

Multi-Tier Inventory Systems With Space Constraints

A Thesis
Presented to
The Academic Faculty

by

Stephanie A. Jernigan

In Partial Fulfillment
of the Requirements for the Degree
Doctor of Philosophy

School of Industrial and Systems Engineering
Georgia Institute of Technology
January 2004

Multi-Tier Inventory Systems With Space Constraints

Approved by:

John J. Bartholdi, III, Advisor

Steven T. Hackman

Gunter P. Sharp, Co-Advisor

John Vande Vate

Mark Ferguson

Date Approved: _____

ACKNOWLEDGEMENTS

I could not have completed this degree without the support and love of my family. It seems almost ridiculous to try to put into words all they have done for me. The biggest help of all was knowing that anywhere they were, there would be side-splitting laughter and people who would do anything for each other. My parents, Tom and Sandy Jernigan, have helped put graduate school in much-needed perspective at times when I was practically foaming at the mouth. My brothers, Mike and Doug Jernigan, and my sister-in-law, Sarah Thomas Jernigan, have been constant cheerleaders, each in unique, creative ways.

There are a couple of almost-family members to whom I am also indebted. Without the friendship and encouragement of Eva Regnier, I would not have made it through school, and her witty, creative personality made everything much more fun. I also have also been very lucky to have Sam Ransbotham in my life. His support during the most hectic part of my degree program has meant so much to me. And our impromptu brainstorming sessions remind me why I wanted a Ph.D. in the first place. I am only sorry the graduate office would not accept his idea of a pop-up dissertation.

I also want to express my appreciation for my academic supporters. My advisor, John Bartholdi, has served as a model of a clear and elegant thinker. If I have absorbed even a small fraction of that clarity and elegance, my time under his advisement will have been well-spent. I am immensely grateful for the many opportunities he has given me to expand my capabilities. My co-advisor, Gunter Sharp, cheerfully made time to help me clarify very rough research ideas, which was no small task. My other committee members—Mark Ferguson, Steven Hackman, and John Vande Vate—as well as Jane Ammons, have provided valuable advice and encouragement. Thanks also go to Cristina Gigola for helpful discussions early in the research process. Very belated thanks are due to Benita Albert and Mark Krusemeyer, who helped nurture my love for mathematics long before this thesis was even a twinkle in my eye.

I was able to devote myself to this research through the generous financial support of Ford Motor Company.

TABLE OF CONTENTS

ACKNOWLEDGEMENTS	iii
LIST OF TABLES	viii
LIST OF FIGURES	ix
SUMMARY	xi
1 INTRODUCTION	1
1.1 A multi-tier inventory system in the Avon warehouse	3
1.1.1 Manual lines	4
1.1.2 Cart pick	6
1.1.3 A-frame	7
1.2 Why do multi-tier inventory systems exist?	10
1.3 Why study multi-tier inventory systems?	12
1.4 Focus of this research	12
1.5 Model	13
1.5.1 Storage modes	13
1.5.2 Skus	14
1.5.3 How to characterize a slotting	14
1.5.4 Restocking protocol	16
1.6 Literature Review	19
1.6.1 Research on multi-tier inventory systems	19
1.6.2 Research on the multi-period forward-reserve problem	23
1.7 Organization of results	24
2 SLOTTING A THREE-TIER INVENTORY SYSTEM TO MINIMIZE RESTOCKING COSTS	25
2.1 Allotting space to skus to minimize restock costs	26
2.1.1 Allotting space to skus in Mode I	26
2.1.2 Allotting space to skus in Mode F	26
2.2 Assigning skus to flowpaths	28

3	SLOTING THE AVON INVENTORY SYSTEM TO MINIMIZE PICKING AND RESTOCKING COSTS	31
3.1	Equivalence to a two-tier inventory system	32
3.2	Optimal allocation of storage resources	34
4	MINIMIZING PICKING AND RESTOCKING COSTS IN GENERAL MULTI-TIER INVENTORY SYSTEMS	37
4.1	A multi-tier inventory system is equivalent to a multi-mode inventory system	37
4.2	Finding the best well-ranked slotting	39
5	APPLICATION: SLOTING THE AVON INVENTORY SYSTEM .	42
5.1	Estimating model parameters	42
5.1.1	Computing the volume of each storage mode	44
5.1.2	Computing picking and restocking costs	44
5.1.3	Computing picks and flow	45
5.2	Results	45
5.3	Recommendations for slotting the order fulfillment area	46
5.4	Allocation of storage resources in the Avon warehouse	50
6	SLOTING A FORWARD-RESERVE SYSTEM OVER MULTIPLE PERIODS	51
6.1	Model	52
6.2	Heuristics from the single-period forward-reserve problem	54
6.3	How to evaluate a multi-period assignment	55
6.3.1	Intuitive justification of ranking by sequence viscosity	56
6.3.2	Theoretical justification of ranking by sequence viscosity	59
6.4	Multiperiod heuristics	61
6.4.1	The T -period heuristic	62
6.4.2	The best-viscosity heuristic	65
6.4.3	The best-viscosity-improvement heuristic	66
6.4.4	The best sequence heuristic	68
6.5	Bounding the solution of the multi-period forward-reserve problem	72
6.6	Case Study	72
6.6.1	Profile of the Avon warehouse	74

6.6.2	Comparison of heuristics	74
6.6.3	Implications for the Avon warehouse	78
6.7	Future Research	79
6.7.1	Incorporating different types of reassignment costs	79
6.7.2	Considering multi-mode inventory systems over multiple periods . .	80
7	CONCLUSIONS	81
	REFERENCES	83

LIST OF TABLES

Table 5.1	Profile of campaign 4 in the Avon warehouse	42
Table 5.2	Actual and optimal solutions for campaign 4	45
Table 6.1	Example sku to show how sequence viscosity is computed	56
Table 6.2	Example sku for intuitive justification of sequence viscosity	57
Table 6.3	Characteristics of heuristics for the multi-period forward-reserve problem	63
Table 6.4	Example skus for the T -period heuristic and best-viscosity heuristic . . .	64
Table 6.5	Sequence viscosities of skus in table 6.4	66
Table 6.6	Example skus for the best-viscosity-improvement heuristic	67
Table 6.7	Example sku for the best sequence heuristic	70
Table 6.8	Sequences for sku a : Iteration 1	70
Table 6.9	Sequences for sku a : Iteration 2	70
Table 6.10	Pick savings in an average period for heuristics	77
Table 6.11	Skus moved to the A-frame in an average period	77
Table 6.12	Maximum error of the solution of the best sequence heuristic	77
Table 6.13	Cost savings per campaign if future forecasts are known	79

