Utilizing the Principles and Implications of the Base Stock Model to Improve Supply Chain Performance

by
Brian E. Black

B.S. Electrical and Computer Engineering, Brigham Young University 1991

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

Master of Science in Management
and
Master of Science in Electrical Engineering and Computer Science

in conjunction with the
Leaders for Manufacturing Program

at the
Massachusetts Institute of Technology
June, 1998

Signature of Author ___

Sloan School of Management
Department of Electrical Engineering and Computer Science
May 8, 1998

Certified by ________________

Alvin W. Drake, Thesis Advisor
Professor of Electrical Engineering and Computer Science

Certified by ________________

Stephen C. Graves, Thesis Advisor
Professor of Management Science

Accepted by ________________

Arthur C. Smith, Chairman, Committee on Graduate Studies
Department of Electrical Engineering and Computer Science

Accepted by ________________

Larry Abeln, Director of the Master's Program
Sloan School of Management

This page is intentionally blank.

Utilizing the Principles and Implications of the Base Stock Model to Improve Supply Chain Performance

by
Brian E. Black

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

Master of Science in Management and
Master of Science in Electrical Engineering and Computer Science

Abstract

Manufacturers of low and medium-volume electronic products face great challenges in configuring their supply chains and production environments to meet customer demand in an effective manner. Specific challenges include variable component yield, long component lead time, expensive components, technological obsolescence, poor forecasting abilities, variation in customer demand, and the possibility of very large orders with short lead time requirements. In the face of these challenges, manufacturers must answer these fundamental questions:

What is the appropriate level of inventory in a supply chain to meet customer demand?
Where should this inventory be held?
Once these inventory targets have been set, how can they be met and maintained?
What are the relative benefits of improving cycle time, yield, and yield variation?

This thesis explores these questions within the context of a low volume manufacturing supply chain. It is based on research done at one of Eastman Kodak Company's professional digital camera supply chains. My goal there was to provide a framework through which Kodak could improve its customer service and supply chain inventory levels.

The tools used for this research are a base stock inventory model, basic statistics, the MIT Strategic Inventory Placement Model, and simulation using historical data. By applying the framework described in this thesis, the manufacturing and supply chain team was able to reduce inventory investment by approximately $100,000 at the distribution center and final assembly, and to identify additional savings opportunities worth over $600,000 in supply chain inventory. Additional long-term benefits will include improved customer service, improved inventory control and simplified management processes, global rather than local optimization, and identification and valuation of improvement opportunities.

Thesis Advisors: Alvin W. Drake, Professor of Electrical Engineering and Computer Science
Stephen C. Graves, Professor of Management Science

Acknowledgements

I would like to thank all those who made this research such a great learning experience:

I wish to acknowledge the Leaders for Manufacturing Program for its support of this work.

I am very grateful for the support, advice, guidance and encouragement from my advisors Steve Graves and Al Drake, and from my unofficial advisor Sean Willems whose work and insights made this a great experience indeed.

Thanks and best wishes to my friends at Kodak who helped in this work, particularly: Kathy McFiggens, Bill Carleton, Gary Smith, Vince Andrews, Gerry Edd, Don May, Darren Johnson, Tom Huss, Tom Bridge, Mike O'Brien, Ginny Spielman, and Beth Galloway.

Thanks also to the LFM alumni at Kodak, particularly Mark Zeni, Brian Parks and Charlie DeWitt who provided guidance and helped make my time at the company enjoyable.

DEDICATION

To my wife Stephanie,

whose love, patience, and encouragement
made this work possible.

And to our three children

Amy,
Shauna,
&
Stephen,

who bring joy to my life and constantly remind me
of what life is really all about.

Table of Contents

1 Introduction

Outline

This thesis has three main sections. Section 1 includes background material describing the setting for this research. Section 2 describes the framework that can be used to set appropriate inventory targets and the application of this framework at Kodak's professional digital camera supply chain. Section 3 provides conclusions and recommendations. For reasons of confidentiality, data (including volumes, forecasts, etc.) in this thesis has been disguised.

Setting

I conducted the research for this thesis at the Eastman Kodak Company's Kodak Equipment Manufacturing Division (KEMD) from June through December, 1997. KEMD manufactures several products, including thermal printers, digital cameras, and photofinishing equipment. Kodak currently divides its digital camera production into two product lines: the higher-volume consumer-level cameras, and the lower-volume professional-level cameras. I focused my research on the Kodak Professional Digital Camera supply chain. Kodak's professional digital cameras are used by journalists and other professional photographers, as well as government, educators, and others who have higher resolution needs than the typical amateur digital photographer.

Product overview

The Kodak professional digital cameras are assembled from three major components: an SLR (Single Lens Reflex) camera body, a CCD (Charge-Coupled Device) imager, and a circuit board. The CCD imager sets the performance (resolution and effective speed) of the camera. Three basic CCD imagers are available: a standard imager, a higher-speed imager, and a higher-resolution imager. Each of these is available on either a Nikon or a Canon body. Many of these cameras are available in color, monochrome, color infrared, and monochrome infrared varieties. Some are also made in different branding variations (e.g. Canon brand and Kodak brand). This results in approximately 25 possible camera varieties, though some are more frequently ordered than others. The table below shows the more commonly ordered models. The number in each box indicates the number of brand varieties/versions for each configuration. Boxes with no numbers represent models that either do not exist, or are ordered extremely rarely.

Table 1-1: Number of Model Types

Body Type	Sensitivity	Imager Performance		
		Standard	High Resolution	High Speed
Canon	Color	2	2	2
	Monochrome	2		2
	Color-IR			
	Mono-IR			2
Nikon	Color	4	2	2
	Monochrome	1	1	
	Color-IR	1		
	Mono-IR	1	1	

Process flow

The camera supply chain is shown below:

Figure 1-1: Camera Supply Chain

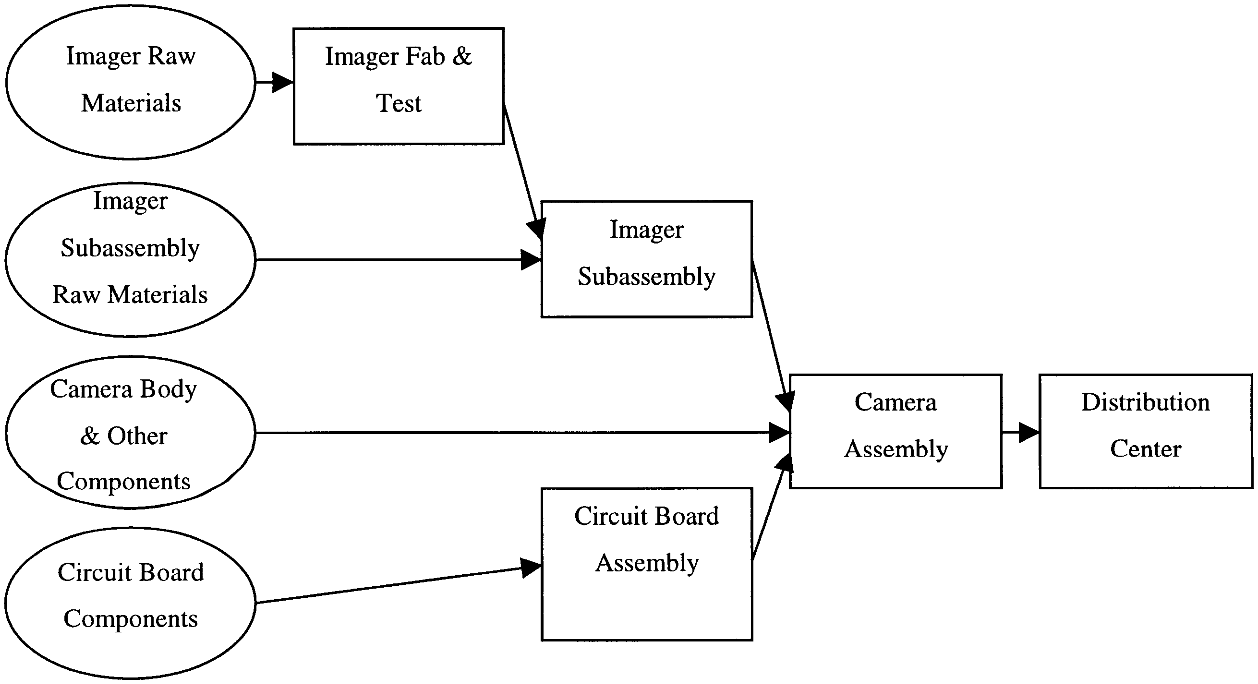

While the circuit board and CCD imager are manufactured internally, the camera body and minor components are sourced externally. The CCD imager is manufactured at the microelectronics imager fab (IF) and then assembled into an optical-mechanical subassembly at another department (IS). The circuit board is manufactured internally at a circuit board assembly division (CBA). The circuit board, imager sub-assembly, camera body, and other parts are then assembled at a final camera assembly (FCA) area. The finished cameras are then sent to a distribution center (DC) where they are held until they are sent to a domestic customer or overseas distribution center. Currently, the imager fab (IF), imager subassembly (IS), circuit board assembly (CBA), final camera assembly (FCA), and the Distribution Center (DC) organizations are located in different buildings, and all have separate reporting structures.

Microelectronics Imager Fab (IF)

Charge-Coupled Device (CCD) imagers are manufactured at Kodak's microelectronics imager fab (IF) for both internal and external use. The three basic imager types produced for cameras in this supply chain represent a minority of the imagers produced at the imager fab, in terms of both unit volume and revenue. The imager fab (IF) CCD process is shown below.

Figure 1-2: CCD Process

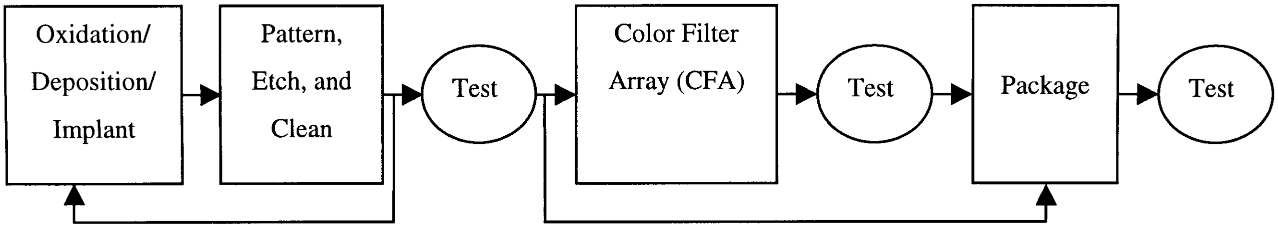

The CCD process is similar in many respects to VLSI integrated circuit fabrication. Bare silicon wafers are sent through recurring oxidation/implant/deposition-pattern-etch cycles to create die with arrays of light-sensitive pixels and support circuitry. One of the final steps before packaging is Color Filter Array (CFA). Imagers that go through CFA become color-sensitive imagers. Imagers that skip this step remain monochrome devices. Additionally, imagers are packaged with cover-glass filters that determine whether they will be infrared sensitive or insensitive devices. Production scheduling occurs weekly, and is currently based on forecasts of customer orders. Prior to the time that I conducted this research, the imager fab (IF) did not explicitly have any stocking targets. This is largely due to the fact that historically, its primary mission was research and development. Recently, however, the imager fab (IF) has been changing to become more production-focused. Accordingly, it has made plans to stock imagers at a few key locations within the IF manufacturing process.

