Optimizing Inventories and Standardizing Planning Procedure in a Multipart Manufacturing System

by

Yu Hua

Bachelor of Science in Mechanical Engineering
Tsinghua University, 2015

Submitted to the Department of Mechanical Engineering
in partial fulfillment of the requirements for the degree of

MASTER OF ENGINEERING IN MANUFACTURING

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

September 2016

Author..Signature redacted..........
Yu Hua
Department of Mechanical Engineering
August 10, 2016

Certified by.....................Signature redacted.....................
Stephen C. Graves
Abraham Siegel Professor of Management Science
Thesis Supervisor

Accepted by.............Signature redacted.........................
Rohan Abeyaratne
Quentin Berg Professor of Mechanics
Chairman, Committee of Graduate Students

Optimizing Inventories and Standardizing Planning Procedure in a Multipart Manufacturing System

by

Yu Hua
Bachelor of Science in Mechanical Engineering
Tsinghua University, 2015

Submitted to the Department of Mechanical Engineering on August 10, 2016, in partial fulfillment of the requirements for the degree of Master of Engineering in Manufacturing at the Massachusetts Institute of Technology

Abstract

This work illustrates how to apply a multi-echelon periodic review base-stock model to a real manufacturing company, Waters Corporation, for their major product Analytical Columns. At Waters, the Analytical Column supply chain ranges from raw materials to the final delivery to customers, and covers US, Europe, and Asia. The goal of this work is to find the best locations to hold safety stock along the supply chain so as to minimize the inventory holding cost for the whole company. To do this analysis, a 5 stage multi-echelon supply chain model is constructed. All the stage costs are measured and standardized based on data from the "SAP" system. To estimate the demand variability, we utilize an adjustment method that accounts for the aggregate bias in the forecast. The final optimal solution will reduce the safety stock level by 67% for the supply chain. We also find a near-optimal solution that is easier to implement; this solution would reduce the safety stocks by 59%. Finally, we argue that the implementation of this model and its assumed operating policies can improve internal communications within the company, leading to better integration across operating units.

Thesis Supervisor: Stephen C. Graves
Abraham Siegel Professor of Management Science

This page is intentionally left blank.

Acknowledgements

I would like to thank our advisor, Stephen C. Graves, for his patient guidance, enthusiastic encouragement, and useful advice throughout this project. I would like to also thank Sean Willems for his assistance and suggestions.

Thank you to Gabriel Kelly and Greg Puszko for the constant guidance at Waters Corporation over the course of this project.

Thank you to everyone at Waters Corporation for their assistance throughout this project, including David Terricciano and Jim McPherson for making this collaboration between MIT and Waters possible.

Thank you to my project partners, Nelson Lee and Yan Han, for their collaboration in producing a result we can all be proud of.

Thank you to David Hardt and Jose Pacheco for giving me the chance to be a member of the Masters of Engineering in Manufacturing program, and for their leadership throughout the past 12 months. I also want to thank all of the other students in the MEngM program for making this past year so memorable.

And finally, thank you to my family for their continuous support throughout my whole year here at MIT.

This page is intentionally left blank.

Table of Contents

List of Figures and Tables

1 Introduction

This thesis illustrates how to apply a multi-echelon periodic review base-stock model to a real manufacturing company, aiming to find the best place to hold safety stock along the supply chain so as to minimize the inventory holding cost for the whole supply chain. Also by implementing this model, we hope that some of the policies like guaranteed service time and forecast-based review will improve the company's planning integration and internal communication.

The thesis is based on a project sponsored by Waters Corporation, an analytical instrumentation company located in Milford, Massachusetts. Waters mainly provides liquid chromatography machines and the consumable analytical columns inside the machine. The analytical column business is their major revenue driver and has undergone huge demand growth in recent years. Just for quarter 1 in 2016, the demand growth rate was as high as 9.5% relative to one year before. This huge demand growth brings extra burden to Waters' current supply chain and operations. The VP of operations in Waters wants to investigate ways to improve the performance of their analytical column supply chain. Therefore, the purpose of this project is to apply a multi-echelon model to Waters' supply chain, and develop recommendations for improvement based on comparisons with a theoretically optimal solution.

Waters maintains a long supply chain, starting from raw materials to the final finished goods, and manufactures most of the products on their own. Geographically, the supply chain is also a global one covering US, Europe, and Asia. So this supply chain is complicated both functionally and geographically, involving many stakeholders. After many interviews and onsite research with different stakeholders, we were able to map out the complete supply chain and the internal planning procedure. We discovered some issues for further investigation, including the lack of cooperation between facilities, holding too much safety stock, and too many build-to-order products.

Particularly, this thesis is focused on using a theoretical model to determine the right amount of safety stock and a reasonable ordering policy. Section 2 will dive deep into the supply chain details at Waters Corporation, including the description of the stakeholders involved in the supply chain, the current planning system and planning procedures, and the current forecasting method. With these details some problems are discussed in the end of section 2. Section 3 introduces the multi-echelon model and some related theories like evolving forecasts and adjustment of forecasting bias. Section 4 goes through the process of implementing the multi-echelon model, stating the assumptions and determining all the necessary parameters. Section 5 presents an optimal supply chain solution as a benchmark to assess Waters supply chain performance. Also for the purpose of an easy starting point, a near-optimal solution is proposed. Finally, section 6 gives the conclusion for the whole project and identifies some of the limitations for future work.

2 Background and Statement of Purpose

The purpose of this thesis is to develop a new supply chain strategy so as to realize systematic improvements on the supply chain of Waters Corporation in Milford, Massachusetts.

This thesis is based on a project conducted by Nelson Lee, Yan Han, and the author as a team at Waters Corporation between January and August of 2016. Part of this chapter will share parts of the text with the other two theses because of the collaboration throughout the project.

In order to develop a system level strategy, we first need to map out how the supply chain functions as a system. Then we can identify points of inefficiency along the supply chain to do further investigation of the causes. So our project started with intensive interviews with different stakeholders along the supply chain to gain a basic understanding of how different stages connect into a system and why inefficiencies and problems happen.

2.1 Background on Waters Corporation

Waters Corporation is a leader in the chemical analytical industry, providing test instruments and consumables to the market. Their products are widely used in pharmaceutical industry, food industry, and research laboratories. Our project is focused on the consumables product line, represented by analytical columns, which is the main revenue driver for Waters Corporation.

An analytical column (called as column in the remaining part) is a piece of metal tube filled with powder, and it is installed into the test instrument. During the test, instruments apply high pressure to force sample fluids to flow through the column, then due to different affinities to the powder, the different components in sample fluids will separate; then we can further investigate what are the components and their quantities are in a sample. The affinities of the powder inside a column will gradually decrease as the number of tests increases. So generally a column will be replaced by a new one after several rounds of tests.

2.2 Breakdown of an Analytical Column

To understand the supply chain of the analytical column, we firstly need to determine the major components of an analytical column. Columns may have different size with their inside diameter varying from 2 mm to 150 mm, and length varying from 100 mm to 250 mm, as shown in figure 2.1.

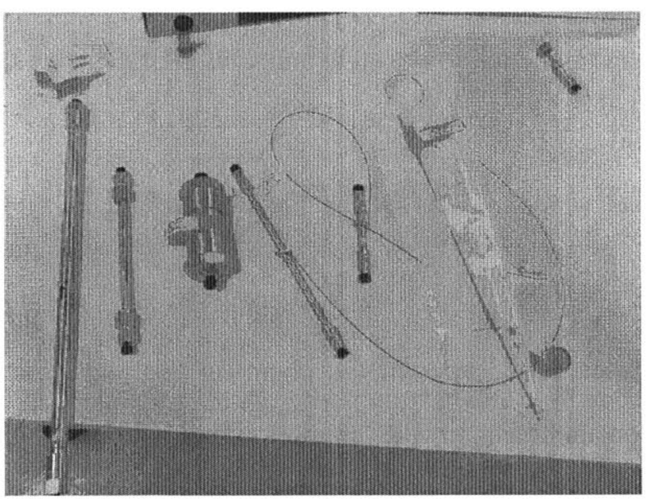

Figure 2.1: Analytical Columns in Different Sizes

In spite of the different sizes, the basic components of different columns are the same.

Generally, a finished column is manufactured at Waters Ireland facility by filling an assembly from Waters Milford facility with a chemical powder from Waters Taunton facility. The assembly is produced by assembling a set of machined parts at Waters Milford facility. Figure 2.2 is a picture showing the breakdown of an assembly with inside diameter of 2 mm.

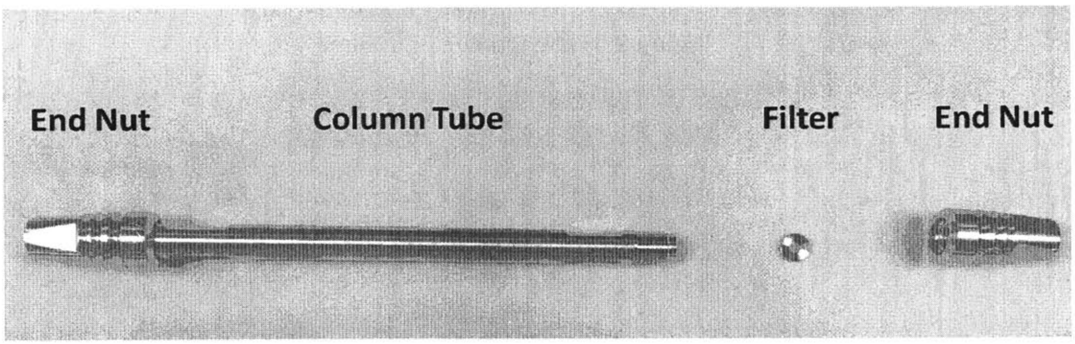

Figure 2.2: Breakdown of an Analytical Column Assembly

From the picture we can see that an assembly is assembled from three major components. The first one is called the column tube, which is manufactured at Waters Milford machine shop. The second one is called end nut. Currently Waters is using a hybrid sourcing strategy for the end nut, which means some of the end nuts are manufactured at Waters Milford machine shop, while some are outsourced to external suppliers. The third part is the filter, which is assembled at Waters Milford machine shop from lower level components, as shown in figure 2.3. For these lower level components, the frit and disk are totally purchased from external suppliers, while the filter house is partly self-manufactured and partly outsourced. A Column Assembly will have one column tube, two end nuts, and two filters.

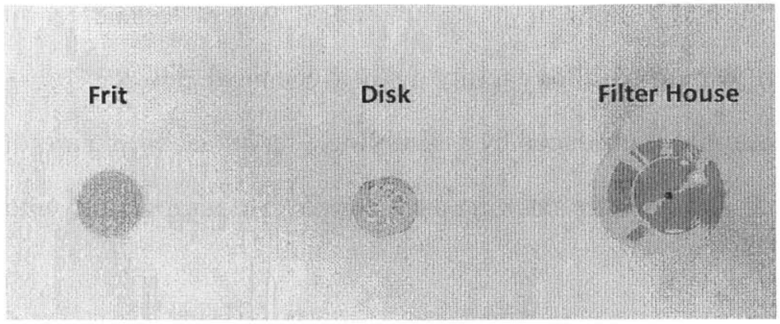

Figure 2.3: Breakdown of a Filter

2.3 Flow of Column Production Along the Supply Chain

After figuring out the detailed breakdown of a column, we can draw a supply chain map for the material flow of an Analytical Column, as shown in figure 2.4. This thesis is mainly focused on the supply chain for the Column Assembly division, which is from the component production in Milford to final sales representatives. We don't include Taunton facility, the chemical powder division, in our analysis.

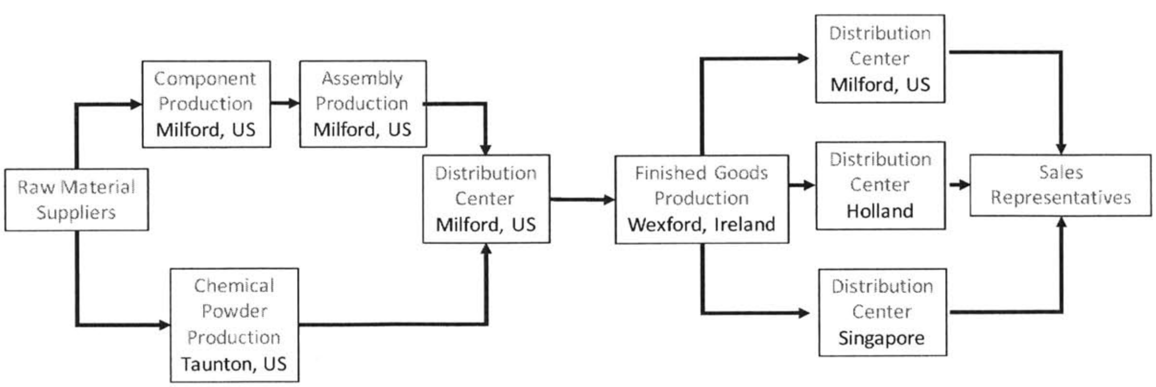

Figure 2.4: Supply Chain for Analytical Column

2.3.1 Waters Milford Facility

The Waters Milford facility, also called The Waters Advanced Manufacturing Center, houses a 50,000 square-foot Machining Center of Excellence which produces 2.7 million parts annually covering 28,000 SKUs, a 29,000 square-foot Advanced Instrument Assembly & Accessory Kitting Operation facility which produces over 130,000 finished goods assemblies, spare parts, and accessory kits, and an 8,500 square-foot Class 10,000 Clean Room for optics, micro valves, and critical parts. The project of interest focused on the operations of two departments, one is the Machining Center of Excellence (referred to as the "Machining Center" or "machining center"), which produces precision-machined metal components for the column assembly like column tube, end nut, and filter house. Another department we are interested in for this thesis is the column assembly line, which gets supplies from the machining center, and further assembles them into the column assemblies.

The Machining Center operates for 24 hours per day, 6 days per week, 52 weeks per year, producing 2.7 million parts annually covering 1500 unique SKUs. The Center is divided up into four main departments: NC Turning, NC Milling, Valve Cell, and Column Cell. Our focused area, The Column Cell, is laid out as a production cell, with a variety of different machines arranged next to each other that complete different operations on the same part.

2.3.2 Distribution Center

The distribution center at Milford plays a significant role in the whole supply chain. On one hand, it functions as a shipping department for the Waters Milford facility and Waters Taunton facility. Currently, these two facilities will send what are required by the

Ireland facility to the distribution center in Milford; then every Wednesday and Friday, the distribution center uses air package to ship all these components to Ireland facility for final assembly. On the other hand, this distribution center also functions as a warehouse for the finished goods. Sales representatives will directly place orders from customers on the distribution center. So the distribution center will pull products from its inventory to fulfill the order, and at the same time place orders on the Ireland facility to get replenishment for its inventory. Currently Waters has three distribution centers around the world, one in Milford, US, one in Netherlands, and the last one in Singapore. This thesis only focuses on the US distribution center because of the two roles it plays.

2.3.3 Flow of Demand and Materials Along the Supply Chain

Before diving into details, we first can have a general overview of the flow process between different facilities, shown in figure 2.5. The demands coming from customers are always placed on the distribution center. For build-to-stock products, the distribution center should have keep a certain amount of inventory in the warehouse so that once an order comes in, the distribution center can immediately pull the product from the warehouse, pack and ship it to the customer. The distribution center also can look at the demand forecast of 18 months into the future. Based on the demand forecast, the distribution center will generate a replenishment plan placed on the Ireland facility, to replenish its warehouse inventory into the future. These orders are then the forecasted demand for the Ireland facility. Based on this demand forecast, the Ireland facility will generate its own production plan and the corresponding demand orders for column assemblies placed on the Milford distribution center. In a similar way, the distribution center passes its demand forecast along to the Milford facility, which further develops its production plan.

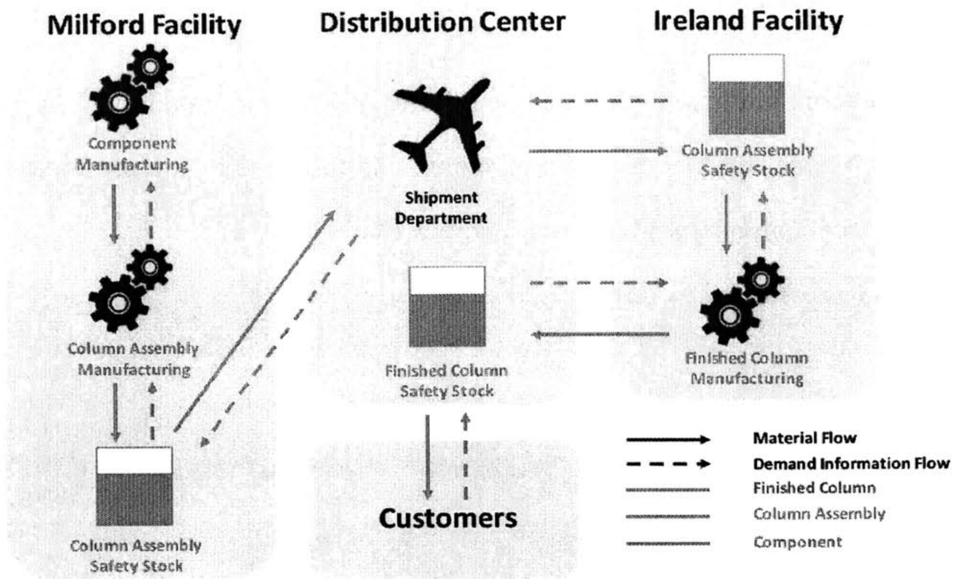

Figure 2.5: Overview of Demand Information and Material Flow Process

This flow of demand information and materials between different facilities are completely managed by Waters' Schedule and Planning (SAP) system. By entering different Material Recourses Planning (MRP) group's code, we can see the SAP managing panel for different facilities.

In each facility's managing panel, people can easily see the status of any material in this facility, including the current inventory level, current safety stock level, current and future production plan, current and future demand from downstream, and current and future demand placed on upstream.

When one facility makes some changes on its managing panel, like placing an order or issuing a new production plan, the changes will be automatically transferred to all the other facilities through related materials. In this way, the SAP system realizes the function of connecting different facilities together. And each facility only needs to focus on its own SAP managing panel as its guidance for production planning.

To give a more intuitive understanding of how does SAP system functions, figure 2.6 shows the SAP panel for US distribution center, and the material we are looking at is a finished analytical column, with material number 186006937.