LIST OF FIGURES

Figure 1.1 A multi-tier inventory system.	1
Figure 1.2 A forward-reserve inventory system	2
Figure 1.3 A multi-mode inventory system	3
Figure 1.4 Ways to restock the manual lines area	5
Figure 1.5 Pick-to-light station	5
Figure 1.6 Ways to restock the cart pick area	6
Figure 1.7 Cross-section view of A-frame, with safety rail.	7
Figure 1.8 Dispenser channels on A-frame	8
Figure 1.9 Ways to restock the A-frame area	9
Figure 1.10 Storage modes in the Avon warehouse and ways they are restocked	10
Figure 1.11 A three-tier inventory system with costs.	14
Figure 1.12 Skus assigned to the four possible flowpaths in a three-tier system.	15
Figure 1.13 Decision process when tier q receives a restock request	17
Figure 1.14 Ways to restock the A-frame area	18
Figure 1.15 The Bartholdi-Hackman multi-mode inventory system.	23
Figure 2.1 The inventory system for which we minimize restock costs	25
Figure 2.2 The space in Mode F divided between flowgroups.	27
Figure 3.1 The simplified Avon inventory system.	31
Figure 3.2 Example of equivalent inventory systems	32
Figure 3.3 Two-tier inventory system equivalent to the simplified Avon inventory system	33
Figure 3.4 Avon system where all flowrack is considered one mode.	34
Figure 3.5 A two-mode system equivalent to the Avon system with one tier of flowrack.	35
Figure 4.1 Flowpath with Q tiers.	38
Figure 4.2 Multi-tier inventory system where space in each tier has been allotted to flowgroups.	38
Figure 4.3 Multi-mode inventory system equivalent to system in Figure 4.2.	39
Figure 5.1 Piece forecast of the top 500 skus, in descending order of sales.	43
Figure 5.2 Configuration of Avon's warehouse in Suwanee, GA.	43
Figure 5.3 A cost breakdown of the actual and optimal solutions.	47

Figure 5.4 Total cost of picking and restocking if more skus were dispensable.	49
Figure 6.1 Forward-reserve inventory system in the multi-period forward-reserve problem.	52
Figure 6.2 Space occupied by sku i in the forward mode over three periods.	58
Figure 6.3 Configuration of Avon's warehouse in Suwanee, GA.	73
Figure 6.4 Number of skus in the Avon warehouse by period	75
Figure 6.5 Pieces sold from the Avon warehouse by period	75
Figure 6.6 Distribution of skus by pieces sales in a typical campaign	75
Figure 6.7 Cost savings per campaign from several heuristics	76
Figure 6.8 Cost savings per campaign from multi-period heuristics	77

SUMMARY

In the warehouse of a large cosmetics company, a mechanized order picker is restocked from nearby shelving, and the shelving is restocked from bulk storage, forming a three-tier inventory system. We consider such multi-tier inventory systems and determine the storage areas to which to assign items, and the quantities in which to store them, in order to minimize the total cost of picking items and restocking storage locations. With this research, we increase the number of inventory systems for which simple search algorithms find a provably near-optimal solution. The model and method were tested on data from the Avon Products distribution center outside Atlanta; the solution identified by the model reduced picking and restocking costs there by 20%.

The sales forecasts of items stored in the warehouse may change, however, and new items will be introduced into the inventory system and others removed. To account for these changes, some warehouses may periodically reassign items to storage areas and recompute their storage quantities. These reassignment activities account for additional costs in the warehouse. The second focus of this research examines these costs over several time periods in a simple multi-tier inventory system. We develop heuristics to determine the storage areas to which to assign items and the quantities in which to store them in each time period, in order to minimize the total cost of picking items, restocking storage locations, and reassigning skus over multiple periods.

CHAPTER 1

INTRODUCTION

The most basic type of inventory system a warehouse can have is a single storage mode: each stock-keeping unit, or *sku* (pronounced “skew”), is picked from that mode, and replenishment stock is stored there after it is received into the warehouse. An example is a warehouse that has only pallet rack—even if a sku has a only single case present in the warehouse, the case is placed on a pallet and stored in a pallet location.

A warehouse that picks some skus with an automatic picking device may have a more complicated inventory system, as shown in figure 1.1. The automatic device holds a small quantity of each sku, and this supply is replenished from a supply stored in case flow rack located nearby. We say that the automatic picker is *restocked* from the case flow rack. The supply in the case flow rack is replenished from bulk storage: this path is shown by the solid arrows in figure 1.1. Skus in the warehouse that are not picked from the automatic picking device are picked from shelving that is restocked from bulk storage, as shown by the dotted arrows in figure 1.1. This is an example of a multi-tier inventory system. Observe that no skus are picked from the case flow rack or bulk storage; these storage modes serve only to restock the A-frame.




Figure 1.1: A multi-tier inventory system.




Figure 1.2: A forward-reserve inventory system

Definition 1.1 *A multi-tier inventory system is a set of storage modes in a warehouse where associated with each storage mode is*

- *a subset of modes that replenish it*
- *a designation of whether or not skus can be picked from the mode.*

In the inventory system in figure 1.1, for example, the storage mode consisting of case flow rack is replenished from bulk storage only, and skus can be picked from the mode. The commonly known *forward-reserve* inventory system, shown in figure 1.2, is another example of a multi-tier inventory system. In this case, there are two storage modes: a reserve mode where skus are typically stored in bulk quantities, and a forward mode where skus are typically stored in smaller quantities. Skus may be picked from either the forward mode or the reserve mode, and the supply of skus in the forward mode is replenished from the reserve mode.

The *multi-mode* inventory system, an extension of the forward-reserve system with several forward modes, is also a multi-tier inventory system. As shown in figure 1.3, skus may be picked from either the reserve mode or one of several forward modes, and the supply of skus in each forward mode is replenished from the reserve mode.

Storage modes in both the forward-reserve and multi-mode systems are restocked directly from a bulk storage mode, but this is not the case in more complex multi-tier inventory systems.

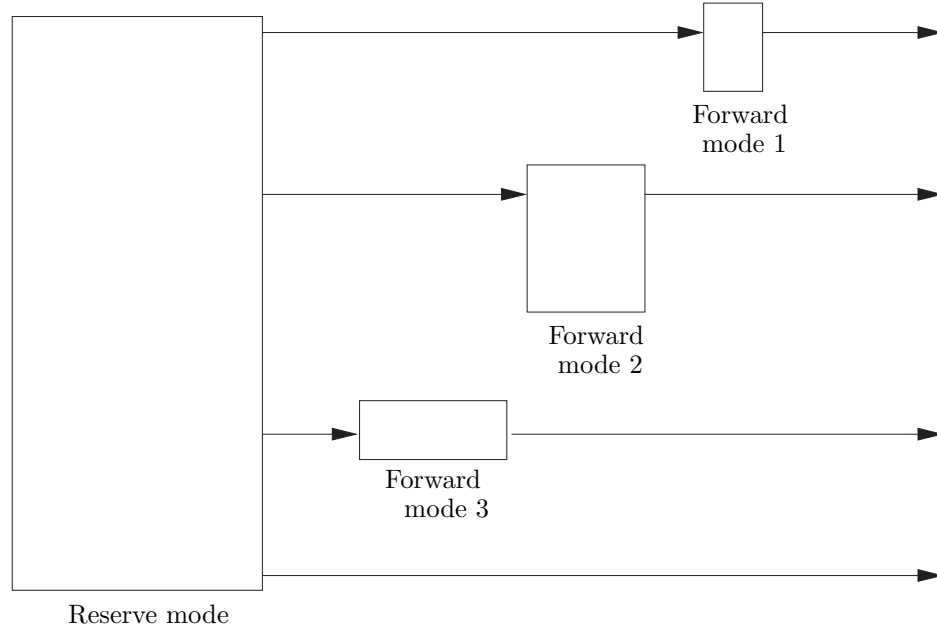


Figure 1.3: A multi-mode inventory system

1.1 A multi-tier inventory system in the Avon warehouse

The warehouse for Avon Products, Inc. outside Atlanta has a multi-tier inventory system. We describe the system below for two reasons: to highlight the issues that must be addressed to minimize operations costs, and to show why a warehouse may establish a multi-tier inventory system. We will use the Avon inventory system as an example to illustrate results throughout this research.

