Holding Costs

In order to estimate the inventory holding costs by product line, we will proceed as we did on Chapter 4 for the aggregated demand: first, we need to know the relationship between safety stock and service level by product line, and then estimate the inventory holding cost as a function of safety stock and average variable cost.

A mentioned in Chapter 4, Levesque (2004) estimated the inventory targets required to meet different service levels. This work was based on a two-stage base-stock model and considered the demand at the aggregated level only. Chow (2004) extended Levesque's model (described in Section 4) to calculate inventory targets needed to meet a certain level of service for each of Intel's product lines. (Table 6 – actual results not shown for confidentiality).

	Fill Rate	**70%**	**75%**	**80%**	**85%**	**90%**	**95%**	**96%**	**97%**	**98%**	**99%**	**99.0%**
Weeks of Safety Stock	**Desktop**	WD_{70}	WD_{75}	WD_{80}	WD_{85}	WD_{90}	WD_{95}	WD_{96}	WD_{97}	WD_{98}	WD_{99}	$WD_{99.9}$
	Mobile	WM_{70}	WM_{75}	WM_{80}	WM_{85}	WM_{90}	WM_{95}	WM_{96}	WM_{97}	WM_{98}	WM_{99}	$WM_{99.9}$
	Server	WS_{70}	WS_{75}	WS_{80}	WS_{85}	WS_{90}	WS_{95}	WS_{96}	WS_{97}	WS_{98}	WS_{99}	$WS_{99.9}$

Table 6

Since the demand for the different product lines is known, we can calculate the number of units of safety stock for each service level.

Safety Stock [units] = Weekly Demand * Safety Stock [weeks]

As mentioned in Section 4.2, the inventory holding cost at Intel was modeled by Bridge (2004) as a function of the variable cost of building inventory. Applying Bridge's formula,

Inventory Holding Cost [$] = I% * 100 *Average Variable Cost * Safety Stock [units]

Data such as variable cost and demand cannot be disclosed due to confidentiality issues, but is readily available at Intel, so the inventory holding cost for each service level and product line can be easily calculated.

6.3. The Scrap Cost

In order to complement the inventory holding cost aforementioned, it necessary to calculate the scrap cost due to excess inventory. Using Intel's historical data, such cost can be estimated as a percentage of the inventory volume for the aggregated demand. Unfortunately, there is no historical data available by product line, so we will use the aggregated rate as an approximation.

Scrap Cost [$] = Average Variable Cost * Scrap Rate * Safety Stock [units]

Data such as variable cost and demand by product line cannot be disclosed due to confidentiality issues, but is readily available at Intel, so the scrap cost for each service level can be easily calculated. Adding this cost to the inventory holding cost calculated in 6.2. we get the total inventory cost associated with each service level by product line.

6.4. The Optimal Service Level

Like we did on Chapter 4 for the aggregated level, now we can add together the stockout cost and the inventory cost in order to get the total cost associated with each service level in order to find the optimal service level for each product line.

6.4.1. Desktop Product Line

Graphs 13 to 15 show how stockout cost, inventory cost and total cost vary with service level for the different Desktop products. In all cases, we can observe that as service level goes up, the stockout cost decreases and the inventory cost increases. However, the stockout cost decreasing rate is faster than the inventory cost increasing rate, causing the optimal service level to shift towards the top values for Mainstream and Performance products. Despite the small difference in the optimal service levels, the shapes of the total cost curves are rather different. In the case of Performance products, we can see that the total cost curve keeps decreasing with service level and reaches its minimum for a service level close to 100%, and only starts to slightly pick up beyond that point. This is due to the high selling prices; in case of a stockout, the loss is so high that easily offsets the inventory costs. For Mainstream products, the results are somewhat different. Due to the much lower prices, the total cost reaches its minimum at a service level slightly lower and picks up quickly beyond that point.

For Seam & Value chips, the curves are almost identical than for Mainstream. This is because their price is slightly lower but they face a tougher customer response (more lost sales); and these effects tend to balance each other.