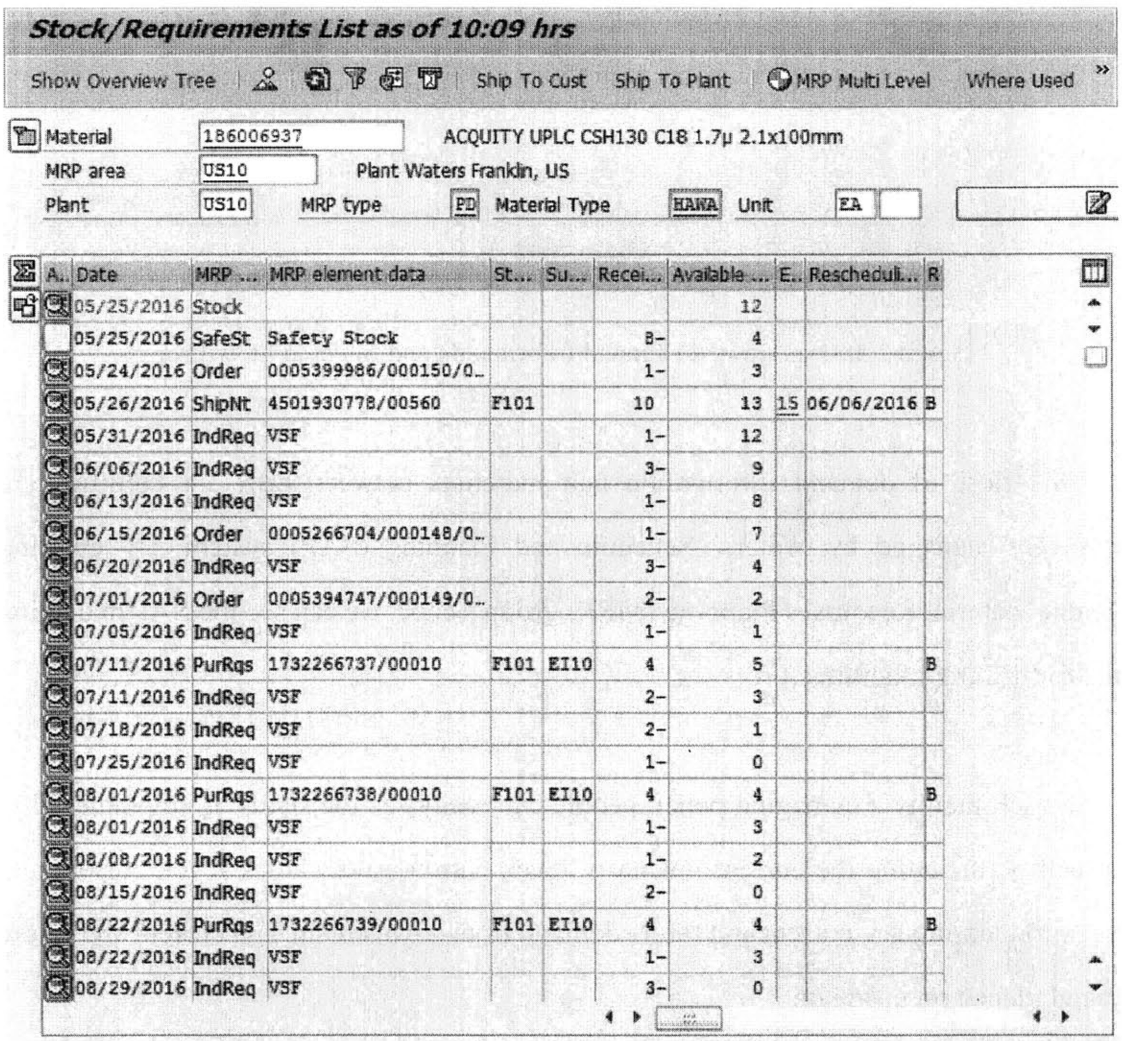
Stock/Requirements List as of 10:09 hrs

Show Overview Tree | Ship To Cust | Ship To Plant | MRP Multi Level | Where Used

Material 186006937 ACQUITY UPLC CSH130 C18 1.7µ 2.1x100mm
MRP area US10 Plant Waters Franklin, US
Plant US10 MRP type PD Material Type HAWA Unit EA

A..	Date	MRP ...	MRP element data	St...	Su...	Recei...	Available ...	E..	Rescheduli...	R
	05/25/2016	Stock					12			
	05/25/2016	SafeSt	Safety Stock			8-	4			
	05/24/2016	Order	0005399986/000150/0..			1-	3			
	05/26/2016	ShipNt	4501930778/00560	F101		10	13	15	06/06/2016	B
	05/31/2016	IndReq	VSF			1-	12			
	06/06/2016	IndReq	VSF			3-	9			
	06/13/2016	IndReq	VSF			1-	8			
	06/15/2016	Order	0005266704/000148/0..			1-	7			
	06/20/2016	IndReq	VSF			3-	4			
	07/01/2016	Order	0005394747/000149/0..			2-	2			
	07/05/2016	IndReq	VSF			1-	1			
	07/11/2016	PurRqs	1732266737/00010	F101	EI10	4	5			B
	07/11/2016	IndReq	VSF			2-	3			
	07/18/2016	IndReq	VSF			2-	1			
	07/25/2016	IndReq	VSF			1-	0			
	08/01/2016	PurRqs	1732266738/00010	F101	EI10	4	4			B
	08/01/2016	IndReq	VSF			1-	3			
	08/08/2016	IndReq	VSF			1-	2			
	08/15/2016	IndReq	VSF			2-	0			
	08/22/2016	PurRqs	1732266739/00010	F101	EI10	4	4			B
	08/22/2016	IndReq	VSF			1-	3			
	08/29/2016	IndReq	VSF			3-	0			

Figure 2.6: SAP Panel of US Distribution Center for a Finished Column

On the panel there are many rows and columns. Each row indicates a particular action related to this SKU. The "Date" column shows the due date for each action. The "MRP" column indicates which type of action it is. The "Receipt" column shows how many units are needed for this action, and the "Available" column shows the expected inventory level after each action.

In a SAP managing panel, the first row always shows the current date and current inventory level of this material in this facility, from which we know currently there are 12 units of this material in the US distribution center. The second row means in the US distribution center, the safety stock level for this material is 8 units. The third row marked with "Order" in "MRP" column means the distribution center currently has an order request from a customer requiring 1 unit of this SKU, which is due on May 24th. And the forth row marked with "ShipNt" in "MRP" column means currently there are 10 units of this SKU in shipment to the distribution center, and will be delivered on May 26th. The rows marked with "IndReq" are forecast demand for this SKU, each with an expected due date. Finally, the rows marked with "PurRqs" mean the future replenishments this distribution center is planning to get from Ireland for this SKU, based on the demand forecast to avoid stock out.

Those "PurRqs" transactions are automatically transferred to the upstream stage, the Ireland manufacturing facility. Figure 2.7 is the SAP panel of Ireland facility for the same SKU. This panel has many similar attributes. We note that for finished columns, the Ireland facility doesn't have safety stock on hand. All the "TrRes" rows indicate the forecasted demands from downstream distribution centers. We can see that by June 24^{th}, July 15^{th}, and August 5^{th}, Ireland facility each has a "TrRes" planned to ship to US distribution center. And referring to Figure 2.6 again, these shipments correspond to "PurRqs" rows in DC SAP panel with deadline as July 11^{th}, August 1^{st}, August 22^{nd}. This means the planned lead time for shipment from Ireland facility to US distribution center is 17 days. To satisfy these demands from distribution centers, SAP will generate a production plan marked as "PlOrd", and when the production plan begins, it will change status from "PlOrd" to "PrdOrd".

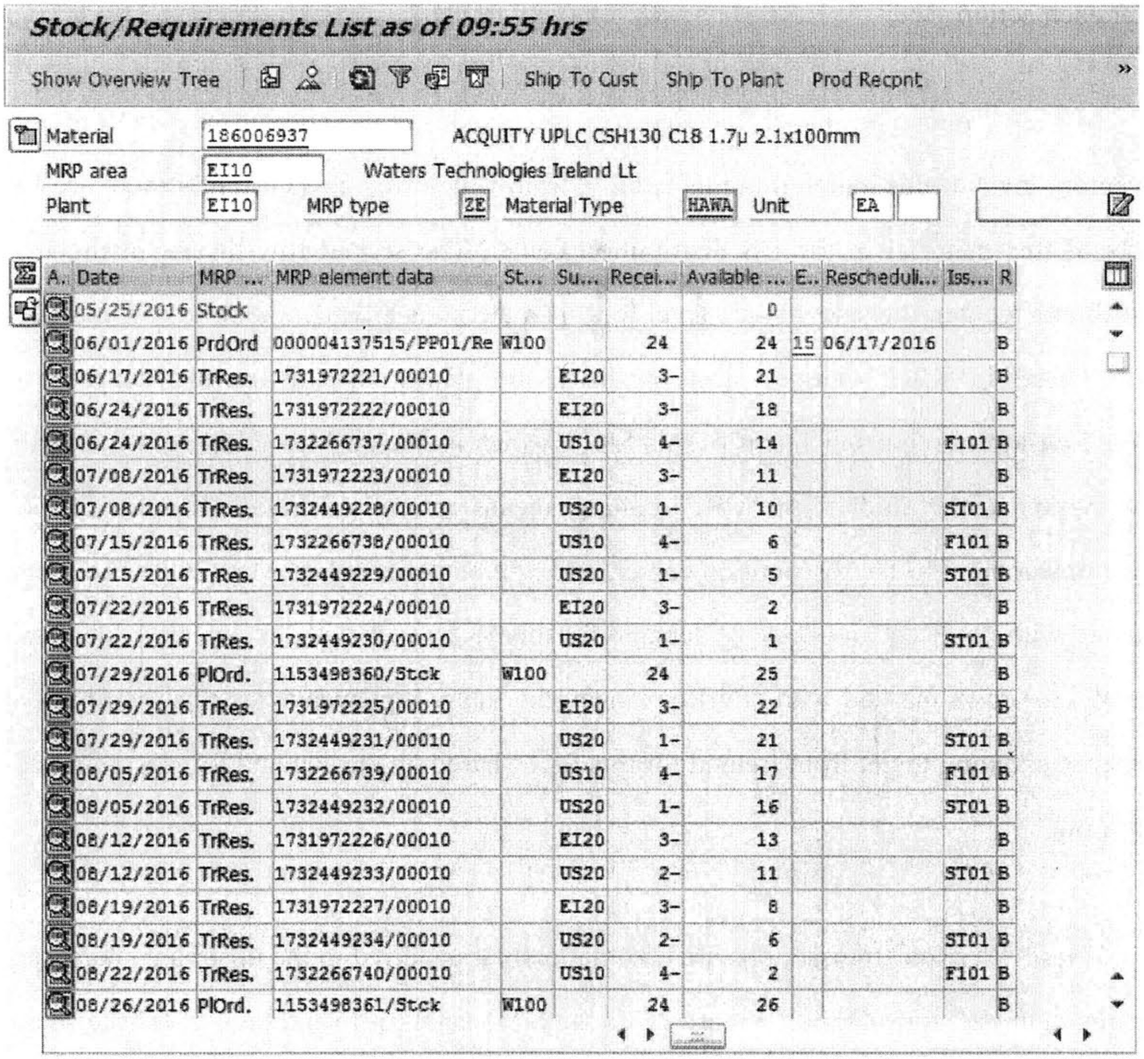

Stock/Requirements List as of 09:55 hrs

Show Overview Tree | Ship To Cust | Ship To Plant | Prod Recpnt

Material 186006937 ACQUITY UPLC CSH130 C18 1.7µ 2.1x100mm

MRP area EI10 Waters Technologies Ireland Lt

Plant EI10 MRP type ZE Material Type HAWA Unit EA

A.	Date	MRP ...	MRP element data	St...	Su...	Recei...	Available ...	E..	Rescheduli...	Iss...	R
	05/25/2016	Stock					0				
	06/01/2016	PrdOrd	000004137515/PP01/Re	W100		24	24	15	06/17/2016		B
	06/17/2016	TrRes.	1731972221/00010		EI20	3-	21				B
	06/24/2016	TrRes.	1731972222/00010		EI20	3-	18				B
	06/24/2016	TrRes.	1732266737/00010		US10	4-	14			F101	B
	07/08/2016	TrRes.	1731972223/00010		EI20	3-	11				B
	07/08/2016	TrRes.	1732449228/00010		US20	1-	10			ST01	B
	07/15/2016	TrRes.	1732266738/00010		US10	4-	6			F101	B
	07/15/2016	TrRes.	1732449229/00010		US20	1-	5			ST01	B
	07/22/2016	TrRes.	1731972224/00010		EI20	3-	2				B
	07/22/2016	TrRes.	1732449230/00010		US20	1-	1			ST01	B
	07/29/2016	PlOrd.	1153498360/Stck	W100		24	25				B
	07/29/2016	TrRes.	1731972225/00010		EI20	3-	22				B
	07/29/2016	TrRes.	1732449231/00010		US20	1-	21			ST01	B
	08/05/2016	TrRes.	1732266739/00010		US10	4-	17			F101	B
	08/05/2016	TrRes.	1732449232/00010		US20	1-	16			ST01	B
	08/12/2016	TrRes.	1731972226/00010		EI20	3-	13				B
	08/12/2016	TrRes.	1732449233/00010		US20	2-	11			ST01	B
	08/19/2016	TrRes.	1731972227/00010		EI20	3-	8				B
	08/19/2016	TrRes.	1732449234/00010		US20	2-	6			ST01	B
	08/22/2016	TrRes.	1732266740/00010		US10	4-	2			F101	B
	08/26/2016	PlOrd.	1153498361/Stck	W100		24	26				B

Figure 2.7: SAP Panel of Ireland Facility for a Finished Column

The Ireland facility also holds inventories for Column Assemblies, so once they begin a production plan, they can directly pull Column Assemblies out from the warehouse. The warehouse will get replenishment of Column Assemblies from the Milford facility. Ireland facility's SAP panel of Column Assemblies clearly shows the transactions in Figure 2.8.

Stock/Requirements List as of 09:13 hrs

Show Overview Tree | Ship To Cust | Ship To Plant | Prod Recpnt

Material: 289001874 — ASSY, COLUMN, 2.1 X 100MM, HP
MRP area: EI10 — Waters Technologies Ireland Lt
Plant: EI10 — MRP type: P3 — Material Type: HAWA — Unit: EA

A.	Date	MRP ...	MRP element data	St...	Su...	Recei...	Available ...	E..	Rescheduli...	R
	05/25/2016	Stock					1,512	96		
	05/25/2016	SafeSt	Safety Stock			2,000-	488-			
	04/05/2016	DepReq	186002878	W200		96-	584-			M
	04/26/2016	DepReq	186008381	W200		1-	585-			M
	04/26/2016	DepReq	186003533	W200		96-	681-			M
	05/09/2016	OrdRes	186004496	W200		3-	684-			M
	05/09/2016	OrdRes	186007095	W200		8-	692-			M
	05/09/2016	DepReq	186007095	W200		56-	748-			N
	05/10/2016	OrdRes	186005297	W200		16-	764-			M
	05/25/2016	OrdRes	186006937	W200		24-	1,576-			N
	05/25/2016	OrdRes	289003434	W200		32-	1,608-			N
	05/25/2016	OrdRes	186005407	W200		96-	1,704-			M
	05/25/2016	OrdRes	186005407	W200		2-	1,706-			N
	05/25/2016	DepReq	186002352	W200		96-	1,802-			N
	05/25/2016	DepReq	186002352	W200		96-	1,898-			N
	05/26/2016	DepReq	186008381	W200		2-	1,900-			N
	05/30/2016	DepReq	186003539	W200		96-	1,996-			N
	05/31/2016	DepReq	186008381	W200		2-	1,998-			N
	05/31/2016	DepReq	186002352	W200		96-	2,094-			N
	05/31/2016	DepReq	186002352	W200		96-	2,190-			N
	05/31/2016	DepReq	186002878	W200		96-	2,286-			N
	05/31/2016	DepReq	186005988	W200		20-	2,306-			N
	06/01/2016	SchLne	4501922368/00020	W200		720	1,586-	10	05/25/2016	M
	06/03/2016	DepReq	186002854	W200		96-	1,682-			N
	06/03/2016	DepReq	186007095	W200		56-	1,738-			N
	06/03/2016	DepReq	186003539	W200		96-	1,834-			N
	06/08/2016	SchLne	4501927611/00020	W200		720	1,114-	10	05/09/2016	M
	06/08/2016	DepReq	186002352	W200		96-	1,210-			N

Figure 2.8: SAP Panel of Ireland Facility for a Column Assembly

In this panel, all the "OrdRes" rows are demand requirements corresponding to the "PrdOrd" rows in the Finished Column SAP panel, while all the "DepReq" rows are expected demand requirements corresponding to the "PlOrd" rows in the Finished Column SAP panel. The "SchLne" rows are planned replenishment to get from the US distribution center.

As mentioned above, the Column Assemblies are also shipped to Ireland facility from the US distribution center. So actually Ireland facility's replenishment requirements are transferred to the US distribution center's SAP panel for Column Assemblies, as shown in Figure 2.9.

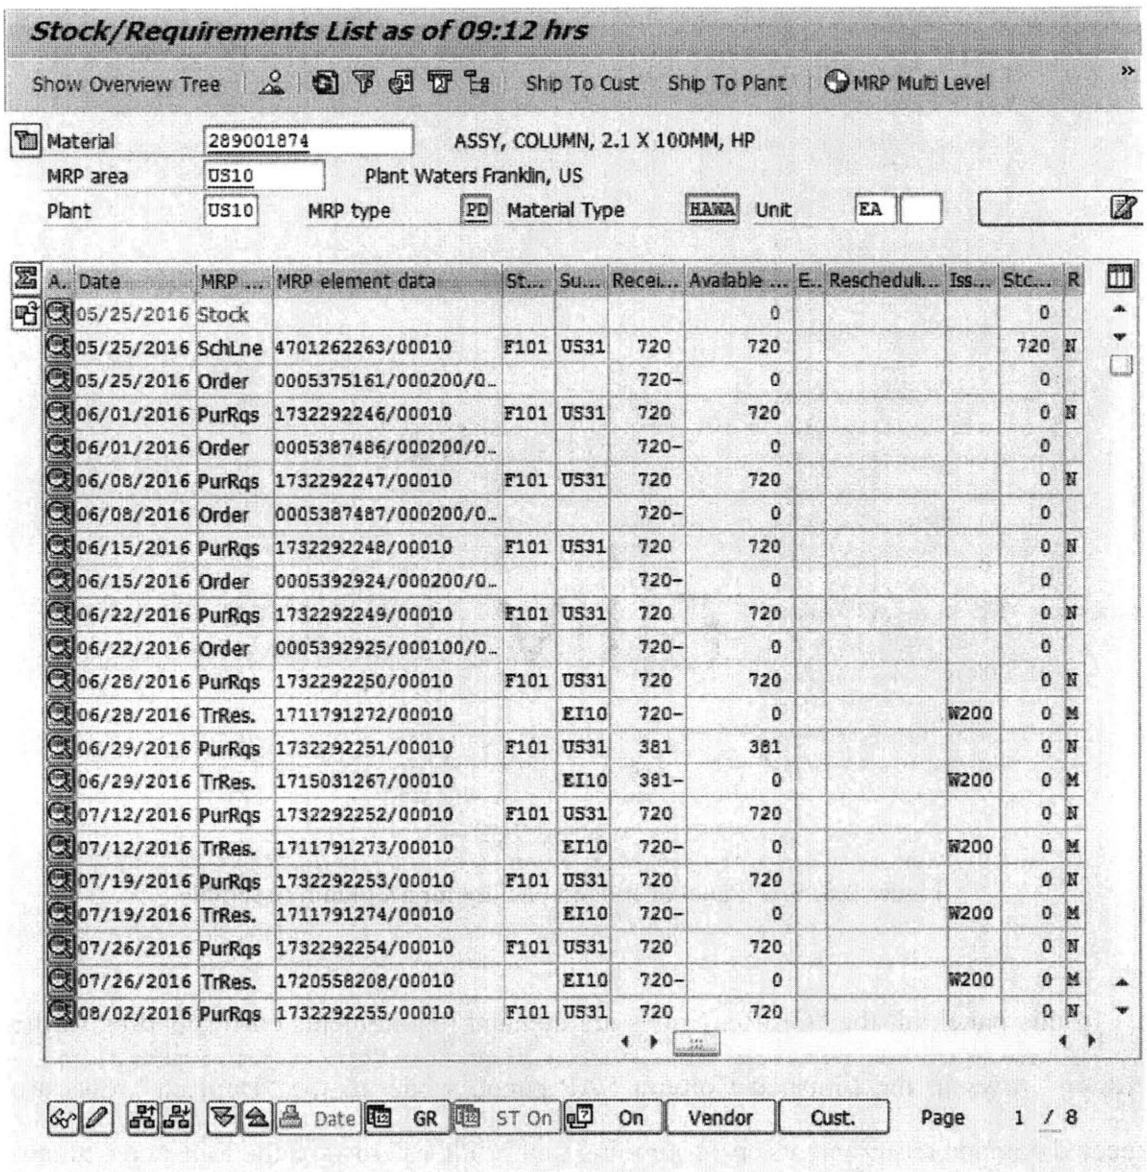

A..	Date	MRP ...	MRP element data	St...	Su...	Recei...	Available ...	E..	Rescheduli...	Iss...	Stc...	R
	05/25/2016	Stock					0				0	
	05/25/2016	SchLne	4701262263/00010	F101	US31	720	720				720	N
	05/25/2016	Order	0005375161/000200/0..			720-	0				0	
	06/01/2016	PurRqs	1732292246/00010	F101	US31	720	720				0	N
	06/01/2016	Order	0005387486/000200/0..			720-	0				0	
	06/08/2016	PurRqs	1732292247/00010	F101	US31	720	720				0	N
	06/08/2016	Order	0005387487/000200/0..			720-	0				0	
	06/15/2016	PurRqs	1732292248/00010	F101	US31	720	720				0	N
	06/15/2016	Order	0005392924/000200/0..			720-	0				0	
	06/22/2016	PurRqs	1732292249/00010	F101	US31	720	720				0	N
	06/22/2016	Order	0005392925/000100/0..			720-	0				0	
	06/28/2016	PurRqs	1732292250/00010	F101	US31	720	720				0	N
	06/28/2016	TrRes.	1711791272/00010		EI10	720-	0			W200	0	M
	06/29/2016	PurRqs	1732292251/00010	F101	US31	381	381				0	N
	06/29/2016	TrRes.	1715031267/00010		EI10	381-	0			W200	0	M
	07/12/2016	PurRqs	1732292252/00010	F101	US31	720	720				0	N
	07/12/2016	TrRes.	1711791273/00010		EI10	720-	0			W200	0	M
	07/19/2016	PurRqs	1732292253/00010	F101	US31	720	720				0	N
	07/19/2016	TrRes.	1711791274/00010		EI10	720-	0			W200	0	M
	07/26/2016	PurRqs	1732292254/00010	F101	US31	720	720				0	N
	07/26/2016	TrRes.	1720558208/00010		EI10	720-	0			W200	0	M
	08/02/2016	PurRqs	1732292255/00010	F101	US31	720	720				0	N

Figure 2.9: SAP Panel of US Distribution Center for a Column Assembly

We can see that the third row called "Order" shows that the DC has received a confirmation from Ireland that they should ship out 720 assemblies to Ireland by May 25th. Given 6 days' lead time, this matches with Ireland panel's "SchLne" row, which will have 720 units delivered by June 1st. From talking with people in warehouse, we know that currently the DC in US ship assemblies to Ireland on every Wednesday and Friday. So the "SchLne" row will keep in this state until Wednesday or Friday, and once it is shipped, it will change from "SchLne" into "Deliv", meaning it is now in transit . Rows called "TrRes" are planed shipments of this assemblies to Ireland, which have not been confirmed by Ireland and once confirmed they will become "Order". However, we know that distribution center doesn't produce anything, which means to ship these assemblies to Ireland, the DC first needs to get these assemblies from Milford facility which manufactures the assemblies. The second row marked with "SchLne" means Milford already has the 720 units ready for the DC to ship. The other rows called "PurRqs" mean that the DC plans to receive things from Milford facility, but the Milford facility doesn't have products ready for shipping yet.