Avon Products, Inc. sells cosmetics and gift items; in 2002, they were the fifth biggest presence in the cosmetics industry by sales. They sell items through a network of 3.4 million sales representatives in 139 countries, each of whom collects the orders of their customers and places an order to her assigned warehouse.

The warehouse outside Atlanta fills orders for approximately 150,000 sales representatives in the southeastern United States. Sales representatives transmit their orders to the warehouse every two weeks, a period known as a *campaign*. Most representatives place only one order per campaign, but some with high sales volumes may place two or three. Between 150,000 and 170,000 orders will be processed in a typical two-week period. Each order requires 60 pieces, on average, meaning that over 10 million pieces are sold every two

weeks. Because each representative usually requests only one or two of each sku they order, these pieces represent almost the same number of picks. Piece picking is therefore a very labor-intensive activity in the Avon warehouse, and it is a priority of the management to reduce costs for this as much as possible.

Orders are picked in a section of the warehouse called, appropriately, the *order fulfillment area*. The order fulfillment area has a footprint of approximately 120,000 sq. ft., and must have a picking location for 6,000–8,000 skus. Because there is limited storage space, most skus have additional supply in the bulk storage area of the warehouse, known as the *back warehouse*, and storage locations in the order fulfillment area are restocked from the back warehouse as needed. Restocks happen constantly as skus are being picked; one-third of the labor force in the order fulfillment area is dedicated to restocking skus.

In 2002, only about 64% of the orders fulfilled in the Atlanta-area warehouse contained all the items the customer ordered; some of this was due to manufacturing backorders, but some was due to stockouts in the order fulfillment area. For this reason, Avon management wants to ensure that each sku has sufficient supply to avoid stockouts, and to use restocking labor as efficiently as possible.

Skus can be picked from one of three different storage modes: a section of pick-to-light flowrack called *manual lines*; a section of shelving and flow rack called *cart pick*; and from an automated picking machine called the *A-frame*.

1.1.1 Manual lines

The manual lines area consists of six conveyor belts where each belt is bordered by eight U-shaped stations of pick-to-light case flowrack, as shown in figure 1.4. A picker stands in each station waiting for totes to flow along the belt. When a tote stops at a picker's station, she places the items needed from her station in the tote and waits for the next tote to arrive. A picture of the station from the picker's point of view is shown in figure 1.5. In the Avon warehouse, each station has different skus from another station in the same line, but the six lines of stations all contain the same skus.

The flowrack in the manual lines area is restocked in one of two ways. Some skus

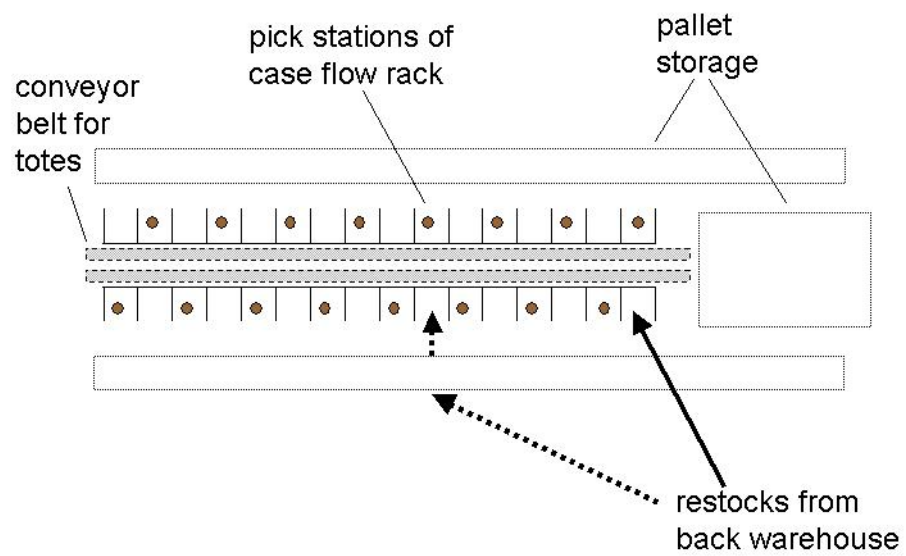


Figure 1.4: Ways to restock the manual lines area



Figure 1.5: Pick-to-light station




Figure 1.6: Ways to restock the cart pick area

are restocked directly from the back warehouse, as shown by the solid arrow in figure 1.4. Other skus have pallet locations located near the flowrack; the supply of these skus in the flowrack is replenished from the pallet locations, and the pallets are replenished from the back warehouse, as shown by the dotted arrow in figure 1.4.

1.1.2 Cart pick

The cart pick area consists of three different types of storage locations: shelving, case flow rack, and Barton rack, which holds bins approximately the size of a shoe box. Pickers in the cart pick area work in teams of two to push carts that hold 100 order totes from section to section, darting into the shelving to collect needed skus, and returning to the cart to place them in the correct tote. More labor per pick is required to pick from the cart pick area than any of the other areas.

All storage locations in cart pick are replenished directly from the back warehouse, as shown in figure 1.6.



Figure 1.7: Cross-section view of A-frame, with safety rail.

1.1.3 A-frame

The A-frame is an automated order picking machine. On each side of the A-frame is a row of dispenser channels, where each channel holds a supply of at most one sku. A conveyor belt passes underneath the row of channels, and each order is allocated an amount of space on the belt for one pass through the machine. As the space corresponding to an order passes underneath each channel, if the sku in the channel is needed for the order, the appropriate amount of the sku is dispensed onto the belt. The photo in figure 1.7 shows a cross-section view, and the photo in figure 1.8 shows dispenser channels on one side. The A-frame in the Avon warehouse has 2,432 channels and can pick up to 20 orders a minute. The channels are 52 inches high on one side of the machine and 72 inches on the other.

As seen in figure 1.8, each channel in the A-frame can hold a limited amount of each sku—typically 40-100 eaches. The limited amount of space available, together with the



Figure 1.8: Dispenser channels on A-frame




Figure 1.9: Ways to restock the A-frame area

speed with which the A-frame picks orders, means that the supply of items in the A-frame can be quickly depleted and must be restocked frequently. Because of this, each sku in the A-frame is restocked from a supply held in case flow rack located next to each side of the A-frame, as illustrated in figure 1.9. When a sku in the A-frame needs to be restocked, a replenisher need only to turn around to find a supply of the sku. Some skus are restocked in the flowrack directly from the back warehouse, as shown by the dashed arrow in figure 1.9. Other skus have pallet locations near this flowrack; these skus are restocked in the flowrack from the pallet locations, which are then replenished from the back warehouse, shown by the solid arrow in figure 1.9.

Picks in the A-frame require very little labor, and so it is natural that the warehouse management would like to store a lot of skus there. If too many skus are stored there, however, the most popular ones will have a small supply, meaning that replenishers will have to hover over that sku to keep it from stocking out. In that case, the cost of restocking the A-frame might outweigh the pick savings achieved.