Across the industry, the production of CCD imagers has inherently low and variable yields. Unlike digital integrated circuits, which depend on circuits being either on or off, analog CCD imagers must exhibit specific response curves to a range of light inputs. Since these imagers have a long manufacturing lead time, the effect of this yield variation is that supply is variable and uncertain. This results in stockouts and expediting when yield is lower, and excess inventory when yield is higher. The information feedback of yield and inventory is inherently delayed, which leads to oscillations in inventory levels.

Imager Subassembly (IS)

Once an imager has been manufactured and tested at the fab (IF), it is sent to the imager subassembly (IS) department, where it is tested, aligned, and mounted to an optical-mechanical-electronic assembly. Due to the nature of the assembly and test process, there can be additional fallout at IS.

Figure 1-3: Imager Fallout; when an imager fails at test, it is sent back to the preceding operation for root cause analysis

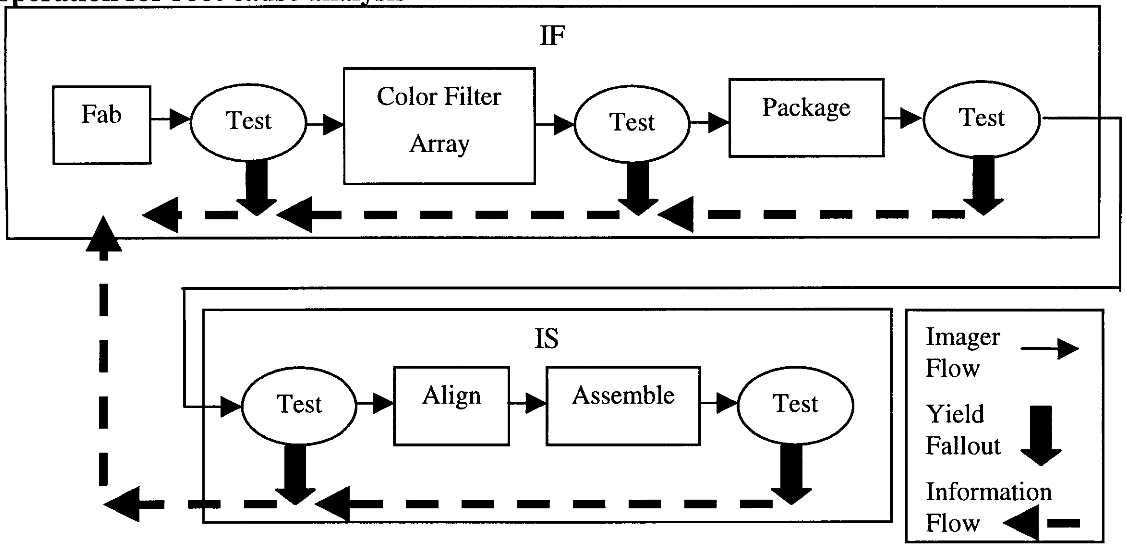

The imager subassembly department (IS) uses highly specialized equipment to perform the alignment task. This multiple axis system is used to align the imager to a frame. Once aligned, the imager is then mounted to the frame, and a flexible circuit cable is attached. The flex cable allows the imager to communicate with the circuit board. Production is scheduled weekly, based on the production schedule of final camera assembly (FCA). Inventories of higher volume imager varieties are held both as raw materials (unmounted imagers) and finished goods (mounted imager subassemblies).

Circuit Board Assembly (CBA)

Only about 3% of the board types assembled at Kodak's Circuit Board Assembly (CBA) area are made for this supply chain's cameras. Inventories of higher-volume professional camera circuit boards, as well as raw components are held at CBA. Production is scheduled weekly based on demand forecasts, finished goods inventory targets, and actual levels. Finished goods inventory targets are determined by an algorithm within the CBA MRP system, which takes into account average demand as well as past forecast accuracy. Although circuit board yield is quite high, board revisions caused by end-user quality issues can result in board rework or scrap. CBA has achieved major lead time reductions, and is close to achieving a manufacturing lead time of one day.

In my research, I decided to focus on the imager portion of the camera supply chain and not on circuit boards because the opportunities for improvement appeared to be greater with imagers. The same principles that I used in the imager supply chain may be applied across the entire camera supply chain.

Final Camera Assembly and Distribution Center

At final camera assembly (FCA), the CCD imager assembly, the circuit board, the camera body, and other externally sourced parts are assembled into a final product. They are then tested, packed, and shipped to a local distribution center. From the distribution center, the camera is typically shipped to a domestic camera dealer or to a stocking location in another country or geographic region. Individual dealers carry little or no inventory and place orders based on end customer demand. Country and regional stocking centers will typically carry a limited amount of stock.

People from manufacturing, supply chain, and purchasing functions, meet weekly to schedule final camera assembly production. At the time that I began this research, the production schedule was based on finished goods inventory, forecasted demand, and parts availability. The goal of scheduling was to keep finished goods inventories relatively constant. To do this, the schedulers looked at current finished goods levels, subtracted average demand over the production lead time, and added work in process. They then compared this projected finished goods inventory with finished goods inventory targets. Since there are many end products and these calculations were done mentally during the scheduling meeting, schedulers sometimes took shortcuts or made mistakes. One shortcut was to assume that current work in process was approximately equal to average demand. Schedulers would then use the current finished goods inventory level rather than projected finished goods inventory as an indicator of whether new production was needed. Once the schedulers determined the desired production numbers, they compared these to the loaded MRP schedule. Any difference between desired production and the MRP schedule led to negotiations between the schedulers and the purchasing representative over how much they would deviate from the original plan (and if so what should be done about the MRP shortage/excess). After the production scheduling meeting, the purchasing representative would load the resulting schedule changes into the MRP system. This updated MRP schedule then served as the forecast for the rest of the supply chain, including the other two imager departments, imager subassembly (IS) and imager fab (IF). The weekly changes to the forecast led to system nervousness, resulting in numerous expedite and de-expedite notices. In addition, long lead times and inaccurate forecasts led to production and inventory oscillations over a longer term.

Part of my work was to help final camera assembly to transition to a replenishment scheduling process. In such a process, production would be scheduled to equal demand over the previous review period. One of my challenges was overcoming resistance from the purchasing department that came when deviating from the MRP schedule, since such changes might result in expedite and de-expedite flags.

Challenges and opportunities of inventory management

At Kodak, the major inventory management challenge was establishing a formal process by which appropriate inventory levels could be set and achieved. At the time that I conducted this research, each department acted independently, and communication between them was limited. Most of the inter-departmental communications consisted of updated demand forecasts or quality information. Typically, the departments didn't look upstream or downstream to assess total supply chain excesses or shortages, but focused on their own department inventory (unless there was a shortage). This local department focus, coupled with forecast inaccuracy caused oscillating inventory. While Kodak makes great efforts to improve operating performance and to reduce inventory levels and other costs, these inventory reductions are sometimes made without first using formal methods to identify the proper inventory targets for meeting customer demand. Instead, inventory targets were based on rules of thumb developed through years of intuition. As a result, inventories may be higher or lower than needed. Without formal methods and models, it also becomes very difficult to assess the impact of operating improvements on the appropriate inventory levels.

2 Application of framework at Kodak

While at Kodak I developed a set of tools that are appropriate for setting inventory targets throughout the imager supply chain. My long-term goal is to embed these tools into the supply chain's production and inventory management processes. This set of processes would then constitute an effective framework for managing the supply chain. This section discusses the progress that we made while I was at Kodak, and highlights some of the work that remains to be done. The steps that we followed to implement the framework are:

1. Understand demand and supply parameters which affect inventory
2. Evaluate and design the supply chain
3. Choose an appropriate inventory model and use it to establish the appropriate inventory targets
    - Use the base stock model base stock model to determine the appropriate inventory targets for imager subassembly (IS) and final camera assembly (FCA & DC)
    - Use a base stock model, adapted to account for yield variation, to set targets for the Imager Fab (IF)
    - Develop an anticipation stocking strategy for large orders
4. Identify and make necessary changes to meet inventory targets
5. Identify improvement opportunities
6. Establish a regular review process

The first six sections correspond to these steps. The final section describes the results achieved, as well as projected results.

2.1 Understand Demand and Supply Parameters Which Affect Inventory

Inventory is a function of several factors, including demand (mean, variation), supply (mean, variation), time (manufacturing and component lead times, review period, and customer lead time), and customer service level targets. The first step in choosing an appropriate inventory model is to understand the environment, including customer demand and product supply.

Since a manufacturer's goal is to match product supply to customer demand, it is important to properly characterize the demand seen by the manufacturing system. This demand may have an underlying distribution (e.g. normal or Poisson), but it may also depend on product lifecycle phases, on seasonal factors, and on large infrequent orders. These factors should be considered when choosing a model and using it to set inventory target levels. The following figure illustrates the different types of demand that exist for Kodak's professional digital cameras.