Next the information is transferred to the Milford facility, as shown in Figure 2.10. Corresponding to the "PurRqs" rows in distribution center, the Milford facility has "TrRes" rows with a 2-day lead time. So the "TrRes" rows due by May 31st and June 7th are planned to arrive at US distribution center and get shipped on June 1st and June 8th. Assemblies are made from machined parts inside Milford facility, so Milford also has plans to replenish the assemblies, shown as the "PlOrd" row. And once the Milford facility decides that they do need the replenishment, they will confirm the plan to make this "PlOrd" into a task called "PrdOrd". Then the machine shop will perform this task to replenish the inventory of this assembly in Milford.

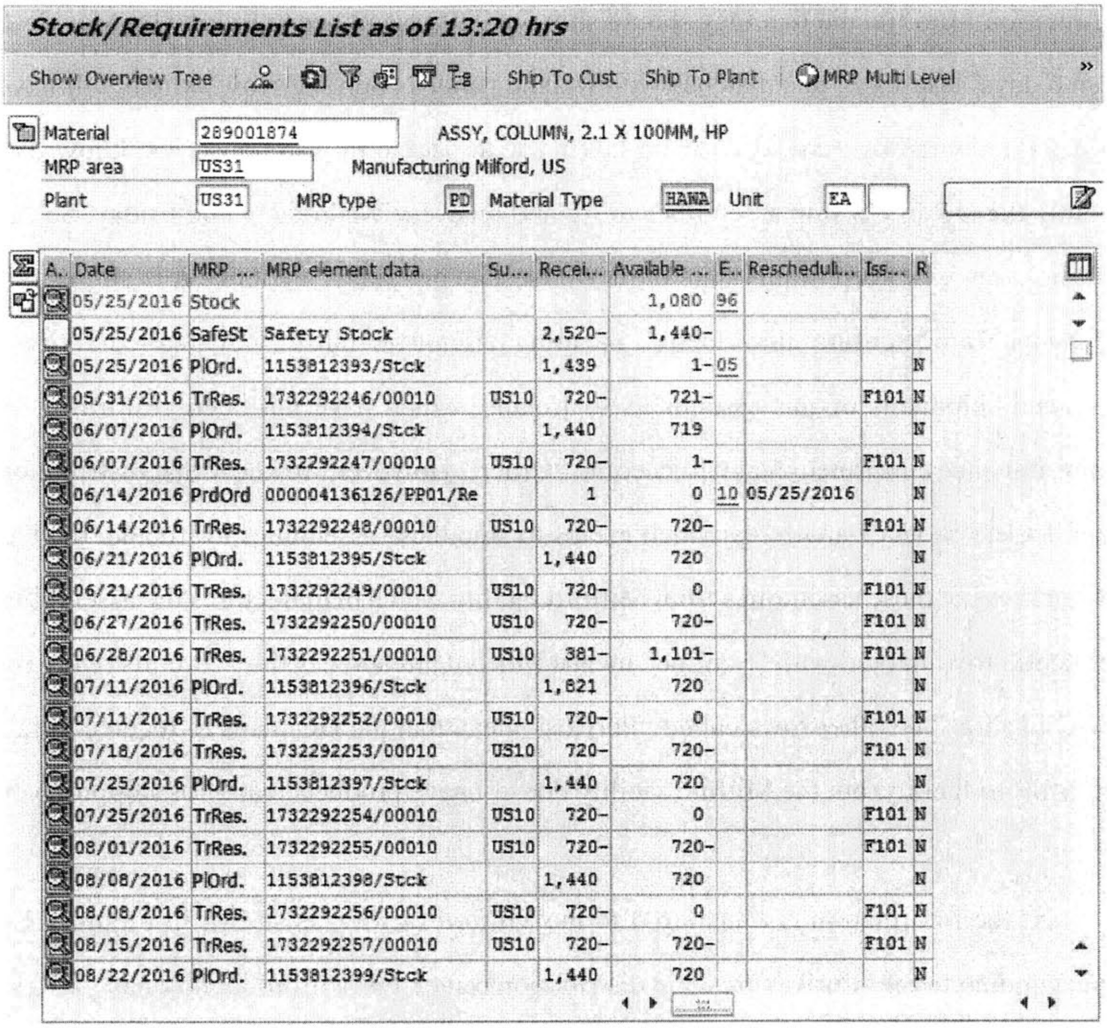

Stock/Requirements List as of 13:20 hrs

Show Overview Tree | Ship To Cust | Ship To Plant | MRP Multi Level

Material 289001874 ASSY, COLUMN, 2.1 X 100MM, HP
MRP area US31 Manufacturing Milford, US
Plant US31 MRP type PD Material Type HAWA Unit EA

A..	Date	MRP ...	MRP element data	Su...	Recei...	Available ...	E..	Rescheduli...	Iss...	R
	05/25/2016	Stock				1,080	96			
	05/25/2016	SafeSt	Safety Stock		2,520-	1,440-				
	05/25/2016	PlOrd.	1153812393/Stck		1,439	1-	05			N
	05/31/2016	TrRes.	1732292246/00010	US10	720-	721-			F101	N
	06/07/2016	PlOrd.	1153812394/Stck		1,440	719				N
	06/07/2016	TrRes.	1732292247/00010	US10	720-	1-			F101	N
	06/14/2016	PrdOrd	000004136126/PP01/Re		1	0	10	05/25/2016		N
	06/14/2016	TrRes.	1732292248/00010	US10	720-	720-			F101	N
	06/21/2016	PlOrd.	1153812395/Stck		1,440	720				N
	06/21/2016	TrRes.	1732292249/00010	US10	720-	0			F101	N
	06/27/2016	TrRes.	1732292250/00010	US10	720-	720-			F101	N
	06/28/2016	TrRes.	1732292251/00010	US10	381-	1,101-			F101	N
	07/11/2016	PlOrd.	1153812396/Stck		1,821	720				N
	07/11/2016	TrRes.	1732292252/00010	US10	720-	0			F101	N
	07/18/2016	TrRes.	1732292253/00010	US10	720-	720-			F101	N
	07/25/2016	PlOrd.	1153812397/Stck		1,440	720				N
	07/25/2016	TrRes.	1732292254/00010	US10	720-	0			F101	N
	08/01/2016	TrRes.	1732292255/00010	US10	720-	720-			F101	N
	08/08/2016	PlOrd.	1153812398/Stck		1,440	720				N
	08/08/2016	TrRes.	1732292256/00010	US10	720-	0			F101	N
	08/15/2016	TrRes.	1732292257/00010	US10	720-	720-			F101	N
	08/22/2016	PlOrd.	1153812399/Stck		1,440	720				N

Figure 2.10: SAP Panel of Milford Facility for a Column Assembly

The upstream stage of this whole supply chain is the manufacturing of components in the Milford facility. The SAP panel for components are shown in Figure 2.11. For most of the components, the Milford facility doesn't have safety stock. And the panel shows an 8-10 days' lead time for Components to be assembled into Column Assemblies. For example, in Figure 2.10, the Milford facility plans to produce 1439 units of Column Assemblies by May 25th, while in Figure 2.11, 1439 units of required Components will be produced and ready for being assembled by May 17th.

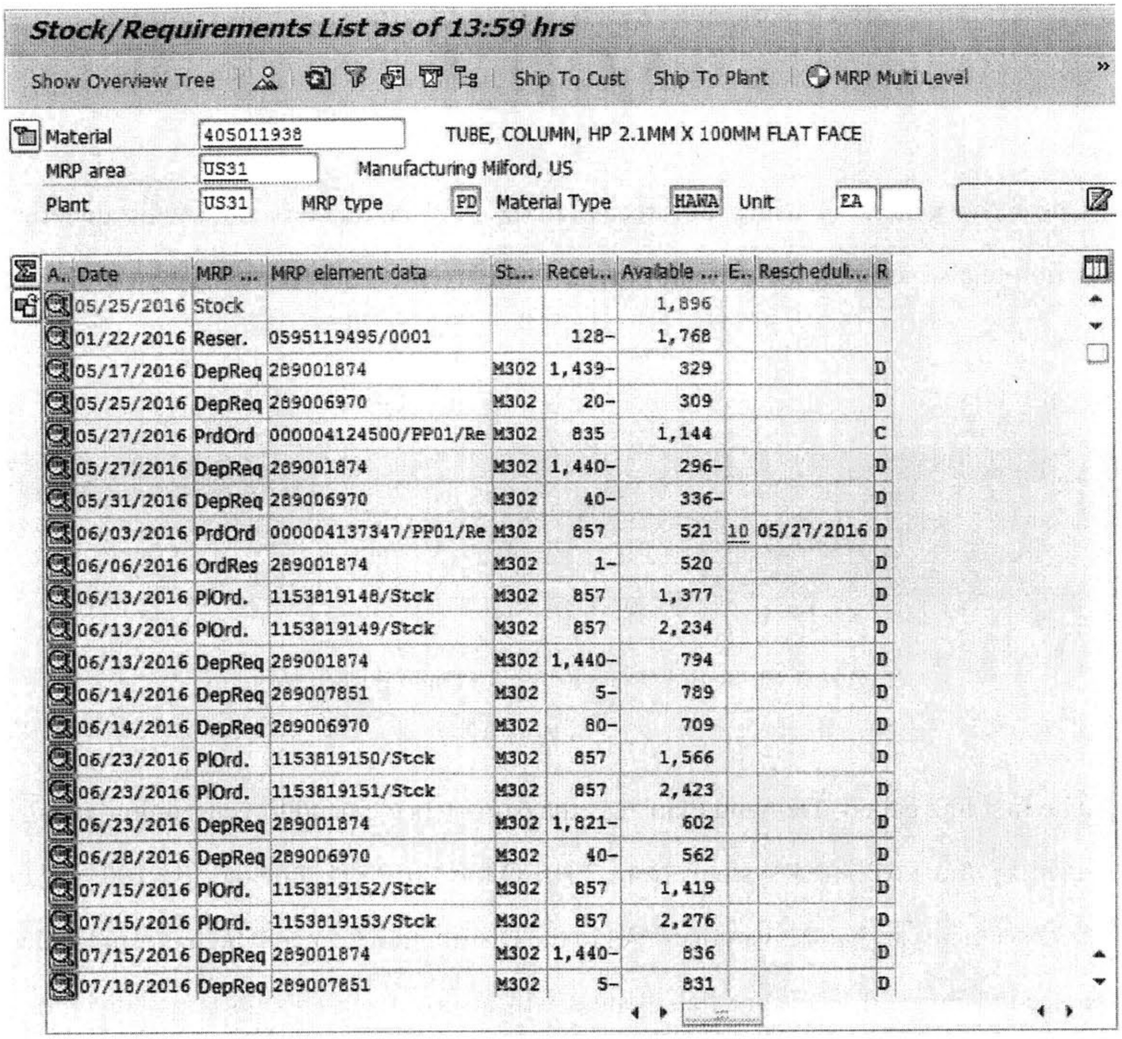

Stock/Requirements List as of 13:59 hrs

Show Overview Tree | Ship To Cust | Ship To Plant | MRP Multi Level

Material: 405011938 — TUBE, COLUMN, HP 2.1MM X 100MM FLAT FACE
MRP area: US31 — Manufacturing Milford, US
Plant: US31 | MRP type: PD | Material Type: HAWA | Unit: EA

A..	Date	MRP ...	MRP element data	St...	Recei...	Available ...	E..	Rescheduli...	R
	05/25/2016	Stock				1,896			
	01/22/2016	Reser.	0595119495/0001		128-	1,768			
	05/17/2016	DepReq	289001874	M302	1,439-	329			D
	05/25/2016	DepReq	289006970	M302	20-	309			D
	05/27/2016	PrdOrd	000004124500/PP01/Re	M302	835	1,144			C
	05/27/2016	DepReq	289001874	M302	1,440-	296-			D
	05/31/2016	DepReq	289006970	M302	40-	336-			D
	06/03/2016	PrdOrd	000004137347/PP01/Re	M302	857	521	10	05/27/2016	D
	06/06/2016	OrdRes	289001874	M302	1-	520			D
	06/13/2016	PlOrd.	1153819148/Stck	M302	857	1,377			D
	06/13/2016	PlOrd.	1153819149/Stck	M302	857	2,234			D
	06/13/2016	DepReq	289001874	M302	1,440-	794			D
	06/14/2016	DepReq	289007851	M302	5-	789			D
	06/14/2016	DepReq	289006970	M302	80-	709			D
	06/23/2016	PlOrd.	1153819150/Stck	M302	857	1,566			D
	06/23/2016	PlOrd.	1153819151/Stck	M302	857	2,423			D
	06/23/2016	DepReq	289001874	M302	1,821-	602			D
	06/28/2016	DepReq	289006970	M302	40-	562			D
	07/15/2016	PlOrd.	1153819152/Stck	M302	857	1,419			D
	07/15/2016	PlOrd.	1153819153/Stck	M302	857	2,276			D
	07/15/2016	DepReq	289001874	M302	1,440-	836			D
	07/18/2016	DepReq	289007851	M302	5-	831			D

Figure 2.11: SAP Panel of Milford Facility for a Component

2.4 Forecasting Mechanism

Currently Waters is using a forecast horizon of 6 quarters, and will generate production plans based on the demand forecast and inventory targets, as shown in Figure 2.12.

PN	Plan Type	Q3 2016	Q4 2016	Q1-17	Q2-17	Q3-17	Q4-17
186000494	Demand Plan	377	377	377	377	374	378
186000494	Planned Ending INV	226	233	240	247	257	263
186000494	Proposed Production	384	384	384	384	384	384

Figure 2.12: Sample Forecasting for One Finished Column

The first line called "Demand Plan" has the 6 quarters' demand forecast automatically generated by the SAP APO system, from 2016 Q3 to 2017 Q4. For the top 100 SKUs, Waters has a target on-hand inventory level in the Distribution Center, which is represented in the second line. So to ensure that Waters will always have enough target inventory, combining demand forecast and current inventory together, the planner at DC will generate a proposed production plan for each quarter, which is shown in third line.

Then this proposed production plan is sent to the Ireland facility, as their corresponding forecasted demand for the following 6 quarters. However, every month the planner at DC will make some revisions for the current quarter's plan, based on the past months' sales. For instance, suppose at at the beginning of this quarter, the planner planned that Ireland should produce 300 units of a SKU. Then after one month at Week 5, the planner will observe the past month's demand. If the past month's demand is 150, then he may make some changes to this quarter's production plan; for example, he might increase the plan from 300 units to 400 units. Once Ireland facility get the update, they will update the production plan in system by increasing 100 units (400-300). So we can simulate a simple procedure shown in Table 2.1 below.

Table 2.1: Simulation of Production Plan in Ireland Facility

	Actions	Change in amount due	Amount due
Week 1	Production Plan	300	300
Week 1	1 batch finished	-24	276
Week 2	1 batch finished	-24	252
Week 3	1 batch finished	-24	228
Week 4	1 batch finished	-24	204
Week 5	Revision made	100	304
Week 5	2 batch finished	-48	256
Week 6	2 batch finished	-48	208
Week 7	2 batch finished	-48	160
Week 8	2 batch finished	-48	112
Week 9	Revision made	0	112
Week 10	1 batch finished	-24	88
Week 11	1 batch finished	-24	64
Week 12	1 batch finished	-24	40
Week 13	2 batch finished	-48	-8
Week 14	Production Plan	300	292

At week 1, the Ireland facility sees the initial forecast production plan and produces accordingly. At week 5, the planner at the DC decides that we may need 100 more for this quarter. So 100 is added to the Ireland facility's amount due. Then based on this new revised production plan, the Ireland facility may increase its production from 1 batch per week to 2 batches per week (Assuming a batch has 24 units). At week 9, the planner at DC changes nothing. So for the last months, the Ireland facility produces according to its original plan. And finally, since the Ireland facility only produces in integer batches, the extra 8 units of products will go to next quarter's credits.

2.5 Issues with Planning and Forecasting System

Sections 2.2-2.5 present the results of a detailed supply chain analysis, including the product breakdown, product features, material planning, and demand forecasting. This analysis is based on the investigation of SAP system and from conducting interviews with each stakeholder along the supply chain. After diving deep into details about the results, some areas in planning and forecasting system with potential opportunities for improvement become clear.

2.5.1 Ambiguous Safety Stock Policy

Waters currently doesn't have a data-driven standard method to determine its safety stock policy. This problem arises in two ways.

Firstly, from the overview of supply chain, it's obvious to see that Waters doesn't have a consensus of where to put safety stock in the supply chain. Each facility decides whether to hold safety stock and how much to hold based on different incentives. For facilities seeking for a "Lean Production", they don't have safety stock for their products, while facilities focused on "On Time Delivery Rate" will hold many safety stocks on hand. This leads to potential redundant safety stock in the supply chain as well as material starvation in some places. For example, in Figure 2.5, there are no safety stocks for Components and Finished Columns in manufacturing facilities, while Column Assemblies have safety stock in two places, in Milford and in Wexford.

Secondly, through talking with managers and planners in different facilities, we discovered that there is not a formal calculation for determining the amount of safety stock.

When it comes to determining safety stock, there are generally two methods adopted. The first one is based on intuition and past experience. For the top 100 SKUs, planners will determine the safety stock level to be 4 weeks of demands. The second one is based on performance. If a product doesn't meet the on time delivery requirements, then the supervisors will increase the safety stock level. If it performs better next time, then this safety stock level will be maintained until next time when starvation happens.

There are two factors contributing to this problem. The first one is lack of global level's cooperation and optimization. Each facility in the supply chain is doing optimization locally, without considering the effects on other facilities. The second reason is people don't have a correct understanding of safety stock and its function.

2.5.2 Poor Planning Procedure

Through interviews with planners and deeply diving into SAP system to track particular orders, we find that there is no standardized procedure for planning in Waters, which means it is possible for arbitrary changes to the planning parameters. Firstly, for lead time there is an apparent discrepancy between reality and the SAP system. In reality, it takes 3 days to ship Column Assemblies from US distribution center to Ireland facility. It takes 1 week for Ireland facility to finish a production order. And finally it takes 5 days to ship Finished Columns from Ireland to US distribution center. However, in the SAP system, these numbers are determined and input manually, and each facility's planners have the authority to change them. People tend to increase these numbers so that they may have enough time as a buffer against late deliveries. So in SAP system, these numbers are 7 days for Colum Assemblies to be shipped to Ireland, 7-15 days for Finished Columns to be manufactured, and 15-18 days for Finished Columns to be shipped from Ireland to distribution center. These are shown in the Figure 2.13 below.

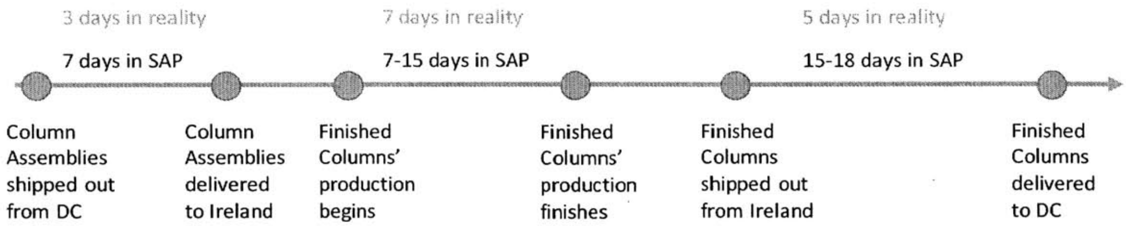

Figure 2.13: Illustration of Discrepancy in Lead Time between Reality and SAP System

This discrepancy is always adjusted manually by distribution center's planners. The planners are in charge of all the replenishment in distribution centers, and they usually look only 2 weeks ahead. Every day the planners will check based on SAP planning and forecasting, what finished goods should be delivered to the distribution center 2 weeks later and then confirm it. When they are confirming the planned orders, they will manually change the lead time from 15-18 days to 7 days. Here is an example to explain this procedure.

