

DEVELOPING STRATEGIES FOR SYSTEM ASSEMBLY, FLEXIBLE
LABOR, AND INVENTORY IN THE ELECTRONIC MANUFACTURING
SERVICES INDUSTRY

by

Timothy Brian Frederick

B.S. Mechanical Engineering, University of California, San Diego, 1995

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

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

in Conjunction with the
Leaders for Manufacturing Program
at the
Massachusetts Institute of Technology
June 2000

© 2000 Massachusetts Institute of Technology
All Rights Reserved

Signature of Author _____

Department of Mechanical Engineering
Sloan School of Management
May 5, 2000

Certified by _____

Dr. Stanley B. Gershwin, Thesis Advisor
Department of Mechanical Engineering

Certified by _____


Professor Lawrence M. Wein, Thesis Advisor
Sloan School of Management

Accepted by _____

Ain Sonin, Chairman, Departmental Committee on Graduate Studies
Department of Mechanical Engineering

Accepted by _____

Margaret Andrews, Director of Master's Program
Sloan School of Management



Developing Strategies for System Assembly, Flexible Labor, and Inventory in the Electronic Manufacturing Services Industry

by

Timothy Brian Frederick

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

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

Abstract

Contract Electronic Manufacturing is an industry that has dramatically reshaped manufacturing in the past ten years by proving that it can manufacture and distribute products better than the companies that design an market those products. Toronto-based Celestica, Inc. competes in this industry, and like its competitors, continually searches for new methods of reducing costs and increasing customer service.

The research described in this thesis follows the challenges faced by Celestica's Exeter, New Hampshire site during the second half of 1999. It falls into three topics: assembly system design, workforce composition and the use of flexible labor, and a production system design for a high-mix computer accessory packing area.

With every new product, Celestica struggles with the decision of whether to assemble the system *progressively*, as on a traditional assembly line, or at a *single-station*, with one assembler performing all operations. A framework is presented to help guide the decision based on product- and factory-specific criteria.

Like many labor-intensive operations, Celestica has a dual workforce consisting partly of regular permanent employees and partly of flexible labor hired on a temporary basis. A single-period quantitative model is developed to provide an optimal workforce size and ratio of permanent to temporary employees in the face of stochastic demand.

Finally, a production system is designed and implemented for a pick-and-pack assembly environment. A hybrid build-to-stock and build-to-order production system is developed which simultaneously lowers inventory, improves customer service, and improves productivity by buffering demand variation.

Thesis Advisors

Dr. Stanley B. Gershwin, Department of Mechanical Engineering
Professor Lawrence M. Wein, Sloan School of Management

TABLE OF CONTENTS

Table of Contents	5
List of figures and tables	6
Acknowledgments	7
Chapter 1	9
Introduction to Celestica and Electronic Manufacturing Services	9
1.1 Site History	10
1.2 Site Activities.....	10
1.3 The Contract Manufacturing Industry.....	12
1.4 This Thesis.....	12
Chapter 2	15
Deciding Between Progressive and Single Station Assembly	15
2.1 Background and Problem Definition.....	15
2.2 Electronic Systems Assembly Described	16
2.3 Decision Criteria.....	18
2.4 Hybrid Assembly Systems.....	22
2.5 Conclusion and Suggestions for Future Work	23
Chapter 3	25
Workforce Composition Optimization Model	25
3.1 Background and Problem Definition.....	25
3.2 Cost Optimization Model	27
3.3 Cost Parameters.....	30
3.4 Example: Celestica.....	32
3.5 Potential Cost Savings.....	33
3.6 Intelligent Use of Flexible Labor	36
Chapter 4	37
Design and Implementation of a production system for software and localization kitting .	37
4.1 Background and Problem Definition.....	37
4.2 Existing Planning, Production, and Inventory Management Processes	41
4.3 Proposed Planning, Production, and Inventory Management Processes	44
4.4 Detailed Implementation and Problems Encountered	51
4.5 Realized Benefits	59
4.6 Conclusion and Recommendations	60
References	63
Appendix A: Workforce Model Plot	65
Appendix B: Software Replenishment Ticket	67

LIST OF FIGURES AND TABLES

Figure 1: Progressive and Single Station Assembly.....	15
Figure 2: Hybrid production system for late-stage customization	22
Figure 3: Workforce Model Plotted.....	30
Figure 4: A Possible Workforce Relative Maintenance Cost Difference Diagram.....	31
Figure 5: Monte Carlo cost analysis results for various workforce compositions	34
Figure 6: Celestica's optimal workforce ratio as a function of the standard devaiaton of demand.	35
Figure 7: Software and Localization Demand History.....	38
Figure 8: Shipment histogram	39
Figure 9: Cumulative orders by SKU.....	46
Figure 10: Roles and Responsibilities in a 4-person production system.	50
Figure 11: Existing Software and Localization Line Layout	52
Figure 12: Redesigned Software and Localization Line Layout	54
Table 1: Top Five EMS providers, by sales.....	12
Table 2: Typical component list for a computer.....	16
Table 3: Common products for EMS providers.....	17
Table 4: Example of Finished Goods Inventory System: Three SKU's.....	56

ACKNOWLEDGMENTS

This research would not have been possible without the aid of a great many people at Celestica New England. Thanks specifically to:

Chuck Johnston for supervision and advice, as verbose as it might be.

Toni Chaput for friendship, energy, and insight into Celestica and everything else.

Pat Mongold and Sonja Jauch for everything.

Meg Dunne for proving that Dunkin' Donuts' coffee is a complete breakfast. And lunch.

Chris Reddy, who will be a great inventory manager when he stops buying stuff.

Lanny Meade for showing me how to get things done.

Bill Macanirlan for gambling tips.

Stanley Gershwin and Lawrence Wein provided insight and guidance in their roles as thesis advisors. Their time and efforts are much appreciated.

Steven Graves was gracious enough to lend his time to correct my work while I developed the model in Chapter 3.

The Fellows of the Leaders for Manufacturing Class of 2000 have comprised the greatest collection of teachers I have ever encountered, and I am honored to have spent this short time with them.

My deepest thanks to Mike and Lynn Frederick for their support and for allowing a small bit of their work ethic and devotion to learning rub off on me somewhere along the line. Thanks to Kevin Frederick for showing me what weekends are for.

Simone Miller deserves my undying thanks for providing love and support although she was 3000 miles away and had her own education to worry about.

The author also wishes to thank the Leaders for Manufacturing Program for support of this research.

Chapter 1

INTRODUCTION TO CELESTICA AND ELECTRONIC MANUFACTURING SERVICES

“We expect to do \$10 billion in 2001- not run rate - full year. And we're on track.”

*Eugene V. Polistuk
President and Chief Executive Officer, Celestica*

Celestica is very old for a “new” company. After eighty years as the manufacturing arm of IBM Canada, it was incorporated in 1994 as a wholly owned subsidiary. In 1996, the Onex Corporation purchased the company and began an aggressive expansion fueled mainly by acquisitions. These acquisitions, 17 since 1997, are either other contract manufacturers or divestitures from original equipment manufacturers as Celestica’s customers shed their manufacturing operations. One of the first properties Celestica had its eye on was Hewlett-Packard’s Exeter, NH systems assembly operation. In June of 1997, the Exeter site became Celestica New England in a purchase price of \$187.5 million.

Celestica sees contract electronic manufacturing (CEM) or electronic manufacturing services (EMS) as everything except marketing and basic research. Celestica is involved in design, prototyping, manufacturing and assembly, test, regulatory assurance, distribution, supply chain management, and after-sales support and repair. Its original equipment manufacturer (OEM) customers have made the decision that these services are not core competencies and are best left to the specialists.

In spite of the lofty goal of everything from design to distribution, EMS providers, including Celestica, are best known for two services: printed circuit board assembly (PCB/PCA) and electronic systems assembly (box-build). The Exeter, NH facility is primarily a box-build site due to its long history as a computer assembly factory.

1.1 Site History

The Celestica New England site began its life as one of New England's first high-tech startups, Apollo Computer. Apollo made a name for itself in the early 1980's as a designer and manufacturer of innovative graphics workstations for scientists and engineers. It enjoyed 40% annual revenue growth until 1987, at which point it had secured a 30% market share against competitors like Hewlett-Packard, IBM, and Sun Microsystems.

The tide turned in 1988, when the competition moved to standardize on UNIX-based operating systems while Apollo stuck with its proprietary Domain operating system. During that year, Apollo's market share fell to 13.5% as it lost sales to Sun and a new high-performance, low-cost competitor called Silicon Graphics, Inc. HP was determined to buy market leadership in graphics workstations, and swallowed Apollo in April 1989 for \$476 million.

The 1990's saw HP come to the realization that it was not a manufacturing company, but a R&D and marketing powerhouse. In everything from printers to PC's, HP became a leading user of contract manufacturing, often by selling its manufacturing operations to CEM's. In 1997 it was decided that the Exeter plant was next to go.

A large amount of autonomy was given to the employees of the Exeter site in the selection of their new employer. Some of the suitors offered to buy the plant and leave it to compete with other plants in the same company for business. Celestica was chosen because it presented a vision of a united company leveraging the advantages of a global scope.

Since acquisition by Celestica, HP Workstation Systems Division remains the top customer in sales. The site, however, has expanded its customer scope to include other HP divisions and other computing, telecommunications, networking, and storage systems customers.

1.2 Site Activities

Celestica New England provides four main services to customers in the Exeter site: electronic system assembly and test, consolidation and distribution, return asset management, and printed circuit board assembly.

1.2.1 Electronic System Assembly and Test

Exeter's emphasis since Apollo has always been box-build. This can be described as integrating electronic components inside a chassis. While the desktop computer is the mainstay of this business, the same description can be applied to almost any electronic device destined for the end user. Chapter 2 describes the assembly processes in more depth.

After assembly, most products go through an involved test to insure correct assembly and material quality. Test can take anywhere from only a few minutes, for storage devices, up to multiple days for high-end telecommunications equipment. Test times for computers range from 1 to 24 hours, and involve software-intensive, automated test systems that run a battery of tests and load software on to the computer's hard drive.

1.2.2 Consolidation and Distribution

Most of Celestica's products are shipped directly to end-users. Often, other items manufactured elsewhere accompany the product. Celestica will stock items like manuals, keyboards, monitors, software, and accessories in its distribution warehouse. These items will be consolidated with the manufactured product and sent in the same shipment to the end-user. Chapter 4 documents an improvement effort in one of these distribution areas.

1.2.3 Return Asset Management

For some customers, Celestica operates a program for the refurbishment or value recovery of returned merchandise. Most of the volume in this area are products manufactured at Celestica which have been damaged in shipment, sat unsold past obsolescence, or returned from third parties unused. Celestica will either refurbish for sale or break down the returned products for parts or scrap.

1.2.4 Printed Circuit Board Assembly

Celestica also has a small shop for the assembly of components onto printed circuit boards. The area is still rather new, but it is hoped that it will support systems assembly through the on-site production of PCB's.

1.3 The Contract Manufacturing Industry

In 1999, the global EMS market is estimated at \$90 billion and projected to grow at 25% annually. The top five contract manufacturers control only 31% of this market. The competition for contracts leads to a competition in size. Contract manufacturers fear a coming day when they are evaluated on the scale of their operations. Potential customers want to pick the manufacturing partner that will best help them reach global markets instantly while enjoying huge economies of scale in purchasing and supply chain and distribution logistics.

Table 1: Top Five EMS providers, by sales

<i>Company</i>	<i>1999 Sales (\$ billions)</i>
Solectron	9.8
SCI Systems	7.2
Celestica	5.6
Flextronics International	3.3
Jabil Circuit	2.3

The key metric for contract manufacturers is operating margin. Best in class for high-volume manufacturing is above 5% operating margin, with the top five companies ranging in the 2.2-5.4% range. Getting above 5% requires maximizing utilization of resources. Generating profits in contract manufacturing requires that every square foot of every factory be adding value 24 hours a day every day.

The cost-cutting and efficiency goals can collide with the hunger for growth. New factories are postponed as long as possible to free available capital for acquiring already operating plants. Consequently, production density increases as more product is pushed through a tighter space in already cramped factories.