Figure 2-2: Types of Demand

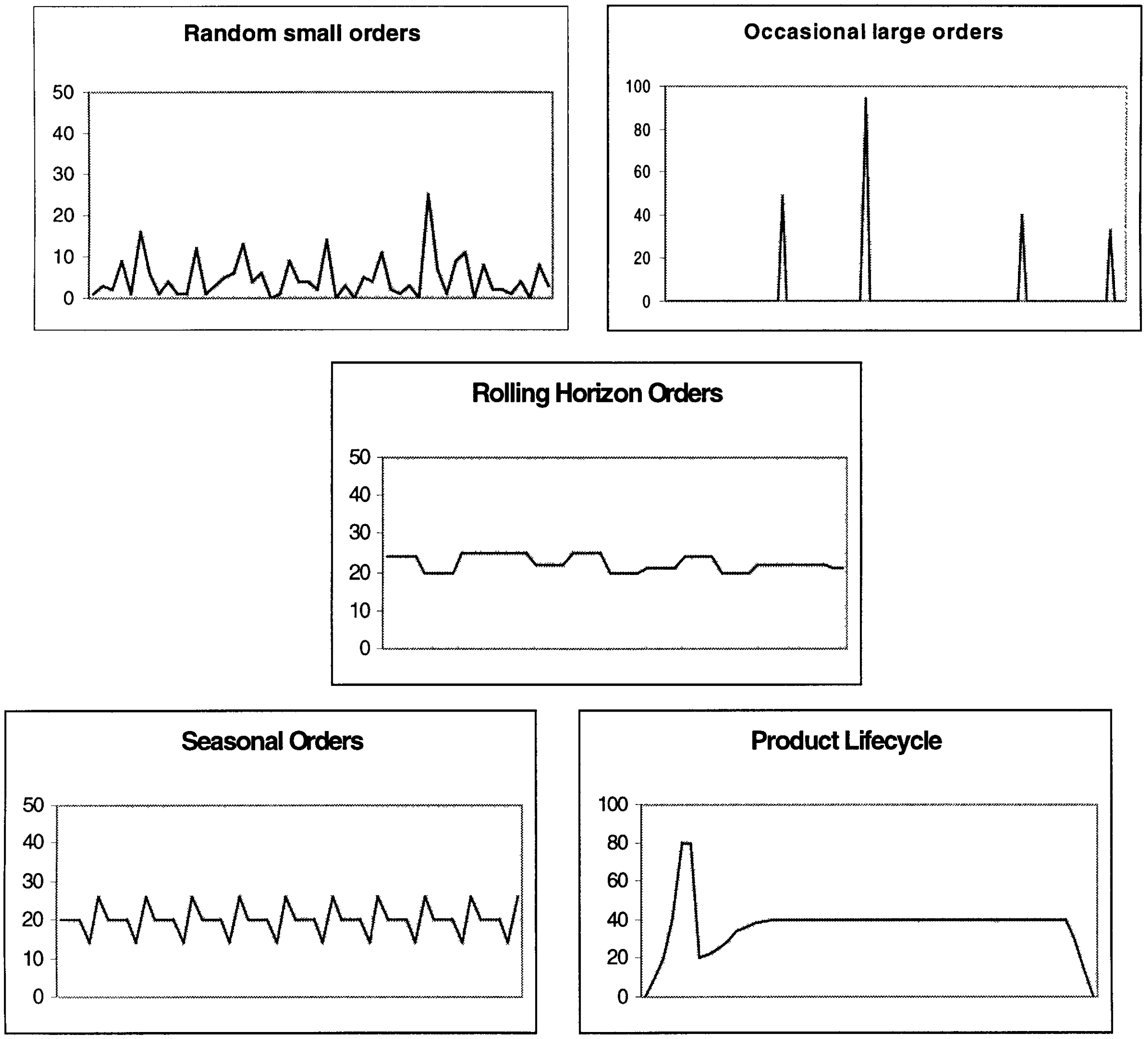

I found it useful in my research to consider demand with respect to the frequency of the orders, the variation in order size, customer lead time, and longer-term trends. The first three graphs show different types of random demand without underlying trends. The last two show trends which may underlie these random orders, including seasonality, and upward and downward ramps related to product lifecycle. Other possible trends, such as ramps due to marketing campaigns and competitive pressures, are not discussed in this thesis.

Order Frequency, Size, and Customer Lead Time

Random Small Orders

One useful measure of order regularity is the coefficient of variation (c.v.), which is defined as the ratio of standard deviation of demand to mean demand (σ/μ). I divided camera models into two types, 'A-type' and 'B-type' based on their coefficient of variation. 'A-type' cameras have a lower c.v. Typically, these are the color version cameras. 'B-type' (typically monochrome and infrared) cameras are ordered so infrequently that when looked at on daily intervals, they do not follow a normal or Poisson distribution. For such products, I recommended a slightly different approach for setting inventory targets. The following graphs show typical examples of 'A-type,' 'B-type,' and aggregate camera demand over ten weeks (including large orders).

Figure 2-3: Camera Demand

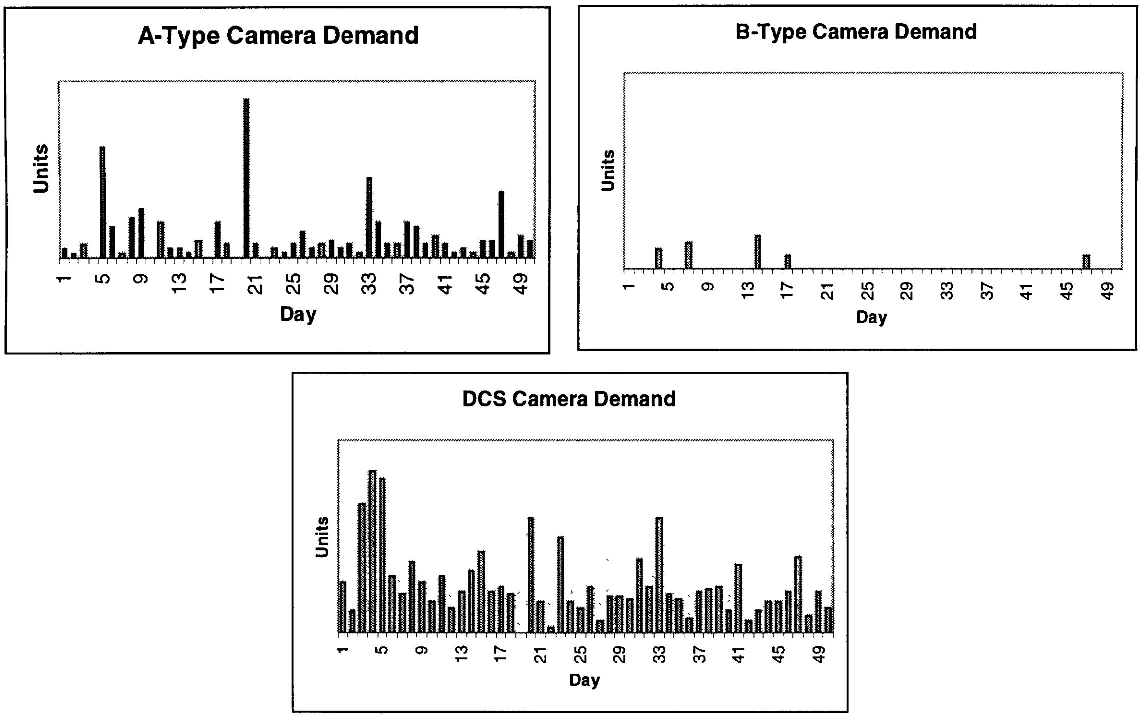

Large orders

Many Kodak professional digital camera products have an infrequent large order demand component in addition to the underlying random small order distribution. At Kodak, these orders are difficult to predict and are typically placed with a due date that is less than the component and manufacturing lead times. Additional inventory should be held in anticipation of these large orders. The manufacturing and scheduling people at final camera assembly try to complete large orders within the short order lead time they are given by the customer. The assumption is that the customers are not willing to take smaller partial shipments. Some managers and schedulers were looking into the possibility of having customers receive smaller partial shipments. If customers would accept partial shipments, then the anticipation stock could be held further back in the supply chain at lower cost.

For the purposes of planning inventory, I separated large order demand from random small order demand. If I had kept the small and large demand components together, I would have calculated inventory targets higher than needed to meet daily random orders, but insufficient to meet the large orders. Additionally, this inventory would likely be held at the wrong location (e.g. finished goods rather than raw materials). To deal effectively with large infrequent orders, I proposed an anticipation stocking strategy (described in section 2.3) that considers stocking location, desired customer lead time, and capacity constraints.

Rolling horizon orders

Another type of demand is what schedulers at Kodak referred to as a "rolling horizon" order. For a small number of models, Kodak has a single customer who places orders well in advance. On a monthly basis, the orders for future months are updated. Orders within a certain time period are "fixed" such that they can not change. This serves to provide visibility into customer demand. Since this demand is known in advance, it does not necessarily require additional safety stock inventory (if demand is known far enough in advance, the forecast is perfect and variation is essentially zero). The farther in advance that this demand is set, the farther back the safety stocks may be held.

Trends

Seasonality

In general, demand often has an underlying seasonal demand component, at the weekly, monthly, quarterly, and even yearly level. Upon examining demand data at Kodak, I found that seasonality exists at the weekly level, with average orders being higher than the grand average on Mondays, and lower than the grand average on Fridays. The following figure shows the relative mean and variation of demand for each day of the week over 35 weeks for all professional digital camera models.

Figure 2-4: Demand Seasonality

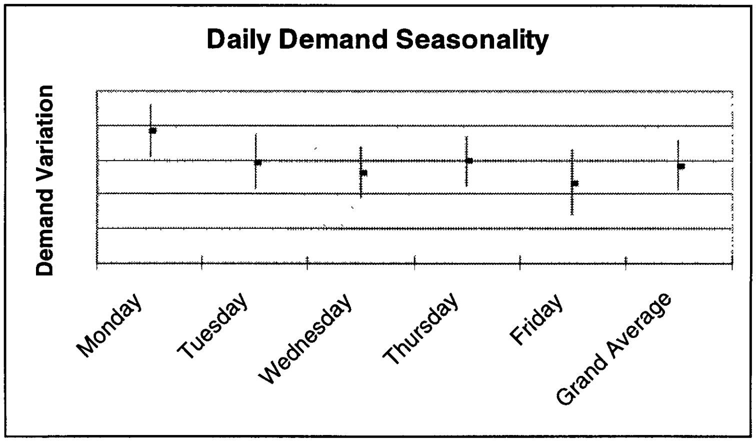

At the time that I did this research, final camera assembly (FCA) scheduled production weekly, so this seasonality did not affect operations. However, some of the improvements that we considered included going to daily review and scheduling (r=0). Consequently, I discussed the implications of this seasonality with the production supervisor. A seasonal demand component requires us to decide whether to chase demand through increasing and decreasing production, or to smooth production by carrying an additional anticipation stock to cover the seasonal demand peaks. To chase demand, final camera assembly would staff up to meet the higher Monday demand and staff down to meet Friday demand. The production supervisor preferred this approach to building ahead of demand on Fridays. He found it helpful to know about this seasonality because it helped him plan training and improvement activities for the end of the week.

While some production and scheduling people felt that there was also seasonality related to the end of month, quarter, and year, I was not able to determine if this was the case because data sufficient for a statistical analysis did not exist.

Lifecycle

Products go through different stages in their lifecycle, including:

- Product introduction and channel filling
- Demand ramp-up
- Midlife demand
- End-of-life ramp-down

A product's current lifecycle stage can help determine the approach to setting inventory targets. During new product introduction, filling the channel, and initial market acceptance, stable historic demand data does not exist, so forecasts should be used to set the desired inventory and production levels. During mid-life, forecasts or historical data can be used with the base stock model to set the appropriate inventory levels. Mid-life represents the bulk of a products lifecycle, and is therefore the focus of this thesis. However, shortening product lifecycles are causing the midlife to shrink, making a smaller percentage of the product's total life. As this trend continues, it increases the need for better methods for dealing with the early and final stages of the product lifecycle.