Suppose that based on SAP planning and forecasting, there should be 100 units of Finished Column "A" delivered to distribution center by July 30th as replenishment. Then based on the system's 17-day lead time here, according to the plan the Ireland facility should ship out the 100 units Finished Columns by July 13th. This means they need to finish the production even before July 13th. But on July 13th, when a planner logs into the SAP system and looks 2 weeks ahead, he will only see what is due upto July 27th. So he won't confirm the shipment of 100 units of Finished Columns "A" until July 16th, when the SAP system will signify that this is due 2 weeks later. And once he confirmed the order, he will change the expected lead time from 17 to 7, which means now the expected time of delivery is changed to July 23rd, one week earlier than forecasted. Although the planner finally changes expected delivery date to 7 days, the manufacturing department has already been producing finished goods according to original plan. This discrepancy leads to a great increase in pipeline inventories, as shown in Figure 2.14.

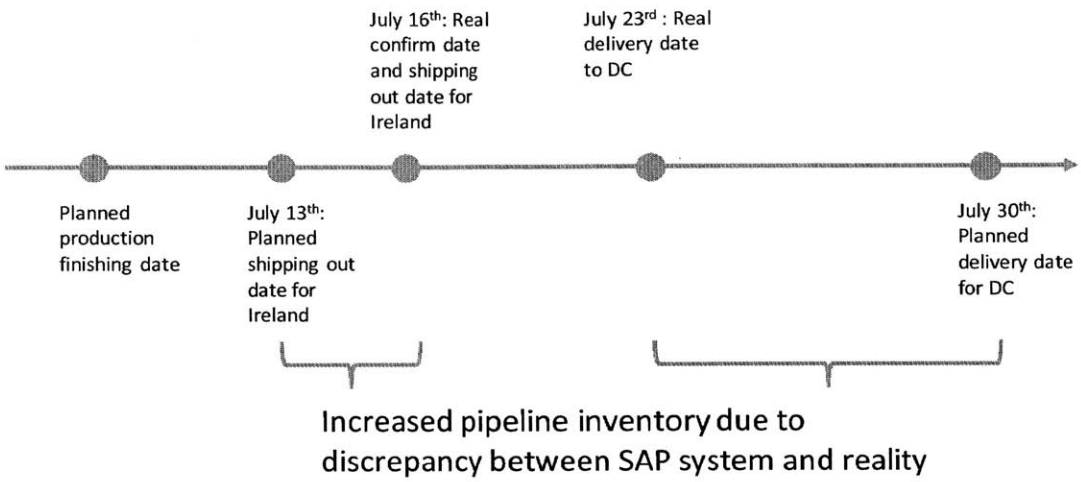

Figure 2.14: Illustration of Increased Pipeline Inventory

This kind of discrepancy also exists in the planning at Milford facility. Another result of this discrepancy is that sometimes planners in manufacturing facilities will intentionally postpone the production of some orders, because they think these orders are not necessary currently. So this heavily involved manual intervention and high level ambiguity will cause wastes in the supply chain, misunderstandings between facilities, and sometimes even late deliveries. This is also a major reason why Waters can hardly have a standardized policy and global optimization in supply chain.

2.6 Statement of Purpose

According the problems identified above, the objectives of this project can be framed as developing a new supply chain strategy to:

- Reduce waste to make the supply chain "lean".
- Have a more competitive response time to the market and customers.

The methods adopted here are:

- Determining the optimal location of inventory across different stages of the supply chain and the optimal safety stock level at each stage.
- Standardizing the current material planning process to reduce lead time.

3 Literature Review

This chapter mainly describes concepts related to inventory management, including the significance of inventory management, basic principles and terminology in inventory management, and explanations of the multi-echelon inventory model which will be used in Chapter 4 and 5 as the major methodology for developing an optimal safety stock strategy. References for this chapter are heavily based on the work done by Simchi-Levi, Kaminsky, S.Graves, T.Schoenmeyr and S.Willems.

3.1 Inventory Management

In a manufacturing company setting, inventory is often one of the dominant costs, since a lot of capital can be tied to the inventory on hand. Generally, the goal for effective inventory management in the supply chain is to have the correct amount of inventory at the right place at the right time to minimize system costs while satisfying customer service requirements.

The reasons for companies holding inventories on hand are:

1) Uncertainty in customers' demands which happens in two ways. The first one is demand from increasing number of new customers, as a result of sales team's promotion efforts and the expansion of total market size. The second one is changes in demand from old customers.
2) Uncertainty in suppliers' quality, quantity, and delivery times.
3) Manufacturing lead time requires companies to hold inventory to achieve a competitive response time to market.
4) Economies of scale in manufacturing and transportation will motivate companies to

purchase or produce large quantities every time.

One way in supply chain management is to construct supply chain models and make analysis. The parameters usually considered in constructing an effective supply chain model are:

1) Scope and granularity. Supply chain is always complex and large for modern manufacturing companies, and it's often not practical to manage the whole supply chain, so scope defines what parts in this supply chain should be focused on. Granularity defines the level of detail and number of SKUs people should look into to represent the supply chain well. It's also not practical for companies with thousands of SKUs to construct supply chain models for each SKU. So it's critical to pick out a manageable number of representative SKUs for analysis, while still ensuring the results are applicable for all SKUs.
2) Customer demand. The characteristics of customers demand directly decide what kind of forecasting tools to use, and what level of uncertainties to deal with.
3) Standard lead time of each facility, which represents the time from an order being placed to the order being shipped out.
4) Costs. There are generally three types of costs in a supply chain setting. The first one is the product cost including the material cost and all the costs needed to manufacture this product. The second one is the transportation cost, which is more important in a global supply chain setting. And the final one is the inventory holding cost, including taxes and insurance on inventories, maintenance costs, obsolescence costs derived from the risk that a product lose its value as market changes, and opportunity costs representing the return on investment if the capital were invested in something else.
5) Target service level. Each company will have its target service level to customers depending on the company's strategy.

It's easy for companies to take safety stock simply as inventory, while in reality inventory have many different categories and safety stock is only one of them. The general classification of inventory that has been useful in practice has the following different types:

1) Anticipatory Stock. It is inventory held on hand for demands derived from a one-time event, which means the demands will either be captured or lost forever. For example, by the end of a year, many companies usually have budget clear out so they will procure large amount of products. This kind of one-time demand greatly increases the uncertainties faced by suppliers, and to suppliers these demands will never be regained once they are lost. Based on this situation, suppliers need to make decisions of whether or not to hold inventory on hand for this kind of demand. The inventory held to meet this one-time large demand is anticipatory stock.
2) Cycle Stock. Cycle stock is the type of inventory as the result of periodic patterns for supply chain operations. In real business world, the supply chain cannot operate continuously in time. For example, there are certain dates in a week that transportation will ship out products, while before these dates some products will just pile up in warehouse. Also, products are always produced in batches to meet demand for a certain period. So there will always be an increase in inventory right after a manufacturing process or transportation inbound, and then the inventory level gradually goes down until next peak shows up.
3) Pipeline Stock. Aside from the inventory stored in warehouse, there is also a significant amount of inventory moving along the supply chain as a result of the lead time between different stages or even within a stage in the supply chain, called pipeline stock. In a broader definition, pipeline stock represents all the inventory moving along the supply chain, including moving under transportation

process and moving through production process. For a company with global supply chain, the pipeline stock is more important due to the longer lead time in transportation.

4) Safety Stock. Safety stock is the type of inventory used as a buffer against uncertainties in the supply chain, including customer demands' uncertainty, manufacturing uncertainty like machine failures, and supply uncertainty like suppliers' disruption.

5) Strategic Stock. It is a type of inventory purchased in large quantities and ahead of real needs, because of some potential benefits and discounts provided by suppliers.

3.2 Multi-echelon Inventory Model

Much research has been done to realize a global optimization for supply chain and avoid local suboptimization that occurs when each step is operating independently with its own metrics and incentives. From this research, a multi-echelon model has proved to be a good framework for modeling a complex supply chain and optimizing the inventories in the supply chain by finding the optimal placement of safety stock. Key assumptions for this model are that each stage of the supply chain is operating under a periodic-review, base-stock policy, and each stage quotes a guaranteed service time to next stages. In this model, a safety stock is the inventory used to decouple the upstream and downstream, so that the downstream can work independently from the upstream, and the uncertainties are pushed to the upstream. The original work for this model was done by Simpson (1958) who determined optimal safety stock for a serial supply chain model. Later Inderfurth (1991) and Minner (1997) extended the model by relaxing the assumptions about demand and internal policies. Graves and Willems (2000) applied this model for spanning-tree networks under a stationary demand real-life case. Schoenmeyr and Graves (2009) further developed the model for the case of forecast evolution, which is more practical in a real world setting where companies usually build products according to forecasts.

3.2.1 Multi-Stage Network

The supply chain is often modeled as a network consisting of nodes and arcs, of which nodes represent stages in the supply chain and arcs represent the supply relationship from an upstream stage to a downstream stage, as shown in Figure 3.1.

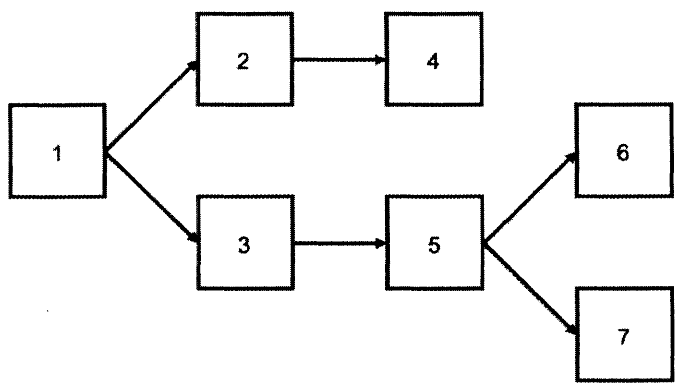

Figure 3.1: Example for a Multi-Stage Network

A stage represents a particular process function in the supply chain, so it could be the procurement of raw materials, the production of components, the transportation of products from one place to another place, and the assembly of components into finished goods. Each stage has a particular material input and material output. The input materials are the output materials from the upstream stages; after the stage's processing, the input materials become the output materials for the downstream stage. For a production stage, the inputs are raw materials and outputs are components, while for a transportation stage, the inputs and outputs are the same. A stage begins when it gets inputs from upstream, and it ends when this stage releases its own outputs. Each stage has a location for holding a safety stock inventory of its outputs at the end of this stage.

The arcs are simple goes-into relationship between two stages, representing the upstream stage sending outputs to the downstream stage. However, in a supply chain network, there will always be two types of stages that have only input or output arcs. The first one is the initial stage like Stage 1 in Figure 3.1, usually denoting the process of changing raw materials into components. Here we usually assume raw materials are already in the company and there are no upstream suppliers. The second type is the final stage that serves customers like Stage 6 and 7, where customers are not supplying anything to any further downstream stages.

For each stage i, there is a deterministic process lead time called stage time and denoted by T_i. A stage time is the time from when all of the inputs for this stage are available until all the outputs of this stage are available to serve downstream stage. In a production stage, the stage time is the time needed to produce the required products, while in a transportation stage, the stage time is the time needed to ship products from one place to another place. We assume that this stage time does not depend on the size of the order; hence we assume that the demands are all bounded and within any existing capacity constraints. And for situations when demands exceed the bound, other actions are required to fix the problem, which is not considered in a multi-stage model.

Each stage j promises a guaranteed service time S_j by which the stage j will meet demands from downstream stages. This means that all the demands coming in at time point t must be 100% filled by time point $t + S_j$. While at the same time, stage j is also getting a 100% guaranteed serviced time from upstream stage i, which is denoted as S_{ij}. Generally S_j is called the outgoing service time for stage j, and S_{ij} is called the incoming service time for stage j, and it's easy to discover that $S_{ij} = S_i$. Actually, the service times within the supply chain network are decision variables for this optimization model, which will be shown in Chapter 4 in detail.

The guaranteed service assumption is very strong. Graves and Willems (2000) have provided reasons for the guaranteed service time assumption. By this assumption, we do not need to model a tradeoff between material starvation costs and inventory holding costs, and are actually defining the problem as how to place safety stocks across supply chain to provide 100% service with least inventory holding cost. There are mainly two benefits for this assumption. The first one is that this strong assumption goes well with real-life practices. In terms of a target service level, managers will always want no stock-outs nor late deliveries, which means 100% is an ideal number for managers as long as the demand is "reasonable." When the demand is extremely high and exceeds the level of "reasonable" demand, it's not unreasonable for demand to not be fully satisfied from inventory. The second benefit is that guaranteed service time greatly facilitates the coordination between different facilities and it's easier to track and determine the safety stocks in supply chain.

3.2.2 Evolving Forecast Model

We adopt a forecast evolution model based on Graves et al. (1986). In period t we denote the forecast for period $t+i$ as $f_t(t+i)$ for $i \in \{1, 2, \dots, H\}$, where H is the forecast horizon. Usually we have $f_t(t) = D_t$ where D_t is the real demand in period t. In each period t we make an initial forecast for the demand in period $t+H$ denoted as $f_t(t+H)$, and make revisions to the other forecasts within the forecast horizon, which is denoted by $\Delta f_t(t+i) = f_t(t+i) - f_{t-1}(t+i)$ for $i \in \{1, 2, \dots, H-1\}$.

Then we can get the relationship as:

$$D_t = f_{t-H}(t) + \sum_{i=1}^{H} \Delta f_{t-H+i}(t)$$

3.2.3 Periodic-Review Forecast-based Order Policy

We assume all stages operate with a periodic review base-stock replenishment policy with a common review period. And based on the supply chain in Waters, here we just adopt a serial supply chain network. For each period the order placed on stage *k*+1 by stage *k* is denoted as $P_k(t)$, and the order is made after observing the forecast revision for period *t*.

We denote the stage directly supplying customer demands as stage 1, and the upstream stages as stage 2, stage 3, until the final stage. As assumed before, each stage *k* guarantees a service time S_k to next stage *k*-1. Usually we assume the outgoing service time to customers S_1 is 0. The replenishment time for a stage *k* is the time for the stage to produce its output for its inventory; it is actually the time to get its inputs from the upstream stage plus the stage time, namely $S_{k+1} + T_k$. We define the net replenishment time for stage *k* denoted as τ_k to be the replenishment time minus its outgoing service time; thus, we have:

$$\tau_k = S_{k+1} + T_k - S_k$$

We define a cumulative lead time for stage *k* as:

$$L_k = S_{k+1} + \sum_{i=1}^{k} T_j = \sum_{i=1}^{k} \tau_i$$

Actually this cumulative lead time represents the shortest time for a production order on stage *k* to reach the customers. It is the sum of the time to get inputs from upstream stage *k*+1 plus all the time needed to process the materials for stage *k* and all the downstream stages. So actually at time point *t*, we will assume that stage *k* will place an order on stage *k+1* in an amount of:

$$P_k(t) = f_t(t + L_k) + \sum_{i=0}^{L_k - 1} \Delta f_t(t + i)$$

This is actually the forecast-based ordering policy, based on the assumption that the forecast of end customer demands is visible to all the stages. When there are no errors in the initial forecasting, then no revisions need to be made in each period; in this case, each stage just orders products according to the initial forecast and then directly pushes this order forward along the supply chain to the customers. However, when there are revisions to the forecast, then the production order needs to be modified as shown above. Here we do not consider the possibility for negative orders, which the revisions are relatively small, compared to the initial forecast. This forecast-based model mimics what companies are doing in practice.

We denote stage *k*'s on-hand inventory at time point *t* to be $I_k(t)$; then we can also have the following relationship:

$$P_k(t) = \sum_{i=1}^{T_k + S_{k+1}} E_t[P_{k-1}(t + i - S_k)] - \sum_{i=1}^{T_k + S_{k+1} - 1} P_k(t - i) - I_k(t) + I_{ss}$$

The first term is the number of products that stage *k* needs to deliver to its downstream stage over its replenishment time. The second term denotes all the materials coming from the upstream stage and being processed at current stage. And I_{ss} means the target safety stock level.

We can also express the inventory $I_k(t)$ as:

$$I_k(t + 1) = I_k(t) - P_{k-1}(t + 1 - S_k) + P_k(t + 1 - S_{k+1} - T_k)$$

$P_{k-1}(t + 1 - S_k)$ is the order placed on stage *k* that needs to be delivered on time point *t+1*. And $P_k(t + 1 - S_{k+1} - T_k)$ is the production order that stage *k* is going to complete at time point $t+1$. Combining this inventory equation with the order equation together we can state another equation:

$$I_k(t + T_k + S_{k+1}) = I_{ss} - \sum_{i=t+1}^{t+\tau_k} \sum_{j=i}^{t+L_k} \Delta f_i(j)$$

This equation clearly shows that the inventory level on hand at stage k is directly related to the forecast revisions over the period of the cumulative lead time.

To make it more intuitive, we can further write the revisions term as:

$$\begin{aligned}
\sum_{i=t+1}^{t+\tau_k} \sum_{j=i}^{t+L_k} \Delta f_i(j) &= \sum_{i=t+1}^{t+L_k} \sum_{j=i}^{t+L_k} \Delta f_i(j) - \sum_{i=t+\tau_k+1}^{t+L_k} \sum_{j=i}^{t+L_k} \Delta f_i(j) \\
&= \sum_{j=t+1}^{t+L_k} \sum_{i=t+1}^{j} \Delta f_i(j) \\
&- \sum_{j=t+\tau_k+1}^{t+L_k} \sum_{i=t+\tau_k+1}^{j} \Delta f_i(j) \\
&= \sum_{j=t+1}^{t+L_k} \left(D_j - f_t(j)\right) - \sum_{j=t+\tau_k+1}^{t+L_k} \left(D_j - f_{t+\tau_k}(j)\right)
\end{aligned}$$

The first term is the cumulative forecast errors made at time t for the next L_k periods, while the second term is the forecast errors made at time $t+\tau_k$ for the next $L_k - \tau_k$ periods. So intuitively the revisions represent the cumulative forecast errors over the next τ_k period. To avoid stock-out, we need to have the inventory on hand not be negative; thus, the safety stock level should be large enough to cover this cumulative forecasting error.

If we assume a bound to the cumulative forecasting errors over the net replenishment time, we can directly set the safety stock to be the bound so that we can ensure 100% service level without excess inventories. As for determining the bound, Schoenmeyr and Graves (2009) have proposed two ways. The first one is to observe the historical data and identify a distribution pattern of the forecast errors so as to further determine the bound.

The second way is to measure the standard deviation of the forecast error and construct a maximum forecast error to find the bound. Suppose that we denote the maximum cumulative forecast error over L time periods to be $F(L)$ and can represent this by:

$$F(L) = z\sigma(L)$$

Where z is the safety factor we assume, and $\sigma(L)$ is the standard deviation of our cumulative forecast error over L periods. If the cumulative forecast error is smaller than $F(L)$, we can always ensure the 100% service level. With the bound function, we can easily determine the corresponding safety stock level to realize the 100% guaranteed service.