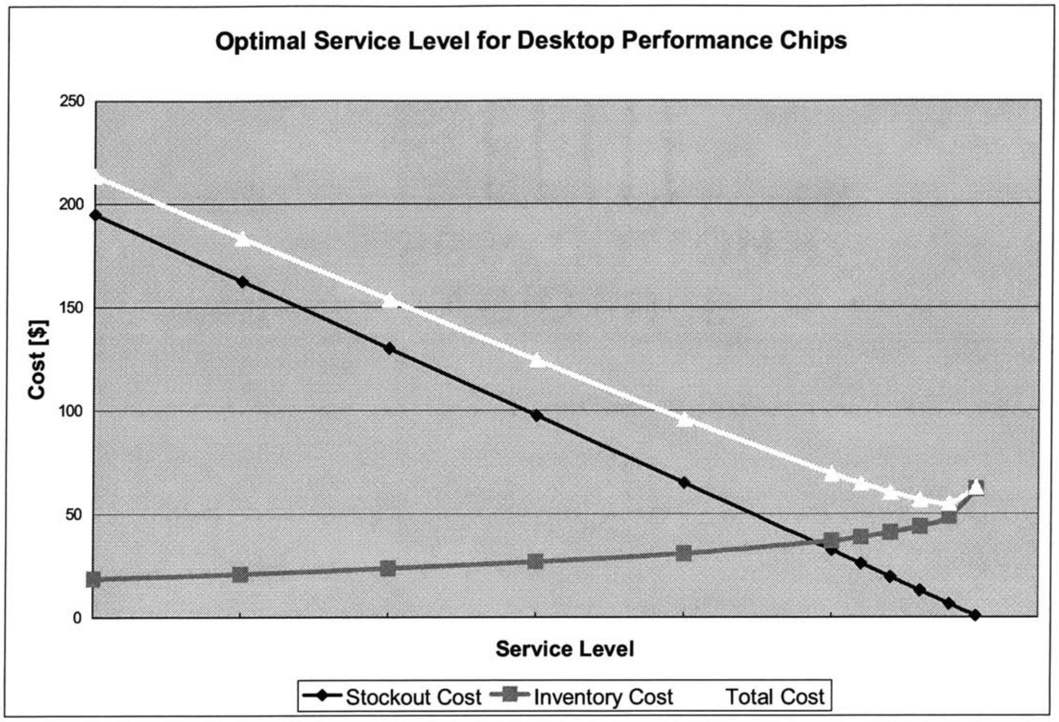

Note: scale normalized to protect confidentiality

Graph 13

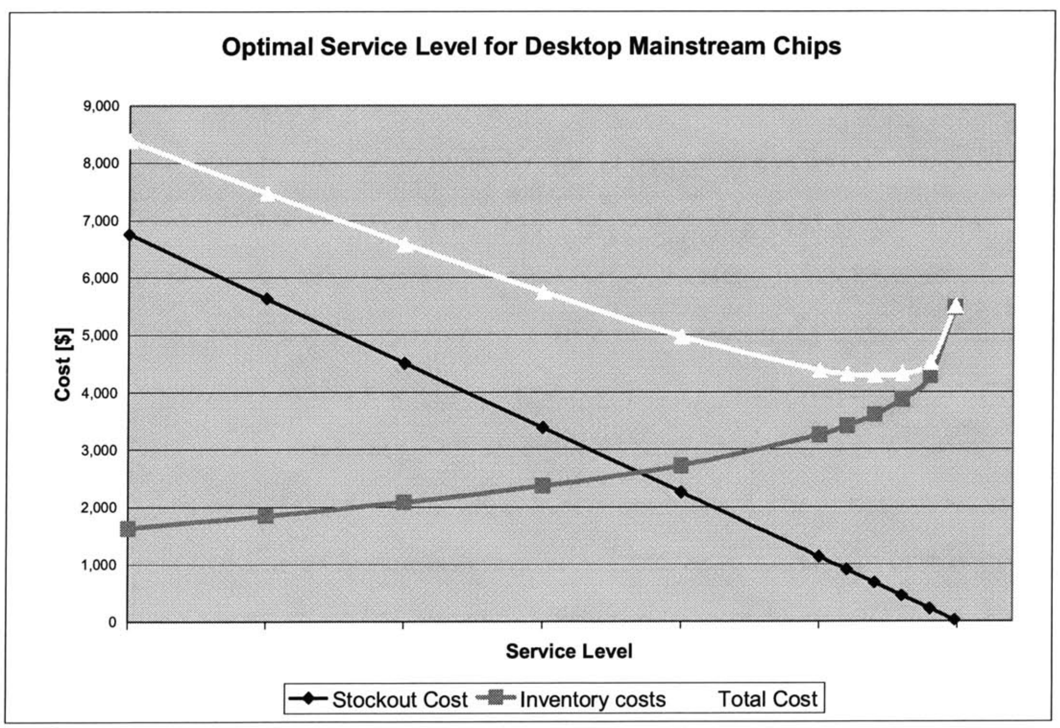

Note: scale normalized to protect confidentiality

Graph 14

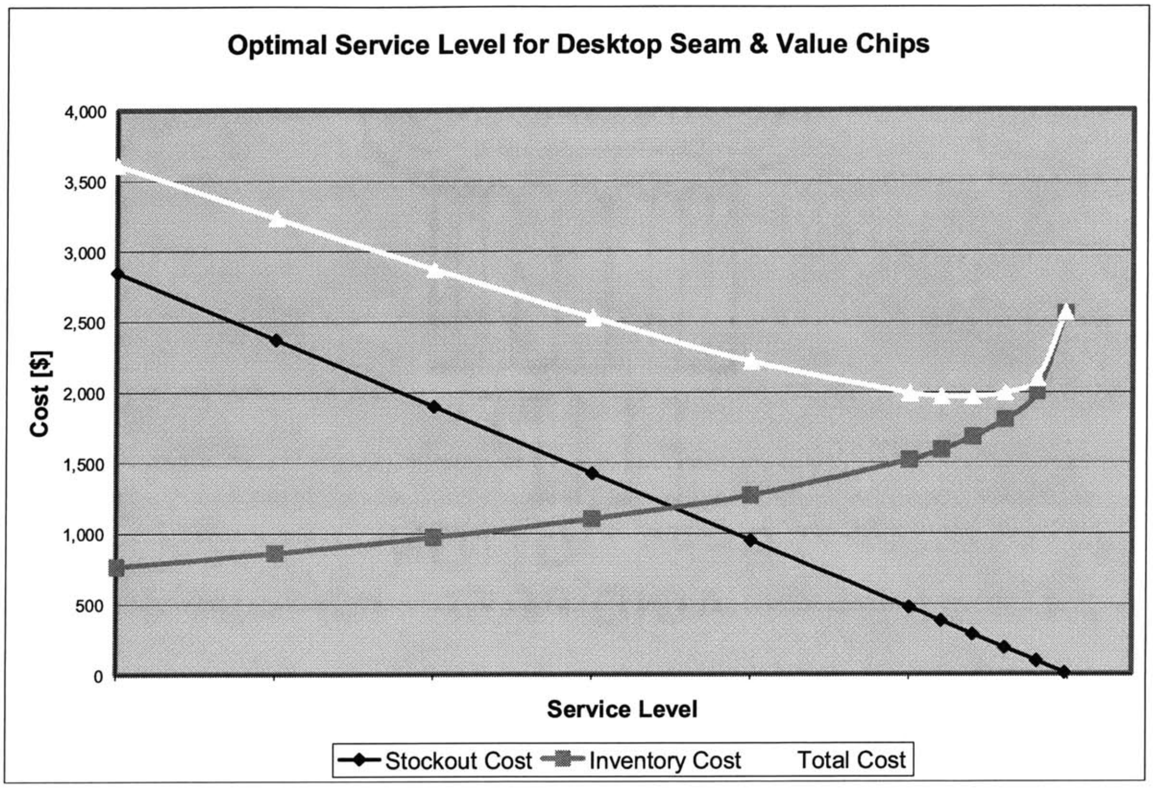

Note: scale normalized to protect confidentiality

Graph 15

6.4.2. Mobile Product Line

Graphs 16 to 18 show how stockout cost, inventory cost and total cost vary with service level for the different Mobile products. As in the case of Desktop chips, we can observe that as service level goes up, the stockout cost decreases and the inventory cost increases; and again, the stockout cost decreasing rate is faster than the inventory cost increasing rate, causing the optimal service level to shift towards the top values. In this case, there is a clear correlation between price and service level; for Performance chips, the optimal service level is the highest, and for Seam & Value is the lowest. The customer response patterns are very similar in all three cases, so there is virtually no effect that offsets the price impact.