1.4 This Thesis

My Leaders for Manufacturing internship at Celestica came at a time when many changes were happening at once in Exeter. The factory was being squeezed into a tighter space as they were being forced out of the Exeter site before a greenfield plant in nearby Portsmouth was ready for occupancy. The conversion to a contract manufacturing operating structure had not completely taken effect, and management was searching for ways to increase productivity and

accept new customers at the same time.

My research followed the needs of the company during my 6½-month tenure. I started by trying to help bring sense to the decision on production system design through a framework I present in Chapter 2. During that investigation, I found that an important input into the framework was workforce composition. Chapter 1 describes a quantitative model to help guide the strategy around the ratio of permanent to temporary employees in the factory. Halfway through the internship I was asked to lead an improvement effort in the Software and Localization Kit Assembly area. That initiative is documented in Chapter 4.

It is my sincere hope that the information I present here will be utilized by Celestica. Many of the recommendations I make are general enough to be useful for a wide range of problems both in Celestica and in industries outside contract manufacturing.

DECIDING BETWEEN PROGRESSIVE AND SINGLE STATION ASSEMBLY

2.1 Background and Problem Definition

Assembly lines at Celestica are of two forms, progressive or single station. A progressive line is the image that the mind conjures when the words “assembly line” are heard: a sequential series of operations performed by different assemblers (people or machines), in which every assembler adds value to every product. A single station line produces the same product, but one assembler will complete the entire operation at a stationary workbench.

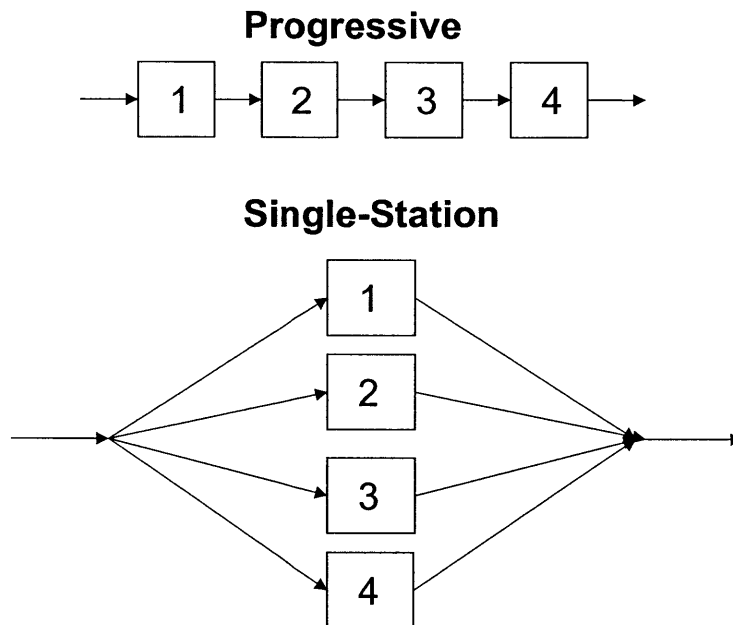


Figure 1: Progressive and Single Station Assembly

For many products, one or two criteria dominate the decision on whether to assemble the project in a parallel or single station manner. On an automobile or airplane, expensive tooling and narrow skills necessitate a progressive process.¹ The other extreme is a disposable pen, which, while most likely automated, is far too simple an assembly to be split between several

¹ Although there have been modern exceptions, like Volvo Car’s Udevalla plant (Womack, Jones, & Roos, 1991).

manual assemblers.

Electronic systems assembly falls somewhere in the middle, and the choice becomes unclear. Most mid- to high-volume electronic products have been designed to be simple enough for one person to assemble with a moderate amount of training. The same product often has a level of complexity that makes it feasible for several workers to progressively build it as well. The choice of assembly system has far reaching effects on areas ranging from inventory cost to customer service.

Celestica is in the business of quickly introducing new products to the factory and executing production in the most effective way possible. Consequently, the question of progressive versus single station assembly needs to be answered repeatedly for a variety of products. This chapter is devoted to providing some structure around the decision. The topic is examined against several criteria of importance to various constituencies affected by the decision.

2.2 Electronic Systems Assembly Described

Electronic systems assembly, often referred to as “Box Build,” is the process of combining electronic components inside a mechanical chassis and testing the assembled system. A common example is a personal computer or graphics workstation. Common PC components are enumerated in Table 2.

Table 2: Typical component list for a computer

<i>Component</i>	<i>Description</i>	<i>Quantity</i>
Chassis	Several sheet metal or injection molded plastic components	1
Motherboard(s)	Large printed circuit board (PCB)	1-2
Memory	Small PCB's for RAM	0-8
Adapter cards	Add-on PCB's to provide graphics, I/O, networking, etc.	0-5
Storage	Non-volatile devices such as hard drives, tape, CD/DVD, removable, etc.	0-5
Power Supply	Self-contained, converts line AC power to DC power for components	1-2
Cooling	Fans to ventilate chassis or provide directed air flow to cool specific components	0-4

Assembly of such a computer would begin with the chassis, then the addition of each component. The computer is typical in that the work is performed in a small space, enabling

only one person at a time to perform assembly. Most components are secured with screws driven either by hand or by powered screwdrivers, although recent designs have shown the benefits of design-for-manufacturing programs and eliminate screws where possible.

A personal computer is considered a worst-case scenario for product variation. Most end users have specific needs and demand custom-configured machines consisting of a collection of components chosen from a long menu of possibilities. In addition, rapid technology change causes constant component and design changes. Table 3 describes some other electronic systems commonly assembled by electronic manufacturing service providers.

Table 3: Common products for EMS providers

<i>Product</i>	<i>Key Characteristics</i>
Computers and storage systems	<ul style="list-style-type: none"> • Highly configurable • Long test (hours)
Telecommunications Equipment	<ul style="list-style-type: none"> • Large chassis • Long, complicated test (hours to days)
Storage devices	<ul style="list-style-type: none"> • Small form factor • Short test (minutes)
Automated Teller Machines (ATM's)	<ul style="list-style-type: none"> • Large form factor • Extensive mechanical assembly
Printers	<ul style="list-style-type: none"> • Extensive tooling requirements • Extensive mechanical assembly
Networking Equipment	<ul style="list-style-type: none"> • Small to medium form factor • Short test (minutes)


Each of the products in Table 3 can be generally described as some form of electronics assembled in a chassis, sometimes with some mechanical assembly required as well. In that definition, however, there is a large range of key characteristics which can drive the choice for assembly system. Eight factors were identified as the most important in the assembly system decision. In the following section these criteria are enumerated and described.

For the personal computer example, fixing the work consists of a set of rules like “Station 2 installs hard drives.” Fixing the work in this way with high product variability makes predicting the behavior of a progressive assembly line extremely difficult, especially when there are small or no buffers allowed between stations. In such a tightly-coupled system, throughput suffers tremendously when there is high variation in station cycle times (Gershwin, 2000).

Fixing the station cycle time leads to more predictable line performance (and higher throughput) but at the expense of the gains in training cost described in section 2.3.1. Now each assembler must have the ability to perform not only their operation but that of the station(s) up- and downstream from their own.

A single station system decouples the workers and the work content, absorbing the variation in the system input queue.

2.3.3 Demand Variability



Every enterprise must decide how it plans to manage variation in the demand for its product. For products with highly configurable features, where finished goods inventory is impractical, demand variability is felt on the factory floor. In general, contract manufacturers see their competitive advantage over OEM manufacturing as the ability to provide flexibility with no additional cost. They feel that variation in demand and product can be tempered by the scope of their operations across different customers and industries.

Realizing this flexible vision requires a production system that can quickly scale and adapt to changes. When operating at capacity, a progressive assembly line must add people in order to increase production rate. This implies adding length to the line and redistributing work elements, both of which require stopping production and upsetting the production environment. Choices involving running multiple products down the same progressive line involve making sure that all products are physically compatible with the existing conveyor system.

Scaling single station assembly systems is trivial: add more stations. Since there is no division

of work elements, no procedures need to change. A single station environment provides flexibility bounded only by physical floor space, tooling, material delivery, and training.

2.3.4 Material Delivery




Imagine assembling cars in a single station system. Every body, every door, every engine, every seat set, every different part would need to be delivered, when needed, to multiple points of use throughout the factory. Compared to progressive assembly lines, with one point of use for every part, the material handling costs would be enormous.

In the EMS world, many products can be delivered as a kit to the assembly line. In a “supermarket,” the parts required for a specific configuration are picked into a tote and delivered on a cart or via conveyor to the assembly line or station. This works equally well for progressive or single station assembly, but is an absolute requirement for single station assembly since holding stocks of large items at every workstation causes inventory to balloon.

Progressive assembly lines have an advantage here, because inventory can be easily stored in or delivered to one location on the line.


2.3.5 Rework Rate



Rework, or repair, is carried out when a product fails a test or an inspection. Progressive assembly environments require a separate repair station to perform the diagnostic and repair work. If the problem is due to an assembly error, it is often difficult to enforce a learning loop back to the assembler.

A single station operator can double as his own rework station, especially if a preliminary turn-on test is performed before the product leaves. The learning takes place immediately, and the operator can consult with other nearby experts or send the unit to central engineering when troubleshooting is difficult.


2.3.6 Tooling Costs



By its nature, EMS usually has few tooling requirements. Exceptions lie in areas like optical networking or inkjet printing where alignment and calibration can require specialized and expensive tooling.


If tooling is required during assembly, it must be replicated at every assembly cell. For single station assembly, this cost could be prohibitive. Progressive assembly only requires tooling at the one station it is used.

2.3.7 Engineering Change Rate



Some products, especially during ramp-up, experience a large volume of Engineering Change Orders or Notices (ECO/ECN's). An ECO it is sometimes as simple as a part replacement with no affect on assembly. However, when ECO's require assembly process changes, assemblers must be trained in the new procedure. Single station assembly requires that all assemblers be notified and trained. Progressive assembly contains the changes to one assembly station.

2.3.8 Handling Time Relative to Takt Time



Takt time is defined as the available minutes of production divided by demand. The ratio r relates this station cycle time to the handling time. When r is small, it indicates that the product is difficult to move from station to station, and is probably better left in place for single station build. Large r indicates that material handling times are inconsequential and a progressive line will bear little penalty for frequent moves.

Indeed the history of progressive assembly lines tells us that r was involved in the development of the modern automotive assembly plant. The moving assembly line resulted, in part, from Henry Ford's willingness to invest in its development to bring this ratio down to make progressive assembly possible.

2.4 Hybrid Assembly Systems

An understanding how the eight factors described in section 2.3 affect the production system decision for pure progressive or single station environments can lead to innovative production cells which capitalize on the factory's and products' unique characteristics. I will describe two examples.

2.4.1 Complex product with late-stage customization

When most factors point to progressive assembly, one alternative is to build "vanilla" products on a progressive line, which are then sent to single stations for configuration and assembly verification.




Figure 2: Hybrid production system for late-stage customization

In Figure 2, stations 1-3 build vanilla base systems while stations 4 and 5 do last stage customization. The ability to delay customization is largely dependent on the design of the system being assembled.

Ideally, a mix of new and experienced would staff the progressive stations and the most experienced permanent employees would do final customization and quality check. This allows the benefits of on-line training with product variability and high skill requirements confined to stations 4 and 5.

2.4.2 Complex Highly Configurable Products: Miniature Assembly Cells

Another hybrid possibility when postponement of configuration is not possible is a collection

of mini-cells consisting of 2-3 people working around an assembly table. Every takt, the workers shift one position. The face-to-face contact enhances training, while the small cell size lessens the effect of product variability on the overall production system.

It is also easy to enforce self-balancing line procedures in this environment. When finished with a product, assembler #3 takes over from #2, who in turn takes over from assembler #1. Assembler #1 pulls a new kit into the cell. This method works only when station shifts are easy and the workers are trained to do much of the work in the cell.

2.5 Conclusion and Suggestions for Future Work

The goal of this chapter was to present some objective criteria for evaluating the production system decision. Every reader will interpret the factors with a unique perspective and weighting system based on their concern at the moment. For example, the shop floor supervisor who can not keep enough trained assemblers in the factory will weight product complexity very high. The production engineer, trying to guess what new product will be introduced next will value the flexibility of single station production.