Figure 1.10 shows a sketch of the different storage modes in the Avon warehouse and how

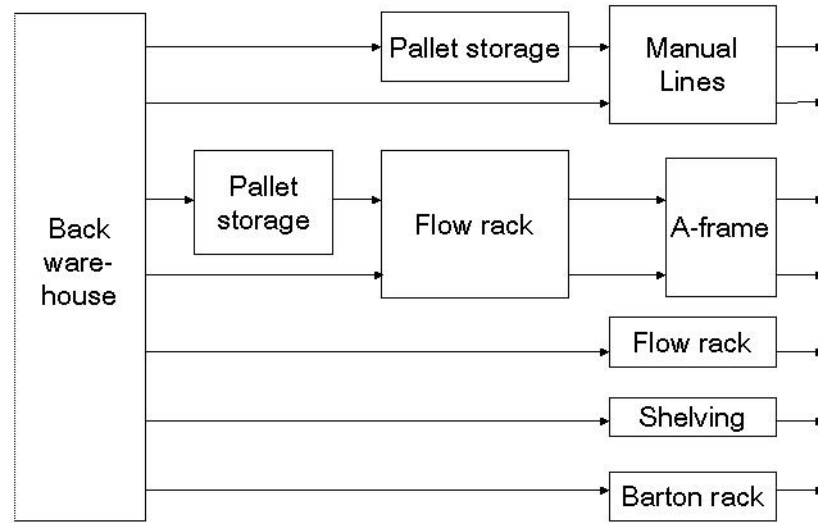


Figure 1.10: Storage modes in the Avon warehouse and ways they are restocked

skus are picked and replenished. For efficient operation of the warehouse, it is not enough to decide which skus are picked from which mode, but we must also investigate how they should be replenished in that mode, and we must balance picking costs with replenishment costs.

1.2 *Why do multi-tier inventory systems exist?*

After examining the multi-tier inventory system in use at the Avon warehouse, it is natural to wonder why such systems were created in the first place. The reason that forward modes are established in inventory systems is fairly easy to understand: they can reduce the cost of picking skus in the inventory system. They hold a small amount of many skus, creating a high pick density, and therefore lowering the cost per pick. The tradeoff is that these modes must be restocked from bulk storage modes. Intermediate modes are established between the forward modes and the bulk storage modes to reduce the cost of restocking the forward modes and to ensure that restocks are close at hand when needed.

Intermediate modes can reduce restock costs with little capital investment The A-frame in the Avon warehouse has a very small amount of storage area—it has only 2% of the storage area of the pick-to-light mode—but the cost per pick there is essentially zero, making it a desirable storage mode for small skus with many picks. Avon wants many skus to take advantage of its low pick savings, and for this to happen, the cost of restocking the machine needs be as low as possible.

The cost of restocking a storage mode can be reduced by either decreasing the cost per restock, or by increasing the storage volume of the mode, as shown in [1]. Adding A-frame storage is a significant capital investment that the warehouse was unable to make, however. Because of this, the warehouse instead reduced the cost per restock by establishing an intermediate mode of case flow rack within arm’s reach of the A-frame. The cost per restock of the A-frame from the case flow rack is very low, since the replenishers need only turn around to find additional supply of each sku. The cost of buying the flowrack is cheaper than adding channels to the A-frame.

Intermediate modes can reduce restock costs without increasing pick costs The manual lines mode in the Avon inventory system consists of 4-foot deep bays of flowrack equipped with pick-to-light sensors. It has a lower pick cost than the cart pick area, but the more skus assigned there, the less space each sku will get, and restock costs will increase. If we reduce the cost of restocking the manual lines mode, we can store more skus there and thus achieve greater pick savings.

One way to decrease the cost of restocking the manual lines mode is to increase the size of the manual lines mode. But increasing the size of the storage mode may decrease the pick savings in the mode: if the flowrack in the manual lines mode were extended and pickers had to travel farther to pick items, the savings per pick would decrease. Instead, Avon has established pallet positions on the floor near the manual lines mode to serve as an intermediate mode between the back warehouse and the manual lines mode. The skus stored in this intermediate mode are restocked in bulk, and their location in the manual lines mode can be restocked from the pallet locations at a very low cost. The total cost of

restocking these skus this way is lower than making frequent trips to the back warehouse to restock these skus in the flowrack.

1.3 Why study multi-tier inventory systems?

Multi-tier inventory systems are commonly established in warehouses to reduce the costs of picking and restocking skus and to ensure that storage locations are restocked promptly. Simple guidelines to make the operations of such systems as efficient as possible could be of benefit to many warehouses.

Multi-tier inventory systems can also be established inadvertently in warehouses. At Avon, workers who replenished the cart pick storage mode would store extra cases of the most popular skus on the floor near the shelving, effectively forming a storage mode between bulk storage and the cart pick mode. They did this to ensure that the supply of those skus would not run out before a replenishment could arrive from the bulk storage mode. Therefore there may be inventory systems for which a multi-tier configuration could improve operations.

1.4 Focus of this research

In a typical warehouse, 70% of the total operating cost is attributed to picking and restocking activities ([4]). To minimize these costs in an existing multi-tier inventory system, the warehouse manager can control which skus are stored in which storage modes, and how much of each sku is stored there. We refer to the process of deciding the storage modes to which to assign a sku and how much space to allot to the sku there as *slotting* the inventory system. Any result of this process is also referred to as a slotting. In common usage, slotting an inventory system means also that each sku is assigned an *exact location* within a mode [4], but since this research does not address assigning skus to exact locations, we use the term more generally.

In this research we first focus on slotting a multi-tier inventory system to minimize picking and restocking costs. We then additionally consider the costs of *reslotting* the warehouse: some warehouses do this regularly, to accomodate changes in sales profiles of

skus, to introduce new skus, and to remove skus that are no longer needed. The process can be quite labor-intensive, and thus when determining a slotting for a given time period, it can be important to consider the cost of future changes to the slotting. The second focus of this research is determining how to slot a warehouse to minimize the costs of picking, restocking, and reslotting.

1.5 Model

1.5.1 Storage modes

Our model of a multi-tier inventory system is based on the model of the forward-reserve system developed by Hackman and Rosenblatt in [7]. We assume that every multi-tier inventory system that we model will have the following characteristics:

1. Upon receipt into the warehouse, all skus are stored in a bulk storage mode which has sufficient supply of each sku to restock all other modes as needed. This mode, called the reserve mode, is denoted mode R .
2. Each storage mode in the system has a known storage capacity, denoted V_m for mode m . The reserve mode can be assumed to have infinite capacity. Our default unit of measurement will be cubic feet.
 - Each sku assigned to a given storage mode will be allotted a fraction of the available storage space. It is assumed that the sku can completely fill whatever space it is allotted. For simplicity we assume that modes are dedicated storage.
3. Associated with each storage mode is a set of storage modes that can supply restocks for the mode, called the *predecessor modes* for the mode.
 - In the example in figure 1.11, the only predecessor mode for the intermediate mode is the reserve mode. The forward mode counts both the intermediate mode and the reserve mode as predecessor modes.