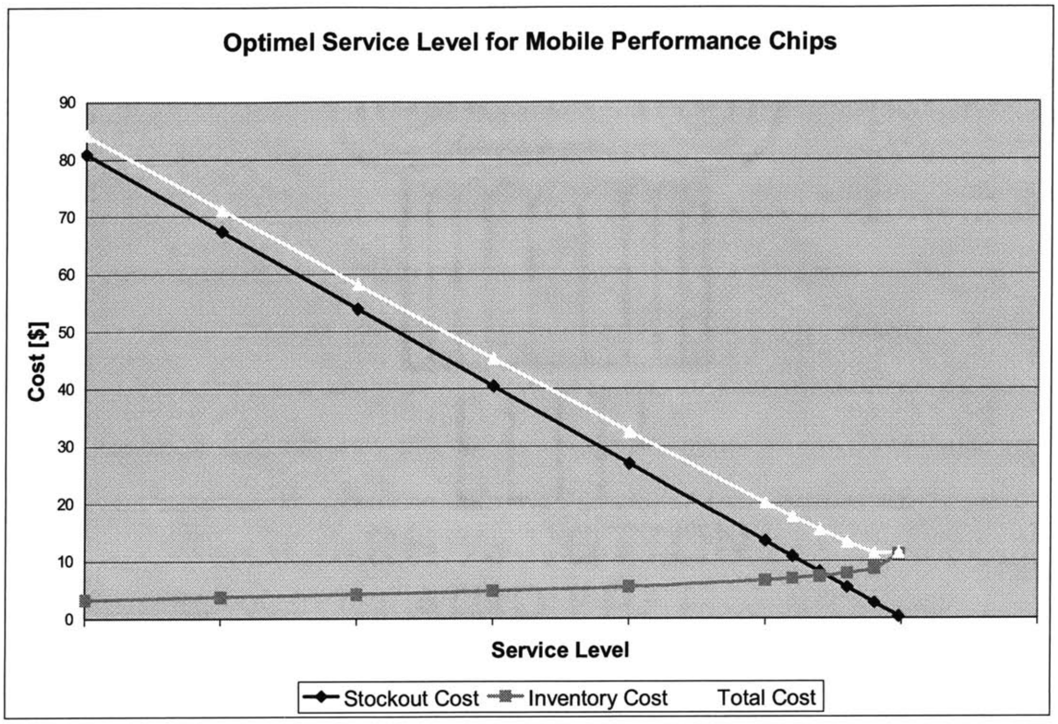

Note: scale normalized to protect confidentiality

Graph 16

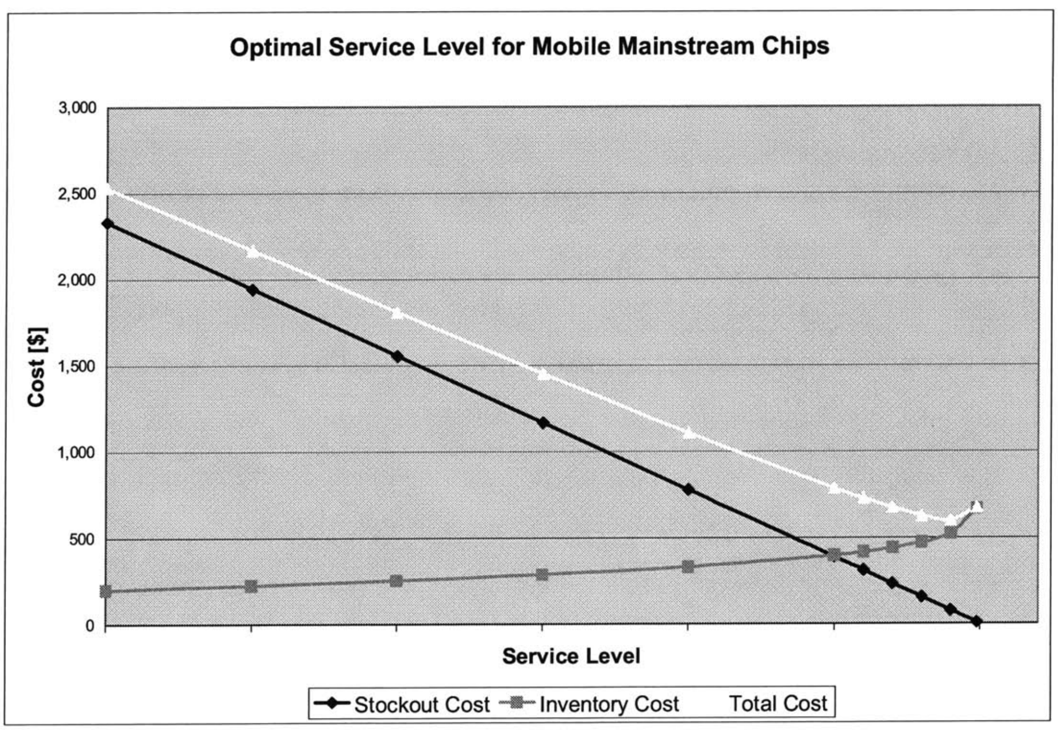

Note: scale normalized to protect confidentiality

Graph 17

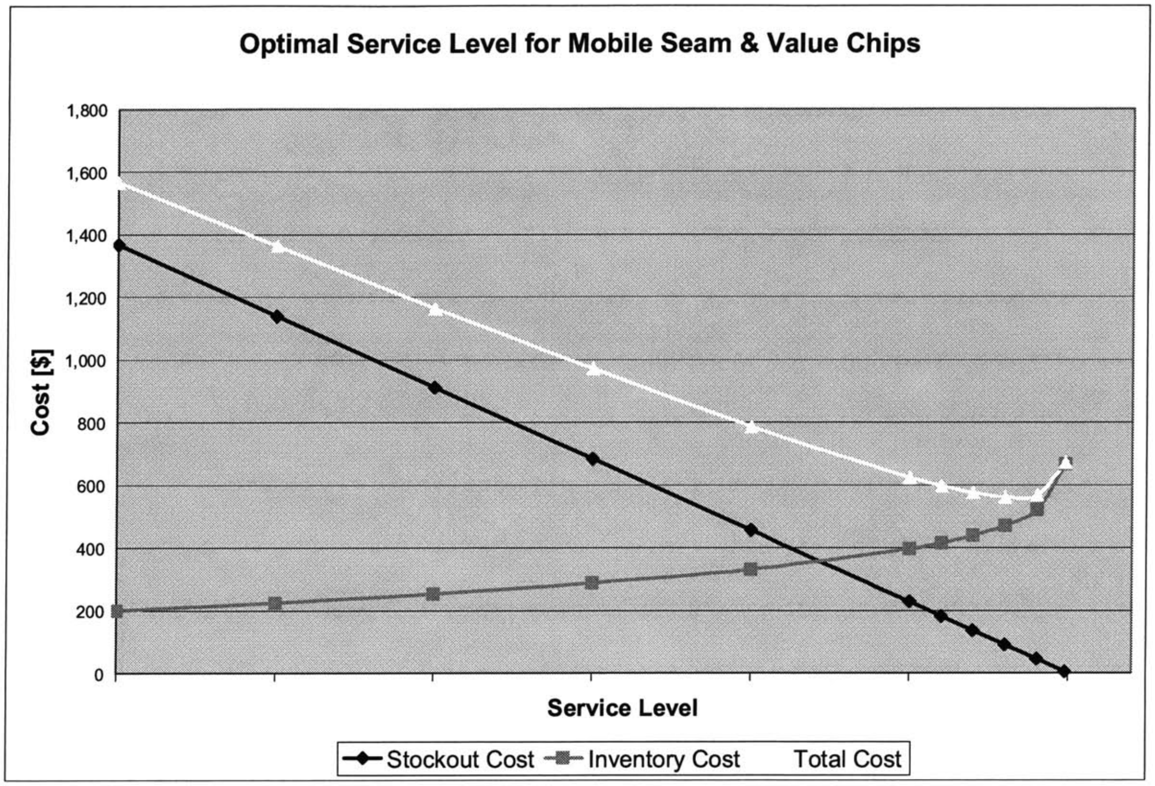

Note: scale normalized to protect confidentiality

Graph 18

6.4.3. Server Product Line

Graphs 19 and 20 show how stockout cost, inventory cost and total cost vary with service level for the two Server products. As in the previous cases, we can observe that as service level goes up, the stockout cost decreases and the inventory cost increases; and again, the stockout cost decreasing rate is faster than the inventory cost increasing rate, causing the optimal service level to shift towards the top values. Like in the case of Mobile products, there is a clear correlation between price and service level; the optimal cost is reached at a 100% service level for MP and slightly lower for DP. Again, the customer response patterns are very similar in both cases, so there is virtually no effect that offsets the price impact.