The eight criteria outlined are in terms of vague notions like high/low, small/large, etc. My intention for this project was to provide concrete metrics for the EMS industry that would quantify these terms. As it stands, it is left up to the reader's judgement and experience to decide where any particular product falls in the given spectra. Future work might include developing metrics for the criteria developed here, and guideline benchmarks based on historical experience and studies of other companies.

WORKFORCE COMPOSITION OPTIMIZATION MODEL

A key finding from the investigation described in Chapter 2 was that training cost had a large influence in optimal production system design. In a stable workforce, training is a secondary concern. However, during 1999 Celestica New England was facing the tightest labor market in recent memory. Like most electronic manufacturers, a significant portion of Celestica's workforce is comprised of *flexible labor*: temporary workers who were technically employees of an outside agency. Celestica contracts with the outside agency for the use of the workers, but shoulders the cost of training new "temps".

In a hot labor market, turnover in the temporary ranks was tremendous as workers were easily lured away to permanent employment elsewhere. Some months, turnover could be as high as 30% *per month*. I was approached during one of these months by two second-line managers who complained of the turnover and management costs associated with flexible labor. They asked me if I could come up with some guidelines around the use of flexible labor, and specifically what fraction, if any, of the workforce should be temporary.

3.1 Background and Problem Definition

In highly manual industries, the demand for labor closely follows the demand for product. Companies like Celestica use flexible labor as a buffer against uncertainty in the demand for labor. "Temps" can be a good tool for managers who wish to protect their permanent employees from demand fluctuations.

For as long as anybody at Celestica could remember, the company used a 50% ratio of permanent employees to total workforce for direct labor. Nobody with whom I spoke had a clear picture about why this number was chosen. Some speculated that it was a simple management rule: in times of growth: for every new temp hired, one temp could be converted to permanent. Partly it appeared to have a cultural history. For several years, Celestica New

England was a division of Hewlett-Packard, a company with a long history of avoiding layoffs of permanent employees at all costs.

In any case, when those two managers approached me, I translated their concerns into an indication that *flexible labor is expensive*. This is a non-obvious conclusion, because from the easiest measure, wages, temps are usually paid less per hour than permanent employees, and Celestica was no exception. However, the difference was not significant—less than 19%, fully burdened—and there were costs not included in hourly wage, such as training (proportional to turnover) and supervision costs. Further investigation, the details of which appear in Section 3.4, revealed that when all costs are taken into account, temps do cost Celestica more than permanent employees.

In fact, it is my assertion that the total cost of maintaining a productive temporary employee is always higher than maintaining a permanent employee. If this were not the case, then we would see many companies eliminating permanent employment altogether. The cynical reader will respond that this does not occur only out of fear that the temps will organize, which is certainly a possibility. I would contend that the higher total cost of temporary labor is a financial instrument, an option to discharge the employee cost-free at some future date, or convert them to permanent if warranted by their performance and the business environment. Thus the two primary uses of flexible labor is capacity flexibility and new hire screening.

The literature presents a dearth of quantitative methods for planning the use of flexible labor. Abraham (1988) makes the assumption that the temporary worker is less productive than a permanent worker and expresses the optimal ratio of permanent to total workforce as a function of this difference. The closest are Herer & Harel (1998), who provide a planning algorithm based in newsboy-style inventory control methods, but assume that temporary labor can be called in at a moment's notice when the actual demand is realized. Looking beyond flexible labor in particular, several authors have presented Markov-chain analyses that forecast movements among different job classes within an organization to predict long-range hiring needs (Rowland & Sovereign, 1969 and Hooper & Catalanello, 1981).

My contribution is a single-period optimal staffing model tailored to the particular realities facing Celestica in 1999. The goal of the following model is not to provide a detailed human

resource planning tool to be used on a continual basis to justify hiring and firing. Instead, I only hope to provide an applicable, quantitative framework for understanding trade-offs and guiding high-level policy.

3.2 Cost Optimization Model

Looking ahead to a period of stochastic demand, a manager needs to estimate the labor required. At Celestica, the direct labor can be separated into the two groups as described above, permanent and temporary. The question the manager needs to answer is “How many of each type of worker should I hire in the face of uncertain demand?” Let L_1 and L_2 represent the number of permanent and temporary workers, respectively, that will be hired to satisfy a near-future demand period. The number of total hires is then $L_1 + L_2$. The length of the period to be modeled is left to the user’s discretion, but one month is used in the example in Section 3.4.

If actual labor demand x , measured in worker-periods, falls within L_1 and $L_1 + L_2$, then the company is liable for an *overage* (excess) of $L_1 + L_2 - x$ temporary workers at a cost C_{ot} per employee. If x falls below L_1 , all temporary workers are subject to C_{ot} and $L_1 - x$ permanent workers set up an overage cost of C_{op} per employee. If x falls above $L_1 + L_2$, an *underrage* (shortage) cost C_u is incurred. An overage cost might be interpreted as a firing or lay-off cost, in which case the assumption is made that all temporary employees will be discharged before the first permanent worker is laid off. An alternate interpretation of overage costs that does not involve releasing workers is described in Section 3.3.2.

The other assumption made in the development of the model is that decisions about temporary labor levels must be made in advance of the demand period in question. That is, once the demand period arrives, and the actual demand is known, it is too late to hire even temporary workers. There were two reasons this assumption was valid for Celestica in 1999. First, the tight labor market made it difficult to fill temporary positions in fewer than 3 weeks. Second, a threshold level of training was required before contributing to production. During the training period, each temporary worker is paired with a more experienced worker and adds no marginal production capability.

Let C_{mp} and C_{mt} denote the cost of maintaining one permanent and one temporary employee for the period in question. Then the cost of the workforce over the period for demand x is:

$$C(x) = \begin{cases} C_{op}(L_1 - x) + C_{mp}x + C_{ot}L_2 & 0 \leq x < L_1 \\ C_{mt}(x - L_1) + C_{ot}(L_1 + L_2 - x) + C_{mp}L_1 & L_1 \leq x < L_1 + L_2 \\ C_u[x - (L_1 + L_2)] + C_{mp}L_1 + C_{mt}L_2 & L_1 + L_2 \leq x < \infty \end{cases} \quad (3.1)$$

If the distribution of possible demands has probability density function $f(x)$, then the expected cost of the workforce Q is:

$$Q(L_1, L_2) = \int C(x)f(x)dx$$

or substituting Equation (3.1) for $C(x)$,

$$\begin{aligned} Q(L_1, L_2) &= \int_0^{L_1} [C_{op}(L_1 - x) + C_{mp}x + C_{ot}L_2]f(x)dx \\ &+ \int_{L_1}^{L_1+L_2} [C_{mt}(x - L_1) + C_{ot}(L_1 + L_2 - x) + C_{mp}L_1]f(x)dx \\ &+ \int_{L_1+L_2}^{\infty} [C_u[x - (L_1 + L_2)] + C_{mp}L_1 + C_{mt}L_2]f(x)dx \end{aligned} \quad (3.2)$$

The first order conditions for a minimum are $\frac{\partial Q}{\partial L_1} = \frac{\partial Q}{\partial L_2} = 0$, which yield:

$$\begin{aligned} \frac{\partial Q}{\partial L_1} &= C_{op}F(L_1) + C_{mp}[1 - F(L_1)] + C_{mt}[F(L_1) - F(L_1 + L_2)] \\ &+ C_{ot}[F(L_1 + L_2) - F(L_1)] + C_u[F(L_1 + L_2) - 1] = 0 \\ \frac{\partial Q}{\partial L_2} &= C_{mt}[1 - F(L_1 + L_2)] + C_{ot}F(L_1 + L_2) + C_u[F(L_1 + L_2) - 1] = 0 \end{aligned} \quad (3.3)$$

in which $F(x)$ is the cumulative density function of demand. Solving Equations (3.3) for $F(L_1)$ and $F(L_1 + L_2)$ present the compact relations

$$F(L_1 + L_2) = \frac{C_u - C_{mt}}{C_{ot} + C_u - C_{mt}} \quad (3.4)$$

and

$$F(L_1) = \frac{C_{mt} - C_{mp}}{C_{mt} + C_{op} - C_{mp} - C_{ot}}. \quad (3.5)$$

Two observations can be earned from inspection of Equations (3.4) and (3.5). Since probability dictates that $F(x)$ is bounded by 0 and 1, setting the numerator of Equation (3.5) greater than the denominator yields $C_{mt} > C_{mp}$. This follows from the previous assumption that temporary workers are more expensive to maintain than permanent workers. The same operation to Equation (3.4) reduces to $C_{ot} > 0$, or that there must be a non-zero overage cost for temporary workers for the model to be valid. This inequality makes sense to the mathematician, who will claim that of course this the true, otherwise there is nothing to prevent the company from asking several hundred extra temps to show up every day, “just in case.” Actually, there are some examples where excess labor is routinely brought in just in case, such as union hiring halls in large cities. Yet I will contend that most companies do incur overage costs for temporary employees. The most obvious in this example is the training and wage costs incurred during the time between when the temp starts work and when the demand is realized. Other costs might include the space, uniforms, and security for idle workers.

Solving Equations (3.4) and (3.5) for L_1 and L_2 involves (A) choosing representative cost parameters, and (B) choosing an appropriate distribution $f(x)$ for the labor demand. Section 3.3 discusses choosing cost parameters, and section 3.4 provides an example based on Celestica’s costs with a normally distributed labor demand.

For normally distributed demand with $\sigma = 0.1\mu$ and $C_u = 3C_{mp}$, the model can be visualized as in Figure 3. Figure 3 plots the optimal permanent fraction of the workforce as a function of the ratio of C_{mp} to C_{mt} for various relative values of the overage costs. The plot shows optimal perm to total workforce ratio falls as the maintenance costs approach parity, rapidly decreasing slope above $C_{mp}/C_{mt}=80\%$. Also, increasing C_{ot}/C_{op} shifts the curve vertically upwards, signifying that if it costs as much to have extra temps as extra perms, there is no reason to hire temps. This behavior agrees with intuition, which tells us that inexpensive temps, in both wages and dismissal costs, will cause us to use them more. A larger version of this plot appears in Appendix A.

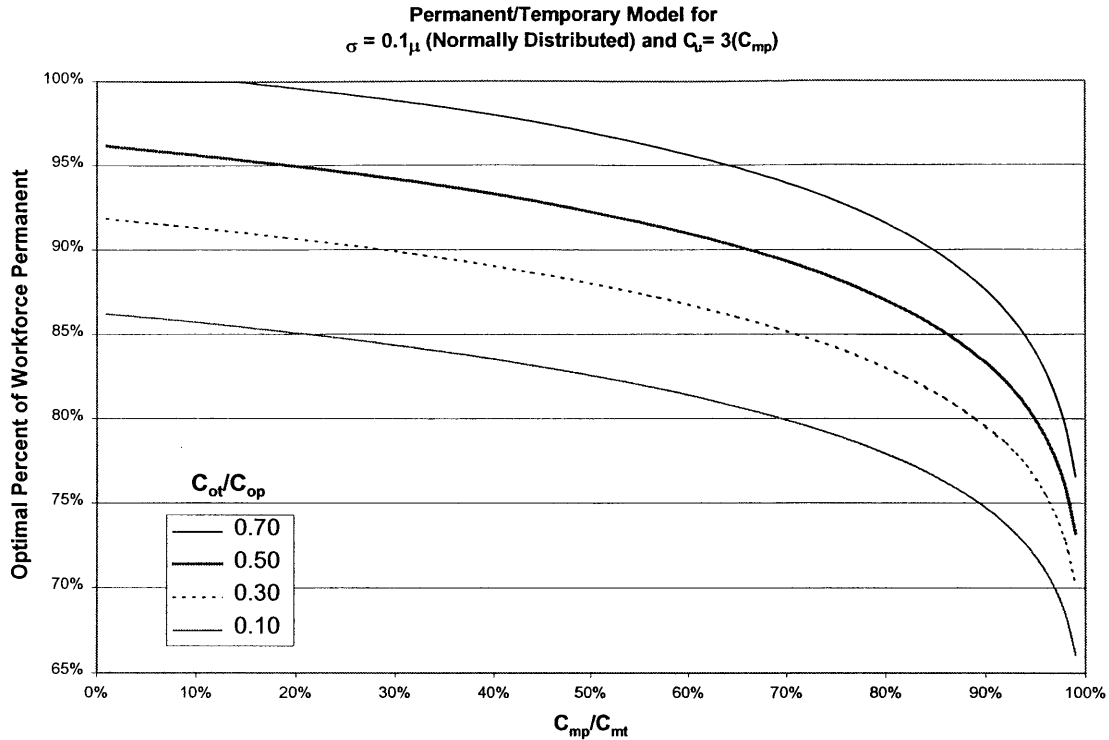


Figure 3: Workforce Model Plotted

3.3 Cost Parameters

The various costs of the labor force must be estimated for Equations (3.4) and (3.5) to be used. The costs can be divided into three groups: maintenance (variable) costs, overage costs, and underage cost.