3.2.4 Safety Stock Optimization

The forecast-based order policy model can be further used to formulate an optimization problem whose goal is to minimize the total safety stock holding costs along the supply chain. We denote the bound for forecasting error at stage k over the net replenishment time to be $B(\tau_k)$. Actually we know τ_k is a function of the service times S_k, so the bound can be written as $B(S_k)$. We also denote inventory holding cost at stage k to be h_k; then the optimization problem is as follows:

$$\min \sum_{k=1}^{N} h_k\, B(S_k),$$

$$s.t.\ \tau_k \geq 0,\ \ \forall k,$$

$$S_k \geq 0, S_1 = S_{N+1} = 0.$$

In general we can define $g(L) = F^2(L) = z^2 var[\sum_{j=t+1}^{t+L}(D_j - f_t(j))]$, so then we can further develop the optimization problem into:

$$\min \sum_{k=1}^{N} h_k \sqrt{g(L_k) - g(L_{k-1})},$$

$$s.t.\ \tau_k \geq 0,\ \ \forall k,$$

$$S_k \geq 0, S_1 = S_{N+1} = 0.$$

We can also add $\sum_{j=1}^{k-1} T_j$ to the two sides of the first constraint, and then the optimization can be expressed in terms of L_k:

$$\min \sum_{k=1}^{N} h_k \sqrt{g(L_k) - g(L_{k-1})},$$

$$s.t.\ L_k \geq L_{k-1},\ \ \forall k,$$

$$L_k \geq \sum_{j=1}^{k} T_j, \forall k,$$

$$L_0 = 0.$$

It is safe to assume that $g(L)$, the variance of forecast errors over the L time period, is a strict increasing function of time L. So the constraints in the optimization problem can be further changed into:

$$\min \sum_{k=1}^{N} h_k \sqrt{g(L_k) - g(L_{k-1})},$$

$$s.t.\ g(L_k) \geq g(L_{k-1}),\ \ \forall k,$$

$$g(L_k) \geq g(\sum_{j=1}^{k} T_j), \forall k,$$

$$g(L_0) = 0.$$

To make it simple in expression, we just use Z_k to denote $g(L_k)$. Since $g(L_k)$ is a strict increasing function, each single Z_k will have an unique $g(L_k)$ and thus unique L_k and S_k, which ensures the solution are constantly the same for all the above optimization equations. So the final expression of the problem is:

$$\min \sum_{k=1}^{N} h_k \sqrt{Z_k - Z_{k-1}},$$

$$s.t.\ Z_k \geq Z_{k-1},\ \ \forall k,$$

$$Z_k \geq g(\sum_{j=1}^{k} T_j), \forall k,$$

$$Z_0 = 0.$$

Simpson (1958) proved that this format of optimization problem has a concavity property that the optimal solution lies in the corner of the feasible region. He further concluded an all-or-nothing property for this optimization problem, which means that for any stage in the network, it either holds no safety stock or keeps a large amount of safety stock to decouple it with the downstream stages.

3.2.5 Forecasting Bias

When it comes to determining the standard deviation of the forecast error, the above model is based on an assumption that the forecasting is unbiased, which means in the long term the expected value of the forecast is equal to demand. However, this assumption isn't necessarily true in real life situations. Actually, M.Manary and S.Willems (2008) have described a situation in Intel, where a biased forecasting has resulted in significantly more inventories. To test for bias, the relative forecast accuracy for each SKU can be calculated using a ratio:

$$\theta = \frac{F}{F + D}$$

Where F means the forecast and D means the actual demand. For unbiased forecasting, the time series of ratio θ should be centered at 0.5.

Ideally the best thing to do for forecast bias is to reengineer the forecasting process in the system to remove the bias. However, this requires lots of efforts and time inputs. So M.Manary and S.Willems (2008) described a mathematical method to adjust the standard deviation of forecast error with forecast bias. To do this adjustment, we consider the distribution of forecast errors based on the percentage of occurrence for each error measurement. A revised approximation of the standard deviation of forecasting errors is:

$$\hat{\sigma} = \frac{\frac{1-\theta_{1-\alpha}}{\theta_{1-\alpha}} - 1}{t_{1-\alpha.df}} \cdot \mu$$

Where α means our target service level, so $1-\alpha$ means the percentage of time we cannot satisfy the demand. $\theta_{1-\alpha}$ is derived from the time series of the ratio θ by finding out the corresponding quantile point of $1-\alpha$. $t_{1-\alpha.df}$ is the student-t distribution with a cumulative density of $1-\alpha$ and degrees of freedom based on the number of historical data points we collected in the time series. And μ denotes the average demand based on the historical data.

If the forecasting is not biased, then as the number of historical data points increases, the $\hat{\sigma}$ will converge to the standard deviation of forecasting error calculated in the usual way:

$$\sigma = \sqrt{\frac{\sum_{i=1}^{n}(F_i - D_i)^2}{n-1}}$$

4 Implementation of Multi-Echelon Model

This chapter mainly describes the implementation process of a periodic-review forecast-based inventory policy in a multi-echelon supply chain setting. The process starts with the clarification of assumptions, determination of model parameters, model construction, and ends with a result of optimal safety stock strategy. The theoretical supports for this chapter all are from chapter 3.

4.1 Assumptions and Parameters

As stated in chapter 3, before constructing a supply chain model and making an analysis, the first step is always determining the parameters that we will use in the model and clearly defining our assumptions. Below are the discussions of important assumptions and parameters for the model.

4.1.1 Assumptions

Before implementing the model, we need to clearly define our assumptions for the model. Since our model cannot be a one-for-all model, we need to make sure the assumptions reveal or go well with most of the real life situations.

The first assumption is deterministic lead time for each stage, which means the stage time is not impacted by the size of the order, and further means that each stage does not have a capacity constraint. For Waters, the transportation stages perfectly follow this assumption since air shipment time has nothing to do with order size. However, for the production stages, this assumption may be a little strong. To make it reasonable we need to combine with another assumption of bounded demand, and use the demand bound to calculate the maximum lead time for a production stage. Then it may be reasonable to use this time as the stage time with a deterministic lead time assumption.

The second assumption is bounded demand. According to historical data we can always find a meaningful upper limit for the demand that can be covered by the safety stock. Once the demand exceeds the upper limit, then it is an extreme situation that is beyond the protection that is provided by the safety stock. This bounded demand assumption also helps us standardize and fix each stage's lead time, so as to apply a guaranteed service time policy.

The third assumption is for the guaranteed service time policy. As stated in Chapter 3, this assumption is strong but greatly improves the integration and cooperation level for all the stages. Especially for Waters, as we know from Chapter 2, one of the biggest problems is the lack of integration and communication, which sometimes results in late deliveries. So the assumption itself of guaranteed service times can be a solution to this problem of lack of integration.

The final assumption is that all stages are using a periodic-review policy and share a common review period. And there is no time delay between review and order placement, which means once the review is finished, each stage will place an order to its upstream stage. For the review period, we assume it to be 7 days, which is commonly adopted by modern manufacturing companies.

4.1.2 Scope and Granularity

The supply chain is typically complex and large for modern manufacturing companies, and it is often not practical to manage the whole supply chain. So people need to define what parts to focus on in this supply chain. As stated in chapter 2, Waters' supply chain network for Analytical Columns is broken down into two major functional divisions: one is the Column Assembly division focused in Milford facility while the other one is the Chemistry Powder division focused in Taunton facility. Compared to Column Assembly, the Chemistry Powder division's supply chain is relatively less complicated and more fixed; so this model will only consider the Column Assembly division. For the downstream part of the supply chain, we choose the stage of shipping products from the Milford distribution center to customers as the final stage. For the upstream of the supply chain, Waters always keeps more than one year of raw material inventory on hand, so we will neglect the raw materials procurement stage and define the most upstream stage as getting raw materials to produce Components.

The fact that Waters has nearly 6100 SKUs makes it rather complicated to construct models for each of these SKUs. As a starting point for the implementation of a forecast-based model, we want to pick out a manageable number of SKUs for analysis, and ensure that these SKUs will have significant effects on the whole supply chain. From analysis we find that the top 80 SKUs in Waters contribute to more than half of the total Analytical Column's sales volume, as shown in Figure 4.1. Also for the top 80 SKUs, Waters is holding a Finished Column safety stock level of 4 weeks' demand, which means the capital tied up in these SKUs is significant enough so that an improvement in these top 80 SKUs' supply chain will also bring observable improvements at the company level. And 80 is a manageable number to do the model analysis. So for the granularity of the model, we focus on the top 80 SKUs out of total 6100 SKUs.

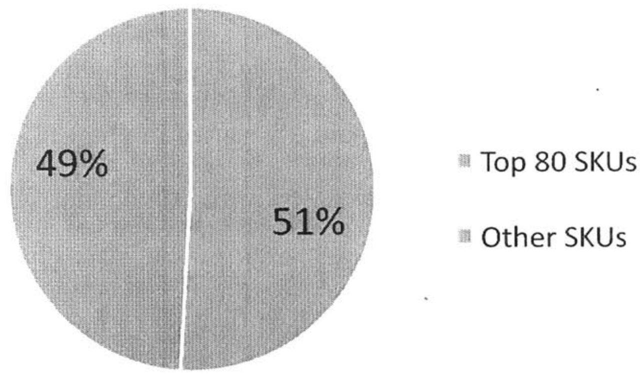

Figure 4.1: Total Sales Volume for Top 80 SKUs and Other SKUs

For the scope that we have chosen, we can break down these 80 SKUs into different kinds of Column Assemblies and further into different kinds of Components. Table 4.1 shows the detailed break down of these 80 SKUs, which are extracted from the SAP system "Single BOM" transaction. All these materials exist in our model and we need to consider all of them in analysis. The table has two sections and each section has five columns. The first column is the list of top 80 Finished Columns. The second column indicates the type of Column Assembly used to produce the Finished Column in first column. For example, SKU "186002350" to "186007114" are all made from Column Assembly "289001855". And similarly, the other columns show the Components used in the corresponding Finished Column.

From Table 4.1 we can easily see that these 80 SKUs only have 24 types of Column Assemblies, 19 types of Column Tubes, 9 types of End Nuts and Filters, which means we have the opportunity to enjoy the benefits of pooling if inventories are held in terms of these components.

Table 4.1: Breakdown of Top 80 SKUs

Finished Column	Column Assembly	Column Tube	End Nut	Filter	Finished Column	Column Assembly	Column Tube	End Nut	Filter
186002350	289001855	405011567	405003600	289004437	186000480	289000491	WAT015927	405000614	WAT015931
186002853					186003033	289000494	WAT091-02		
186002877					186000442	289000495	WAT091-03		
186002884					186000492				
186003538					186003034				
186004495					186003045				
186005296					186003116				
186005965					186003335				
186007093					186005270				
186007114					186000494	289000496	WAT091-04		
186002349	289002583	405011937			186000496				
186002352	289001874	405011938			186003010				
186002854					186003117				
186002878					186003031	289000493	WAT091-06		
186002885					WAT046970	WAT015923	WAT015927	WAT015930	
186003461					WAT046980				
186003533					WAT052885				
186003539					WAT086344				
186003837					WAT027324	WAT015924	WAT015928		
186004496					WAT066224	WAT551-04	WAT091-01		
186004801					WAT066220	WAT551-05	WAT091-02		
186005297					186002559	WAT092-03	WAT091-03		
186007095					WAT045905				
186007116					WAT045995				
186008316					186000180	WAT551-06			
186002353	289002445	405011939			186001342				
186003376					186002554				
186003534					186003729				
186003540					WAT200632				
186004742					186000112	WAT092-04	WAT091-04		
186004802					186001346				
186005298					186002560				
186007096					186003748				
186008315					WAT054270				
186005225	289007997	405012517	405008848	289007723	WAT054275				
186007378					186003975	289004057	405007034	405005625	289004545
186005226	289007998	405012518			186003976				
186002979	289001600	405009466	405006406	289004915	PSS830615	PSS550425	PSS550425	PSS613120	PSS669018
186003028	289000488	WAT200-04	405000614	289002235	PSS831915				

4.1.3 Costs

As stated in chapter 3, we need to determine three types of costs in a supply chain analysis: production costs, transportation costs, and inventory holding costs.

The production costs have two parts. The first part is the materials costs for this product. The second part is the facility operation cost, including labor and machine maintenance. Waters calculates the total annual operation cost, then divides it by total production hours, getting the hourly operation cost. Finally based on the production lead time and hourly operation cost, each material is assigned with a particular operation cost. We note that the production cost of downstream materials includes the material cost of upstream materials. For example, based on the breakdown in Table 4.1, the production cost of Finished Column "186002350" already contains the production cost of Column Assembly "289001855". And the production cost of Column Assembly "289001855" contains the production cost of components like Column Tube "405011567", End Nut "405003600", and Filter "289004437". Figure 4.2 is an illustration of the production cost.

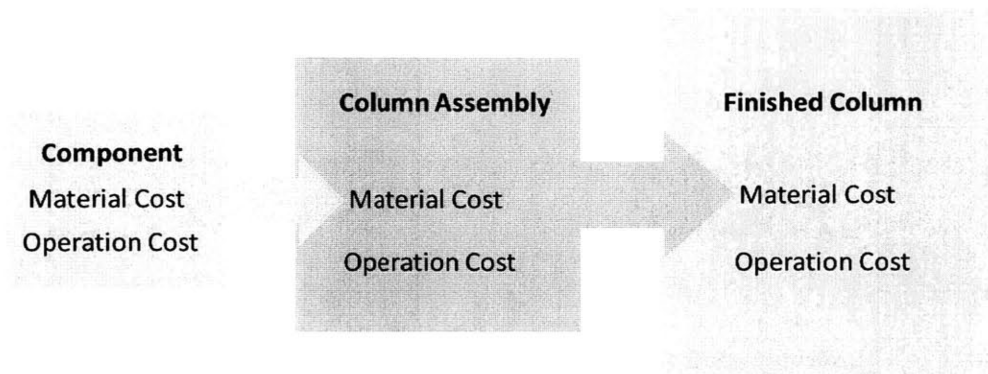

Figure 4.2: Breakdown of Material Production Cost

We can get the production costs from SAP system "Display Material Cost Estimate" transaction. For the sake of confidentiality, each product's production cost is standardized. We denote the production cost of the component Column Tube "405011567" to be 1 in the standardized costs system. Table 4.2 shows all the 141 materials' standardized production costs. The Components section has three colors; yellow indicates Column Tubes, green indicates Filters, and blue indicates End Nuts. We note that for Finished Column, we do not include the cost of Chemistry Powder, because we only consider and optimize the Column Assembly division, assuming that the Chemistry Powder division functions individually and holds enough inventories. So whether the Chemistry Powder is packed in Finished Column or is stored in warehouse, it will incur the same inventory cost.

Table 4.2: Standardized Production Cost for Materials

Finished Column	Production Cost	Finished Column	Production Cost	Column Assembly	Production Cost	Component	Production Cost
186003021	6.91	186008316	8.98	289000481	3.82	405007034	1.24
186003028	7.77	186002353	9.68	289000488	4.68	405009466	11.91
186000480	6.81	186003376	9.68	289000491	3.61	405011567	1.00
186003033	6.37	186003534	12.04	289000493	2.77	405011937	0.90
186000442	6.42	186003540	9.48	289000494	3.26	405011938	1.46
186000492	6.61	186004742	12.20	289000495	3.38	405011939	1.95
186003034	6.66	186004802	10.69	289000496	4.07	405012517	1.64
186003045	6.62	186005298	11.97	289001600	40.14	405012518	2.58
186003116	6.65	186007096	10.53	289001855	4.27	PSS550425	1.65
186003335	6.48	186008315	10.69	289001874	4.73	WAT015927	1.73
186005270	8.87	186003975	16.89	289002445	5.22	WAT015928	2.94
186000494	7.74	186003976	16.91	289002583	4.17	WAT015960	4.48
186000496	7.73	186002979	52.35	289004057	4.29	WAT091-01	1.04
186003010	7.61	186005225	28.59	289007997	6.81	WAT091-02	1.38
186003117	7.78	186007378	9.63	289007998	7.75	WAT091-03	1.50
186003031	5.93	186005226	31.58	PSS550425	4.01	WAT091-04	2.19
186002350	8.37	PSS830615	5.39	WAT015923	3.01	WAT091-06	0.89
186002853	8.73	PSS831915	6.88	WAT015924	4.22	WAT098-03	1.06
186002877	8.24	WAT046970	6.49	WAT015934	7.06	WAT200-04	2.41
186002884	8.31	WAT046980	6.18	WAT092-03	2.78	WAT022978	0.27
186003538	8.38	WAT052885	5.92	WAT092-04	3.47	289002235	0.64
186004495	11.19	WAT086344	6.04	WAT551-04	2.31	WAT015931	0.45
186005296	8.10	WAT027324	7.77	WAT551-05	2.65	289004437	0.30
186005965	20.72	WAT066224	5.28	WAT551-06	2.78	289004545	0.30
186007093	8.90	WAT066220	5.59			289004915	9.20
186007114	8.90	186002559	6.59			289007723	1.12
186002349	8.27	WAT045905	5.93			PSS669018	0.05
186002352	8.98	WAT045995	5.80			WAT015935	0.67
186002854	10.19	186000180	5.69			405000613	1.08
186002878	8.77	186001342	5.92			405000614	0.46
186002885	8.97	186002554	6.60			405003600	1.25
186003461	10.39	186003729	5.97			405005625	1.09
186003533	8.84	WAT200632	6.65			405006406	4.52
186003539	8.98	186000112	6.84			405008848	1.36
186003837	11.24	186001346	6.96			PSS613120	1.13
186004496	11.65	186002560	7.66			WAT015930	0.11
186004801	10.40	186003748	6.88			WAT015961	0.52
186005297	8.84	WAT054270	6.80				
186007095	9.38	WAT054275	6.88				
186007116	4.42	WAT201549	12.45				

Waters uses air shipment between different facilities, so the transportation costs are measured based on the weight of the materials. There are two transportation stages, one is shipment of Column Assemblies from Milford to Ireland, and the other is shipment of Finished Columns from Ireland to Milford Distribution Center. The weight of each product can be found from the SAP system "Display Stock and Requirements" transaction. Again we standardize the transportation in the same way above, and Table 4.3 shows the standardized transportation costs for all the Finished Columns and Column Assemblies. From the table we can easily see that due to the low weight of the Column, the transportation costs are relatively low compared to production costs.

Table 4.3: Standardized Transportation Cost for Different Materials

Finished Column	Trans Cost	Finished Column	Trans Cost	Finished Column	Trans Cost	Finished Column	Trans Cost	Column Assembly	Trans Cost
186003021	0.02	186007093	0.02	186007096	0.02	WAT200632	0.03	289000481	0.00
186003028	0.03	186007114	0.02	186008315	0.02	186000112	0.04	289000488	0.01
186000480	0.07	186002349	0.02	186003975	0.02	186001346	0.17	289000491	0.01
186003033	0.02	186002352	0.02	186003976	0.02	186002560	0.04	289000493	0.00
186000442	0.06	186002854	0.02	186002979	0.14	186003748	0.03	289000494	0.00
186000492	0.07	186002878	0.02	186005225	0.03	WAT054270	0.05	289000495	0.01
186003034	0.03	186002885	0.02	186007378	0.03	WAT054275	0.04	289000496	0.01
186003045	0.03	186003461	0.02	186005226	0.04	WAT201549	0.07	289001600	0.13
186003116	0.03	186003533	0.03	PSS830615	0.03	186003538	0.02	289001855	0.00
186003335	0.04	186003539	0.03	PSS831915	0.03	186004495	0.02	289001874	0.01
186005270	0.03	186003837	0.02	WAT046970	0.03	186005296	0.02	289002445	0.01
186000494	0.03	186004496	0.02	WAT046980	0.03	186005965	0.02	289002583	0.00
186000496	0.04	186004801	0.01	WAT052885	0.03	186003540	0.03	289004057	0.00
186003010	0.04	186005297	0.02	WAT086344	0.03	186004742	0.02	289007997	0.01
186003117	0.04	186007095	0.02	WAT027324	0.04	186004802	0.03	289007998	0.01
186003031	0.02	186007116	0.02	WAT066224	0.03	186005298	0.02	PSS550425	0.01
186002350	0.02	186008316	0.02	WAT066220	0.03	186000180	0.03	WAT015923	0.01
186002853	0.17	186002353	0.05	186002559	0.03	186001342	0.04	WAT015924	0.02
186002877	0.02	186003376	0.04	WAT045905	0.03	186002554	0.04	WAT015934	0.04
186002884	0.02	186003534	0.03	WAT045995	0.03	186003729	0.05	WAT092-03	0.01
								WAT551-06	0.01
								WAT551-05	0.01
								WAT551-04	0.01
								WAT092-04	0.01

From interviews with warehouse managers, we know that Analytical Column is a very special type of product, which never goes obsolete and which occupies very little storage space. This means Waters has nearly no operational costs in holding Columns inventory. So the primary cost in holding inventory at Waters is the opportunity cost of the capital tied up in the inventory. The production costs and transportation costs indicate the capital Waters has spent on the inventory, and this capital directly determines the inventory holding costs at Waters. Here we assume the opportunity cost is 8% annually of the total capital tied up in inventory.

4.1.4 Target Service Level

The Analytical Column Industry is very competitive. If customers cannot get what they want immediately, they may directly go to competitors. So any missed order may cause significant loss to the company in the long term. From talking and discussing with the managers in Waters, we determine that 98% is a reasonable service level target for the top 80 products. This means that for 98% of the time, our supply chain can cover 100% of the demand from customers. In the remaining 2% of the time, the demand exceeds our bound and this problem cannot be handled only by safety stock.

4.1.5 Forecast Errors

Originally we planned to use an evolving forecast model for Waters' supply chain. However, when we continued the project, we found that Waters lack the necessary data to determine the maximum cumulative forecast error $F(L)$. So we still use the Simpson model (1958) but instead of using demand variability, we use the variability of forecasting error to determine the safety stock.