Forecasts

A simple test of forecast accuracy is to compare the forecast error to the error that would have resulted if historical averages were used as the forecast. If the forecast contains useful data the actual forecast error should be similar to or lower than the error associated with historical averages. In the case of the cameras that I studied, the business unit forecast was less accurate as an indicator of demand than was a forecast based on the actual previous demand data. The following table of disguised data compares the forecast error (standard deviation) of three A-type camera models to the error associated with using past averages as a forecast. The disguised data is for monthly demand periods, taken over six months. Each forecast is for the following month.

Table 2-1: Forecast Error

Model	Mean Demand	Forecast Error	Forecast Error using 6-month moving average as Forecast
Model 1	289	227	92
Model 2	573	379	128
Model 3	138	56	24

Large orders contributed to this forecast inaccuracy. The practice at Kodak was to raise the forecast when a large order was anticipated, and lower it if it was clear that an order would not materialize. As a result, forecasts varied greatly from month to month, and parts were expedited or cancelled as expectations of large orders changed. These fluctuations due to big orders increased the degree to which "beer game phenomena"[1] existed.

In summary, the major characteristics of this supply chain which affected inventory and production targets were:

- A mixture of large and small orders
- Different demand types: products with high and low coefficients of variation

- Historic demand-based forecasts that are more accurate than business unit forecasts

2.2 Use SIPModel to Evaluate and Design the Supply Chain

Once I had developed an understanding of the digital camera supply chain's supply and demand parameters, I was prepared to evaluate it and consider alternative configurations. The Strategic Inventory Placement Model (SIPModel)[2] provides a powerful way to evaluate a supply chain and set global inventory targets. Indeed, it allowed me to create a graphical representation of the global digital camera supply chain. Inputs to the SIPModel include customer service time for the final stage, customer service level, cost added at each stage, manufacturing cycle time at each stage, and demand parameters (mean and standard deviation). An algorithm[3] finds the lowest total inventory cost to meet the desired customer service times for the supply chain. The SIPModel places safety stocks at strategic locations within the supply chain to de-couple a process stage from downstream stages. Stages between safety stocks are coupled in such a way that they are linked together in a single pull system. To illustrate this, figure 2-5 shows an example of the SIPModel user interface. Each box represents a process stage. Within each box is either a circle which represents an operation, or a circle and a triangle which represents an operation followed by a de-coupling safety stock. In this example, the stages labeled CFA, pkg/test, subassy, camera, ckt board, build/test/pack, process order, and ship to customer are all grouped together. Production for each of these stages is planned together. Production planning at the stage labeled wafer fab would be done independently of downstream stages, due to its de-coupling safety stock.

Forcing the service time for a given stage to zero allows us to compare alternate configurations. This stage is then fully de-coupled from downstream stages, resulting in the need for an additional safety stock between it and the following stage. The SIPModel can be used to compare the costs of these various configurations. This comparison can be used as a sort of sensitivity analysis of the costs of alternate configurations. The SIPModel analysis may be used to guide supply chain improvements by setting inventory targets and specifying de-coupling stock locations.

The SIPModel assumes a well-behaved demand distribution, no significant supply variation (yield or cycle time), and a continuous review pull system. A major insight of the SIPModel is that by looking at the supply chain in a global fashion, we can reduce the number of stocking locations for safety stock inventory, thus reducing the total safety stock investment.

In the case of the digital camera supply chain, the SIPModel's optimized configuration combined the back end of the imager wafer fab (IF), the circuit board assembly (CBA), the final camera assembly (FCA), and distribution center into one effective department. This means that each step in these combined stages would be fully coupled to the other stages, with no intermediate safety stocks. The figure (2-5) below shows this optimized configuration.

Figure 2-5: SIPModel Optimized Supply Chain Configuration

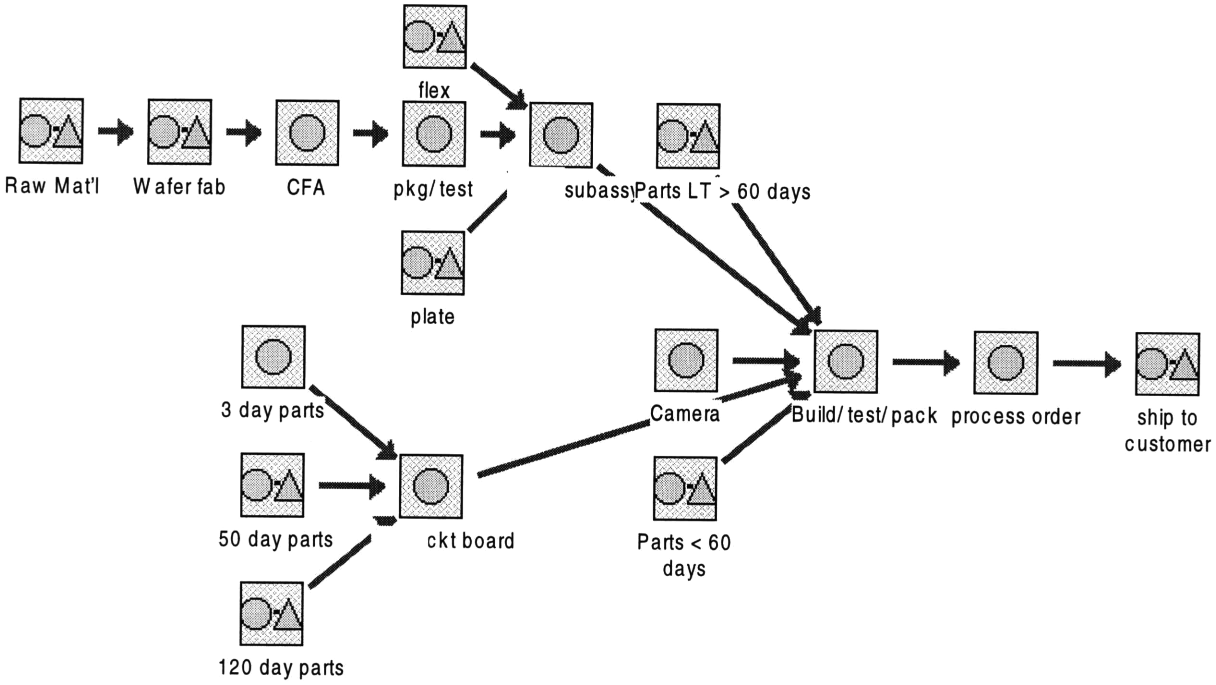

In reality, this optimized configuration is not feasible. Both the imager fab (IF) and circuit board assembly (CBA) supply many other customers, which were not modeled in the digital camera SIPModel. The actual supply chain configuration in operation at the time I conducted my research was better suited to serving these multiple customers. When I compared the model of this actual configuration (figure 2-6) with the SIPModel-optimized configuration (figure 2-5), the difference in predicted inventory costs was a reasonably small 2.1%. The actual supply chain inventory, however, was higher than the levels suggested by the model for the actual configuration (figure 2-6). Part of this difference is due to additional stock needed to compensate for yield variation at the imager fab (IF). But even after accounting for yield variation, there is excess inventory which can be removed from the supply chain.

Figure 2-6: Actual Supply Chain Configuration

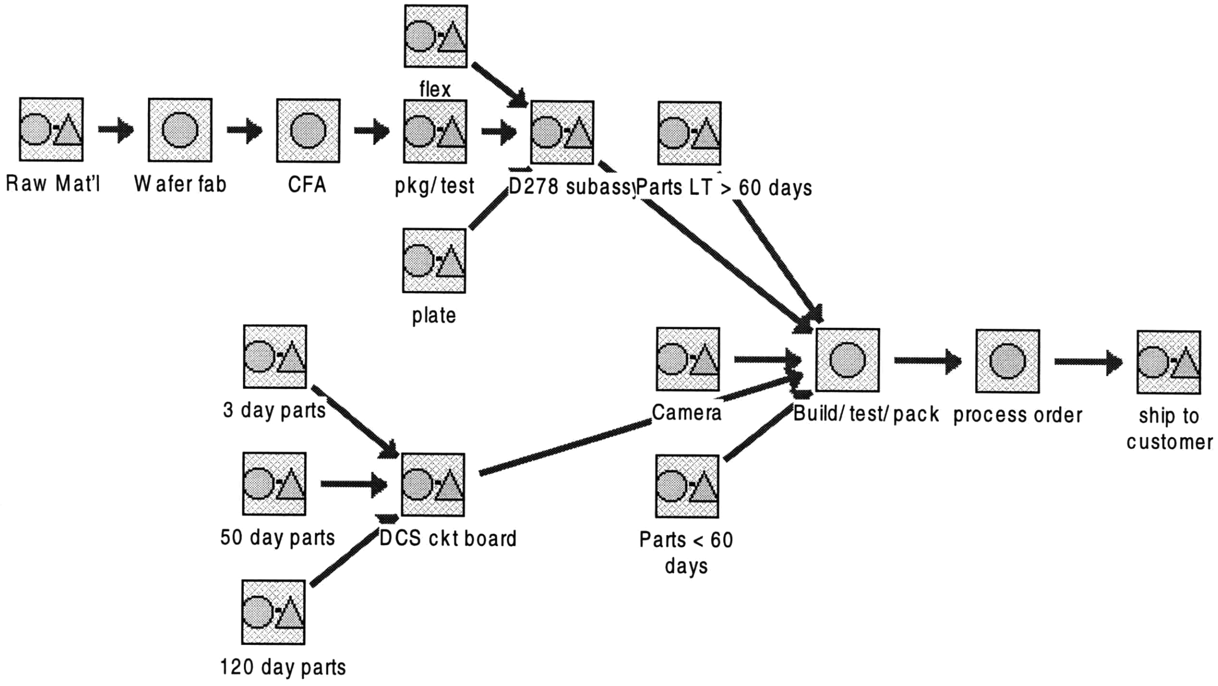

Thus the SIPModel helped to identify opportunities for inventory reduction within the existing configuration. It also established that the current configuration was close to being optimal.

2.3 Choose appropriate inventory models

A key to effective inventory management is the ability to set appropriate inventory targets. Simple models can effectively replace the ad hoc approaches that are often used to set these targets. An appropriate model yields improved targets that do not rely on years of acquired intuition. Manufacturers can use these same models to identify and evaluate improvement opportunities.

As described in the previous section, I used the SIPModel to determine where de-coupling stocks should exist. I then used the base stock model to set inventory targets within the de-coupled IS (imager subassembly) and FCA (final camera assembly) departments. For the imager fab, I used a yield variation inventory model. The base stock model formula is essentially the same formula that is used by the SIPModel to do its calculations and optimizations. While I could have used the SIPModel to determine inventory targets, I chose to use the base stock formula directly. This was primarily a matter of convenience.