3.3.1 Maintenance Costs

The parameters C_{mp} and C_{mt} refer to the cost of maintaining one productive permanent or temporary employee. The largest part of this cost will be wages, including all benefits and other variable costs.

However, as stated before, wages alone do not measure the cost of maintaining an employee. Another sizeable difference between permanent and temporary workers is training cost as a function of turnover. For example, if the turnover rate is r_t per period and the cost of training a worker is C_t , this is reflected as an additional $r_t C_t$ incurred per period per worker. Celestica found that while it cost the same to train both types of workers, the turnover in the temporary ranks was dramatically higher than that in the permanent ranks.

Other cost differences are more difficult to quantify. There are management costs associated with the extra effort required to communicate with an outside agency for temporary workers. Supervision costs can be much greater due to a real or perceived increased likelihood of shirking or theft. Figure 4 displays a possible relative cost structure for the components of maintenance cost.

Additionally, a multiplying factor may be applied to the cost of temporary workers to represent a productivity difference relative to permanent workers as described by Abraham (1988).

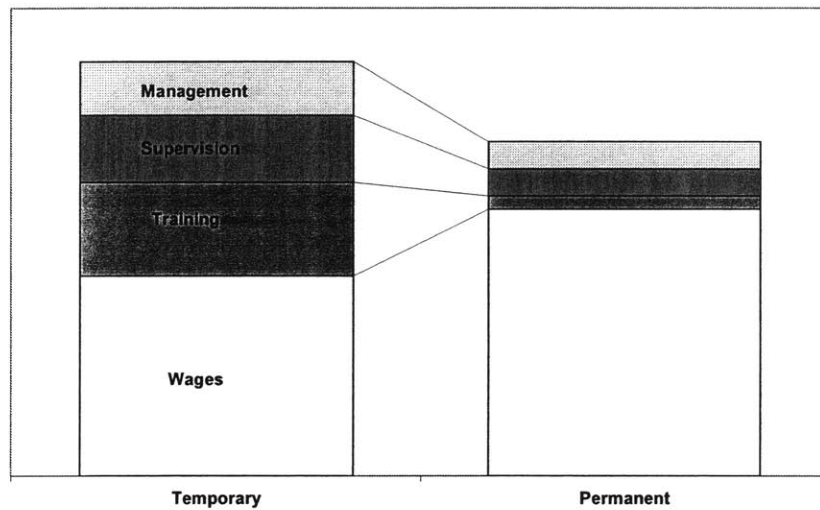


Figure 4: A Possible Workforce Relative Maintenance Cost Difference Diagram. Wages alone do not tell the whole story.

3.3.2 Overage Costs

For the firm that, at the realization of actual demand, lays off all of the extra workforce the overage costs C_{op} and C_{ot} , describe severance costs. A good approximation for C_{op} might be two weeks of pay plus whatever other severance benefits due a permanent employee. For C_{ot} , the conventional wisdom is that it is smaller than C_{op} , but, as previously argued, not zero. For both workers, there is most likely a time lag between the beginning of the demand period and the decision for a lay-off. In this case the overage penalty is the pay for the employee over this time period.

The firm that chooses not to lay off workers, or release only a fraction, may prefer alternate definitions for C_{np} and C_{nr} . Instead of total costs, these can be per worker profit contributions

with a negative cost. The overage penalties C_{op} and C_{ot} could then be considered the costs of keeping a worker on the payroll with no work to be done. In this case, we would expect direct wages to dominate the calculation of overage costs ($C_{op} > C_{ot}$) since there is no supervision, no management, and workers who leave on their own volition are not replaced. There might also be a *salvage value*, involved when a worker is diverted to alternate tasks that generate less profit contribution.

3.3.3 Underage Cost

The cost of not meeting demand depends on the business structure. If the customer can easily go elsewhere or not purchase at all, this is the cost of the lost sale. For a company with market power or high switching costs, unmet demand is more likely backordered. The degradation to customer service might be quantified in lost future sales or increased expedited shipping costs.

Another option, depending on the firm, is overtime. The model provides a good method of understanding the trade-off between planned overtime and planned above-mean capacity in the temporary ranks.

3.4 Example: Celestica

Let us apply this model to Celestica. Celestica's fully loaded labor costs are \$13.65/hr for temps and \$16.80/hr for permanent employees. In addition, the monthly turnover rate is approximately 30% for temps and 1% for permanents, for which training costs are \$1,700² per employee. We conservatively estimate that temps require \$0.50/hr more in supervision costs than perms. For a period of one month (160 work hours):

$$C_{mt} = 160(13.65 + 0.50) + 0.30(1700) = \$2774$$

$$C_{mp} = 160(16.80) + 0.01(1700) = \$2705$$

Should actual demand fall below the as yet undetermined optimal workforce level, overage costs will be encountered. At this point, the temps can be released at no cost other than the wages paid for the first week, but permanent employees require 2 weeks severance pay:

² Source: Previous Celestica study.

$$C_{ot} = \frac{1}{4}(C_{mt}) = \$694$$

$$C_{op} = \frac{1}{4}(C_{mp}) + \frac{160}{2}(16.80) = \$2020$$

If demand is not met, all overage is covered through the use of overtime at twice the rate of permanent labor:

$$C_u = 2(C_{mp}) = \$5376$$

Evaluating Equations (3.4) and (3.5) with these costs gives us $F(L_1 + L_2) = 0.790$ and $F(L_1) = 0.049$. Next, we will assume a normally distributed labor demand forecast with a mean of 100 workers and a standard deviation of 10 workers. I make no claim that a normal distribution is the best model for uncertain demand for labor. In fact, phenomena like material and equipment constraints and contract obligations will limit the spread of the distribution. However, for the purposes of this model, a well reasoned mean and standard deviation on *any* distribution will provide the majority of the benefit, and in the absence of any dramatic evidence otherwise, the normal distribution is a good starting point. Using a table or spreadsheet function to calculate the inverse of the normal distribution, we arrive at the values for L_1 and L_2 :

$$L_1 + L_2 = 108.0$$

$$L_1 = 83.5$$

Or, alternately, that the optimal ratio of permanent to total workforce is 77%.

3.5 Potential Cost Savings

It would be valuable to understand the shape of Q , the expected cost for the workforce, as a function of the fraction of permanent workers. For some labor demand distributions $f(x)$, Equation (3.2) can be solved to provide a closed-form expected cost function. However, a more useful tool is a Monte Carlo analysis. A Monte Carlo simulation has the advantages of being easy to implement in a spreadsheet while at the same time producing an estimate of the variance of the expected cost. Additionally, the effects of various possible distribution shapes, including those that are discontinuous, can be quickly studied.

A Monte Carlo simulation with the cost parameters calculated in Section 3.4 tells us that moving from a 50% permanent ratio to the optimal 77% will save an expected \$14,000 (5%) in cost per 100 man-months of demand. However, the standard deviation of the simulated expected cost is \$26,000 (9%). With this magnitude of possible variation in cost, any savings is unlikely to be noticed in the short term. There is even a good chance that the cost would go up simply from the stochastic nature of the demand.

The simulation can be repeated for many values of L_1 and L_2 while holding their sum constant. Using the same cost parameters calculated for Celestica, such an exercise would generate the plot seen in Figure 5. To better visualize the gentle slope to the left of the minimum, the simulation for this plot used the same sequence of 1000 normally distributed random labor demands to calculate each point on the line, producing the smooth lines seen here.

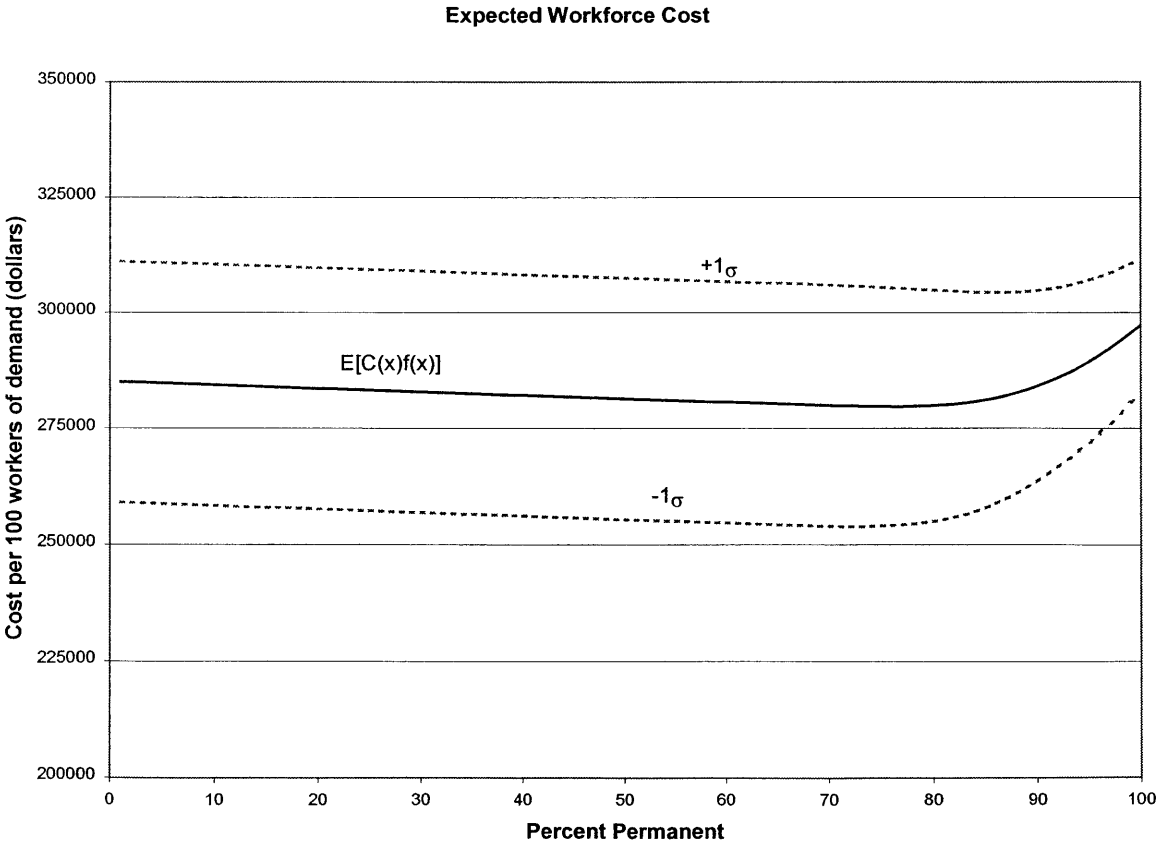


Figure 5: Monte Carlo cost analysis results for various workforce compositions. Cost parameters are as calculated in Section 3.4.

The feature to note Figure 5 is the very moderate slope below the minimum. The slope in this region is dominated by the relative magnitudes of C_{mp} and C_{mt} , which as calculated are almost the same. For larger values of C_{mt}/C_{mp} , this slope would be more dramatic approaching the minimum. At Celestica, it tells us that while there is not an extreme amount of cost savings associated with raising the permanent ratio, there is also no reason why it should not be raised.

To the right of the minimum, the slope increases as more permanent workers are released. This indicates that the standard deviation of the demand has a large effect on the optimal ratio. In fact, the plot in Figure 6 suggests an almost linear relationship for this particular example. A fairly accurate rule of thumb for Celestica might read, "Hire enough permanents to satisfy the expected labor demand less two standard deviations."