The cost per restock of a sku in mode m from predecessor mode ℓ is denoted $c_{\ell m}$. The cost is the same for all skus restocked in mode m from mode ℓ . The restock cost is assumed

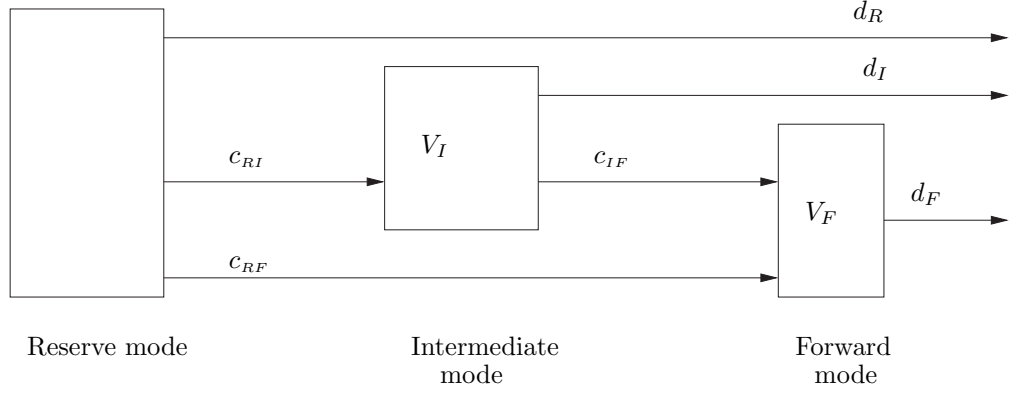


Figure 1.11: A three-tier inventory system with costs.

to be nonnegative and is independent of the size of the restock. We do not consider the cost of restocking the reserve mode in this research.

If skus can be picked from mode m , the cost per pick is denoted d_m . We initially assume that this cost is the same for all skus picked from mode m , but we will note when we relax this assumption.

Figure 1.11 shows a model of a three-tier inventory system.

1.5.2 Skus

Let S represent the set of skus to be stored in the inventory system, and assume that we are analyzing warehouse operations over a fixed period of time. For each sku $i \in S$, we know

- the total number of orders on which the sku appears each period, called the *picks* of sku i and denoted p_i , and
- the total cubic feet sold per period, called the *flow* of sku i and denoted f_i .

We will frequently use the square root of the flow of each sku in computations; we refer to the quantity $\sqrt{f_i}$ as the *rootflow* of sku i .

1.5.3 How to characterize a slotting

In the Hackman-Rosenblatt model of a forward-reserve inventory system, a slotting is characterized by stating which skus are picked from each mode and how much space each sku is allotted in the forward mode. In a multi-tier inventory system, however, this is not sufficient, since two skus picked from the same mode may have been restocked in that mode




Figure 1.12: Skus assigned to the four possible flowpaths in a three-tier system.

from different predecessor modes. In a slotting of a multi-tier inventory system, each sku is restocked in the mode from which it is picked via a path of storage modes, which we will call the *flowpath* of the sku. Each mode in the flowpath restocks its successor in the path.

When it is important to know the component nodes of a flowpath, we will use path notation from graph theory and denote each flowpath as a path of its component modes. Figure 1.12 shows a multi-tier inventory system with four flowpaths; skus a through d are each assigned to a different flowpath.

- Sku a is assigned to flowpath R and is picked from the reserve mode.
- Sku b is assigned to flowpath $R, 2$: it is picked from mode 2, and its supply there is replenished from the reserve mode.
- Sku c is assigned to flowpath $R, 2, 3$: it is picked from mode 3; its supply in mode 3 is replenished from mode 2; and its supply in mode 2 is replenished from the reserve mode.
- Sku d is assigned to flowpath $R, 3$: it is picked from mode 3; its supply in mode 3 is replenished directly from the reserve mode.

We can characterize a slotting of a multi-tier inventory system by stating the assignment of skus to flowpaths and the amount of space that each sku is allotted in each mode in its flowpath. We make the following assumptions about the flowpaths in an inventory system.

1. Each sku is assigned to exactly one flowpath.

2. All flowpaths begin with the reserve mode.

The q th mode in the flowpath is referred to as the q th *tier* of the flowpath. A given storage mode may be the q th tier of one flowpath and the r th tier of another. In Figure 1.12, mode 3 is the third tier of flowpath $R, 2, 3$ but the second tier of flowpath $R, 3$.

A flowpath can be thought of as a sequence of tiers, where tier q restocks tier $q + 1$. The *length* of a flowpath is defined as the number of tiers it comprises, and an inventory system whose longest flowpath has length ℓ is referred to as a ℓ -*tier inventory system*. The forward-reserve and multi-mode inventory systems are both two-tier inventory systems.

The set of all skus assigned to a flowpath is known as a *flowgroup*; the flowgroup that corresponds to flowpath g is denoted $S(g)$.

1.5.4 Restocking protocol

Over the long term, the rate at which each sku is restocked in its picking mode must equal the rate at which the picking mode requests restocks of that sku. For analytical tractability, we make two simplifying assumptions. We assume that restocks of a sku in a mode occur instantaneously from the appropriate predecessor tier, and that a sku is restocked in a storage mode when the supply of the sku there has been entirely depleted, with no allowance for safety stock. If v_{im} is the amount of space allotted to sku i in mode m , then sku i will be restocked in mode m a total of f_i/v_{im} times.

When a picking mode requests a restock of a sku, it may trigger a series of replenishment requests down the flowpath of the sku. The restocking procedure is summarized in the flowchart in figure 1.13. To illustrate the restocking protocol, we consider the flowpath shown in figure 1.14. The figure shows the series of requests triggered the first time mode C stocks out. Over the long run, for every restock of mode C , mode B is restocked $5/8$ times, and mode A is restocked $5/6$ times. If a sku assigned to this flowpath sells 48 cubic feet worth of product each period, then on average mode C will be restocked 9.6 times, mode B will be restocked 6 times, and mode A will be restocked 8 times.