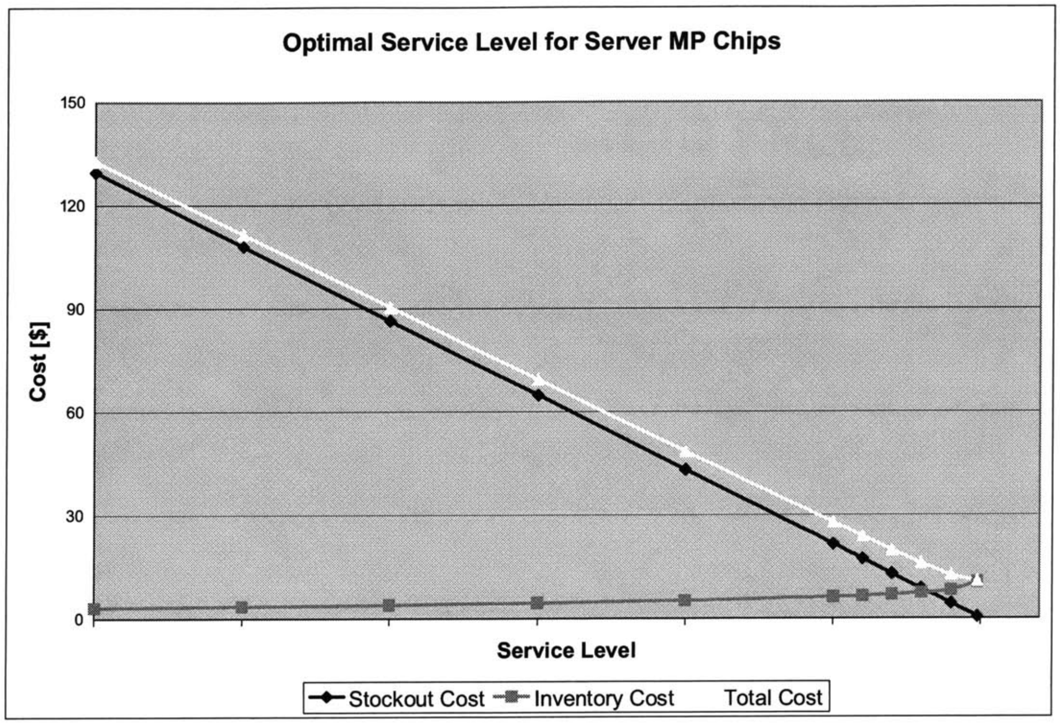

Note: scale normalized to protect confidentiality

Graph 19

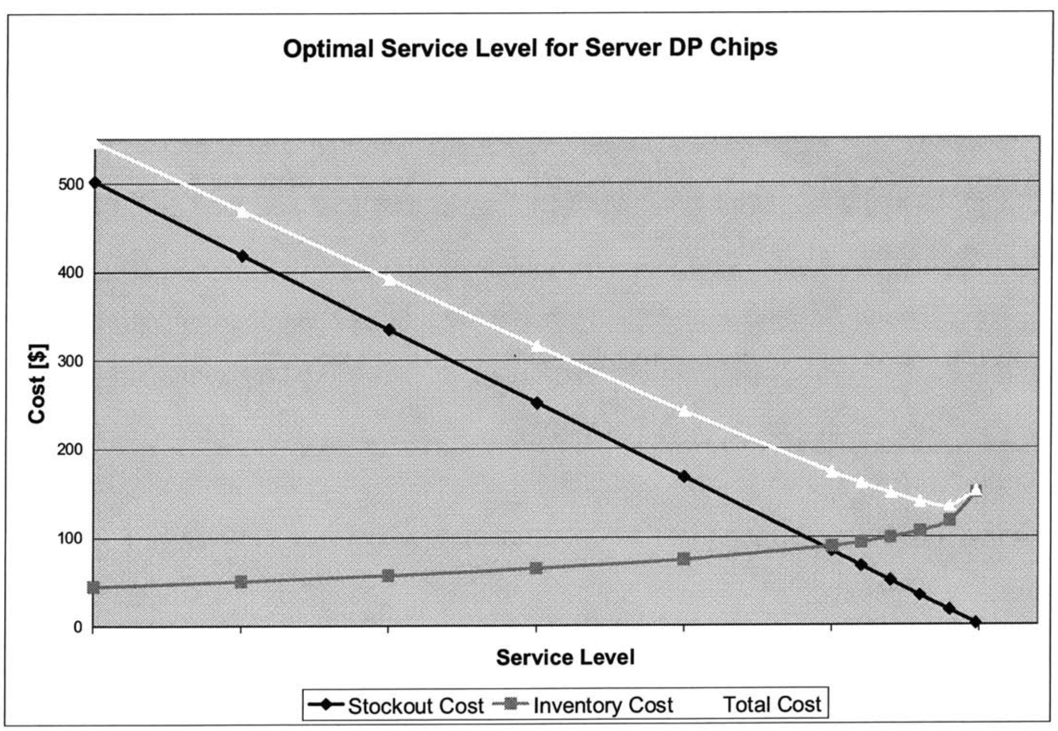

Note: scale normalized to protect confidentiality

Graph 20

Chapter 7 - Organizational Barriers and Change

7.1. Introduction

As mentioned in Chapter 1, inventories play a key role in a company's operations, and have a decisive impact on its performance. Inventories allow companies to buffer fluctuations in production and demand, increase order fill rate and achieve higher customer satisfaction; but on the other hand, keeping inventories involves costs and risks. Due to these advantages and disadvantages, inventory management is commonly affected by behavioral and political conflicts. According to Silver and Peterson (1979), "managing inventory means managing conflict". They summarized the organizational forces pushing towards higher and lower inventory as follows.

(a) Organizational forces pushing for higher inventory

1. Middle to senior management; in general, prefer higher buffer stocks to cover mistakes and inefficiencies in their operations that they have not been able to remove. Their promotion and reward depend on smooth operations.
2. Production management prefers higher inventories because they allow:
   - Lower operating costs
   - Longer production runs
   - More in-process stock
   - Higher raw materials levels
3. Marketing/sales management prefer higher inventories because they make it possible to provide:
   - Better customer service
   - Shorter lead times
   - Higher order fill rates
   - Full product lines
   - More new products
   - More flexibility

(b) Organizational forces pushing for lower inventory

1. When a corporation faces difficult times one of the first actions examined is lower inventory investment—that is, reducing the organizational slack present in the form of buffer stocks.
2. Finance/accounting management are rewarded for:
   - Reducing working capital requirements
   - Demonstrating higher return on investment on the money tied up in inventories
   - Increase profits by reducing carrying costs
   - Keeping better records on managers who may be overly using inventory buffers
   - Diverting money tied up in inventory into other, more profitable investments

Balancing these advantages and disadvantages is a complex task that depends on a wide array of factors such as demand patterns, demand forecast accuracy, manufacturing process, ordering process, service requirements and costs factors that vary widely across industries. In the particular case of the microprocessor industry, demand is highly volatile, accurate forecasting is difficult, and manufacturing processes are extremely complex, making inventory management particularly challenging. On top of that, the results of this research showed that inventories should go up, so some resistance was expected from the organizational stakeholders pushing towards lower inventories. For this reason, we considered interesting to make a "three lenses" (strategic, political, and cultural) analysis of the organizational barriers to change.

7.2. The Stakeholders