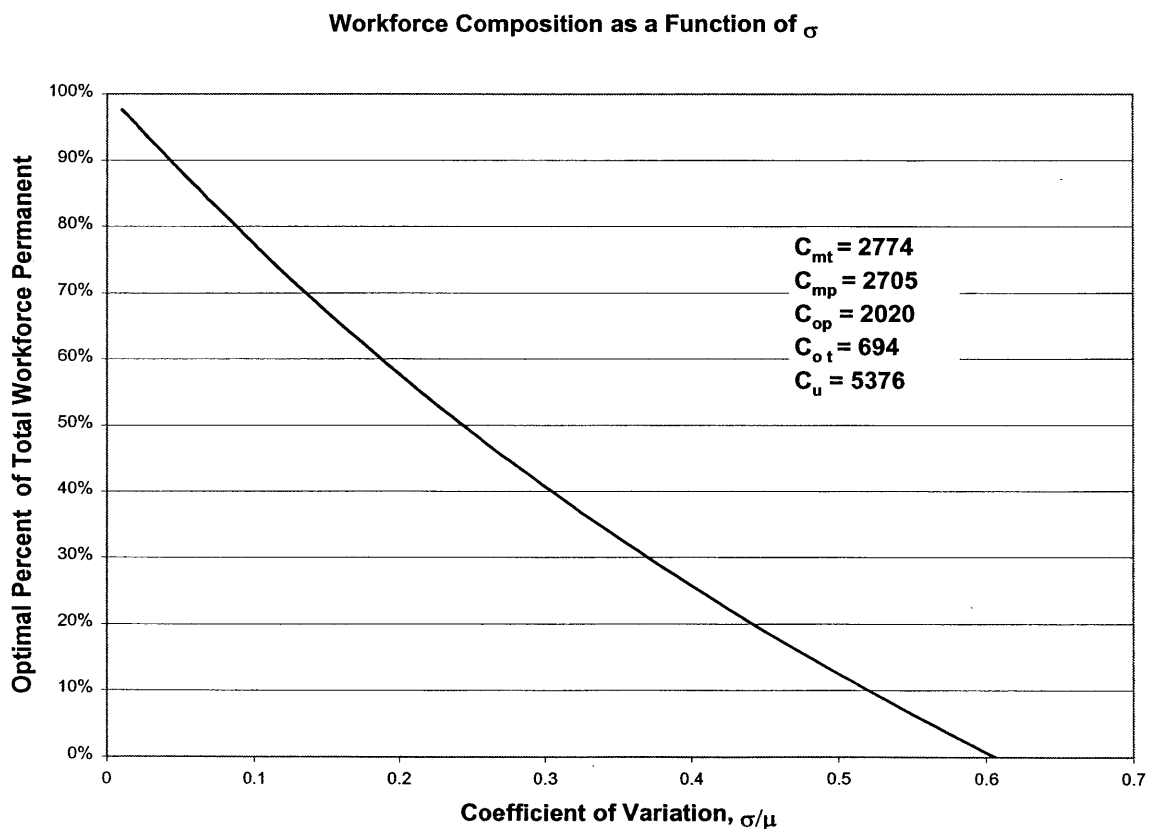


Figure 6: Celestica's optimal workforce ratio as a function of the standard deviation of demand.

3.6 Intelligent Use of Flexible Labor

I will end this chapter with a short opinion on how to best integrate temporary employees into the factory. Often what occurs on the manufacturing shop floor is that permanent employees settle into roles that they are especially good at or enjoy, then the flexible labor fills around them as demand dictates. I claim that the opposite should be encouraged: flexible labor takes on the static roles while the group of highly-trained permanent employees roam to provide capacity flexibility.

The reason for this is that training cost is often under-appreciated as a component of workforce cost. When training is significant, as was shown here, minimizing total labor expenditures is an exercise in reducing the turnover of the trained people. The greatest investment in training should go to the sector of the workforce that turns over the least, the permanent portion. Flexible labor, on the other hand, should be used for jobs for which the training requirement is lowest, and, equivalently, the monetary cost of turnover is lowest.

In any area, when given the choice between cross-training an experienced, permanent person and a temp on a new task, the nod should be given to the permanent person. That is, the risk of the training investment walking out the door or being discharged is much less with the permanent employee. A documentation system to track workforce development is crucial to making the correct choices in these situations, and effective use of such a system can lower the adverse costs of turnover dramatically.

Chapter 4

DESIGN AND IMPLEMENTATION OF A PRODUCTION SYSTEM FOR SOFTWARE AND LOCALIZATION KITTING

Every computer CPU that is assembled at Celestica New England requires an assortment of complementary items before it can be operated by the end user. Examples of these items are keyboards, mice, hardware user manuals, and software media and documentation. This chapter describes an initiative undertaken during the period of October-December 1999 that lowered the cost and improved the performance of the factory area that assembled these items into software and localization kits.

4.1 Background and Problem Definition

One of the services that Celestica provides to its OEM customers is to assemble accessory kits. The kits contain a variety of software media, license contracts, manuals, books, and small hardware items - power cords, keyboards, mice, cables, etc. Celestica handles all the purchasing and inventory management, the assembly of the items into boxes, and the shipment of the completed kits. During late 1999, Celestica was under pressure to improve customer service, reduce space, and reduce costs in this area.

4.1.1 Customers and Customer Expectations

Celestica ships computers and the accompanying software and localization kits to several different customer types. First are end users, who receive and use finished good shipped directly from Celestica. Second are OEM customer sites or other Celestica sites that act as distribution centers for the finished goods. Third are distributors and value added resellers (VAR's) who stock finished goods for retail sale or modification (custom hardware or software enhancements.)

Each of these customers has a different typical order. The end users usually place small orders of quantities in the 1-10 pieces range. The other customers - distributors, OEM's and VAR's

– are more likely to expect larger bulk orders to replenish inventories or service large customers.

This combination of customer types leads to a “lumpy” demand pattern. The demand created by the end users and other small orders is interrupted by spikes of large orders from the OEM, distributors, and VAR’s, as can be seen in Figure 7 on day 31.




Figure 7: Software and Localization Demand History

In spite of this mixed demand picture—or perhaps because of it—the aggregate demand for the kitting area could be demonstrated by a normal distribution, as suggested by Figure 8. For the remainder of the project, the simplifying assumption was made that all demands were normally distributed.

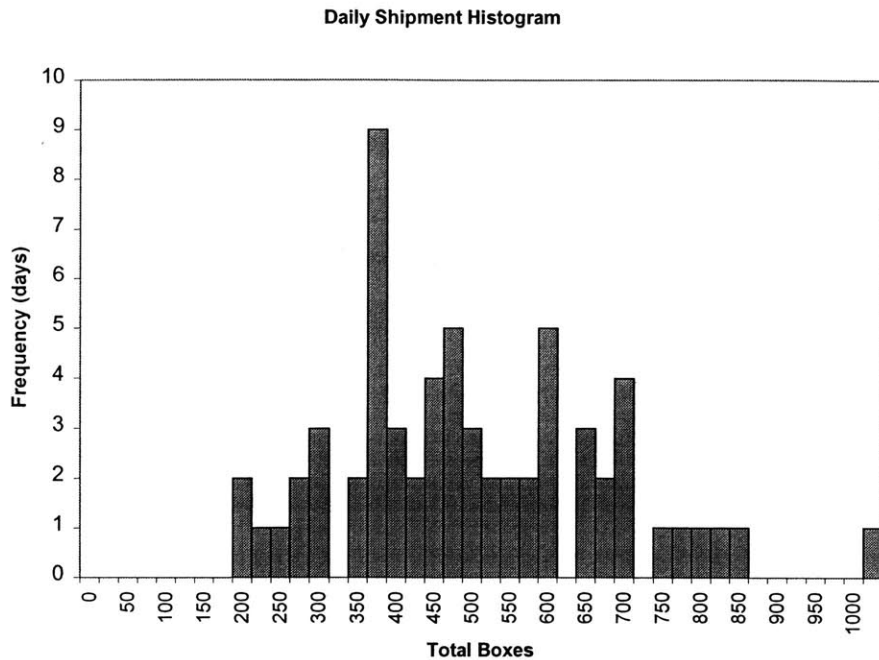


Figure 8: Shipment histogram

Traditionally, Celestica New England’s products have been high-end UNIX workstations and servers. The customers for these products paid premiums for the reliability and performance associated with these products. In addition, Celestica’s OEM customer had proprietary control over the version of UNIX shipped with these machines, and the end users made purchase decisions primarily based on technical suitability for a particular task or compatibility with existing software or networks. In other words, the OEM had some measure of monopoly power in the supply chain since the products were highly differentiated.

Two industry developments have served to lessen this power. First, highly standardized lower-priced Microsoft Windows-based computers have made enormous strides in performance and reliability, eroding the market share of UNIX computers. Second, the growth of the Internet and, more specifically, standard Internet network protocols has reduced users’ dependence on any one hardware-software combination. End users now find it much easier to mix-and-match computing platforms to best suit their needs.

The resulting increased customer power has forced the OEM’s – and their manufacturing partners – to compete in the areas of customer service and cost. These are both determined to a large part by manufacturing lead time: faster deliveries of custom configured hardware and

lower work-in-process inventories and the associated carrying cost.

As a result, Celestica New England has chosen to do whatever it took to reduce the flow time in the factory from 5 days to 1 day. This is called the “24-hour factory” initiative. “24 hours” was a slight misnomer, since the purpose of the initiative is next day shipment. That is, an order received early in the morning but not shipped until late in the evening of the following day qualifies as a next day shipment, although it may consume up to 48 hours of flow time.

The software and localization kitting area (“Kits”) must support the 24 hour factory. In addition to promising next day shipment of any stand-alone kit, it is imperative that any kit which completes the order of a computer be available when or before the rest of the order is available. The worst case scenario is a \$100,000 order of hardware sitting unable to be shipped because it is waiting for a \$10 kit to be completed.

4.1.2 Internal pressures

In addition to the external demands of the customer, Celestica New England (CNE) is under increasing stress resulting from two internal directives associated with space and people.

CNE is in the process of building a new factory that will be located 10 miles from the current site. In the meantime, the current building has been sold to another manufacturing company and Celestica is leasing back the space it needs, slowly vacating the building at a previously determined rate. In addition, new OEM customers are being added to the factory continuously, usually requiring their own production areas. The resulting squeeze has forced every part of the factory to shrink its footprint.

Like any EMS provider, Celestica is also under constant pressure to cut costs. Since a large part of EMS services is labor-intensive, great effort is being put into improving productivity. This would allow variable labor costs to be reduced or, in a tight labor market like New Hampshire in 1999, re-deployed elsewhere in the factory to fuel business growth.

Therefore, the three objectives of the improvement initiative in the kitting area were to reduce cycle time, reduce space, and reduce costs by improving productivity.

4.2 Existing Planning, Production, and Inventory Management Processes

The software and localization kitting area of the Celestica factory is separate from all other manufacturing. The area receives raw materials either from the on-site warehouse or directly from suppliers and returns finished kits to the distribution area.

The area is characterized by a high number of available finished good SKU's. Each SKU corresponds to a different kit configuration. During any three-month period, as many as 140 different SKU's will be assembled in the kitting area. These SKU's are not static: every 6 weeks the kit contents and configurations will be updated, and the SKU numbers will change accordingly.

There are 6 ½ people assigned to the Kits. Six are direct labor and the remaining half resource is a planner/team leader who provides direction regarding what kits to make and when to make them.

4.2.1 Existing Planning and Inventory Processes

Every morning, the planner will look at a paper report, or "PAC384," which shows the daily orders for every SKU for the next 10 days. The planner's goal is to make sure that the line is building orders that have a scheduled ship date 3 days from today. This is called building to "T-3."

Another way to look at building to T-3 is to say that the line uses a fixed 3-day lead-time to buffer demand. Sometimes, during periods of heavy demand, the line would slip to T-2, and conversely when demand would slack they might build up to T-4, but they will always have a day or two of cushion before the order became late.

By generally building only what appears on PAC384 report, the kits area was considered build-to-order. However, build-to-order did not result in low finished goods inventories. In dollar terms the kits area averaged 5.4 days of finished goods inventory. There were three primary reasons for high FGI. First, building ahead to T-3 guarantees that the finished kit will remain in the distribution warehouse for at least three days before it is shipped. Second, three days provided ample opportunity for the order to be cancelled, at which time the finished kit will

remain in finished goods until it can be shipped with a new order. For slow-moving kits, the wait could be months. Third, pressures to support the 24-hour factory resulted in an ad-hoc stocking plan that positioned some safety stock for each kit, even the slow-moving items.