In a make-to-stock environment, forecast errors directly determine the variability we are facing and thus how much safety stock we should keep to cover the variability. As stated in Chapter 3, we can use a forecast ratio to measure forecast error:

$$\theta = \frac{Forecast}{Forecast + Demand}$$

When the ratio is larger than 0.5, it means forecast is larger than real demand, and when the ratio is smaller than 0.5, it means forecast is smaller than real demand.

From Waters we gathered the one-month ahead forecast and real demand data for the past 13 months for top 80 SKUs. For each SKU we have 13 observations of the forecast ratio. Figure 4.3 shows the boxplot of the forecast ratio for all the SKUs.

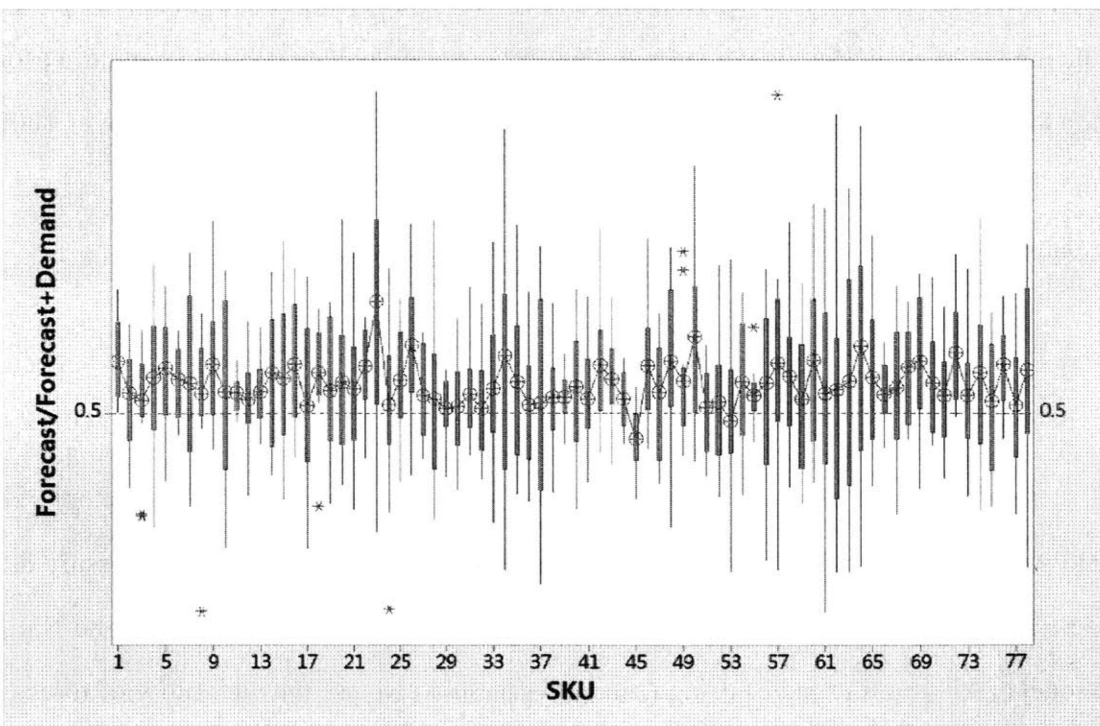

Figure 4.3: Boxplot of Forecasting Ratios

We can make two inferences based on this plot. First, it is apparent that nearly all of the SKUs have a mean forecast ratio larger than 0.5, which means the current forecasting system in Waters has a bias in forecast. So we need to use a mathematical method introduced in Chapter 3 to adjust the bias. The second inference the boxplot gives us is that generally all the SKUs have a similar mean of forecast ratio. Based on this we can make an assumption that Waters has the same forecasting ability for all the top 80 SKUs. Due to the lack of data, we only can make observations for 77 SKUs. To validate this assumption, we can make an ANOVA test based on the total 1008 observations, viewing the 77 SKUs as 77 experiment levels. The results are shown in Table 4.4. From the table we see a large p value, which suggests that our assumption of same forecasting ability is reasonable. This assumption is to determine the quantile point. If we decide this point based on each SKU's data, then because the sample size is small we have to choose the smallest forecast ratio. But it's easy for a single SKU to have extreme situations, so this way is not likely to be accurate.

Table 4.4: ANOVA Test for Different SKUs' Forecast Ratio

Source	DF	Adj SS	Adj MS	P-Value
SKU	77	0.33	0.00	0.40
Error	930	3.90	0.00	
Total	1007	4.23		

Now based on the assumption, we can draw the distribution of all the forecast ratios, and find the corresponding quantile point, which is 2% in our case. Figure 4.4 shows how to find the 2% point in the distribution plot. For the sake of confidentiality, we cannot give the exact number here.

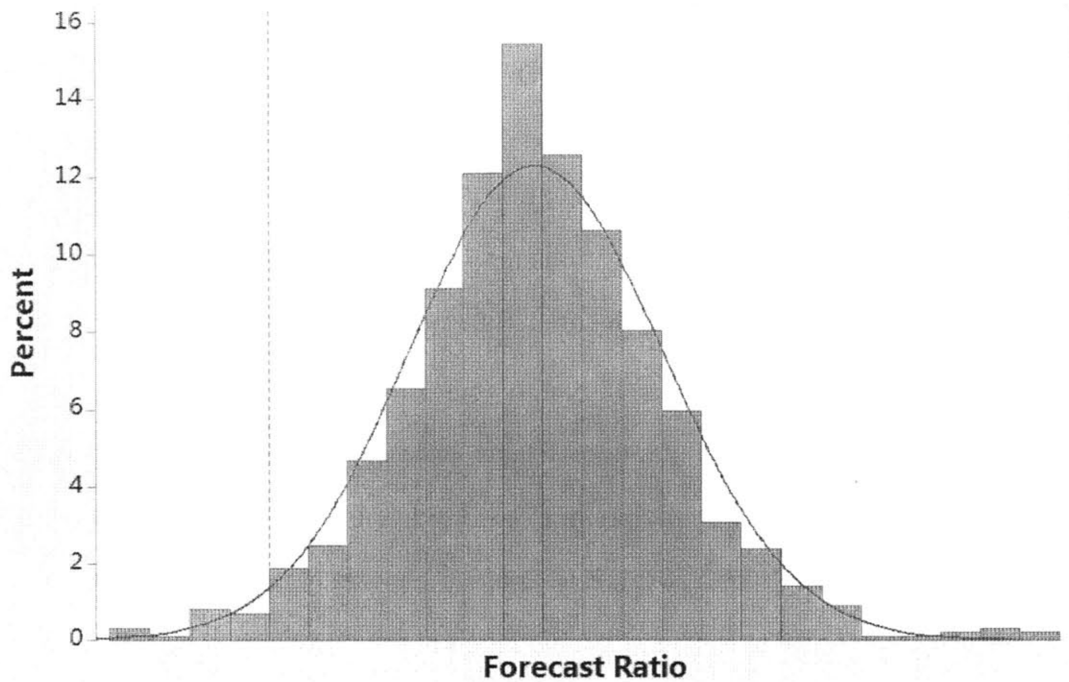

Figure 4.4: Distribution of Forecast Ratios

Now with the $\theta_{0.02}$ from the forecast ratio data, we can use the formula stated in Chapter 3 to calculate adjusted standard deviation of forecast error:

$$\hat{\sigma} = \frac{\frac{1-\theta_{0.02}}{\theta_{0.02}} - 1}{t_{0.02,df}} \cdot \mu$$

One thing we need to notice is that since we are assuming that all the SKUs have the same forecasting ability, it is possible that some of the SKUs' smallest forecast ratios are even larger than $\theta_{0.02}$ here. This suggests that this SKU's forecast may be more accurate and less biased, so in this situation we will just use the smallest forecast ratio in the past 13 months for this SKU.

After determining the variability of the forecast error for each of the Finished Columns, we continue to determine the variability of the forecast error on the Column Assembly and

Component level. At these levels, there are no extreme cases due to the consequences from pooling; thus, we are able to directly use the smallest forecast ratio to calculate the standard deviation of the forecast error. However, the effect of pooling may cause situations in which even the smallest forecast ratio is very large. For instance, when it's larger than 0.47, then the standard deviation is only 6% of the mean demand. In this situation we just use the original standard deviation of forecasting errors, as a conservative estimate.

From the above we will obtain the standard deviation of forecast errors for each SKU, and each assembly and component. However, these standard deviations are based on a monthly data. Since our review period is 1 week, we need to convert the monthly standard deviations into weekly standard deviations. Now we assume that the weekly forecast errors are independent; since there are 4 weeks in a month, we have the following calculations:

$$\begin{aligned} &Variance\ of\ Monthly\ Forecast\ Error \\ &= Sum\ of\ Variance\ of\ Weekly\ Forecast\ Error \\ &= 4 \times Variance\ of\ Weekly\ Forecast\ Error \end{aligned}$$

So we have:

$$Weekly\ Standard\ Deviation\ of\ Forecast\ Error = \frac{Monthly\ Standard\ Deviation\ of\ Forecast\ Error}{2}$$

Table 4.5 lists the weekly standard deviations of forecast error for all the SKUs in different material levels. Notice that the variability for Filter and End Nut is larger, because one Column needs two Filters and two End Nuts, so the variability of Finished Column will be doubled.

Table 4.5: Weekly SDFE in Different Material Levels

SKU	Finished Column	Column Assembly	Tube	Filter	End Nut	SKU	Finished Column	Column Assembly	Tube	Filter	End Nut
186002350	68.9	60.5	60.5	108.8	107.0	186003033	8.2	8.2	5.5	7.5	9.3
186002853	15.2	3.9	3.9	7.0	6.9	186000442	10.2	4.7	5.9	9.8	12.2
186002877	9.5	3.5	3.5	6.2	6.1	186000492	11.8	8.8	11.0	18.2	22.6
186002884	10.1	3.4	3.4	6.2	6.1	186003034	21.3	16.9	21.1	34.9	43.3
186003538	27.8	11.9	11.9	21.3	21.0	186003045	9.6	6.3	7.9	13.0	16.2
186004495	2.2	2.2	2.2	4.0	3.9	186003116	21.3	10.3	12.9	21.3	26.5
186005296	9.3	3.9	3.9	7.0	6.9	186003335	7.4	3.1	3.9	6.5	8.1
186005965	7.5	1.9	1.9	3.5	3.4	186005270	10.4	3.9	4.9	8.1	10.0
186007093	3.1	1.4	1.4	2.6	2.5	186000494	7.9	7.2	6.8	11.6	14.5
186007114	4.2	1.1	1.1	1.9	1.9	186000496	25.5	17.6	16.1	27.7	34.4
186002349	8.0	8.0	8.0	3.7	3.6	186003010	6.6	3.2	3.0	5.1	6.4
186002352	84.0	51.3	51.3	61.4	60.4	186003117	30.4	19.8	18.2	31.2	38.8
186002854	14.8	7.9	7.9	9.4	9.3	186003031	14.7	14.7	14.7	34.9	43.4
186002878	7.8	5.6	5.6	6.7	6.6	WAT046970	10.3	9.3	9.1	11.8	12.5
186002885	10.0	5.7	5.7	6.8	6.7	WAT046980	27.5	25.0	24.3	31.7	33.5
186003461	5.8	3.6	3.6	4.3	4.2	WAT052885	5.4	4.9	4.8	6.2	6.6
186003533	8.8	4.9	4.9	5.9	5.8	WAT086344	26.8	24.4	23.7	30.9	32.7
186003539	23.1	22.5	22.5	27.0	26.5	WAT027324	16.5	16.5	16.5	52.2	55.3
186003837	7.9	4.5	4.5	5.4	5.3	WAT066224	17.2	17.2	17.2	18.5	19.6
186004496	3.8	2.8	2.8	3.3	3.2	WAT066220	7.1	7.1	5.3	7.2	7.6
186004801	7.0	5.6	5.6	6.7	6.5	186002559	10.1	5.5	7.2	11.9	12.6
186005297	9.6	7.6	7.6	9.1	9.0	WAT045905	34.7	15.6	20.3	33.6	35.6
186007095	5.1	2.8	2.8	3.4	3.3	WAT045995	6.9	5.2	6.8	11.2	11.8
186007116	4.5	1.7	1.7	2.1	2.0	186000180	6.0	5.1	4.2	6.9	7.3
186002353	20.3	19.7	19.7	13.8	13.6	186001342	7.5	4.3	3.5	5.8	6.2
186003376	7.8	5.1	5.1	3.6	3.5	186002554	12.4	12.4	10.1	16.8	17.7
186003534	8.6	5.7	5.7	4.0	3.9	186003729	10.2	10.1	8.3	13.7	14.5
186003540	12.4	10.7	10.7	7.5	7.4	WAT200632	16.6	11.9	9.7	16.1	17.0
186004742	10.1	7.1	7.1	5.0	4.9	186000112	11.5	8.2	10.8	18.5	19.6
186004802	5.7	5.5	5.5	3.9	3.8	186001346	4.0	2.8	3.7	6.4	6.8
186005298	11.4	7.5	7.5	5.2	5.2	186002560	11.7	8.4	11.0	18.9	20.0
186007096	1.9	1.9	1.9	1.3	1.3	186003748	5.1	3.1	4.0	6.9	7.3
186005225	12.6	8.1	8.1	19.4	19.4	WAT054270	13.7	7.0	9.2	15.8	16.7
186007378	3.3	1.0	1.0	2.4	2.4	WAT054275	50.7	27.5	36.1	61.9	65.5
186005226	7.6	7.6	7.6	8.5	8.5	186003975	21.7	14.6	14.6	29.2	29.2
186002979	1.7	1.7	1.7	3.4	3.4	186003976	13.6	5.5	5.5	10.9	10.9
186003028	3.6	3.6	3.6	7.1	7.1	PSS830615	14.7	14.2	14.2	28.5	28.5
186000480	6.8	6.8	4.4	5.7	7.1	PSS831915	13.4	13.4	13.4	26.9	26.9

Here we are using multi-echelon analysis for each individual SKU. On the Finished Column level, it's easy to decide the variability; but on the Component level, we need to consider both the pooling effect and the individual SKU. Our way to do this is to split the Components' standard deviation to different SKUs. For example, we assume SKU A and B are using a common Component C, as shown in Figure 4.5. The pooling effect decides that Component C has a forecast ratio of 0.46, so its adjusted standard deviation of forecast error is 17.22. But this number is calculated from the mean demand of 200 units for Component C. So if we use 17.22 as the $\hat{\sigma}$ for A and B to decide the safety stock level on their Component level, we will overestimate the safety stock level. So we use the pooled ratio of 0.46 and the mean demand for individual SKU to decide the SKU's Component safety stock level. In our example, when calculating SKU A, if we have safety stock in the Component stage, then the safety stock needed will be:

$$z \cdot \frac{\frac{1-0.46}{0.46} - 1}{2.05} \cdot 90 \cdot \sqrt{L} = 7.75 \cdot z \cdot \sqrt{L}$$

And similarly, on the Component stage, the safety stock needed for SKU B is $9.47 \cdot z \cdot \sqrt{L}$. SKU A and B together on the Component stage will need safety stock of $7.75 \cdot z \cdot \sqrt{L} + 9.47 \cdot z \cdot \sqrt{L} = 17.22 \cdot z \cdot \sqrt{L}$, which is exactly the same as the number directly calculated from Component C.

SKU	Forecast	Demand	Forecast Ratio	Adjusted SDFE	
A	80	90	0.47	5.49	
B	90	110	0.45	11.92	
Component	**Forecast**	**Demand**	**Forecast Ratio**	**Adjusted SDFE**	**Safety Stock**
C	170	200	0.46	17.22	z*17.22*Sqrt(L)
Component Stage	**Demand**	**Forecast Ratio**	**Adjusted SDFE**	**Safety Stock**	**Summed Safety Stock**
SKU A	90	0.46	7.75	z*7.75*Sqrt(L)	z*17.22*Sqrt(L)
SKU B	110	0.46	9.47	z*9.47*Sqrt(L)	

Figure 4.5: Example of Splitting Components' Safety Stock

4.2 Model Construction

After determining the parameters and stating the assumptions, we can build a multi-echelon supply chain model for Waters.

As stated in section 4.1.1, we only look at the Column Assembly division which begins when Milford Machine Shop producing Components and ends with Finished Columns delivered to Customers. The graphic illustration of the supply chain network is shown in Figure 4.6. Sequentially there are 5 stages. For stage 5 although there are 3 different Components, they are manufactured in the same place and have similar lead time, so they can be viewed as an identical stage.

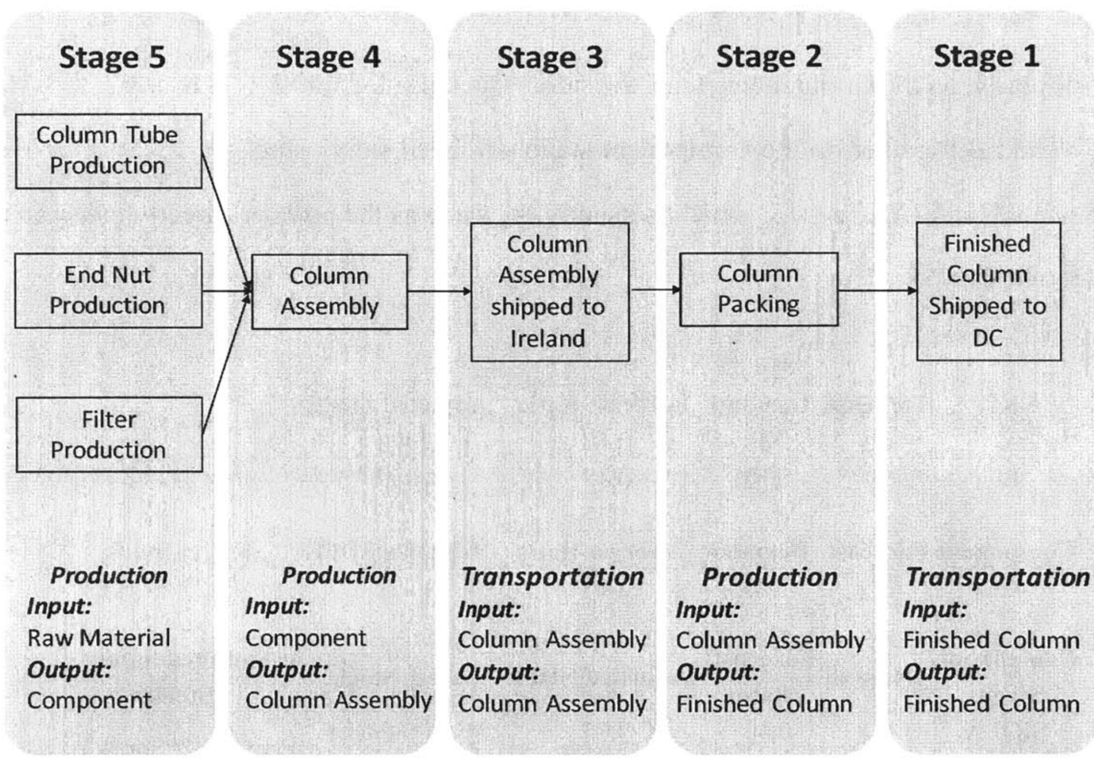

Figure 4.6: Multi-Echelon Network of Waters Supply Chain

Of the 5 stages, there are 3 production stages where the stage time is the lead time to finish production, and there are 2 transportation stages where the stage time is the lead time spent in shipment of the inputs. Table 4.6 shows the time parameters in the supply chain network model.

Table 4.6: Time Parameters for Supply Chain Network

Stage			Stage Time/days	Incoming Service Time/days	Outgoing Service Time/days
1	***End***	***DC***	5	S2	0
	Output	Finished Column			
	Begin	***Ireland Facility***			
	Input	Finished Column			
2	***End***	***Ireland Facility***	14	S3	S2
	Output	Finished Column			
	Begin	***Ireland Facility***			
	Input	Column Assembly			
3	***End***	***Ireland Facility***	3	S4	S3
	Output	Column Assembly			
	Begin	***Milford***			
	Input	Column Assembly			
4	***End***	***Milford***	7	S5	S4
	Output	Column Assembly			
	Begin	***Milford***			
	Input	Component			
5	***End***	***Milford***	14	0	S5
	Output	Component			
	Begin	***Milford***			
	Input	Raw Material			

All the black numbers are deterministic time parameters while all the red numbers are decision parameters we will use in finding the optimal safety stock policy. It is clear that all the outgoing service times are to be determined except for stage 1, where the outgoing service time is preset to be 0. Considering that Waters has enough inventory of raw materials, the incoming service time for stage 5, meaning that whenever the Components production team wants raw materials, they can get them immediately. All the stage times are determined by conducting interviews with stakeholders to learn about the standard production time and transportation time. For all the transportation stages, the stage times are easy to determine, which are quoted by the air shipment companies. For the production stages, the stage time is the maximum lead time to produce the amount of upper demand bound.