Among the many areas supported by the IDM group, there is the Microprocessor Marketing and Business Planning Group, or MMBP, which allocates microprocessor production among Intel's customers based on demand, capacity and current strategy.

MMBP works closely with IDM; some previous LFM interns, Joseph Levesque and Jim Chow, have worked in projects involving both areas. In addition, by the time this internship project started, MMBP was considering reviewing its inventory targets aiming to achieve the optimal service level. For these reasons, MMBP became a main stakeholder of this work right from the start. Due to the wide array of information needed to complete this work, several other areas within the company became involved. The most relevant were Finance, which provided information related to margins and costs; the Geographic Representatives or GEOs, which owned data regarding customer behavior and fill rates, and Logistics, providing historic information on scrap rates related to excess inventories.

From the methodology point of view, the project consisted on gathering the information from related areas, testing hypothesis with different stakeholders and finally making recommendations based on the outcome of the work. In consequence, it involved a high amount of interaction and communication, and would have never been possible without the collaboration of many people from many areas.

7.3. The Organizational Processes

7.3.1. The Strategic Design

The Inventory and Demand Management Group's role is to provide support and advice regarding supply chain optimization to other areas within Intel. Due to the nature of this project, related to fundamental issues such as inventory targets and service level, it was not possible to address changes in a particular area without affecting others. As a result, a change in the supply chain cannot be considered in isolation, since its effects will cascade down many other business areas and may lead to unanticipated side effects. For these reasons, it was necessary to contemplate the effects of the project outcome on all the related areas and to carefully evaluate the pros and cons on each of them.

On the other hand, the fact that IDM is an independent, support area gave this project certain freedom of action, since it was not tied to the needs of a specific group but to suit the interests of the company as a whole. The objective was to determine the optimal service level for the microprocessor business as a whole instead of suiting the needs of a particular area; in other words the project objectives were aligned with the company strategy, even though part of the challenge was to make them fit with the needs of different areas within the organization. In that sense, the formal structure of the company definitely helped: the goal was to gather information, make hypothesis, analyze the results and make recommendations aiming to improve the overall performance of the business, and each of the related areas would have the responsibility to implement them, supported by the well-tuned coordinated systems in place. The real challenge was to put together a strong foundation behind such recommendations to get stakeholders to buy-in and to accurately evaluate the pros and cons in order to make sure that the former out weighted the latter.