Replenishment Process

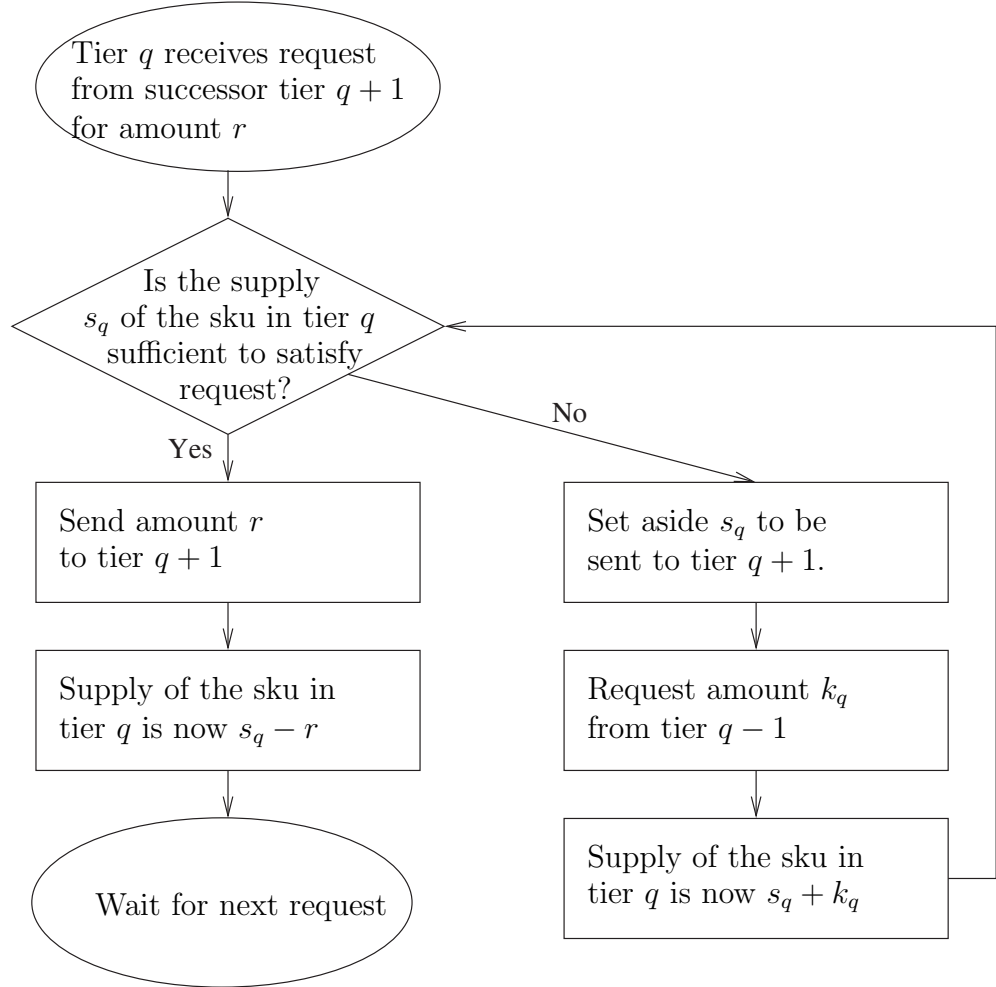




Figure 1.13: Decision process when tier q receives a restock request


The flowpath. All modes have a full supply of the sku.




Mode C stocks out and requests 5 from Mode B.




Mode B sends 5 to Mode C.




Mode C stocks out again and requests 5 from Mode B.




Mode B places 3 in its holding area and requests 8 from Mode A.




Mode A places 6 in its holding area and requests 6 from Mode R.



Mode R sends 6 to Mode A.



Mode A sends 8 to Mode B.



Mode B sends 5 to Mode C.




Figure 1.14: Ways to restock the A-frame area

In formal terms, each tier q follows this set of rules when it receives a request to restock a sku.

- 1: If there is no current, unfilled request from tier $q + 1$, wait.
- 2: If there is a current, unfilled request from tier $q + 1$ for amount r :
 - a: If $r \leq s_q$, withdraw amount r from tier q , mark the request “filled”, send the inventory for this order to tier $q + 1$ and go to Rule 1. (Note that amount $s_q - r$ remains in inventory at tier q .)
 - b: If $r > s_q$, withdraw amount s_q from tier q and hold it toward the request. Reduce the amount requested from r to $r - s_q$ and make a request to tier $q - 1$ for amount k_q . Go to Rule 3.
- 3: On receiving inventory from tier $q - 1$, store it and go to Rule 2.

Under this restocking protocol, a mode can accumulate enough supply of a sku to completely replenish a successor mode, so the amount of space a sku occupies in a mode does not affect the amount of space a sku occupies in any other mode. Under some restock protocols, however, this is not the case. As an example, consider an alternative restock protocol where restocks of a sku in mode m were restricted to be no larger than the maximum supply stored in the predecessor mode for that sku. In this case, the sku would effectively never occupy more space in mode m than in the predecessor mode. In the inventory system in figure 1.14, assume that sku i is allotted three cubic feet in mode A, and four cubic feet in mode B. Each time that mode B requests a restock from mode A, the alternate restocking protocol would dictate that we send only three cubic feet, meaning that no more than 3 cubic feet of the space in mode B would ever be occupied by sku i .

1.6 Literature Review

1.6.1 Research on multi-tier inventory systems

The multi-tier inventory system studied in this research is based on the forward-reserve inventory system presented by Hackman and Rosenblatt in [7], which is pictured in figure 1.2. In that system, skus can be picked from either the forward mode or the reserve

mode, although all skus have a supply in the reserve mode. It costs less to pick a sku from the forward mode than from the reserve mode, but a sku in the forward mode must be restocked from the reserve mode, and there is a cost incurred for each restock. The cost of a restock is assumed to be independent of the size of the restock. In this model, the cost per pick of a sku from the forward mode may vary by sku, as may the cost per restock of a sku in the forward mode. Picking and restocking activities happen concurrently.

Hackman and Rosenblatt focused on slotting the inventory system to minimize the total cost of picking and restocking, given the constraint that the forward mode has a limited amount of space to allot to skus for storage. It is desirable to assign many skus to the forward area to take advantage of the low pick cost there, but the more skus that are assigned there, the less space each gets, meaning that each sku will require more restocks, which drives up the cost of restocking. They do not address issues of the exact storage location in a storage mode to which to assign a sku, or how much supply of a sku to hold in the reserve mode.

Hackman and Rosenblatt first assume an assignment of skus to the forward mode, and they show how to optimally allocate the available space among the skus. Each sku assigned to the forward mode is assumed to completely fill the space it is allotted. If $S(R, F)$ represents the set of skus assigned to the forward mode, then

Theorem 1.1 *In order to minimize the cost of restocking the forward mode, sku i in flow-group $S(R, F)$ should be allotted*

$$\frac{\sqrt{f_i}}{\sum_{j \in S(R, F)} \sqrt{f_j}} V_F \quad (1.1)$$

cubic feet in the forward mode.

They then develop a measure to rank skus by their suitability for placement in the forward mode: we refer to this ranking as the *viscosity*¹ of a sku, where the viscosity of sku i is $p_i/\sqrt{f_i}$. Hackman and Rosenblatt show that a near-optimal assignment of skus to storage locations and allocation of storage space to skus has the characteristic that some

¹called the *economic assignment quotient* in [7]

number of the skus with the highest viscosities are assigned to the forward mode. This means that given a list of skus pre-sorted by viscosity, a near-optimal slotting of a forward-reserve system with n skus can be found by enumeration in $O(n)$ time. They go on to show that if the k highest-viscosity skus are assigned to the forward mode, then the total cost of picking and restocking the inventory system is unimodal as a function of k . This implies that search procedures can be used to reduce the time needed to find a near-optimal assignment to $O(\log n)$.

Frazelle, et al. study a slight variation of the Hackman-Rosenblatt model of the forward-reserve system in [5]. They allow the size of the forward mode to be a variable, and the costs of picking and restocking skus in the forward mode are functions of this variable. In their model, congestion is constrained to be under a certain user-defined limit, where congestion is defined as the number of workers per square foot of walking area. The goal of their research is to find the size of the forward mode and the corresponding slotting of the forward-reserve system where the total cost of picking and restocking is minimized and the congestion constraint is not violated.