Raw material procurement is handled on a commodity-by-commodity basis as determined by lead-time. Long lead-time items – custom-branded keyboards and mice, software media (compact disks and tapes), power cords, and published books – were purchased in bulk as dictated by the MRP system driven by the customer’s forecast. These were held in an on-site warehouse and delivered to the kit assembly line as required.

About 70% of the raw materials used in the kitting area is of a second type: short lead-time print-on-demand (POD) items. A single vendor, located a few hours away, supplies printed material such as cards, pamphlets, licensees, and softbound manuals and books. If they receive an order before noon, they will print and ship the materials for arrival by late afternoon the following day.

POD ordering is somewhat automated. Using an internet-based supply chain tool called webPLAN, the procurement specialist communicates orders to the POD supplier. The input into webPLAN is an “order demand file,” or ODF. The ODF is hand-loaded with the forecast for finished goods for the upcoming planning period. For example, if the forecast calls for 100 of SKU A1234 during the next month, the ODF will show a requirement for 100 pieces on day 1 of the month. When this demand falls into webPLAN’s “look-ahead” window (about 7-10 days), webPLAN explodes the bill of materials, subtracts current inventory and expected arrivals, and communicates an order for all of the required POD items to the vendor. At the same time, webPLAN is monitoring actual orders, and will send alerts when actual demand outstrips the forecast.

4.2.2 Existing Production Process

At T-3 the planner releases a build authorization, or “B/A,” for a SKU. The B/A is a sheet of paper that describes the kit to be assembled in terms of the SKU, the quantity required, and the bill of materials, and the due date. If the daily demand is larger than about 20 kits, there might be several B/A’s for the same SKU in order to break up the production into manageable batches. The B/A is placed in a bin that acts as a queue.

The kit assemblers will then draw the B/A from the bin. One person will check for availability of the parts by searching for their locations in the line-side stores, or “supermarket.” The search is done by looking up the parts one at a time at a computer terminal that accesses an inventory database. If there is a parts shortage, the B/A is placed in an “unbuildable ” pile for review by the planner. If there are no parts shortages, the same person will travel through the parts supermarket with a wheeled cart and gather enough of each part to meet the quantity on the B/A. Gathering a batch of parts is called “kitting,” although here with a different definition that the overall process of creating software kits, also referred to as kitting.

The next step is to fold boxes for the batch of kits. Usually, the boxes are not one of the parts gathered in the previous step. Instead, the flat boxes are stored separately, and retrieved when actual assembly of the kits is about to begin. After the boxes are folded, a SKU label is applied to the open box.

The folded boxes, up to 20, are arranged on a long non-powered roller conveyor that serves as a work surface. For each part in the kit, the assembler walks the length of the conveyor, placing one in each box. When all the boxes are filled, the boxes are closed, then sealed with tape if necessary, and finally stacked on a pallet.

When the pallet is complete, an “in-transit slip” is completed by hand, and the pallet is moved to an in-transit area. The in-transit is a piece of paper describing the pallet’s contents. Here a person working in the distribution area will put away the pallet’s contents into finished goods inventory locations, using the in-transit slip to perform the inventory transaction at a computer terminal.

When the production information system recognizes that all items in an order have arrived into finished goods, a “pick list” is printed in the distribution area. The distribution team uses the pick list to find the line items of the order in the distribution warehouse and consolidate them into an order that can be shipped to the customer.

4.3 Proposed Planning, Production, and Inventory Management

Processes

As it stood, the current processes for the software and localization areas did not meet Celestica's needs for cycle time, headcount, and space reductions. Several changes to the system were required to meet these goals.

When the whole factory was issuing 5-day lead times to the customer, a T-3 release date was more than adequate to buffer the demand. However, the 24-hour factory initiative meant that most orders would begin arriving within the 3-day window and render this buffer useless. Another method of dealing with demand variability was needed.

There are three basic alternatives for dealing with demand variability. The first is capacity. In this case, the production system is sized to handle a maximum daily volume some number of standard deviations above the mean. For example, the mean demand from the area is 496 units with a standard deviation of 202 units. To satisfy the demand 95% of days (1.6 standard deviations above the mean), a capacity of 820 boxes/day would be required. Since capacity in this manual process means adding workers, and thus increasing cost, this was unacceptable.

The second potential method is flexibility. For this to take place, a means would need to be found to add kitting capability to other assembly lines in the factory. After investigation, this was ruled out for fear that production elsewhere would be disrupted.

The third possibility is to add buffers. If the buffer is on the front end of the process, it is a customer quoted lead-time. This time buffer is what is failing now: building to T-3. Otherwise, a buffer appearing after the manufacturing process has begun is physical inventory. It was decided that finished goods inventories would need to be kept to insure customer service.

4.3.1 Finished Goods

Having decided to hold that finished goods inventory, several questions needed to be answered. How much inventory needs to be held, and for which finished goods? How much space will this take? Without orders to trigger assembly, how do we maintain the inventory levels?

To answer these questions, the analysis begins by looking at the customer demand data. For each individual active SKU, daily shipment data for the previous three months was extracted from Celestica's management information systems. At this point the reader may be asking, "How can you be sure that *shipment* data corresponds to *demand* data?" The answer is, of course, that we can not be sure. In fact, there are some mechanisms that will smooth the demand and some that will induce or amplify variation. For example, the mechanism for "slot planning," or negotiating ship date with the customer based on already allocated capacity, will tend to smooth out large demand spikes. On the other hand, a parts shortage may lead to a burst of hastened production and a large rush of shipments when the parts finally arrive. In any case, it was determined that since shipment data was the most readily available, it was close enough.

Since SKU's for the same product will change, some measure of manual cutting and pasting was necessary to combine the data to make it useful. After the data was cleaned up, the mean daily volumes and standard deviations of the demand data were computed, and the SKU's were sorted in order of decreasing volume.

It quickly became apparent that a large part of the volume was in relatively few SKU's. In fact, of the 136 active SKU's only 16 SKU's comprised the top 80% of orders (Figure 9). Some SKU's were shipped only a few times over the three-month period. This led to the conclusion that significant improvements could be made by treating these high-running SKU's differently than the remaining SKU's. As might be guessed from a glance at Figure 9, software and localization kits were excellent candidates for an ABC-type inventory management treatment (Peterson & Silver, 1979).




Figure 9: Cumulative orders by SKU. Note the high concentration of orders in only the top SKU's

Separating the SKU's into a high-running "A" group from a low-running "B" group provides several benefits. First, concentrating management attention on the highest volume items would yield the greatest return on On-Time-Delivery (OTD) metrics. Suddenly, instead of having to actively manage daily orders for 136 SKU's, there would only be 16 or so for which the inventory levels would need to be monitored, and the remaining 120 would be dealt with on an as-needed basis.

A second benefit is that the A/B delineation allows creating a hybrid build-to-stock and build-to-order production system. Maintaining finished goods inventory in *only* A parts provides most of the demand buffer required in a compact space requirement. That is, it is easy to hold large volumes of a small number of items in a small physical space. However, holding only small inventories of a large variety of goods leads to an explosion in space requirements since no two SKU's are allowed to share the same inventory location. The reason for this is that inventory locations in the typical warehouse are generically sized, and rarely matched to the size of the item.

The third benefit resides in the ability to create a self-regulating priority system for the

production line. The first priority for production should always be A items for which there is a stock-out condition, since it is almost guaranteed that these will delay orders if not replenished. Second priority are B items due to be shipped in the next 24 hours, since there are no finished goods stocks from which to draw. Third priority are A inventory replenishment orders, since we are assured that these orders can be shipped from stocks. This input queue should be dynamic, allowing for A stock-out and B orders to be placed at the head of the queue as they arrive.

The inventory system described above, coupled with a production system designed for just enough capacity to keep the input queue manageable, will ensure that all orders are met within 24 hours with little active order management.

4.3.2 Raw Material

The line-side raw material stores were a major concern for Celestica. They required what seemed like inordinate resources in terms of management and space. The line-side stores, or supermarket, contained both POD and long-lead time items. POD items were delivered to the stores directly from the vendor, while everything else was stored in the warehouse and delivered to the supermarket on an as-needed basis.

The assembly crew was responsible for maintaining the supermarket inventory levels for warehouse parts. It was left to their judgement to decide when to place orders and for how many units. They were limited in their decision only by how much shelf space was available in the supermarket and by pan-size³ limitations. The proposed improvement for this process was the implementation of an already-developed information technology solution called WCS⁴ Locator. WCS Locator was a link between the warehouse control system and the shop-floor control system. It monitored the usage out of the supermarket, and replenished the line-side inventory according to an (s,S) replenishment policy (Nahmias, 1997).

The anticipated benefits of the WCS Locator system were threefold. The most important was the automation of the supermarket inventory maintenance, which was hoped to free line resources to concentrate on assembling kits. Also important would be the establishment of

³ The pan-size is the smallest quantity that the warehouse will handle, usually determined by number of units per box.

⁴ WCS, or Warehouse Control System, is the brand name of Celestica's computerized warehouse management system.

consistent inventory policies that would minimize the use of supermarket shelf space, and thereby allow the supermarket footprint to be reduced. The third advantage of automated replenishment is a smoothing of the replenishment requests to the warehouse. Currently, the assembly team surveys their inventory and outstanding orders twice per day and then requests all necessary material from the warehouse at once. The warehouse personnel see this as a flood of replenishment requests and in response scramble to keep up.

The replenishment system for the POD materials was also a target for improvement. The automated webPLAN system provided the functionality required to maintain designed inventory levels with daily periodic review and ordering. Vendor deliveries were already being made to Celestica on a daily basis. With a small amount of analysis, the opportunity for lowered space requirements and improved material availability was at hand. It was proposed that historical demand analysis be reconciled with forecast sales, and the resulting volume be level-loaded in webPLAN's ODF. This would create a periodic review policy in which we could set target inventory levels in terms of days of inventory.⁵

The ABC analysis performed for the finished goods also had an application for raw materials, albeit for a different purpose. Identifying the high running SKU's allowed the raw material locations to be rearranged by volume, with the high volume parts located on or near the line, and for parts commonly shipped together to be grouped in the supermarket. This minimized both the number of parts requiring retrieval from the supermarket as well as the walk distance for the assembler.

4.3.3 Proposed Production Process

The production system faced two hurdles. First, the existing batched production system was marked by a large use of space. For example, a keyboard kit, at 27 inches long, requires at least 45 feet of conveyor for assembling a batch of 20 at one time. This was too long for the new space provided. Since the length of the conveyor drives the overall area footprint, it was decided that the conveyor length, and thus the production system, would need to change to match the available space.

Second, the headcount goal, determined by an across-the-factory uniform target, was a

reduction of 2 direct labor to a total of 4. For a largely manual process, a reduction the labor by one third will, in the absence of process improvements, reduce the production capacity as well.

In other production areas Celestica had experienced good results implementing facets of lean manufacturing. Specifically, quality, throughput, and productivity metrics had increased when batching was eliminated and a “single-piece” flow system was instituted. After data indicated that it would be possible to meet demand with this type of system in software and localization, we began the design of a lean process.

In contrast with the ad-hoc procedures of the 6-person production system, success with a 4-person system would require more discipline and adherence to specified roles, responsibilities, and procedures. As described in Figure 10, four distinct roles were developed. The **Kitter** selects the next B/A in the queue, and gathers the parts in the supermarket onto a cart. The **Folder** gathers the correct boxes for the SKU, folds the first box for the batch, and applies the label. The **Assembler** fills the box from the cart and from line-side materials. Finally, the **Packer** seals the box and delivers the batch to the distribution area.

⁵ Hadley & Whitin (1963) calls this an <R,T> system.