The Base Stock Model

I used the base stock model as a foundation for my work at Kodak. I include here a higher-level explanation of the base stock model. For a more complete and mathematically detailed description, see Graves (1987). A form of the simple base stock formula is:

$$I = \mu_D * (r + l) + z * \sigma_D * \sqrt{r + l} \qquad (1)$$

Where: I = Inventory available over coverage time, or WIP + FGI

$r+l$ = coverage time, or review period + manufacturing lead time

μ_D = single period mean demand

σ_D = single period standard deviation of demand

z = customer service level factor

Also, it is useful to conceptually break the base stock inventory into its two components. Pipeline stock, or average WIP is equal to $\mu_D*(r+l)$. Safety stock is equal to $z*\sigma_D *sqrt(r+l)$. The quantity $r+l$ represents the time between reviews plus the time from review until the product will be manufactured and available for use at the next stage (due date). For example, if scheduling is done every five days (r=5), and production planned today is actually completed in 9 days, then $r+l$ = 14 days, regardless of the actual manufacturing or touch time.

Conceptually, we use the base stock model to balance supply with probable demand. The supply available over the time period $(r+l)$ is given as I. Demand in this case is a random variable, characterized by a single-period mean (μ_D) and a single-period standard deviation (σ_D). Assuming that the demand is independent of supply and that it follows a well-behaved distribution, we can say that demand over the coverage time $(r+l)$ will be less than or equal to:

$$Demand \le \mu_D * (r + l) + z * \sigma_D * \sqrt{r + l} \qquad (2)$$

The customer service level factor (z) determines the probability that demand will be less than or equal to the expression on the right. By setting the total base stock inventory (I) equal to the expression given in equation (1), z corresponds to the probability that supply will be sufficient to meet demand until the next production arrives $(r+l)$.

For example, if demand is normally distributed with a mean of 10 units per day and a daily standard deviation of 5, and $r+l$ is equal to 4 days, we could calculate the desired base stock inventory for a desired 95% coverage (z = 1.64) as:

$I = 10 * 4 + 1.64 * 5 * \sqrt{4} = 56.4$ units.

The base stock model assumes that both lead time and supply variation are negligible and that demand can be modeled as a well-behaved random variable.

The base stock model shows that in order to lower inventory and maintain the same customer service level, either review period, lead time, or demand (average or variation) must be reduced. At manufacturers in general, inventory reduction targets are often set (somewhat arbitrarily) by upper management. The base stock model implies that unless this inventory reduction is matched by an appropriate decrease in review period, lead time, or demand, the result will be lower customer service.

Final Camera Assembly and Distribution Center

In applying the base stock model for final camera assembly and the distribution center, I calculated the production lead time, *l*, to be the time from the production planning event to the time that production is due. This assumes that raw materials are available.

As noted earlier, demand for most professional digital cameras consists of small random orders plus infrequent large orders. I separated these components and used the base stock model to establish inventory targets appropriate for meeting small random orders. Once I removed large orders from the data, I classified models having a coefficient of variation (σ/μ) less than 1 as 'A-type' and models having a coefficient of variation greater than 1 as 'B-type.' I chose to use different customer service level factors (z) for these different types of camera models. The following matrix shows the rule of thumb that I used to assign customer service levels and their associated z values.

Table 2-2: Customer Service Level Values

		Relative Model Cost	
		Low & Mid Range	High End
Demand Type	A-Type	99% (2.33)	90% (1.28)
	B-Type	50% (0)	50% (0)

For all but the most expensive 'A-type' camera models, I chose a customer service level of 99% which corresponds to a z factor of 2.33. Demand for these camera models represented over 75% of total camera unit demand. I set a very high customer service level for these models so that most of the exceptional demand situations would be either large orders or less-frequently ordered models. The exceptions to this rule were a few very expensive models, where the demand planners felt that the inventory savings were worth the customer service tradeoff. One other exception was a new model, which did not have historic demand data. For this model, I used forecasted demand numbers for demand mean and standard deviation, and used a z factor of 1.28. Once demand for the product reaches steady state, I would recommend using historic data. For 'B-type' cameras, I used a very low customer service level (z=0). These models were ordered so infrequently that the demand planners thought it best to have just a few cameras on hand. Their goal was to move such low-volume, high-variation models to a "make-to-order" setting, with no finished goods inventory. This would become more feasible if lead time were reduced to under a week and scheduling occurred daily (r=0).

By using the base stock model, I generated a set of inventory targets that were in line with the expectations of people who were familiar with the products. This similarity between the model-generated targets and the team's intuition helped facilitate acceptance of the base stock model. These targets were, however, lower than actual inventory. By setting targets that included both WIP and FGI, we made the production planning process more focused on replenishing current demand.

Imager Subassembly (IS)

I also used the base stock model to calculate inventory targets for the imager subassembly department. I discussed the base stock model with the IS production scheduler, but we did not have time to implement these new inventory targets. The inventory targets currently in use are fairly close to those that the model calculated.

Extensions to the Base Stock Model to Account for Yield Variation

The base stock model makes a number of simplifying assumptions that may or may not be reasonable for a given manufacturing environment. These assumptions include:

- Capacity is generally unconstrained
- Manufacturing lead time is not highly variable
- Manufacturing yield variation is negligible
- Demand can be approximated as a stationary, well-behaved random variable

When these assumptions are not approximately true, adjustments should be made to the base stock model. For digital camera imagers, yield variation was significant at the imager fab. Therefore I used an extended base stock model that incorporates yield variation for the imager fab (IF).

Like the base stock model, the extended model balances supply to probable demand over the coverage time. We start with a characterization for the demand random variable over the coverage time:

$$E(Demand) = \mu_D * (r + l) \qquad \text{Variance(Demand)} = \sigma_D{}^2 * (r + l) \qquad (3)$$

This is the same expression used to characterize demand in the base stock model. Next we quantify the probable supply that will be available over the coverage time. For the base stock model, this was the current work in process (WIP) plus finished goods inventory (FGI). Now, however, supply is uncertain, due to yield variation. Like demand, supply can be represented as a random variable. If production is done in lots or batches, each lot having a mean (μ_S) and standard deviation (σ_S), then we can characterize the in-process supply as:

$$E(WIP) = \mu_S * Q \qquad \text{Variance(WIP)} = \sigma_S{}^2 * Q \qquad (4)$$

where Q is the number of lots in process. By adding the current finished goods (FGI) which are known to be good, we arrive at the following characterization for supply over the coverage time:

$$E(Supply) = FGI + \mu_S * Q \qquad \text{Variance(Supply)} = \sigma_S{}^2 * Q \qquad (5)$$

By combining equations (3) and (5) such that Supply – Demand > 0, we arrive at the following expression:

$$E(Supply - Demand) = FGI + \mu_S * Q - \mu_D * (r + l)$$
$$\text{Variance(Supply - Demand)} = \sigma^2{}_D * (r + l) + \sigma^2{}_S * Q \qquad (6)$$

This development makes the assumptions that supply and demand can be modeled as independent random variables. To assure that supply available over the coverage time exceeds demand over the coverage time, we need the following equation to hold:

$$FGI + \mu_S * Q - \mu_D * (r + l) \text{ - z} * \sqrt{\sigma^2{}_D * (r + l) + \sigma^2{}_S * Q} \geq 0 \qquad (7)$$

Where z denotes a safety factor that determines the probability that supply will be greater than demand. This equation can be used to plan wafer starts at the imager fab by solving for Q. To make it easier to use, I embedded equation (7) into a Microsoft® Excel spreadsheet. Using Visual Basic macros, I created a tool that calculates desired wafer lot starts based on current yield, demand, and supply chain inventory data. This tool can also solve equation (7) for z. This allows the user to predict customer service levels for various wafer start scenarios. An example of the user interface of this spreadsheet program is shown below (data is disguised):

Figure 2-7: Imager Model Spreadsheet

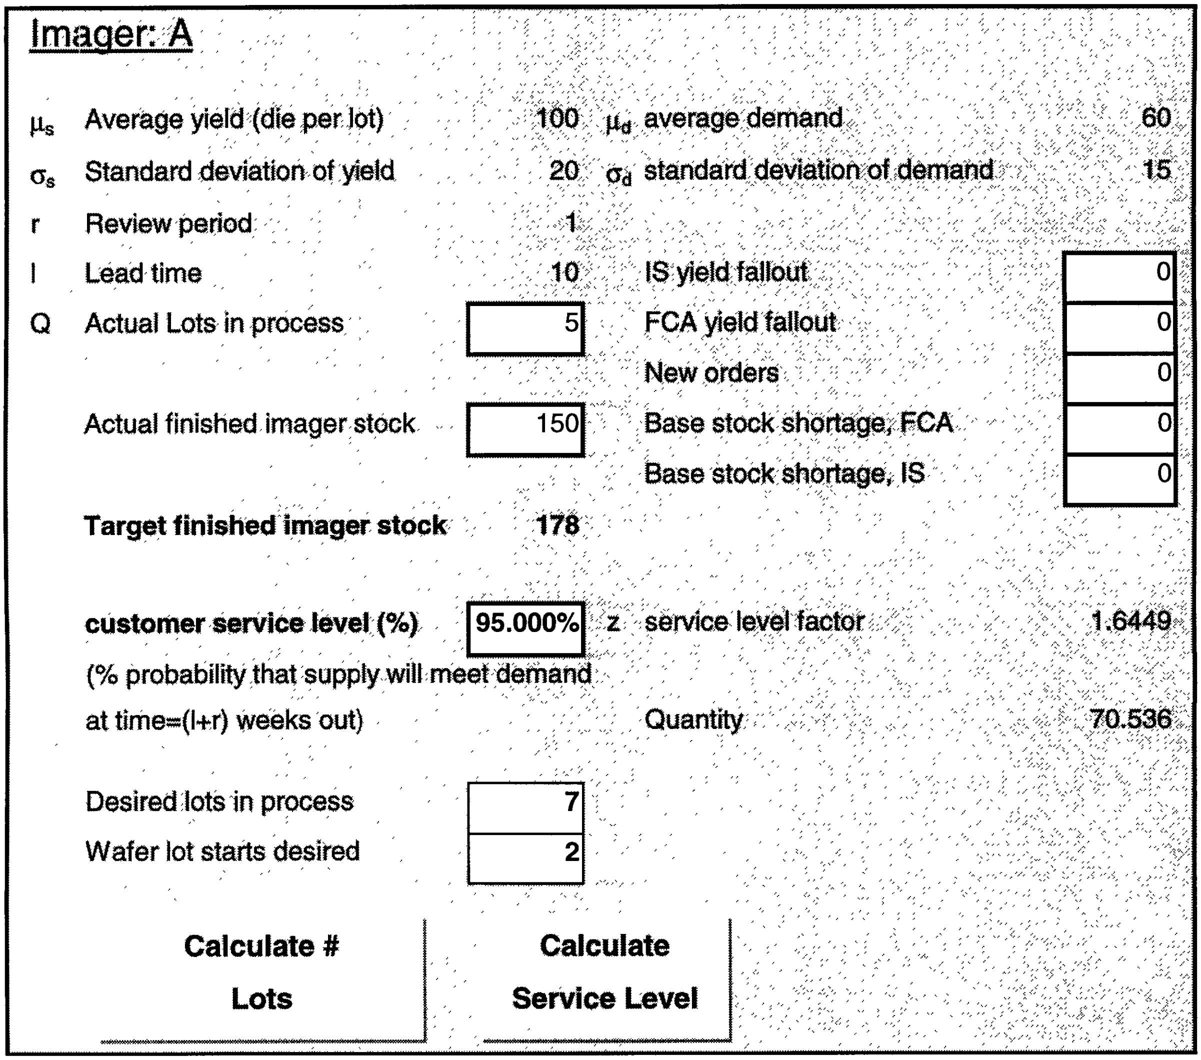

Imager: A

μ_s	Average yield (die per lot)	100	μ_d	average demand	60
σ_s	Standard deviation of yield	20	σ_d	standard deviation of demand	15
r	Review period	1			
l	Lead time	10		IS yield fallout	0
Q	Actual Lots in process	5		FCA yield fallout	0
				New orders	0
	Actual finished imager stock	150		Base stock shortage, FCA	0
				Base stock shortage, IS	0
	Target finished imager stock	**178**			
	customer service level (%)	95.000%	z	service level factor	1.6449
	(% probability that supply will meet demand at time=(l+r) weeks out)			Quantity	70.536
	Desired lots in process	7			
	Wafer lot starts desired	2			

Given recent demand and yield data, I used this model to calculate targets for the imager fab (IF), expressed in desired lots in process and desired finished goods inventory. Since yield is variable, these targets will vary according to the actual level of finished goods inventory.