7.3.2. The Political Aspect

As mentioned in 7.1, the Microprocessor Marketing and Business Planning Group (MMBP) rapidly became a main –and positive– stakeholder of this work; but due to the scope of project, many other areas within Intel such as Finance, Logistics and the Geographic (Commercial) Representatives became involved later.
Some negative stakeholders were expected, especially at the early stages of the project, when Intel's stock plummeted due to excess inventories, which are normally associated with high service levels. However, such negative reaction never happened.

The Finance group also became one of the main stakeholders of the project, especially after overcoming the initial confidentiality issues. By the end of the internship the computer models supporting the project were handed off to this group, which is planning to apply to future projects involving inventory optimization.

7.3.3. The Cultural Aspect

Intel is very involved with the Leaders for Manufacturing (LFM) Program, with a strong alumni presence, offering several internships every year; this creates a very LFM friendly environment, where most people are aware of the value of LFM projects. The combination of these factors makes the adaptation process definitely easier. However, the project ramp-up still offered many challenges, from getting familiar with a new industry to adapting to a very particular corporate culture. In addition, there is always a challenge associated to resistance to change, in particular if the project outcome is controversial. As described in 7.1., decisions related to inventory management tend to be conflictive; and those related to increase inventories are normally resisted by the finance and accounting management. For this reason, some resistance to our recommendations was expected.

At the beginning of the internship, the project was broken down into two main steps: the first one was to put together a strong foundation in order to achieve reliable results; and the second one was to affect change by getting stakeholders to buy in the outcome of my work. The success of the second stage was largely based on the success of the first one, which consisted in building a robust foundation for the project. This meant gathering reliable data, putting together consistent mathematical models and double-checking assumptions and hypothesis. In order to achieve all three of them, communication was critical, since it helped understand the company's formal and informal structure, reach the people that could provide an interesting insight into the project and get to know who owned the necessary data.

At the same time, communication was the biggest cultural challenge, in particular due the diversity of audiences, the language barrier, and the predominance of virtual meetings. Due to the scope of the project, it was necessary to interact with different audiences, explaining clearly and fast what the project goals are and why the outcome could be useful for each of them; hence it was important to figure out what they could find interesting, and how to explain the technicalities they were not familiar with. If effective communication itself is already quite challenging, and gets still more complicated when needs to be done in a foreign language. This is particularly visible at Intel, due to its remarkably diverse and international workforce. Besides, Intel's communication is plagued with an ample and widespread variety of acronyms, and needless to say it is necessary to learn them in order to communicate effectively. To make communication still more challenging, most meetings are over the phone and internet, making it difficult to give a persuasive speech or presentation without being able to gain the benefit of direct communication.

7.4. Leading the Change Process

As mentioned beforehand, the project was broken down into two main steps: to put together a strong foundation in order to achieve reliable results; and to affect change by getting stakeholders to buy in to the outcome of this work. I anticipated the latter to be rather controversial because our final recommendation was to increase inventories in order to achieve an optimal service level, and as mentioned in 7.1., inventory increases tend to be resisted by certain organizational forces. But even though this appeared to be the biggest leadership challenge, it turned out to be all the opposite. Despite the counterintuitive project results, no conflicts of interest or cultural issues happened. The project was embraced by a cross-functional team working on inventory optimization with positive results; the model was handed off to the Finance area, which will own it and use it to evaluate inventory policies in the future; and further research building on this work is going to be conducted in future internships. The success of the second stage was based on the success of the first one—which consisted in building a robust foundation for the project, or in other words, effective sensemaking. In order to achieve it, the first step was to gather the necessary data. Such data was dispersed throughout the company, so it was essential to find who owned it and to persuade the owners to release the information. Due to the extent of the data needed and the confidentiality issues, it was necessary to attract the stakeholders' attention and to persuade them to buy-in the project objectives. Once data was gathered, the next step was to build a mathematical model, so in order to assure its consistency, it was necessary to brainstorm with people from different backgrounds to double-check assumptions and hypothesis. For such reasons, relating was critical. And, for the reasons mentioned in section 7.3. (diversity of audiences, language barrier and virtual meetings), communication could easily become an obstacle, making it the real challenge of the project. On the other hand, it was a relatively easy problem to address—it just depended on communication skills.

Chapter 8 – Recommendations

The primary objective of this internship project was to determine the cost of stockouts and its impact in determining the optimal service level; accordingly, the expected deliverables can be summarized as follows.

- To quantify the financial impact of stockouts aiming to develop optimal inventory strategies
- To make recommendations for changes in the current service levels based on those findings
- To provide a framework for evaluating Intel's inventory policies in the future.

The results of this project show that stockouts have a high financial impact in the microprocessor business. Due to the high margins, each sale lost to the competition means losing a large dollar amount, which may easily outweigh the inventory-related costs of that product. The quantification of the financial impact of the different stockout situations vs. the related inventory costs for different service levels shows that the optimal service level is extremely high. This conclusion applies to both the aggregated microprocessor demand and the individual product lines. Stockouts of top of the line products cause high losses due to the high margins, but this effect is balanced by the weak competition in that market segment. Conversely, stockouts of value products cause relatively low losses due to the lower margins, but this effect is balanced by the high number of sales lost to the aggressive competition existing in that segment.