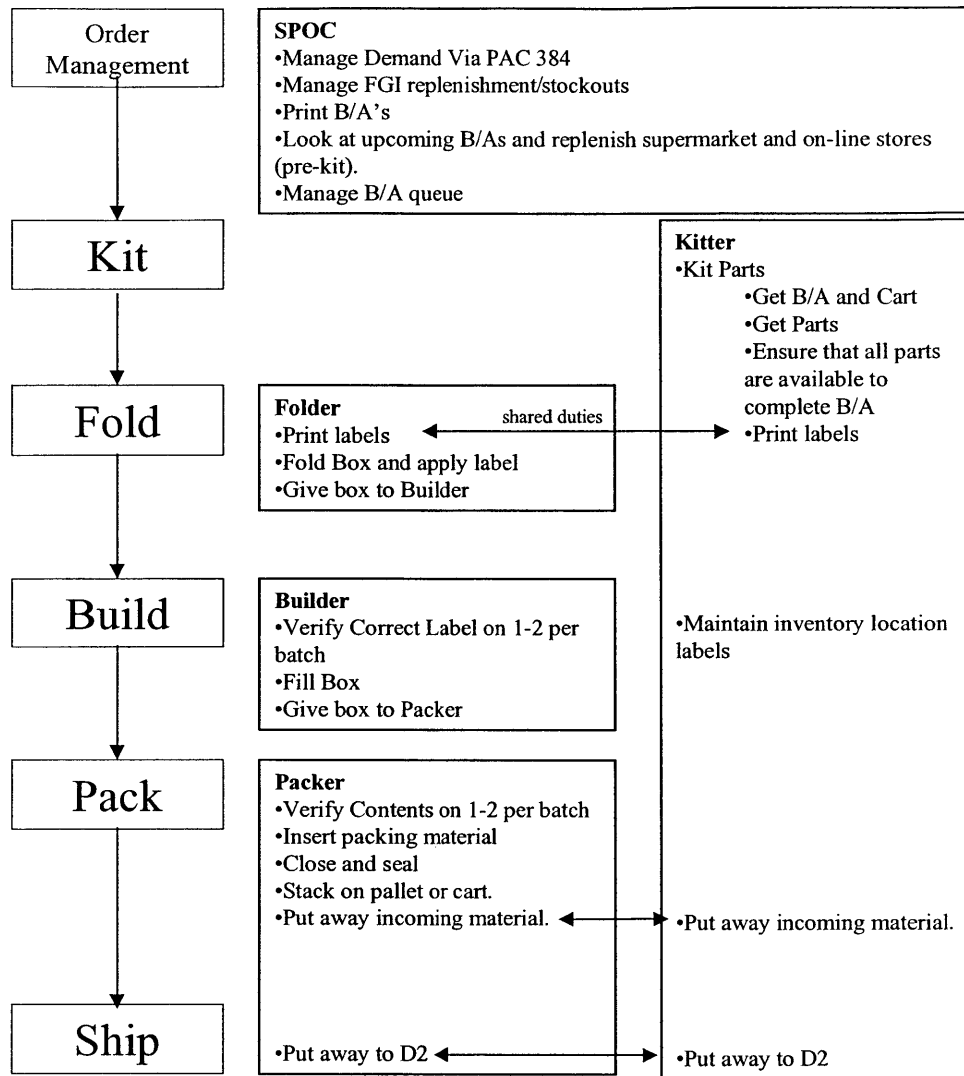


Figure 10: Roles and Responsibilities in a 4-person production system (SPOC = planner). The process steps appear in the first column and the Roles and Responsibilities are on the left.

Experiments were conducted to confirm the feasibility of meeting demand in a 4-person production system. Untrained workers (the author and a co-worker) were able to consistently complete the fold-pack cycle on the most complicated kit in 1.5 minutes. It was assumed at this point that the folder, builder, and packer could find a way to balance the cycle time to a production rate of 1 unit every 30 seconds. For a 450 minute working shift, this line has a theoretical service rate of 900 units/shift. Daily mean demand of approximately 500 units equates to a utilization of 56%. It was felt that this utilization was low enough to account for down time and queuing inefficiencies, especially when combined with a trained workforce and

a mean kit complexity of less than the test case.

4.4 Detailed Implementation and Problems Encountered

The implementation path followed the following milestones:

1. Shrink and move production area and supermarket
2. Establish finished goods location and replenishment system
3. Establish raw material replenishment system
4. Institute the new 4-Person build procedure
5. Implement IT changes

Some of these efforts were more successful than others were, and each had unique challenges for implementation.

4.4.1 Shrink and Move Production Area and Supermarket

Celestica was on the verge of ceding a large amount of floor space to the building's other occupant, requiring the size reduction and relocation of the software and kitting area. The existing layout appears in Figure 11. The available space was approximately half of the existing footprint.

Both the raw material supermarket and the assembly line had to shrink. The first question when looking at the supermarket was "How much inventory do we need on the line?" The first source of data to answer this question was usage (demand) data for the raw materials. However, several different systems were used to track usage and ordering of raw materials. Rectifying these data into a single format would be impossible given the time constraints under which the improvement team was operating. Therefore, finished good demand data was used to compute the demand data for the raw materials.

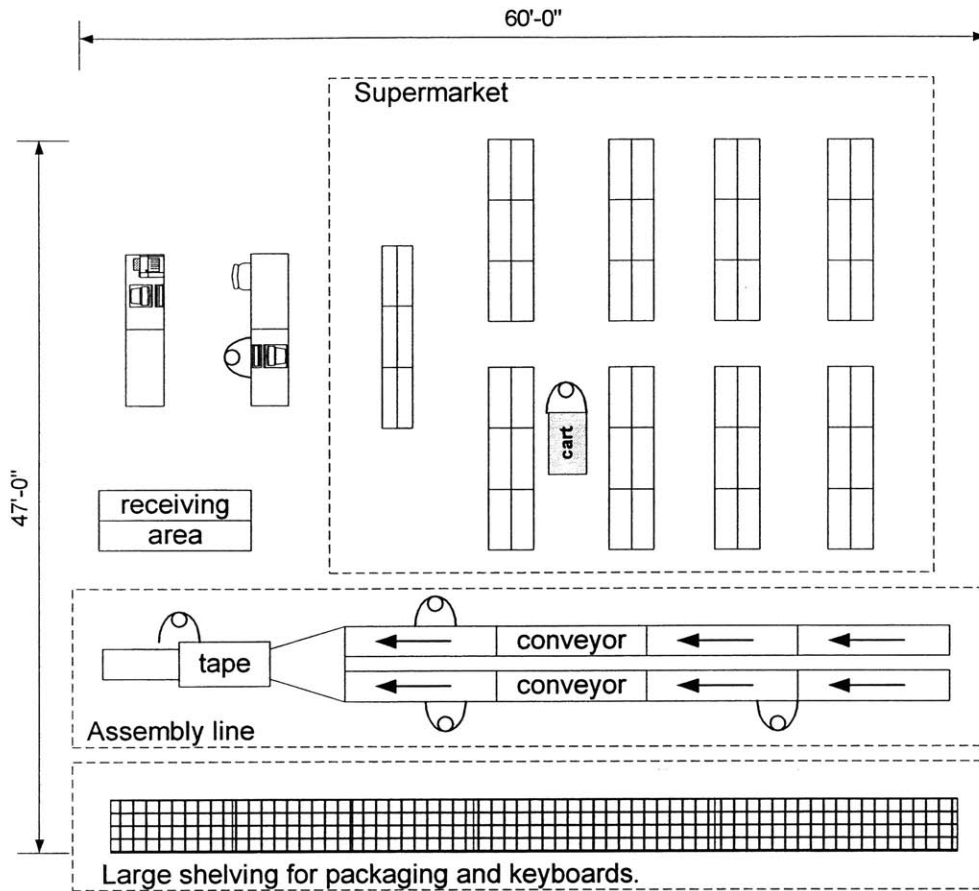


Figure 11: Existing Software and Localization Line Layout

This was not as straightforward as one might imagine, since both mean and variance values are needed to compute inventory levels. Often several finished good SKU's shared common raw materials. Thus, for every SKU the bill of materials was exploded twice, once to for the mean demand for each part and once for the variance⁶. Then the means and variances were summed for each part. That is, let there be N SKU's that share part A, and let μ_i and σ_i^2 be the mean and variance of the demand of SKU i , for each $i \leq N$. Then, if SKU i requires n_i of each part A:

$$\mu_A = \sum_{i=1}^N n_i \mu_i \quad \text{and} \quad \sigma_A^2 = \sum_{i=1}^N n_i \sigma_i^2 \quad (4.1)$$

Spreadsheet macros were created that manipulated the SKU demands and bills of materials to compute Equations (4.1).

⁶ The variance, σ^2 , is the standard deviation squared.

As described below in section 4.4.3, a periodic replenishment system was designed. For mean demand μ with variance σ^2 , replenishment period r , replenishment lead-time L , and customer lead-time l , an appropriate base stock level would be

$$B = \mu(r + L - l) + z\sigma\sqrt{r + L - l}. \quad (4.2)$$

Equation (4.2) is used as follows: Once every period r , orders are placed to replenish stocks up to B units. The safety factor z is a customer service level in terms of standard deviations from μ . For the purposes of line layout, z was chosen to be 2.3, or a 99% confidence of inventory availability (assuming a normal distribution of inventory.) The review period was chosen to be once per day. Examination of Equation (4.2) confirms that if lead time to the customer, l' , is equal to or greater than the time required to review and replenish the inventory, then the system is truly “build-to-order” and requires no inventory.

The lead-times L and l were assumed deterministic. For POD parts, L_{POD} was determined to be 2 days during one shift per day operation. This is because POD parts ordered during the morning of day 1 would arrive at the end of first shift on day 2, for possible use on day 3. For warehouse-stocked parts, a delay of up to 4 hours could arise between when a signal is sent and when replenishment arrives, so $L_{WH} = 0.5$ days.

Customer lead time, l , deserves some discussion. As described earlier, Celestica’s customers have demanded the possibility of 24-hour turn-around on orders, although this option is not yet widely used. For the purposes of sizing the supermarket it was assumed that most, if not all, customers would eventually demand 24-hour shipment and l was set to 0. However, it is useful to note that if a significant fraction of the customer base is content with longer lead times, raw material inventories could be reduced accordingly.

At this point base stock levels are compared to current inventory in order to determine the required supermarket shelf space. After adjusting stocking levels and scrapping obsolete parts, it was determined that the supermarket shelf space could be reduced by half.

A new layout was designed that minimized the supermarket area and conveyor length subject

⁷This assumes l does not include manufacturing lead time.

to the constraint of raw material inventory space (Figure 12). The total footprint was reduced by 1770 ft², or 63%. The conveyor length of 25 feet is actually longer than desired. Ideally, it would just be long enough to assemble one box at a time and thereby forcing single-piece flow. Instead, the length of the conveyor is driven by the length of pigeonhole shelving above the line.




Figure 12: Redesigned Software and Localization Line Layout

In addition to the reduced footprint, there are other advantages of this layout. The smaller supermarket reduces the walk distance for the kitter. The pigeonhole shelves and the end shelves behind the assembler provide locations for most of the raw materials required for A SKU's, which eliminates the kitting step entirely for some high-running goods.

4.4.2 Establish Finished Goods Location and Replenishment System

At a first pass, it was recommended that finished goods storage be located next to the assembly line. The intention was to eliminate physical or responsibility division between the assembly line and the finished goods, making the line accountable for its own inventory and customer service on a visual basis.

Co-location of inventory and assembly proved to be impossible given the space constraints, and instead the finished goods inventory are located in the distribution warehouse, approximately 200 feet and behind a wall from the assembly line. This seemingly small distance turned out to be a significant hurdle, since it requires a formalized replenishment

system to notify the assembly line when to build.

As described in section 4.3.1, finished goods would only be held in SKU's designated as A SKU's. In order to qualify as an A SKU, two criteria must be met. First, it must be considered a high-runner. That is, demand for the SKU must comprise part of the top 80% of orders shipped. Second, the variability in demand for the SKU must be manageable, to avoid holding too much inventory that rarely moves. The threshold for acceptable variability was a coefficient of variation (σ/μ) not exceeding 2.

The replenishment system is manual, for which success depends on simplicity. Thus, the simplest type of inventory system was chosen. This is a continuously reviewed reorder point-reorder quantity, or (Q, R), replenishment system. For any SKU, when inventory level R is reached, a fixed quantity Q is ordered from the assembly line.