Anticipation stock

Thus far, I have described the models that I used to set inventory targets sufficient for meeting small random orders. In this section I propose an anticipatory stocking strategy for meeting large orders. Anticipation stocks are used to counteract demand seasonality, capacity constraints, and unanticipated large orders. Such stocks help smooth production. Generally speaking, a large order could be defined as any single order that is an order of magnitude above an average order. For the purposes of the digital cameras, I define a large order as any single order for 10 or more cameras. There are two elements to meeting big order demand: parts availability and capacity.

Parts Availability

To set an anticipation stocking strategy that is appropriate to meet large digital camera orders, I explored three questions. First, what is the time frame that the anticipation stock should cover? Second, how much anticipation stock is appropriate? Third, how should capacity affect the stocking strategy?

To determine the appropriate time frame, I considered the response time from various points in the supply chain. The figure below illustrates this principle using disguised lead times.

Figure 2-8: Anticipation Stock

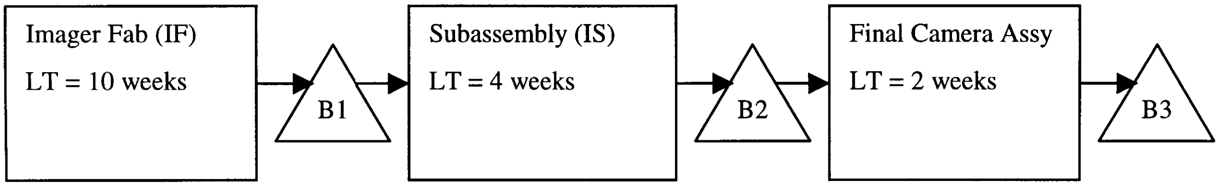

Stock held at B3 is available immediately to fill customer demand. It therefore has a response time of zero. Assuming that capacity is not an issue, then anticipation stock held at B2 has a response time equal to the lead time of FCA, or 2 weeks. For stock at location B1, the response time is equal to the cumulative lead time of IS and FCA, or 6 weeks. If we assume further that the imager fab has raw materials available, the worst case response time is the cumulative lead time of the entire supply chain, or 16 weeks. A large order that exceeds the combined stocking levels of B1, B2, and B3 would therefore be available within 16 weeks.

To determine an appropriate stocking level, I did two things. First, I reviewed historic large order demand. In doing so, I used a sliding window the size of the cumulative lead time (in the above example this would be 16 weeks) to look at historic large order demand for a given imager type. I then established the maximum historic demand over this cumulative lead time as a baseline. The table below shows an example of large order demand over such a sliding window.

Table 2-3: Large Order Example

Starting week	Ending week	Historic large orders over time interval from starting week to ending week (cumulative lead time)
1	16	47
2	17	62
3	18	27

In the above table, the baseline number for large orders is 62 (in actuality, we would want to look at more than three intervals). Once I had calculated a baseline, I used it as a starting point in my discussions with people familiar with the products. Together we determined an inventory target based on expectations of future large orders. This anticipation stock target reflected the maximum expected large order demand over future cumulative lead time intervals. Once we determined the anticipation stocking levels, we discussed where in the supply chain the stock should be held. To do this, we considered the response time from various points in the supply chain, as shown in figure 2-8. Based on historic large order demand and discussions with demand planners, I recommended the following relative anticipation stocking strategy for the three imager types used in professional digital cameras:

Table 2-4: Proposed Anticipation Stock Targets (relative units)

Imager	IF (B1)	IS (B2)	FCA (B3)	Total
A	7	3	0	10
B	28	12	0	40
C	4	2	0	6

Since it is less costly to hold inventory earlier in the process (i.e. fewer value-added and differentiating operations have been performed), there is a tradeoff between cost and customer service. The farther back the inventory is held, the less expensive it is, but the longer it will take to complete the order. The ultimate solution is to decrease cycle time.

This proposed approach to anticipation stock differed from the way large orders were handled in the past. Rather than having separate anticipatory stocking targets, large orders were factored into and out of the forecast as they were anticipated. This led to inventory levels that sometimes appeared to be justified and sometimes did not. As a result, the level of expedite and de-expedite activity was greater than the actual large orders should have warranted. Due to time and resource constraints, I did not succeed in implementing my proposal while I was at Kodak. Nevertheless I believe that such a strategy will benefit the supply chain. First, it should reduce the degree to which inventory oscillates, since a stable target would replace greatly varying forecasts. Second, it allows the business unit to make a quantifiable tradeoff between maximum large orders that can be supported and inventory investment.

Capacity

The strategy described above assumes that if parts are available, then a large order can be processed according to normal manufacturing lead time. Lead time, however, is often related to capacity and demand. In cases where a large order arrives, extraordinary activities such as overtime can be used to process the additional units. In these circumstances, linear programs and optimizations may also be helpful to prioritize work. An additional tool that I created is also well suited to the digital camera manufacturing environment. Since big orders are known a few weeks in advance with some confidence, we can forecast the amount of time it will take to assemble the additional units at some probability corresponding to the safety factor z. This calculation tool works by solving the equations for t_n and t_o:

$$S_n * t_n - O - \mu_D * t_n - z * \sigma_D \sqrt{t_n} \geq 0 \qquad (8)$$

$$S_o * t_o - O - \mu_D * t_o - z * \sigma_D \sqrt{t_o} \geq 0 \qquad (9)$$

Where S_n and S_o = the daily capacity without and with overtime, respectively

O = the size of the large order

t_n = the time needed to make the order at a confidence factor of z

μ_d and σ_d = the daily demand mean and standard deviation excluding the large order

The first equation solves for the time needed without using overtime, and the second solves for overtime conditions.

Figure 2-9: Big Order Calculator Spreadsheet

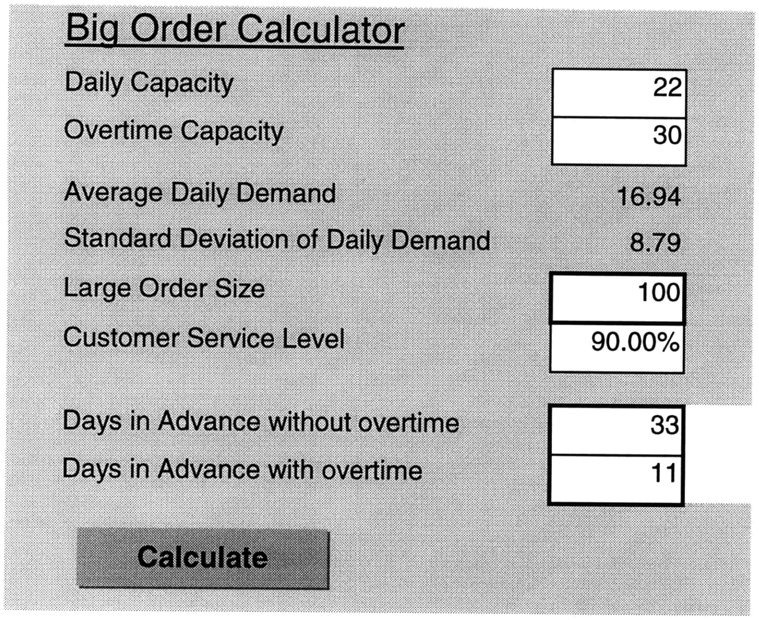

Big Order Calculator

Daily Capacity	22
Overtime Capacity	30
Average Daily Demand	16.94
Standard Deviation of Daily Demand	8.79
Large Order Size	100
Customer Service Level	90.00%
Days in Advance without overtime	33
Days in Advance with overtime	11

Summary

By utilizing the base stock model I was able to set inventory targets for the imager supply chain that are appropriate for meeting small random orders. The anticipation stocking strategy that I proposed can be used to set appropriate inventory levels for meeting large orders, and in quantifying customer service/supply chain inventory cost trade-offs.

2.4 Identify and make necessary changes to meet inventory targets

For the inventory targets to be effective in achieving the desired customer service levels, operations processes need to support the underlying assumptions of the model. This section describes our implementation progress and the challenges we encountered.

Implementation progress

The goal of my research and work at Kodak was to provide a framework for setting and reviewing global inventory targets, and to begin implementing this framework throughout the imager supply chain. While we developed the framework and identified the operational changes needed to support them, implementation was still in the early stages by the end of my research assignment. The implementation was carried out at two levels: within the departments (specifically focused on final camera assembly and the distribution center), and over the integrated imager supply chain (IF, IS, FCA and DC).

Final Camera Assembly

The first implementation phase was within the final camera assembly department. Our goal was to implement a replenishment-based pull system that was more tightly coupled to actual demand than was the current system. Previously, production was based more on inaccurate forecasts than on actual orders. To help tie production to demand, we created a worksheet which could be used each review period to plan production. Essentially, it listed for each product the desired base stock level, and the difference between desired and actual from the previous period. To this shortage, we added orders received over the previous review period to arrive at a desired production figure. Based on capacity and parts availability, we would then schedule as much of the desired production as possible, resulting in a new base stock shortage for the following review period.

We started using this new worksheet in September. The new base stock targets were significantly lower than the actual final camera assembly and distribution center base stock inventory in September. By the time my research ended, we were well on our way to meeting the desired target, which when completed will represent a total final camera assembly and distribution center inventory reduction of approximately $240,000. We determined that additional savings could be achieved by reducing lead time and review period (see section 2.7). At the time that we were exploring these possibilities, other projects of higher priority took the resources needed to make these changes. Nevertheless, the same base stock model and replenishment system principles can be applied throughout the supply chain.