The sensitivity analysis proves that the hypothesis hold even in extreme situations. When a stockout happens, the main cost drivers are margin and customer reaction (whether they choose to postpone the sale, buy from the competition, etc.). From the inventory level point of view, the main cost driver is the inventory holding cost. But even in extreme situations such as doubling the inventory holding cost or cutting the margins by half, the optimal service levels are still above 90%. This validates the assumption that the long-term effects of stockouts such as loss of goodwill have very little impact on the optimal service level, since considering the short-term effects only, the optimal service level is so high that adding an even a small additional cost would bring the optimal service level to 100%.

The final outcome of this work is a recommendation for Intel's microprocessor business is to increase their current service level, and will be a valuable input into further supply chain optimization analysis including adjustment of finished good inventory levels.

One of the key outcomes of this work is a model in Excel that captures the impact of stockout and inventory related costs at different service levels, for the aggregated demand and also by product line. This model will be owned by Intel's ISNG Finance division and used to evaluate inventory policies in the future. Intel's MMBP (Microprocessor Marketing Business Planning) area has already used this model for reviewing the revenue optimal service level in the microprocessor business. The dollar amount impact of this project will depend on the inventory level adopted.

In order to transfer the knowledge, the IDM (Inventory and Demand Management) staff has been educated on the outcome of this work and new internships are being planned in order to build up on it; in fact, this work relies on many assumptions than can be refined by further research.

Our assumption that the long term effect of stockouts (loss of goodwill) can be disregarded is fine from the inventory management point of view, but it would be interesting to study this issue with more detail for strategic reasons.

Other data that could be refined are those related to customer behavior. As described in Section 3.3, we used data gathered by Intel's MMBP area through two spot surveys; hence, the data reflects the reality of two specific instants in time and is not adjusted for effects such as seasonality. Intel could get to know the behavior of its customers much better if such data were gathered on a continuous basis.

Another interesting issue to research would be the effect of product interdependency, which was out of the scope of this work. Product interdependency means that chips of a given specification cannot be produced independently of the other specifications produced. For example, in order to produce a quantity X of spec "A" chips, it is necessary to also produce a quantity Y of spec "B" chips. Hence, if the demand for chips A increases at a higher rate than that of chips B, an excess of chips B will be produced, increasing the scrap rate. Since the scrap rate is related to the inventory cost, it would be interesting to analyze the effect of product interdependence on service levels.

Finally, much of the outcome of this work is based on the two-node inventory model proposed by Levesque (2004). However, there are more than two nodes in Intel's supply chain, so follow-on research should include the development of a multi-node network optimization as per Levesque's recommendation.

References

Bridge, Richard, Intel Corporation, ISNG Finance, Inventory Cost Modeling, Presentation, 2004.

Black, Brian E., "LFM Masters Thesis: Utilizing the Principles and Implications of the Base-Stock Model to Improve Supply Chain Performance", MIT Sloan School of Management and MIT Department of Electrical Engineering and Computer Science, 1998.

Cheslek, Eric, "LFM Masters Thesis: Lean Supplier Relationships in the United Kingdom", MIT Sloan School of Management and MIT Department of Chemical Engineering, 1998.

Chow, Jim, "LFM Masters Thesis: Analysis of New Approaches to Improve the Customer Responsiveness of Intel's Microprocessor Supply Chain", MIT Sloan School of Management and MIT Department of Civil and Environmental Engineering, 2004.

Fine, Charles H., Clockspeed, Perseus Books, 1998.

Hayes, Robert, Pisano, Gary, Upton, David, and Wheelwright, Steven, Operations, Strategy, and Technology – Pursuing the Competitive Edge, John Wiley & Sons, 2005.

Hopp, Wallace J. and Spearman, Mark L., Factory Physics: Foundations of Manufacturing Management, Richard D. Irwin, 1996.

Intel Corporation Website, www.intel.com

Jensen, Arne, "Stockout Costs in Distribution Systems for Spare Parts", International Journal of Physical Distribution and Logistics Management, Bradford, 1992, Vol. 22, Issue 1.

Levesque, Joseph C., "LFM Masters Thesis: Analysis of Variability in the Semiconductor Supply Chain", MIT Sloan School of Management and MIT Department of Chemical Engineering, 2004.

Nahmias, Steven, "Production and Operations Analysis", McGraw-Hill, 4th Edition, 2001.

Oral, M., Salvador, S., Reisman, A., Dean, B., "On the Evaluation of Shortage Costs for Inventory Control of Finished Goods", Management Science, Vol. 18, No. 6, February 1972.

Rosenfield, Donald B., "Disposal of Excess Inventory", Operations Research, May/June 1989.

Savage, Sam L., Decision Making with Insight, Brooks/Cole, 2003.

Silver, Edward A. and Peterson, Rein, "Decision Systems for Inventory Management and Production Planning", John Wiley & Sons, 2nd Edition, 1985.

Schwartz, Benjamin L., "A New Approach to Stockout Penalties", Management Science, Vol. 12, No. 12, Series B, Managerial, August 1966.

Sterman, John D., Business Dynamics – Systems Thinking and Modeling for a Complex World, McGraw Hill, 2000.

Toyota Motor Corporation website, www.toyota.co.jp

Zimmerman, Hans-Jurgen and Sovereign, Michael G. "Quantitative Models for Production Management", Prentice-Hall, Inc., Chapter 7, 1974.