Then we need to determine different stages' inventory holding cost for each of the top 80 SKUs. Given that all the stages are using 8% as the annual inventory holding cost rate, we just need to determine the cumulative capital tied up in the materials in each stage for each SKU. When the products are shipped to customers from the distribution center, it means customers are managing the inventories by themselves, and this holding cost is not considered by us. Table 4.7 shows the cumulative stage costs for each of the Top 80 SKUs.

Table 4.7: Cumulative Stage Costs for Top 80 SKUs

Finished Column	Stage							Finished Column	Stage						
	5			4	3	2	1		5			4	3	2	1
	T	F	N						T	F	N				
186002350	1.00	0.30	1.25	4.27	4.29	8.40	8.42	186000480	1.73	0.45	0.46	3.61	3.67	6.88	6.94
186002853	1.00	0.30	1.25	4.27	4.44	8.90	9.07	186003033	1.38	0.45	0.46	3.26	3.28	6.40	6.42
186002877	1.00	0.30	1.25	4.27	4.29	8.26	8.28	186000442	1.50	0.45	0.46	3.38	3.44	6.49	6.55
186002884	1.00	0.30	1.25	4.27	4.29	8.33	8.35	186000492	1.50	0.45	0.46	3.38	3.44	6.67	6.74
186003538	1.00	0.30	1.25	4.27	4.29	8.40	8.43	186003034	1.50	0.45	0.46	3.38	3.41	6.69	6.72
186004495	1.00	0.30	1.25	4.27	4.29	11.21	11.23	186003045	1.50	0.45	0.46	3.38	3.41	6.65	6.67
186005296	1.00	0.30	1.25	4.27	4.29	8.12	8.14	186003116	1.50	0.45	0.46	3.38	3.41	6.68	6.71
186005965	1.00	0.30	1.25	4.27	4.29	20.74	20.76	186003335	1.50	0.45	0.46	3.38	3.42	6.52	6.57
186007093	1.00	0.30	1.25	4.27	4.29	8.92	8.94	186005270	1.50	0.45	0.46	3.38	3.41	8.90	8.93
186007114	1.00	0.30	1.25	4.27	4.29	8.92	8.94	186000494	2.19	0.45	0.46	4.07	4.10	7.78	7.81
186002349	0.90	0.30	1.25	4.17	4.19	8.30	8.32	186000496	2.19	0.45	0.46	4.07	4.10	7.76	7.80
186002352	1.46	0.30	1.25	4.73	4.75	9.00	9.02	186003010	2.19	0.45	0.46	4.07	4.10	7.64	7.68
186002854	1.46	0.30	1.25	4.73	4.76	10.22	10.24	186003117	2.19	0.45	0.46	4.07	4.10	7.81	7.85
186002878	1.46	0.30	1.25	4.73	4.76	8.80	8.82	186003031	0.89	0.45	0.46	2.77	2.79	5.95	5.97
186002885	1.46	0.30	1.25	4.73	4.76	9.00	9.02	WAT046970	1.73	0.45	0.11	3.01	3.04	6.53	6.56
186003461	1.46	0.30	1.25	4.73	4.76	10.42	10.44	WAT046980	1.73	0.45	0.11	3.01	3.03	6.21	6.24
186003533	1.46	0.30	1.25	4.73	4.76	8.87	8.90	WAT052885	1.73	0.45	0.11	3.01	3.03	5.95	5.98
186003539	1.46	0.30	1.25	4.73	4.76	9.01	9.03	WAT086344	1.73	0.45	0.11	3.01	3.03	6.07	6.09
186003837	1.46	0.30	1.25	4.73	4.76	11.26	11.29	WAT027324	2.94	0.45	0.11	4.22	4.26	7.81	7.86
186004496	1.46	0.30	1.25	4.73	4.75	11.67	11.70	WAT066224	1.04	0.45	0.11	2.31	2.34	5.31	5.34
186004801	1.46	0.30	1.25	4.73	4.74	10.40	10.41	WAT066220	1.38	0.45	0.11	2.65	2.68	5.62	5.65
186005297	1.46	0.30	1.25	4.73	4.75	8.86	8.88	186002559	1.50	0.45	0.11	2.78	2.81	6.62	6.65
186007095	1.46	0.30	1.25	4.73	4.75	9.40	9.42	WAT045905	1.50	0.45	0.11	2.78	2.81	5.96	5.99
186007116	1.46	0.30	1.25	4.73	4.75	4.44	4.46	WAT045995	1.50	0.45	0.11	2.78	2.81	5.83	5.86
186008316	1.46	0.30	1.25	4.73	4.75	9.00	9.02	186000180	1.50	0.45	0.11	2.78	2.81	5.73	5.76
186002353	1.95	0.30	1.25	5.22	5.27	9.72	9.77	186001342	1.50	0.45	0.11	2.78	2.82	5.96	6.01
186003376	1.95	0.30	1.25	5.22	5.26	9.72	9.76	186002554	1.50	0.45	0.11	2.78	2.82	6.64	6.68
186003534	1.95	0.30	1.25	5.22	5.25	12.07	12.09	186003729	1.50	0.45	0.11	2.78	2.82	6.01	6.06
186003540	1.95	0.30	1.25	5.22	5.25	9.51	9.54	WAT200632	1.50	0.45	0.11	2.78	2.81	6.68	6.72
186004742	1.95	0.30	1.25	5.22	5.24	12.23	12.25	186000112	2.19	0.45	0.11	3.47	3.51	6.88	6.93
186004802	1.95	0.30	1.25	5.22	5.26	10.73	10.76	186001346	2.19	0.45	0.11	3.47	3.63	7.13	7.29
186005298	1.95	0.30	1.25	5.22	5.24	11.99	12.01	186002560	2.19	0.45	0.11	3.47	3.51	7.70	7.75
186007096	1.95	0.30	1.25	5.22	5.24	10.55	10.57	186003748	2.19	0.45	0.11	3.47	3.50	6.91	6.94
186008315	1.95	0.30	1.25	5.22	5.24	10.71	10.74	WAT054270	2.19	0.45	0.11	3.47	3.51	6.84	6.89
186005225	1.64	1.12	1.36	6.81	6.84	28.61	28.64	WAT054275	2.19	0.45	0.11	3.47	3.51	6.93	6.97
186007378	1.64	1.12	1.36	6.81	6.84	9.66	9.69	186003975	1.24	0.30	1.09	4.29	4.31	16.91	16.93
186005226	2.58	1.12	1.36	7.75	7.79	31.62	31.66	186003976	1.24	0.30	1.09	4.29	4.31	16.93	16.96
186002979	11.91	9.20	4.52	40.14	40.27	52.49	52.62	PSS830615	1.65	0.05	1.13	4.01	4.04	5.42	5.46
186003028	2.41	0.64	0.46	4.68	4.71	7.80	7.83	PSS831915	1.65	0.05	1.13	4.01	4.04	6.91	6.95

And finally, with all these model parameters determined, we can set up a calculation model in Microsoft Excel, as shown in Figure 4.7. We change the numbers in red and have results as blue numbers. With Excel Solver function, we can easily get the optimal safety stock placement and level for each of the SKU. One thing we need to mention here is that we assume no correlation between different stages' review time. So we are counting each stage's review time into its net replenishment time. This assumption is more conservative in determining the inventory level.

Time Parameters					
Stage	**Review Period/days**	**Stage Time/days**	**Incoming Service Time/days**	**Outgoing Service Time/days**	**Net Replenishment Time/days**
1	7	5	21	0	33
2	7	14	0	21	0
3	7	3	35	0	45
4	7	7	21	35	0
5	7	14	0	21	0

Variability Parameters					
SKU	**Tube**	**Filter**	**End Nut**	**Column Assembly**	**Finished Column**
186002350	60.47	108.83	106.96	60.47	68.91

Cost Parameters							
Stage	**5**			**4**	**3**	**2**	**1**
	T	**F**	**N**				
186002350	1.00	0.30	1.25	4.27	4.29	8.40	8.42

	Safety Stock Level							Total Cost
SKU	**Stage5 Tube**	**Stage5 Filter**	**Stage5 Nut**	**Stage4**	**Stage3**	**Stage2**	**Stage1**	
186002350	0.00	0.00	0.00	0.00	314.29	0.00	306.74	3931.24

Figure 4.7: Optimization Model for Safety Stock Calculation

5 Results and Discussion

This chapter mainly describes the computational results of the model constructed in Chapter 4. An optimal solution is first presented and compared to the current situation in Waters; some observations are then made based on the comparisons. Then the author introduces a sub-optimal solution that is easier to implement as a starting point. The sub-optimal solution has a relatively mild change to the current system so that stakeholders will undergo less difficulty in implementing it.

5.1 Optimal Solution

From Table 4.7 it is clear that for all the Analytical Column products, stage 2 is the most value-add stage, while all the other stages have similar and small value-add. This identical cost pattern leads to identical optimal safety stock solution for each of the SKUs. Table 5.1 shows the optimal safety stock solution calculated for all the top 80 SKUs, and it is clear that for all the SKUs, it is optimal to hold safety stocks in stage 3 and stage 1, while keep no safety stocks in other stages. This means we only have safety stock of Column Assemblies in Ireland facility and safety stock of Finished Columns in the Milford distribution center. Due to lack of necessary data, some of the top 80 SKUs can not be calculated, so the final list only has 76 SKUs.

To interpret this result, we need to know that the optimal location of safety stocks depends primarily on how the lead times and holding costs are distributed across the supply chain. For Waters' case, the cumulative lead time for the supply chain is 78 days, which represents the longest time for raw materials to go through the whole process and fulfill an order. And the safety stock is used to cover the uncertainties faced in the 78 days' lead time. So the problem is actually to decide how to distribute the 78 days to 5 stages, namely determining how many days each stage should cover.

Table 5.1: Optimal Safety Stock Solution for Top 80 SKUs

SKU	Safety Stock Level							SKU	Safety Stock Level						
	T	F	N	4	3	2	1		T	F	N	4	3	2	1
186002350	0	0	0	0	314	0	307	186003033	0	0	0	0	43	0	36
186002853	0	0	0	0	20	0	68	186000442	0	0	0	0	25	0	46
186002877	0	0	0	0	18	0	42	186000492	0	0	0	0	46	0	53
186002884	0	0	0	0	18	0	45	186003034	0	0	0	0	88	0	95
186003538	0	0	0	0	62	0	124	186003045	0	0	0	0	33	0	43
186004495	0	0	0	0	12	0	10	186003116	0	0	0	0	54	0	95
186005296	0	0	0	0	20	0	42	186003335	0	0	0	0	16	0	33
186005965	0	0	0	0	10	0	33	186005270	0	0	0	0	20	0	46
186007093	0	0	0	0	7	0	14	186000494	0	0	0	0	38	0	35
186007114	0	0	0	0	6	0	19	186000496	0	0	0	0	91	0	114
186002349	0	0	0	0	41	0	35	186003010	0	0	0	0	17	0	29
186002352	0	0	0	0	267	0	374	186003117	0	0	0	0	103	0	135
186002854	0	0	0	0	41	0	66	186003031	0	0	0	0	76	0	65
186002878	0	0	0	0	29	0	35	WAT046970	0	0	0	0	48	0	46
186002885	0	0	0	0	30	0	44	WAT046980	0	0	0	0	130	0	122
186003461	0	0	0	0	19	0	26	WAT052885	0	0	0	0	26	0	24
186003533	0	0	0	0	26	0	39	WAT086344	0	0	0	0	127	0	119
186003539	0	0	0	0	117	0	103	WAT027324	0	0	0	0	86	0	74
186003837	0	0	0	0	23	0	35	WAT066224	0	0	0	0	90	0	77
186004496	0	0	0	0	14	0	17	WAT066220	0	0	0	0	37	0	31
186004801	0	0	0	0	29	0	31	186002559	0	0	0	0	29	0	45
186005297	0	0	0	0	40	0	43	WAT045905	0	0	0	0	81	0	155
186007095	0	0	0	0	15	0	23	WAT045995	0	0	0	0	27	0	31
186007116	0	0	0	0	9	0	20	186000180	0	0	0	0	26	0	27
186002353	0	0	0	0	102	0	90	186001342	0	0	0	0	22	0	33
186003376	0	0	0	0	27	0	35	186002554	0	0	0	0	64	0	55
186003534	0	0	0	0	29	0	38	186003729	0	0	0	0	53	0	45
186003540	0	0	0	0	56	0	55	WAT200632	0	0	0	0	62	0	74
186004742	0	0	0	0	37	0	45	186000112	0	0	0	0	43	0	51
186004802	0	0	0	0	29	0	25	186001346	0	0	0	0	15	0	18
186005298	0	0	0	0	39	0	51	186002560	0	0	0	0	44	0	52
186007096	0	0	0	0	10	0	9	186003748	0	0	0	0	16	0	23
186005225	0	0	0	0	42	0	56	WAT054270	0	0	0	0	37	0	61
186007378	0	0	0	0	5	0	15	WAT054275	0	0	0	0	143	0	226
186005226	0	0	0	0	40	0	34	186003975	0	0	0	0	76	0	96
186002979	0	0	0	0	9	0	8	186003976	0	0	0	0	28	0	60
186003028	0	0	0	0	19	0	16	PSS830615	0	0	0	0	74	0	66
186000480	0	0	0	0	36	0	30	PSS831915	0	0	0	0	70	0	60

On the uncertainty side, it is favorable to use one stage to cover all the lead time, because the uncertainty increases with the square root of the lead time. For example, considering a two-stage supply chain with cumulative lead time of 100 periods, and the standard deviation of forecasting for one period is σ. If covered by one stage, the variability needs to be covered is 10σ, while if equally covered by two stages, the total variability for the supply chain is $\sqrt{50}\sigma + \sqrt{50}\sigma \approx 14\sigma$.

On the cost side, it is clear that all the safety stock should be held in the lowest cost stage. However, the structure of supply chain decides that stage cost is increasing from upstream to downstream. Also, upstream stage can push all inventories to downstream stage by setting the outgoing service time equal to it's replenishment time; but downstream stage cannot push all inventories to upstream stage because incoming service time is nonnegative.

With these two observations, we intuitively know that if adjacent stages only have a small difference in stage cost, it's better to put all the safety stock into the downstream stage, because the reduction in uncertainty overrides the increased unit holding costs. However, when two stages have a large discrepancy in stage costs, then this strategy is not favorable because the benefits of reduced uncertainty may not be enough to cover the increased unit holding cost.

In our situation, since stage 5, stage 4, and stage 3 have very small cost increases, the optimal choice is to push all the safety stock to stage 3. Stage 2 is the most value-add stage, so safety stock in stage 3 shouldn't be pushed to stage 2. Stage 2 and stage 1 again have small cost difference, so the safety stock in stage 2 is pushed to stage 1. Thus, the optimal solution is that we only have safety stock in stage 3 and stage 1.

To determine how is Waters performing in the supply chain, we can make comparisons between the optimal solution and Waters' current situation. For finished columns, it's easy to directly make comparisons. For Column Assemblies, since they are also possibly used in other SKUs, we need to modify the numbers in stage 3 by the scale of demand volume. For example, for the top 80 SKUs, the ones using Column Assembly "289007997" are "186005225" and "186007378". From Table 5.1, the safety stock level for "289007997" is the sum of safety stock level in stage 3 for "186005225" and "186007378": $42 + 5 = 47$ units. If the annual demand for "186005225" and "186007378" together is 100 units, and the annual usage for "289007997" is 200 units, then we get the final modified safety stock level for "289007997": $47 \times \frac{200}{100} = 94$.

One thing to notice is that near the end of our thesis project, Waters also launched a "SAP-APO" project, by which they were implementing a more advanced Material Resource Planning (MRP) system instead of the old SAP system. The "SAP-APO" system itself is able to do some improvements on the current safety stock policy. So we are also able to provide comparisons of the supply chain before APO, after APO, and the optimal situation. Unfortunately, the installation of "SAP-APO" system covered over the historical data, so for any data that we didn't gather before installation of "SAP-APO", we cannot have them from anywhere and have to use "N/A" to indicate the data missing.

Table 5.2 shows the comparison of safety stock levels in distribution centers (Stage 1) for top 80 SKUs. It is clear that Waters originally has too much safety stock lying in distribution centers and the implementation of "SAP-APO" system already can cut down a lot of redundant safety stocks. However, compared to the most economical solution, Waters still has much room for further safety stock reduction. According to the value of the Finished Columns, the capital savings in our standard value system are 59000 after implementing "SAP-APO". If the optimal solution is adopted, then we can save an additional 96431 beyond the savings from the "SAP-APO" situation.

For Ireland facility (Stage 2), Waters always follows a no safety stock policy, so it's already the same as optimal results.

Table 5.2: Distribution Center's Safety Stock Comparison for Top 80 SKUs

SKU	Before APO	After APO	Optimal	SKU	Before APO	After APO	Optimal
186002350	1779	1334	307	**186003033**	70	53	36
186002853	118	89	68	**186000442**	101	76	46
186002877	104	78	42	**186000492**	174	131	53
186002884	107	80	45	**186003034**	339	254	95
186003538	331	249	124	**186003045**	122	91	43
186004495	60	46	10	**186003116**	198	149	95
186005296	113	85	42	**186003335**	61	47	33
186005965	61	47	33	**186005270**	65	49	46
186007093	52	39	14	**186000494**	122	92	35
186007114	33	25	19	**186000496**	258	194	114
186002349	51	38	35	**186003010**	53	40	29
186002352	1014	761	374	**186003117**	312	234	135
186002854	147	111	66	**186003031**	322	242	65
186002878	109	83	35	**WAT046970**	108	82	46
186002885	119	90	44	**WAT046980**	281	212	122
186003461	90	68	26	**WAT052885**	58	44	24
186003533	90	69	39	**WAT086344**	290	218	119
186003539	446	336	103	**WAT027324**	506	380	74
186003837	88	66	35	**WAT066224**	173	130	77
186004496	53	40	17	**WAT066220**	66	49	31
186004801	118	89	31	**186002559**	117	88	45
186005297	134	101	43	**WAT045905**	322	242	155
186007095	57	43	23	**WAT045995**	122	92	31
186007116	34	26	20	**186000180**	65	50	27
186002353	216	162	90	**186001342**	52	40	33
186003376	43	33	35	**186002554**	149	111	55
186003534	50	38	38	**186003729**	124	93	45
186003540	118	89	55	**WAT200632**	156	118	74
186004742	84	63	45	**186000112**	168	126	51
186004802	69	52	25	**186001346**	61	46	18
186005298	77	59	51	**186002560**	182	138	52
186007096	20	16	9	**186003748**	67	51	23
186005225	224	168	56	**WAT054270**	154	116	61
186007378	25	19	15	**WAT054275**	566	425	226
186005226	99	75	34	**186003975**	431	324	96
186002979	15	12	8	**186003976**	145	109	60
186003028	65	49	16	**PSS830615**	87	66	66
186000480	52	40	30	**PSS831915**	95	71	60
Total Standard Capital Saving						*59000*	*96431*

Table 5.3 shows the comparison of safety stock levels in Ireland facility (Stage 3) for the Column Assemblies used in the top 80 SKUs. But we lack the data before implementing the "SAP-APO" system. Again Waters still has excessive safety stocks even after implementation of "SAP-APO", compared to the optimal solution. The capital savings we can get, after normalized in our standard value system, is 34251 from this stage.

Table 5.3: Ireland Facility's Safety Stock Comparison for Column Assemblies

Material	Before APO	After APO	Optimal
289000488	N/A	230	50
289000491		100	57
289000493		500	130
289000494		350	167
289000495		1000	422
289000496		800	386
289001600		60	16
289001855		1000	539
289001874		2000	780
289002445		580	536
289002583		200	74
289004057		1500	648
289007997		450	120
289007998		100	65
PSS550425		600	411
WAT015923		1000	446
WAT015924		800	129
WAT092-03		800	229
WAT092-04		1000	374
WAT551-04		400	131
WAT551-05		300	123
WAT551-06		700	327
Total Standard Capital Saving			*34251*

Table 5.4 shows the comparison of safety stock levels in Milford facility (Stage 4) for

the Column Assemblies used in the top 80 SKUs. We lack some of the data before implementing the "SAP-APO" system. For this data, it's very interesting that there are huge reductions in the safety stock after implementing "SAP-APO" systems, summing up to a reduction of 40000 standard capital cost and 8800 units of materials. But according to the interviews with the planners and managers at Milford facility, we learned that this big change has caused lots of chaos to them, and it seems that everyone on the manufacturing floor is in mild panic about this change, because the safety stock they relied on before suddenly disappeared.