R must cover the probabilistic demand over the manufacturing lead time L with some customer service level z . Therefore, similar to Equation (4.2):

$$R = \mu L + z\sigma\sqrt{L} \quad (4.3)$$

The second term in this equation is often called the safety stock, or the inventory level expected to be reached when a replenishment order arrives:

$$S = z\sigma\sqrt{L} \quad (4.4)$$

Q is often calculated as an economic order quantity (EOQ⁸). In this case, the data required to determine EOQ is fuzzy as best, and a more informative question was asked: "How often would you like to replenish any given SKU?" The answer, depending on who was asked, is once every day or once every two days, so a value of 1.5 days of inventory was chosen as a first pass for Q .

It turns out, however, that physical limitations prevent a pure (Q, R) model from being implemented. The ideal inventory holding configuration for a manually maintained (Q, R) system is a tall stack of inventory with a replenishment signal (piece of paper) inserted R units

⁸ See Nahmias (1997) pp. 222-233.

from the bottom of the stack. Unfortunately, most of the SKU's require several inventory locations to accommodate the maximum expected $(S+Q)$ inventory level. For this reason, Q and S are determined according to Equation (4.4), but then $(S+Q)$ is split up into N "bins". N depends on the relative physical sizes of the SKU's box and the inventory location. When any of the bins is emptied, a replenishment signal for one bin quantity is sent to the assembly line. For example, a traditional 2-bin kanban replenishment system is a (Q,R) system broken into bins with $N=2$ and $Q=R$.

Table 4: Example of Finished Goods Inventory System: Three SKU's

<i>SKU</i>	μ	σ	σ/μ	<i>Cum %</i>	S	Q	$S+Q$	<i>Bins</i>	<i>Bin Size</i>	<i>Bins * Bin Size</i>
A4199-70003	129.9	96.2	0.74	20%	193	200	393	7	55	385
A4983-70016	90.1	45.5	0.50	38%	91	140	231	4	64	256
025988	72.1	47.4	0.66	54%	95	110	205	5	42	210

For each bin, a replenishment ticket (kanban card) is created. When a bin is emptied, the replenishment ticket is placed in a folder in the distribution area. Periodically, the tickets are delivered to the assembly line, which triggers replenishment for that bin. The tickets include all the information required for the assembly line to generate a new B/A. An example replenishment ticket appears in Appendix B.

The system is not static. The list of A SKU's and replenishment parameters changes as new SKU's are added or replaced, and as demand rises and falls. To aid maintenance of the system, two software tools were developed. The first was a PC-based automated spreadsheet. The macro programs in the spreadsheet acquire demand data from readily available sources, perform the A/B split, and calculate the replenishment system parameters seen in Table 4. The second was an automated word processing document which created the replenishment tickets (Appendix B) from the parameters generated in the spreadsheet.

4.4.3 Establish Raw Material Replenishment System

For Print-On-Demand parts, an Internet-based replenishment system called webPLAN was currently being used at Celestica. There were two data sets used by webPLAN to replenish POD materials: the forecast and actual unfilled orders. Both were used to determine order quantities.

The most important webPLAN parameter for determining inventory levels is the look-ahead window, T . Every review period r , webPLAN looks out T periods in both forecast demand F and actual orders received O . It then looks at current inventory I (including expected arrivals), unallocated backlog K , and places the order for Q to the following rule:

$$Q = \max\left(\sum_{i=1}^T F_i, \sum_{i=1}^T O_i\right) - I + K \quad (4.5)$$

The index i on F and O can take on integer values from 1 to ∞ , where 1 refers to the current period. When most orders arrive well within T , $\sum F$ will dominate and only in the event of unexpected demand spikes will $\sum O$ cause webPLAN to order above forecast.

If F is constant, the normal behavior of this replenishment algorithm is represented by a simple base-stock periodic replenishment model. That is, if

$$B = \sum_{i=1}^T F_i = TF \quad (4.6)$$

then every review period webPLAN will order up to B . Values for B (and T) are determined from Equation (4.2), and this computation is performed in the same spreadsheet that calculates the inventory parameters for the finished goods. For most parts the base inventory was calculated to be 3 to 5 days ($3 \leq T \leq 5$).

Since webPLAN accepts forecast and order data in terms of finished good SKU's, and then automatically explodes the bills of materials and orders parts, it will tend to hold too much inventory in parts that are shared among several SKU's. In other words, webPLAN lacks the intelligence to automatically sum demands and variances as in Equations (4.1) and compute inventory levels. For this reason, using webPLAN as described here is not a good idea for high-value raw materials shared by several finished goods. In this particular context of a large number of low-cost parts, the difference is not judged significant.

For long lead-time parts, replenishment was from stores located in the warehouse. For this an automatic replenishment system called WCS Locator was implemented. WCS Locator used a continuous review, base-stock replenishment system. For each part, two parameters were kept: R and $(S+Q)$. When R was reached, $(S+Q) - R$ parts were ordered. The parameters

were computed once using Equations (4.3) and (4.4), then maintenance was left to the assembly line's collective discretion based on usage and replenishment performance (lead-time from the warehouse).

4.4.4 Institute the New Four-Person Build Procedure

As could be expected, changing work habits proved to be the second most difficult challenge associated with this initiative (the most difficult being Section 4.4.5: Implement IT Changes, below). At the end of the project (and my internship tenure), the assembly line had started experimenting for short periods with the build process described in Figure 10. Usually, the line operated in the same mode that they had before the area was moved. They were hampered by the smaller conveyor on which they could not batch assemble as many kits as they would prefer.

The assembly crew was used to a high level of autonomy and lack of structure. In this environment, they had settled into work habits dramatically different from those proposed. There was a high level of skepticism that the new build procedure would be as "productive" as the batching process previously used.

Overcoming this skepticism and resistance to change was a leadership challenge that was not effectively met. The improvement team usually worked outside this area, and had no source of past authority or credibility to draw upon. Additionally, there was no opportunity to establish a leadership foundation under the project's extreme schedule pressure.

In these circumstances, the best possible solution would be cooperation with a full-time supervisor who could provide the guidance and source of authority for the changes. Unfortunately, there was no full-time supervisor dedicated to the software and localization area.

Implementing the new four-person build procedure was, in the author's opinion, a failure attributable to a lack of leadership on behalf the implementation team (of which the author was a member and shares full responsibility.) Therefore, at this point in time it is impossible to compare the results of the new build procedure to that of the old.

4.4.5 Implement IT Changes

There were several information technology improvements requested by the improvement team. Some were relatively minor, like small improvements that allowed the assembly team to put away their own finished goods into distribution without needing to hand it off to distribution personnel. These were quickly implemented.

Some changes, however, are more involved and complex. This translates to resource intensive, or requiring a large amount of time from programmers and system administrators. There were 2 items that fall into this realm:

1. Installing Manufacturing Backplane (MBP), Celestica's proprietary shop-floor control system.
2. Implementing WCS Locator, which depends on MBP for usage tracking.

The MBP improvements comprise automatic B/A generation and prioritization. Customer orders will automatically be delivered to the B/A queue, and replenishment orders for A finished goods would be entered at a terminal in the kitting area. The prioritization of the queue would then take place automatically and dynamically: A stockouts, B customer orders, then A replenishment orders. A customer orders would be weeded out, since these would be serviced from inventory.

During the last quarter of 1999 IT resources were scarce. Most of the IT department was occupied with rooting out and correcting bugs associated with the upcoming Year 2000 rollover. In fact, extensive changes the information systems were frozen during November 1999. The MBP and WCS Locator improvement would have to wait until 2000.

4.5 Realized Benefits

The author's tenure at Celestica ended before rigorous quantitative measures could be made of the improvement effort in the Kitting area. However, with a few plausible assumptions, the results can be estimated. Any improvements listed below are in addition to the success of meeting customer demands for reduced lead-time and contractual demands for space reduction.

4.5.1 Inventory

As mentioned above, the previous “build-to-order” system did not produce a small finished goods inventory due to the delay between kitting and shipping. The inventory savings as a result of moving to the A/B FGI strategy was conservatively \$58,000.

While there was considerable space and inventory reduction in raw materials, the value change was negligible due to the low-cost nature of the parts.

4.5.2 Human Resources

Personnel resource savings were one of the explicit goals of the improvement effort. Two direct laborers were freed, and half of a direct laborer in distribution was freed from putting away the Kitting area finished goods. It is anticipated that when the IT improvements are completed, the planner, whose time is split between this and another area, will be unnecessary.

These headcount savings can be quantified at \$6000 per month immediately, with another \$1600 to follow the IT changes.

4.5.3 Process

Benefits from process improvements are difficult to quantify, but important to consider as reduction in management costs and improvements in customer service. Establishing a clear strategy for absorbing demand variability in finished goods inventory provides a simple method of measuring customer service and troubleshooting process problems. Eliminating excess “touches” to inventory reduces the opportunities for failure. Enforcing standard operating procedures and defined roles improves quality and lowers training costs.

4.5.4 Overall savings

The quantifiable sources of savings, inventory and headcount, can be shown to have a significant impact. Assuming an annual cost of capital of 15%, the total return from this 3-month effort is a net present value of \$240,000 over three years.

4.6 Conclusion and Recommendations

The improvement effort in the software and localization kitting area was a highly pressurized program completed to the level described in only 3 months. When evaluated on the return on

investment, it was definitely successful. When evaluated compared to what was possible, the effort has room for improvement.

The undertaking of an aggressive change program that involves several functional areas requires a cross-functional team with the skills and authority to make the changes. In two areas, IT and the shop floor crew, the team involved did not have the influence to complete the changes required. Remedies for these deficiencies would have been an IT resource person and a full-time shop floor supervisor.

It is my sincere hope that Celestica completes the last few elements of this initiative to fully realize the potential benefits. The results so far have been impressive, but I feel the true power of this exercise is as an example. An example, that is, of looking at a production system from beginning to end, eliminating wasted steps and resources, and focusing on the elements that provide value to the customer.

The A/B split of finished goods inventory can be applied to any production system that produces a high variety of products with probabilistic demand. It focuses attention on the high-volume products while still insuring a mechanism for fulfillment of B orders. A secondary benefit is simplified management.


Simple, yet highly effective inventory policies were implemented for both finished goods and raw materials. The use of statistical data analysis in this manner is rarely used but with a small amount of instruction can be a powerful tool to manage inventories. There are many potential applications of intelligent inventory policies throughout Celestica.

REFERENCES

- Abraham, Katharine G. (1988). Flexible Staffing Arrangements and Employers' Short-Term Adjustment Strategies. *Employment, Unemployment, and Labour Utilization*, 288-311.
- Gershwin, S. B. (2000). *Manufacturing Systems Engineering*. Cambridge: Massachusetts Institute of Technology.
- Hadley, G. & Whitin, T. M. (1963). *Analysis of Inventory Systems*. Englewood Cliffs: Prentice-Hall, Inc.
- Herer, Y. T. & Harel, G. H. (1998). Determining the size of the temporary workforce— An inventory modeling approach. *Human Resource Planning*, 21(2), 20-32.
- Hooper, A. & Catalavello, J. (1981). Markov Analysis Applied to Forecasting Technical Personnel. *Human Resource Planning*, 4, 41-54.
- Nahmias, Steven (1997). *Production and Operations Analysis*. Chicago: Irwin.
- Peterson, R., & Silver, E. A. (1979). *Decision Systems for Inventory Management and Production Planning*. New York: John Wiley & Sons.
- Rowland, K. M. & Sovereign, M.G. (1969). Markov-Chain Analysis of Internal Manpower Supply. *Industrial Relations*, 9, 88-99.
- Womack, J.P., Jones, D.T., & Roos, D. (1991). *The Machine that Changed the World*. New York: HarperPerennial.

APPENDIX A: WORKFORCE MODEL PLOT







Permanent/Temporary Model for
 $\sigma = 0.1\mu$ (Normally Distributed) and $C_u = 3(C_{mp})$



APPENDIX B: SOFTWARE REPLENISHMENT TICKET

Software FGI Replenishment Ticket

Instructions: Place this ticket in the bin at the end of the racks **NOW**. If you have taken the last 025961 in D2, please attach a red STOCKOUT sticker with today's date!

Workcell	CD	
Assembly	025961	
DESCRIPTION	KIT REL NOTES HP aC++	
QTY 1	15	
QTY 2	10	
QTY 3		
QTY 4		
QTY 5		
QTY 6		
DEST.	D2	
Total required	25	
Stockout stickers	STOCKOUT	
Replenishment system	2 bins of 25 pcs for a max of 50	