Imager Subassembly, Imager Fab, and the Integrated Supply Chain

The second implementation phase was to integrate the base stock framework across the imager supply chain including imager subassembly (IS) and imager fab (IF). Using the model results, I suggested target inventory levels. For imager subassembly (IS), these inventory targets were compatible with the targets that had evolved through time based on experience. At the imager fab, there had been no explicit stocking targets in the past. By creating a useable model in the form of a spreadsheet, I hoped to simplify the transition to explicit inventory targets. However, as of December when I finished my research, the model was not being used to plan production. The following three charts show that there is a sizeable opportunity for using such a model. For the three digital camera imagers, I show the actual supply chain inventory in terms of imager units at each department for a recent 12 month period. I also show the model-based target as a line across the chart.

Figure 2-10: Imager A Actual and Target Inventory

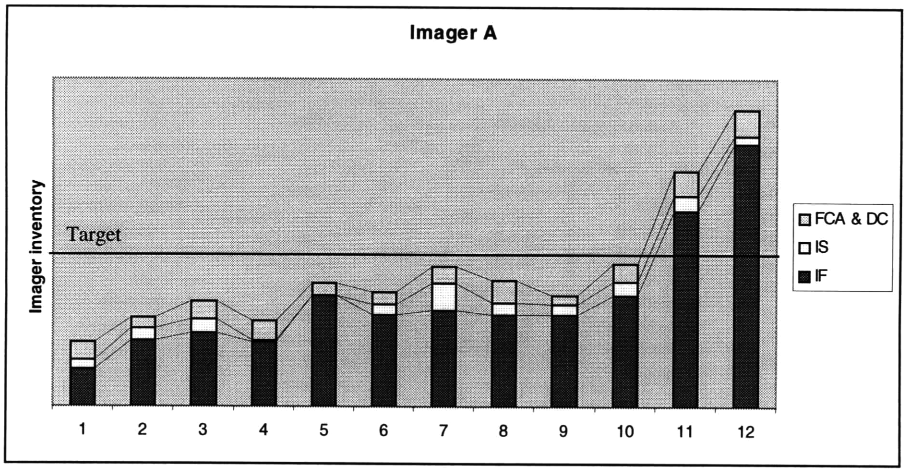

Imager A had too little inventory in the supply chain until month 11, when inventory began to rise rapidly. This would suggest that during the first ten months part shortages kept customer service below the desired level. The inventory buildup during the last two months is the result of a decision at IF to build ahead of demand to protect against shortages during a new manufacturing process ramp-up.

Figure 2-11: Imager B Actual and Target Inventory

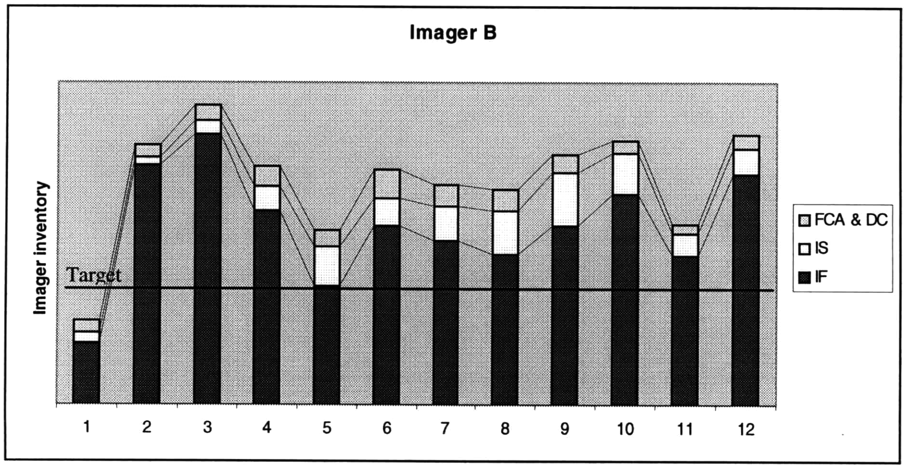

Except for month 1, Imager B inventory was consistently higher than the target. Consequently, parts shortages for this product were extremely rare, and there is a potential for significant inventory savings.

Figure 2-12: Imager C Actual and Target Inventory

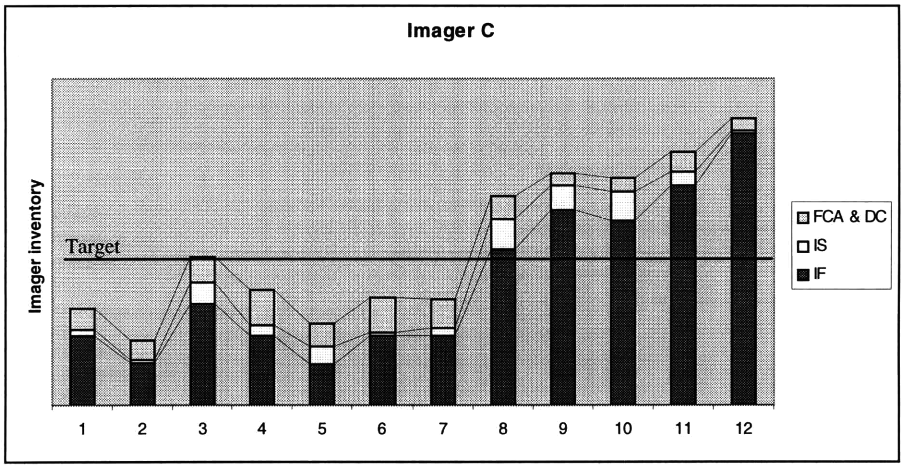

Imager C was understocked for the first seven months. Since this is the highest cost imager, IF tried to keep its investment in imager C as low as possible. After chronic shortages, IF raised imager C inventory. It appears that in the absence of formal modeling tools and techniques the inventory was raised higher than needed. Part of this inventory buildup is also due to the new manufacturing process ramp-up.

Outstanding Issues

Training

To use the framework I have described, it is important that the people involved in setting inventory targets and identifying operations improvements understand the underlying principles. To this end, I developed and taught a series of informal classes on the principles behind the models used, and the processes which would support them. One of the major points of the training was to teach the importance of looking beyond local finished goods inventory when considering inventory targets and production scheduling. Pipeline inventory excesses or shortages should also be considered, as should inventory status in upstream and downstream departments.

Systems

If the desired change is not embedded into current processes, or into new processes, the change will probably not endure. At Kodak, some processes were formalized and relied on computer systems, such as MRP II, while others were less formal and did not rely on an explicit information system. Appropriately embedding the models within the systems and processes that people in the supply chain use remains an important task that would help ensure the success of this project.

Metrics & Organizational Barriers

Often, the metrics by which an individual's success is measured focus on local performance. As such, they often do not align the people of different departments within a supply chain to find and achieve a more global optimum. For example, basing performance primarily on local inventory creates a powerful incentive for departments to shift inventory upstream or downstream to another department. This can lead to inventory oscillations and periods of product or component shortages.

I believe that the reporting structure at Kodak presents another challenge to achieving a global focus. The company's manufacturing groups seem to be dominated by functional organizations rather than product-flow focused organizations. As a result, it becomes more difficult to negotiate the proper metrics and goals, since there is no single authority to help align the different parties that are involved in meeting customer demand. One possible solution is to have a supply chain champion whose mission and authority allows him or her to align the various functions and departments to set appropriate inventory goals, and to develop and implement the processes that will achieve them.

A third specific challenge is that imagers used at final camera assembly (FCA) are a small percentage of imager fab (IF) revenue and volume. This can make it difficult to get the resources needed to implement this framework across the entire supply chain. However, this portion of IF's product line can be seen as a pilot program, with savings and improvements being readily transferable to other product lines. Again, a key part of the success of this program is identifying an individual to drive these efforts as supply chain champion. The search for this champion is still underway at the time of this writing.

2.5 Identify improvement opportunities

The models used to establish targets can also be used to quantify the inventory according to the reasons for holding it. They can then be used to help target improvement activities. For example, we can quantify potential savings of decreasing review period, manufacturing lead time, and demand variation. The yield variation-enhanced model allows us to further quantify the inventory-related savings of yield improvements. The following charts show inventory for each imager (A, B, and C) grouped by location, and by reasons for holding it. Qualitatively, these charts make an important point: the cycle time and yield variation at imager fab (IF) are the primary drivers for holding inventory. By improving these two components, total inventory can be greatly reduced.

Figure 2-13: Base Stock Targets by Supply Chain Location

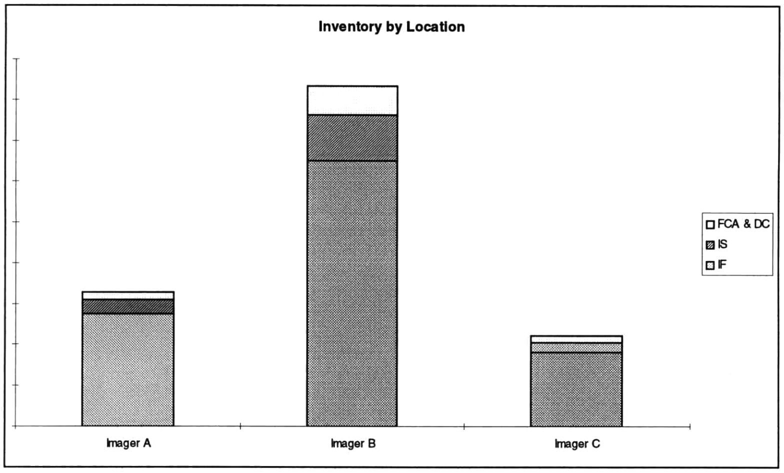

Figure 2-14: Base Stock Targets by Cause

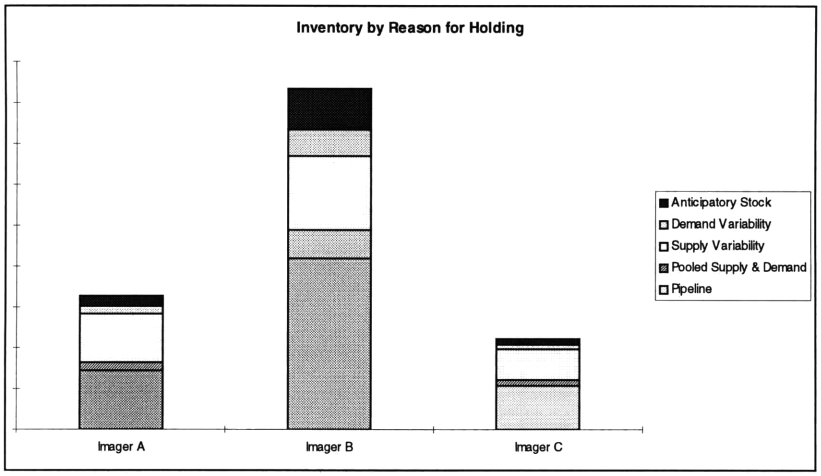

2.6 Establish regular review process

It is important to regularly review performance and goals for at least two reasons. First, by reviewing performance relative to inventory and customer service, we can learn how we are doing compared to the model and find areas in which to improve execution. Second, by reviewing inventory targets and the underlying parameters, we can make adjustments as conditions such as demand or lead time change.