Table 5.4: Milford Facility's Safety Stock Comparison for Column Assemblies

Material	Before APO	After APO	Optimal
289000488	0	0	0
289000491	0	0	0
289000493	288	0	0
289000494	0	0	0
289000495	1440	0	0
289000496	1056	0	0
289001600	0	0	0
289001855	2880	0	0
289001874	2520	0	0
289002445	576	0	0
289002583	80	0	0
289004057	N/A	0	0
289007997	0	0	0
289007998	0	0	0
PSS550425	N/A	0	0
WAT015923	N/A	0	0
WAT015924	N/A	0	0
WAT092-03	N/A	0	0
WAT092-04	N/A	0	0
WAT551-04	N/A	0	0
WAT551-05	N/A	0	0
WAT551-06	N/A	0	0

Finally, Table 5.5 shows the comparison of safety stock levels in Milford facility (Stage 5) for the Components used in the top 80 SKUs. It is clear that the "SAP-APO" system hasn't made any improvements on this part, since after the implementation of "SAP-APO", all the safety stock levels for these components still remain the same. In our optimal solution, this stage doesn't need to hold any safety stock, which also brings very huge capital savings. The total capital saved in the amount of standard cost is 94091.

Based on all of the above analysis, we estimate that beyond the implementation of "SAP-APO", Waters still has room for improvements in safety stock levels across the supply chain. The additional expected total standard capital savings are 224773.

Table 5.5: Milford Facility's Safety Stock Comparison for Components

Column Tube	Before APO	After APO	Optimal	Filter	Before APO	After APO	Optimal
405007034	1350	1350	0	WAT022978	100	100	0
405009466	30	30	0	289002235	100	100	0
405011567	0	0	0	WAT015931	2500	2500	0
405011937	0	0	0	289004437	2500	2500	0
405011938	0	0	0	289004545	1350	1350	0
405011939	0	0	0	289004915	0	0	0
405012517	0	0	0	289007723	0	0	0
405012518	0	0	0	PSS669018	2500	2500	0
PSS550425	0	0	0	WAT015935	0	0	0
WAT015927	300	300	0	**End Nut**	**Before APO**	**After APO**	**Optimal**
WAT015928	300	300	0	405000613	0	0	0
WAT015960	0	0	0	405000614	0	0	0
WAT091-01	200	200	0	405003600	3000	3000	0
WAT091-02	0	0	0	405005625	1550	1550	0
WAT091-03	0	0	0	405006406	100	100	0
WAT091-04	250	250	0	405008848	500	500	0
WAT091-06	200	200	0	PSS613120	2000	2000	0
WAT098-03	0	0	0	WAT015930	1000	1000	0
WAT200-04	100	100	0	WAT015961	0	0	0
Total Standard Capital Saving							*94091*

Figure 5.1 shows an overview of the comparisons of the supply chain situations for top 80 SKUs in Waters before "SAP-APO" system, after "SAP-APO" system, and optimal solution.

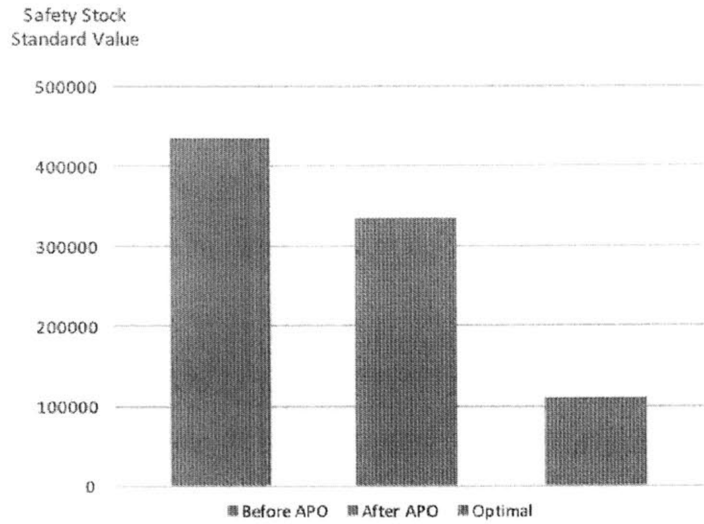

Figure 5.1: Overview of Capital Tied up in Safety Stock

5.2 Sub-Optimal Solution

From section 5.1, it is clear that compared to current situation, the suggested optimal solution is very radical not only in reductions of the amount of safety stock but also in change of place for safety stock. Originally there are 4 stages holding safety stock, while the optimal solution suggests that only 2 stages should hold safety stock. Although the potential benefits are huge for adopting the optimal solution, generally companies won't directly implement such a big change. On one hand, the managers who lose their safety stock may feel unsafe and may not be willing to take the change. On the other hand, different facilities have their own inventory budgets, so how to financially cooperate with each other is also a big problem. Due to these concerns, a sub-optimal solution may be preferable because it would be easier for everyone to buy in as a starting point for changes. It functions as a test for changes, and also helps people to be familiar with working under the new base-stock guaranteed service time supply chain policy. For 1 to 2 months' period, Waters can record its performance under the new policy, and if it works, bigger changes can be further implemented and gradually reach the optimal point.

In the sub-optimal solution, we constrain that stage 3 and stage 4 all quote 0 outgoing service time to the next stage, which means they are all covering their own lead time. And we only do a multi-echelon optimization for stage 5 and stage 4. Also, to be more conservative, we increase the safety factor from 2.05 to 2.33 (service level from 98% to 99%). Table 5.6 shows the list of sub-optimal results. We can see that in the sub-optimal solution, stage 4 and stage 3 all hold safety stocks, while stage 5 and stage 2 hold no safety stocks.

Table 5.6: Sub-Optimal Safety Stock Solution for Top 80 SKUs

SKU	Safety Stock Level							SKU	Safety Stock Level						
	T	F	N	4	3	2	1		T	F	N	4	3	2	1
186002350	0	0	0	315	168	0	349	186003033	0	0	0	43	23	0	41
186002853	0	0	0	20	11	0	77	186000442	0	0	0	25	13	0	52
186002877	0	0	0	18	10	0	48	186000492	0	0	0	46	24	0	60
186002884	0	0	0	18	10	0	51	186003034	0	0	0	88	47	0	108
186003538	0	0	0	62	33	0	141	186003045	0	0	0	33	18	0	49
186004495	0	0	0	12	6	0	11	186003116	0	0	0	54	29	0	108
186005296	0	0	0	20	11	0	47	186003335	0	0	0	16	9	0	38
186005965	0	0	0	10	5	0	38	186005270	0	0	0	20	11	0	52
186007093	0	0	0	7	4	0	16	186000494	0	0	0	38	20	0	40
186007114	0	0	0	6	3	0	21	186000496	0	0	0	91	49	0	129
186002349	0	0	0	41	22	0	40	186003010	0	0	0	17	9	0	33
186002352	0	0	0	267	143	0	425	186003117	0	0	0	103	55	0	154
186002854	0	0	0	41	22	0	75	186003031	0	0	0	77	41	0	74
186002878	0	0	0	29	16	0	40	WAT046970	0	0	0	49	26	0	52
186002885	0	0	0	30	16	0	51	WAT046980	0	0	0	130	70	0	139
186003461	0	0	0	19	10	0	29	WAT052885	0	0	0	26	14	0	27
186003533	0	0	0	26	14	0	45	WAT086344	0	0	0	127	68	0	136
186003539	0	0	0	117	63	0	117	WAT027324	0	0	0	86	46	0	84
186003837	0	0	0	23	13	0	40	WAT066224	0	0	0	90	48	0	87
186004496	0	0	0	14	8	0	19	WAT066220	0	0	0	37	20	0	36
186004801	0	0	0	29	15	0	35	186002559	0	0	0	29	15	0	51
186005297	0	0	0	40	21	0	48	WAT045905	0	0	0	81	43	0	176
186007095	0	0	0	15	8	0	26	WAT045995	0	0	0	27	14	0	35
186007116	0	0	0	9	5	0	23	186000180	0	0	0	26	14	0	30
186002353	0	0	0	102	55	0	103	186001342	0	0	0	22	12	0	38
186003376	0	0	0	27	14	0	40	186002554	0	0	0	64	34	0	63
186003534	0	0	0	30	16	0	44	186003729	0	0	0	53	28	0	51
186003540	0	0	0	56	30	0	63	WAT200632	0	0	0	62	33	0	84
186004742	0	0	0	37	20	0	51	186000112	0	0	0	43	23	0	58
186004802	0	0	0	29	15	0	29	186001346	0	0	0	15	8	0	20
186005298	0	0	0	39	21	0	57	186002560	0	0	0	44	23	0	59
186007096	0	0	0	10	5	0	10	186003748	0	0	0	16	9	0	26
186005225	0	0	0	42	23	0	64	WAT054270	0	0	0	37	20	0	69
186007378	0	0	0	5	3	0	17	WAT054275	0	0	0	143	77	0	256
186005226	0	0	0	40	21	0	39	186003975	0	0	0	76	41	0	110
186002979	0	0	0	9	5	0	9	186003976	0	0	0	29	15	0	69
186003028	0	0	0	19	10	0	18	PSS830615	0	0	0	74	40	0	75
186000480	0	0	0	36	19	0	35	PSS831915	0	0	0	70	37	0	68

Table 5.7 shows the comparisons of safety stock in stage 1. From the table we can see that for relatively small volume SKUs, the sub-optimal solution is very close to the APO number and some of them are even larger than APO number. However, we can still get significant capital savings from high volume SKUs. The capital savings in standard value from sub-optimal solution is 84986.

Table 5.7: Distribution Center's Sub-Optimal Solution for Top 80 SKUs

SKU	Before APO	After APO	Sub-Optimal	SKU	Before APO	After APO	Sub-Optimal
186002350	1779	1334	349	**186003033**	70	53	41
186002853	118	89	77	**186000442**	101	76	52
186002877	104	78	48	**186000492**	174	131	60
186002884	107	80	51	**186003034**	339	254	108
186003538	331	249	141	**186003045**	122	91	49
186004495	60	46	11	**186003116**	198	149	108
186005296	113	85	47	**186003335**	61	47	38
186005965	61	47	38	**186005270**	65	49	52
186007093	52	39	16	**186000494**	122	92	40
186007114	33	25	21	**186000496**	258	194	129
186002349	51	38	40	**186003010**	53	40	33
186002352	1014	761	425	**186003117**	312	234	154
186002854	147	111	75	**186003031**	322	242	74
186002878	109	83	40	**WAT046970**	108	82	52
186002885	119	90	51	**WAT046980**	281	212	139
186003461	90	68	29	**WAT052885**	58	44	27
186003533	90	69	45	**WAT086344**	290	218	136
186003539	446	336	117	**WAT027324**	506	380	84
186003837	88	66	40	**WAT066224**	173	130	87
186004496	53	40	19	**WAT066220**	66	49	36
186004801	118	89	35	**186002559**	117	88	51
186005297	134	101	48	**WAT045905**	322	242	176
186007095	57	43	26	**WAT045995**	122	92	35
186007116	34	26	23	**186000180**	65	50	30
186002353	216	162	103	**186001342**	52	40	38
186003376	43	33	40	**186002554**	149	111	63
186003534	50	38	44	**186003729**	124	93	51
186003540	118	89	63	**WAT200632**	156	118	84
186004742	84	63	51	**186000112**	168	126	58
186004802	69	52	29	**186001346**	61	46	20
186005298	77	59	57	**186002560**	182	138	59
186007096	20	16	10	**186003748**	67	51	26
186005225	224	168	64	**WAT054270**	154	116	69
186007378	25	19	17	**WAT054275**	566	425	256
186005226	99	75	39	**186003975**	431	324	110
186002979	15	12	9	**186003976**	145	109	69
186003028	65	49	18	**PSS830615**	87	66	75
186000480	52	40	35	**PSS831915**	95	71	68
Total Standard Capital Saving						*59000*	*84986*

Table 5.8 shows the comparisons of safety stock in stage 3. Because stage 4 now covers part of the safety stock for stage 3, the sub-optimal solution even leads to less safety stock in stage 3 than the optimal solution, bringing 47139 capital savings.

Table 5.8: Ireland Facility's Sub-Optimal Solution for Top 80 SKUs

Material	Before APO	After APO	Sub-Optimal
289000488	N/A	230	27
289000491		100	31
289000493		500	70
289000494		350	89
289000495		1000	226
289000496		800	207
289001600		60	9
289001855		1000	289
289001874		2000	418
289002445		580	287
289002583		200	39
289004057		1500	347
289007997		450	64
289007998		100	35
PSS550425		600	220
WAT015923		1000	239
WAT015924		800	69
WAT092-03		800	123
WAT092-04		1000	200
WAT551-04		400	70
WAT551-05		300	66
WAT551-06		700	175
Total Standard Capital Saving			*47139*

Table 5.9 shows the comparisons of safety stock in stage 4. Compared to the solution provided by "SAP-APO" system and the optimal solution, the sub-optimal solution offers a mild change. Actually compared to the safety stock situation before "SAP-APO" system,

the sub-optimal solution even increased some of the materials' safety stock level. This probably suggests that Waters' original safety stock policy was a little biased. Although the total amount of safety stock is large, most of it is assigned to a few materials, while some small volume materials have no safety stock and might have problems if demand exceeds forecast.

Table 5.9: Milford Facility's Sub-Optimal Solution for Top 80 SKUs

Material	Before APO	After APO	Sub-Optimal
289000488	0	0	50
289000491	0	0	57
289000493	288	0	131
289000494	0	0	167
289000495	1440	0	423
289000496	1056	0	387
289001600	0	0	16
289001855	2880	0	541
289001874	2520	0	782
289002445	576	0	537
289002583	80	0	74
289004057	N/A	0	650
289007997	0	0	121
289007998	0	0	65
PSS550425	N/A	0	412
WAT015923	N/A	0	447
WAT015924	N/A	0	129
WAT092-03	N/A	0	229
WAT092-04	N/A	0	374
WAT551-04	N/A	0	131
WAT551-05	N/A	0	123
WAT551-06	N/A	0	327

Figure 5.2 shows an overview of the comparisons of the 4 supply chain situations for top 80 SKUs.

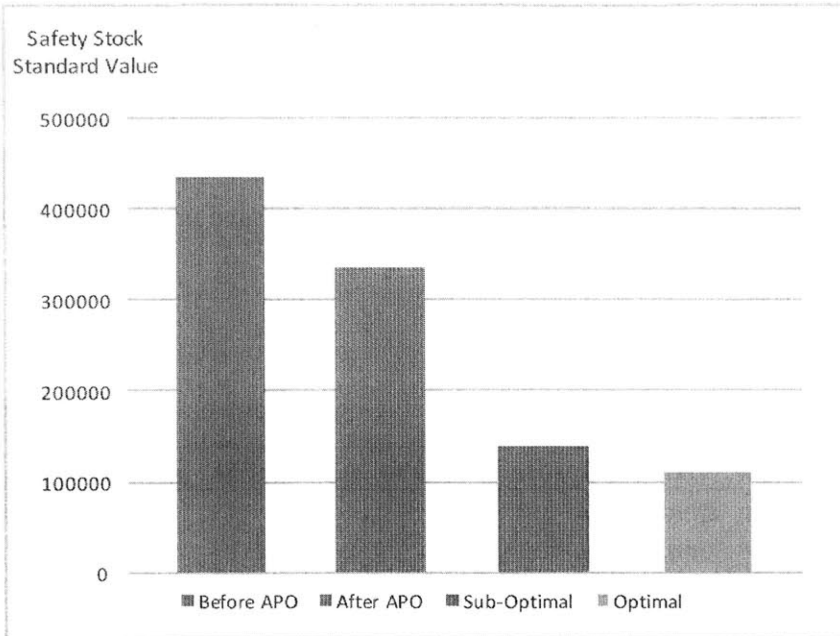

Figure 5.2: Overview of Capital Tied on Safety Stock for Four Situations

6 Conclusions and Future Work

This thesis describes the method of applying a multi-echelon supply chain model for positioning safety stock in a multiple part manufacturing company, Waters Corporation. We reduce the Waters supply chain into a serial 5-stage supply chain model for the Column Assembly division. The basic assumption is that each stage follows a common review period for a base-stock policy, and instead of replenishing for what has been consumed, here each stage makes orders based on forecasting of next period. We assume that each stage is quoting a guaranteed service time to its next stage that within the time limit 100% service will be provided. We also assume that the demands from customers are bounded so that the forecast error for each period will not be too large, and in our case a safety factor of 2.05 is applied, suggesting a 98% service level.

In the implementation step, we show how to set the parameters in our model including the granularity and stage cost. Most importantly, through analysis it is revealed that Waters' forecasting system has a strong bias so that most of the time the forecast will overestimate the demand. We then discuss how to apply the forecast bias adjustment method developed by M.Manary and S.Willems (2008) to our situation.

Finally, we present the results of applying the model to Waters Corporation. We make comparisons between Waters' original supply chain, Waters' improved supply chain, and our optimal solution. We find that although certain improvements have been made on Waters supply chain, they can still further have significant capital savings from inventories. We also observe the problems caused by the implementation of an improved "SAP-APO" systems. Based on this a sub-optimal solution is suggested, aiming to be a test and starting point for Waters to apply changes to its supply chain.

However, as with any research, we end up with some new issues and limitations, and based on them we point out some of the future work that can be done.

Compared to the safety stock positioning problem, it is more important for Waters to streamline its internal planning procedure to be more similar to our assumptions. Actually our assumptions of guaranteed service time and base-stock review are very important and useful in applications, because they represent a coordinated and well-organized internal operation. For example, currently Waters is doing forecasting but is not using forecasting as guidance for production. The production departments are mainly responding to observed demands, which potentially increases the responsive time. Also, departments lack internal cooperation and communications. Instead of 100% fulfill the orders from downstream, upstream stages can evaluate the demands from downstream and postpone the orders they think unnecessary. Finally, the 7 days' review period is also a little challenging for Waters since currently they are planning on monthly period. So before implementing the safety stock policy provided in this thesis, it is more important for Waters to build a more integrated internal operation logistic. Otherwise just reducing inventories without having a corresponding operation system will only cause more serious problems.

In the thesis we only consider the serial supply chain under the Column Assembly division. So this work only optimizes part of the Waters supply chain, assuming the Chemistry Powder division still remains the same. Adding the Chemistry Powder division into our current model not only increases the number of materials we need to consider, but also makes the serial supply chain model into a more complex model. So we leave this more complicated problem for future research.

In our assumptions we apply no constraints on manufacturing facilities' capacities. However, this may not be a true story in reality. Although the managers in manufacturing facilities declare that they have no capacity issues, the fact that sometimes they need to postpone some orders may suggest that they may have capacity constraints. Actually, to make the supply chain model more applicable and credible, capacity limits should be added into the model. But this may make the model too mathematically complicated. So at this time we just assume infinite capacities for manufacturing facilities.

In our current frame for the supply chain model, we don't consider two important features in ordering: lot size and scrap rate. It is always economical for manufacturing facilities to make products in certain lot size, and a lot won't deliver the same amount of products every time. But this highly relies on real situations and can hardly be incorporated into a model. A feasible way to deal with this concern is that in ordering, always round the ordering quantity up so as to be multiples of the lot size. For the scrap rate problem, we can incorporate it into the uncertainties to be covered by safety stock.

Reference

[1] Simchi-Levi, D., Kaminsky, P., and Simchi-Levi, E., 2008, *Designing and Managing the Supply Chain: Concepts, Strategies, and Case Studies*, 3rd Ed., McGraw-Hill/Irwin, Boston, MA.

[2] Graves, S. C. 1988. Safety stocks in manufacturing systems. *J. of Manufacturing and Oper. Management* **1** 67–101.

[3] Graves, S. C., S. P. Willems. 2000. Optimizing strategic safety stock placement in supply chains. *Manufacturing Service Oper. Management* 2(1) 68–83.

[4] S. P. Willems. 1996. Strategic safety stock placement in supply chains. *Proceedings of the 1996 MSOM Conference*, Hanover, NH.

[5] Manary, M. P. and S. P. Willems. Setting Safety Stock Targets at Intel in the Presence of Forecast Bias. *Interfaces*, Mar-Apr 2008, Vol. 38, No. 2, pp. 112-122.

[6] Graves, S. C., S. P. Willems. 2008. Strategic inventory placement in supply chains: Non-stationary demand. *Manufacturing Service Oper. Management* 10(2) 278–287.

[7] Schoenmeyr and Graves. Strategic Safety Stocks in Supply Chains with Evolving Forecasts. *Manufacturing & Service Operations Management* 11(4), pp. 657–673