The researchers show that the skus with the highest viscosities have the strongest claim to the forward mode, and they then present two approaches to solving this problem. They first consider the case where there are relatively few feasible sizes for the forward mode, such as when the forward mode would consist of an integral number of bays of shelving. They iteratively consider a forward mode of each size and assign skus to the forward mode in descending order of viscosity, stopping when either the marginal cost of adding the next sku would outweigh the marginal pick savings or when the congestion constraint is violated, whichever occurs first.

If the forward mode could be one of many possible sizes, however, so that the choice of sizes is close to continuous, the researchers take a different approach. They iteratively consider assigning the k highest-viscosity skus to the forward mode and then determine the size of the forward mode that minimizes picking and restocking costs for those skus.

This research makes a useful modification of the Hackman-Rosenblatt model of a forward-reserve system by allowing the size of a storage mode to be a variable, since it provides

guidance for establishing a forward-reserve inventory system in a warehouse where one may not already exist.

Other researchers have extended the Hackman-Rosenblatt model of a forward-reserve system to a multi-mode inventory system. Hackman and Platzman discuss a general model in [6] that can be applied to find a near-minimum cost slotting of the multi-mode inventory system. Their model allows the cost savings associated with assigning a sku to a forward mode to be any continuous function, and it also allows for the case when the space in a storage mode needs to be allocated in discrete chunks rather than as a fraction of the total available space. They develop a procedure that finds a near-minimum cost slotting of the multi-mode inventory system with a very tight bound, but a great deal of computation and coding is necessary to implement the procedure, making it impractical for use in many warehouses.

To make the problem more tractable, in [2] Bartholdi and Hackman consider a multi-mode inventory system with a more restrictive cost function: the cost per pick from a given storage mode is the same for all skus stored there, as is the cost per restock of a given forward storage mode. This model is illustrated in figure 1.15. As in the Hackman-Rosenblatt forward-reserve system, skus are allotted a fraction of the total space in the forward modes, and it is assumed that they can fully occupy the space.

Definition 1.2 *We say that a slotting is well-ranked for each sku i , no sku with viscosity lower than sku i is picked from a mode with a lower pick cost than the mode from which sku i is picked.*

Theorem 1.2 *A provably near-minimum cost slotting of a Bartholdi-Hackman multi-mode system can be found by considering only well-ranked slottings.*

Using this theorem, a near-minimum cost slotting can be found using enumeration in $O(n)^M$ time. They then derive properties of the objective function that allow simple search methods to be used for a “continuous-time representation” of the problem, reducing the worst-case time to find a near-minimum cost slotting of a multi-mode system to $O((\log n)^M)$ time, assuming a list of skus sorted by viscosity.

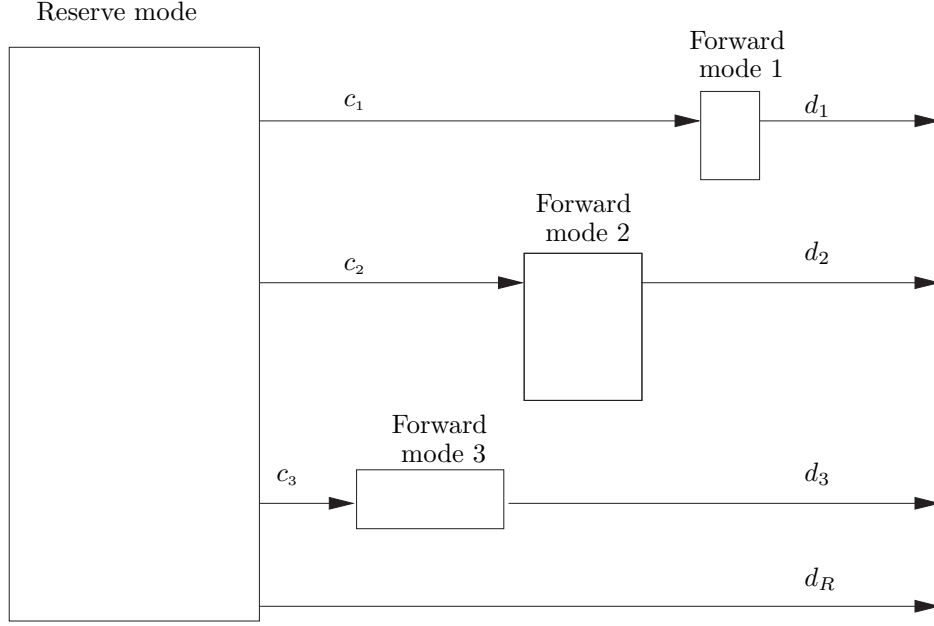


Figure 1.15: The Bartholdi-Hackman multi-mode inventory system.

The existing research on the multi-mode problem has addressed how to minimize picking and restocking costs in those warehouses with multiple forward modes, but it is assumed that all forward modes are restocked directly from a reserve mode. In our research, we will consider inventory systems where the forward modes can be restocked from storage modes other than the reserve mode. The cost of restocking a sku in a mode depends on the mode providing the restock.

1.6.2 Research on the multi-period forward-reserve problem

The research we have examined to this point is concerned with assigning skus to storage modes and allocating them space there over a single period. In [7], however, Hackman and Rosenblatt observed that it can be valuable in some warehouses to consider multiple periods when assigning skus to a forward-reserve system. The skus assigned to the forward mode may change each period, but there is a reassignment cost that is charged each time a sku is reassigned to a different storage mode. The goal, then, is to slot the forward-reserve inventory system in each period to minimize the total picking, restocking, and reassignment costs over all periods. We call this the multi-period forward-reserve problem.

Little research has been published about the multi-period forward-reserve problem since

it was identified in [7]. Sadiq et al. addressed the problem of slotting a forward picking area to minimize time dedicated to picking and moving skus over several periods in [8], but they do not include restock costs in their analysis. Furthermore, the heuristic they develop appears to be too complicated to summarize in the article. In this research, we will develop heuristics for the multi-period forward-reserve problem that warehouse planners can easily execute on commonly available software packages.

1.7 Organization of results

When presenting our results on how to slot a multi-tier inventory system, we will take the approach of introducing one cost to the system at a time. In chapter 2 we focus on minimizing the cost of restocking a multi-tier inventory system, and in chapters 3 and 4 we determine how to minimize the cost of picking and restocking a multi-tier inventory system. We then report on a case study of our results when applied to an inventory system in use at the Avon warehouse in chapter 5. In chapter 6 we consider the question of how best to slot a multi-tier inventory system over multiple time periods when the costs of reslotting the warehouse are considered. We summarize our findings in chapter 7.

CHAPTER 2

SLOTTING A THREE-TIER INVENTORY SYSTEM TO MINIMIZE RESTOCKING COSTS

In the next two chapters, we present results to slot a multi-tier inventory system with the goal of minimizing only restock costs. Besides serving as a foundation for later research, these results apply to multi-tier inventory systems where all forward modes have the same cost per pick.

In this chapter we examine the three-tier inventory system in figure 2.1 to develop the principles of minimizing restock costs. This system has two modes in addition to the reserve mode: an intermediate mode I and a forward mode F . All skus are picked from mode F , but a sku can be restocked in mode F via one of two flowpaths: R, F or R, I, F .

We can affect the total cost of restocking this inventory system in two ways, by

- assigning skus to flowgroups and
- allocating space to skus in each storage mode.

We first assume that skus have been assigned to each flowgroup and we discuss how to allocate space to skus in each flowgroup. We then state a simple procedure for assigning skus to flowpaths to minimize total restock costs.