Metrics and performance measurement

While there may be no ideal set of metrics that will automatically drive the right behavior across an entire supply chain, thoughtfully chosen metrics can help create organizational alignment. Metrics which show total inventory vs. desired inventory (such as are shown in figures 2-10 - 2-12), customer service levels (on-time delivery), and manufacturing cycle time are useful in assessing performance and identifying improvement opportunities.

Model maintenance

Since we are modeling a dynamic manufacturing system, with products whose demand characteristics change over time due to competitive forces and product lifecycle issues, it is important that we regularly review our model parameters and assumptions. At Kodak, we proposed a monthly process, where we would review the following:

- Previous month base stock targets
- Performance vs. targets, with root cause analysis
- Demand mean and standard deviation
- Review Period and manufacturing cycle time
- Customer service level
- Anticipation stock targets
- Product Lifecycle: New and End-of-Life Products

Once these parameters have been reviewed, new base stock targets can be generated from the model for the next month.

2.7 Results

While implementation was not complete by the time my research ended, the following results were either achieved or projected from further implementation of the framework.

Lower inventory

We reduced inventory at final camera assembly and the distribution center by just over $100,000 between September 29 and December 23, when comparing the same camera models. Inventory could be decreased by a total of $240,000 below the September level without further lead time improvements. We projected that this inventory could be cut by 62% of September levels by going to a daily review period and changing the lead time from 9 days to 5 days.

I estimated potential imager supply chain inventory reductions of $626,000 compared with October levels. This estimate makes the conservative assumption that we will need to increase stocking levels for non-imager components to support my proposed anticipation stocking policy. In actuality much of this inventory already exists.

Table 2-5: Potential Inventory Savings ($K)

Savings ($K)	Imager A	Imager B	Imager C	Total
IF	-31	152	260	381
IS	-8	84	112	188
FCA & DC (Imager Only)	46	12	32	90
Other Components	58	-92	1	-33
Total	65	156	405	626

Improved customer service

While it was still too early to know how customer service will change as a result of implementing this framework, I believe that it will probably dip before stabilizing at our stated customer service level target. This is due to the fact that by taking excess inventory out of the system, process inefficiencies are exposed. This is desirable as it forces the manufacturing system to deal with problems that have been hidden in the past.

Improved inventory control

Adopting a replenishment production system throughout the supply chain will reduce or eliminate oscillations due to forecast error. While production will vary as orders vary, these variations are smaller than the actual production variations experienced by building to the inaccurate forecasts that the imager supply chain has been using.

Simplified management processes

This base stock model and replenishment system framework simplifies inventory management for a number of reasons. First, inventory targets across the supply chain are easily calculated and understood. By using parameters of supply and demand these models create more accurate targets while requiring less effort and direct product experience. These targets are set to meet explicit customer service goals. Second, deciding what to build becomes a simple function of demand. There will always be exceptions, such as quality problems, excessive demand, and parts shortages. Even so, the scheduling process under the base stock framework is much more straightforward than the current system of negotiating desired build vs. forecasted schedule.

When exceptions occur, they will be more visible than if there were excess stock, and they can be more effectively prevented in the future.

Global rather than local optimization

When each department sets inventory targets in a consistent manner and considers downstream inventory and customer demand when scheduling production, inventory can be globally optimized. Previously there were occurrences at Kodak in which one department would stop ordering from the upstream department in order to meet a local inventory goal. This cutback in orders would then be perceived as a drop in end-customer demand. The effects of one department seeking a local optimum are inventory and production oscillations that would be amplified as they travel up the supply chain.

Identification and valuation of improvement opportunities

As described in section 2.5, using the framework helped to highlight the relative inventory reduction opportunities by quantifying potential inventory savings associated with reducing lead time or review period. Generally speaking, reducing the review period and manufacturing lead time will help reduce both pipeline stock and safety stock. Also, in the case of the imager fab (IF) we found that by reducing yield variation, we could reduce inventory (since we need less safety stock to protect against this uncertainty).

3 Conclusion

Summary

By using appropriate models that consider sources of variation, time requirements, and customer service goals, we can set inventory targets throughout a supply chain which are appropriate for meeting customer demand. By implementing the necessary supporting operating procedures, we can improve our performance relative to these inventory and customer service targets. This process of modeling and implementation can help to uncover and quantify further improvements. At Kodak, we were able to identify models that are applicable to the imager supply chain. By using these models we were able to quantify inventory levels appropriate to meeting customer service goals. At Kodak prior to my research, goals were set on an ad hoc basis for finished goods levels in the distribution center. As a result, the actual service levels varied for different products, and finished goods levels tended to oscillate since the inventory targets did not explicitly include work in process. By applying the base stock model, two goals were achieved. First, a mathematically based approach to setting inventory targets simplified the process of setting the inventory targets. By standardizing this process, it became easier for a planner without years of experience with these products to set appropriate inventory targets. Also, production scheduling became a replenishment-based system which smoothed production, simplified the scheduling process, and helped in achieving the inventory targets. While we started to make the changes that would allow the supply chain to meet these new inventory targets, much remains to be done.

Recommendations

Based on my work and research, I have made the following recommendations to Kodak.

Imager Supply Chain Recommendations

Establish Supply Chain Champion to Drive Process Development

Implementing this framework is an ongoing project. To be successful, it should be viewed as an opportunity to develop, document, and improve the underlying parts procurement, production scheduling, and inventory goal setting processes. For this to happen, I feel that we must have someone own this at the supply chain level. This "champion" would be able to pull in the appropriate resources from IF, IS, and FCA, and provide alignment and focus to the group's efforts. This person could then help propagate this framework to the rest of the camera supply chain (e.g. international finished goods inventory, circuit boards and other components).

Integrate Imager Subassembly (IS) into Imager Fab (IF)

Because the SIPModel does not account for yield and capacity variation, I used it to perform a first-pass analysis only. Later, when we considered the effects on quality of combining the subassembly operation (IS) with the fab (IF), it became clear that there was much to be gained by combining these two departments. For example, by integrating the subassembly function into the wafer fab department, yield fallout feedback improves, and duplicate tests can be eliminated. By integrating the IS subassembly operation with the imager fab (IF), we would move closer to the optimal configuration determined by the SIPModel. The improvement in feedback of quality data between IS and IF could significantly reduce the yield variation, leading to a decrease in both scrap and safety stock inventory.

Organizational Recommendations

The organizational structure at the time of my research was more weighted towards functional organizations than product flow. By changing the organizational structure to focus on product flow, I believe that interfaces and relationships between people critical to meeting customer demand will be clearer and more effective. For example, rather than having strong reporting hierarchies of purchasing, supply chain, and production functions, those who these resources might more effectively report into a flow manager.

Final Camera Assembly Recommendations

Process Development

Within each of the departments, processes should be continually refined to support the production and inventory framework. The main components of the monthly process that I proposed for final camera assembly (FCA) to maintain its inventory and production modeling are listed below.

1. Characterize demand and forecast accuracy
    - Calculate current products' demand averages and standard deviations with and without big orders
    - Compare average demand with forecast, and parts schedule
    - Investigate significant discrepancies (e.g. demand = 20, but forecast & parts = 10)

    - Adjust demand parameters (μ, σ), forecasts, and parts schedules as necessary

2. Set Base Stock Levels
    - Calculate base stock levels for continuing products
    - Revise parts schedules as necessary to establish/change base stock levels
3. Set Anticipation Stock Levels
    - Identify max demand for large orders over total lead time (use 6 to 12 months of data, if possible)
    - Create proposal for anticipation stock levels and locations
    - Propose anticipation stock strategy to business unit, revise and get buy-in
    - Revise parts schedules as necessary to establish/change anticipation stock levels
4. Set Parts Schedules
    - Adjust safety stock levels based on revised forecasts and anticipation stock levels set in production planning meeting
    - Adjust schedules based on production planning meeting
5. Set Production Capacity
    - Adjust capacity levels based on revised forecast

Final Camera Assembly: Visual Inventory System

One of the primary barriers to fully implementing the framework at FCA was the way in which parts procurement affected production scheduling. As mentioned earlier, the purchasing department typically controlled the production scheduling meetings since the MRP schedule was the de facto production schedule. Any change to this MRP schedule due to differences between forecast and demand resulted in a negotiation over whether parts would be available, or what to do with excess parts. An experienced manager pointed out to me the wisdom of having a visual inventory and production system. In fact, without such a system, the implementation of my work would be seriously limited. What is needed is a simple system in which it is visually clear to purchasing and production personnel whether there is sufficient inventory to meet demand. The Toyota Production System is an example of this type of a go/no-go system. Kodak has also had success in implementing such systems in its Digital Systems and Scanners production. While I did not recognize the importance of implementing such a system in the camera supply chain soon enough to help implement it, I recommend that this be a priority for the final camera assembly department.

Endnotes

---

[1] Senge p.27-54

[2] Developed by Stephen C. Graves, Sean Willems, and John Ruark. The SIPModel is available at http://web.mit.edu/lfmrg3/www/

[3] Developed by Stephen C. Graves and Sean Willems

Bibliography

DeWitt, Charles, Internal Kodak Publication and Personal Correspondence

Gilpin, Brian C. LFM Master's Thesis: Management of the Supply Chain in a Rapid Product Development Environment. MIT Department of Electrical Engineering and Computer Science, MIT Sloan School of Management 1995

Graves, Stephen C. "Safety Stocks in Manufacturing Systems" Journal of Manufacturing and Operations Management, 1988, Vol. 1, No. 1, pp. 67-101.

Graves, Stephen C. and Sean P. Willems "Optimizing Strategic Safety Stock Placement in Supply Chains" Working Paper, January 1998

Hogg, R. V. and J. Ledolter Applied Statistics for Engineers and Physical Scientists, New York: Macmillan Publishing Company, 1992.

Lee, Hau L. and Corey Billington "Managing Supply Chain Inventory: Pitfalls and Opportunities" Sloan Management Review, Spring 1992.

Miller, Michael P. LFM Masters Thesis: Business System Improvements Through Recognition of Process Variation. MIT Sloan School of Management, MIT Department of Chemical Engineering 1997

Nahmias, S. Production and Operations Analysis, Burr Ridge, IL: Richard D. Irwin, Inc., 1993.

Senge, Peter M. The Fifth Discipline, New York, NY: Currency Doubleday, 1990