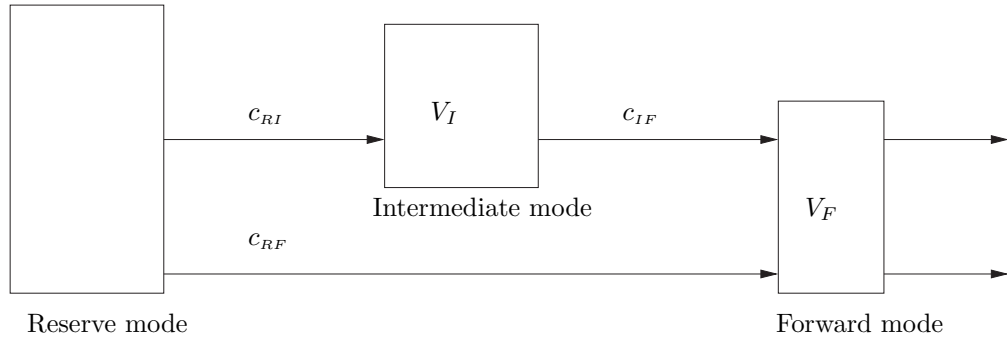


Figure 2.1: The inventory system for which we minimize restock costs

2.1 Allotting space to skus to minimize restock costs

In our model, as in the Hackman-Rosenblatt model of a forward-reserve system, minimizing the cost of restocking a sku in a mode is equivalent to minimizing the number of times the sku is restocked in the mode. Assuming that skus have been assigned to flowpaths in a multi-tier inventory system, we minimize the total cost of restocking all skus in the system if we minimize the cost of restocking each mode in the system, independent of all other modes.

2.1.1 Allotting space to skus in Mode I

The total cost of restocking skus in Mode I can be expressed as

$$\sum_{j \in S(R, I, F)} c_{RI} \frac{f_j}{v_{jI}}.$$

By results in [7] (theorem 1.1), to minimize this cost, sku $i \in S(R, I, F)$ should be allotted volume v_{iI} so that

$$v_{iI} = \frac{\sqrt{f_i}}{\sum_{j \in S(R, I, F)} \sqrt{f_j}} V_I$$

The total cost of restocking Mode I is then

$$c_{RI} \frac{\left(\sum_{j \in S(R, I, F)} \sqrt{f_j} \right)^2}{V_I}. \quad (2.1)$$

2.1.2 Allotting space to skus in Mode F

The total cost of restocking skus in Mode F can be expressed as

$$\sum_{j \in S(R, I, F)} c_{IF} \frac{f_j}{v_{jF}} + \sum_{j \in S(R, F)} c_{RF} \frac{f_j}{v_{jF}}.$$

When determining the values of v_{iF} for each sku i , we cannot directly apply results from [7] since two flowpaths share the mode and each has a different restock cost. Let α be the proportion of space in Mode F to be occupied by skus in flowpath R, I, F , as illustrated in figure 2.2. Then

$$v_{iF} = \begin{cases} \frac{\sqrt{f_i}}{\sum_{j \in S(R, I, F)} \sqrt{f_j}} \alpha V_F & i \in S(R, I, F) \\ \frac{\sqrt{f_i}}{\sum_{j \in S(R, F)} \sqrt{f_j}} (1 - \alpha) V_F & i \in S(R, F) \end{cases}.$$

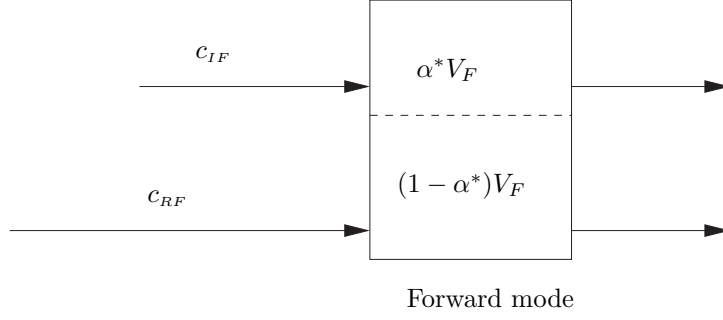


Figure 2.2: The space in Mode F divided between flowgroups.

The cost of restocking the forward mode is then

$$\frac{c_{IF} \left(\sum_{j \in S(R,I,F)} \sqrt{f_j} \right)^2}{\alpha V_F} + \frac{c_{RF} \left(\sum_{j \in S(R,F)} \sqrt{f_j} \right)^2}{(1 - \alpha) V_F} \quad (2.2)$$

We must now find the value of α that minimizes the above expression; we will denote this value α^* .

Theorem 2.1 *To minimize the cost of restocking Mode F , the proportion of mode F to allot to flowgroup $S(R, I, F)$ is*

$$\alpha^* = \frac{\sqrt{c_{IF}} \left(\sum_{i \in S(R,I,F)} \sqrt{f_i} \right)}{\sqrt{c_{IF}} \left(\sum_{i \in S(R,I,F)} \sqrt{f_i} \right) + \sqrt{c_{RF}} \left(\sum_{i \in S(R,F)} \sqrt{f_i} \right)}$$

Proof Take the derivative of expression 2.2 with respect to α ; set this equal to 0 and solve for α^* . \square

Assuming that skus are allocated space in Mode F to minimize the cost of restocking the mode, we can substitute the value of α^* into expression 2.2 to write the cost of restocking Mode F as

$$\frac{\left(\sqrt{c_{IF}} \sum_{j \in S(R,I,F)} \sqrt{f_j} + \sqrt{c_{RF}} \sum_{j \in S(R,F)} \sqrt{f_j} \right)^2}{V_F}. \quad (2.3)$$

The previous results can be extended to an arbitrary mode shared by several flowgroups, as proved in [7].

Theorem 2.2 *Assume that G flowgroups share a mode of volume V . Skus have been assigned to flowgroups, and the total rootflow in flowgroup g is ϕ_g . The cost to restock flowgroup g in the mode from the previous mode on the flowpath is c_g . Let the proportion of*

space in the mode allotted to flowgroup g be denoted α_g . Then to minimize the cost of restocking the mode,

$$\alpha_g^* = \frac{\sqrt{c_g} \phi_g}{\sum_{h=1}^G \sqrt{c_h} \phi_h}$$

The total cost of restocking the mode shared by the G flowgroups is

$$\frac{\left(\sum_{h=1}^G \sqrt{c_h} \phi_h\right)^2}{V}$$

2.2 Assigning skus to flowpaths

In the previous section we assumed that skus were assigned to flowgroups, and we showed how to allocate space to skus to minimize restock costs. With this result in mind, we now consider how best to assign skus to flowgroups to minimize restock costs.

For a given assignment of skus to flowgroups, we add Expressions 2.1 and 2.3 to state the total cost of restocking the inventory system as

$$\frac{\left(\sqrt{c_{RI}} \sum_{j \in S(R,I,F)} \sqrt{f_j}\right)^2}{V_I} + \frac{\left(\sqrt{c_{IF}} \sum_{j \in S(R,I,F)} \sqrt{f_j} + \sqrt{c_{RF}} \sum_{j \in S(R,F)} \sqrt{f_j}\right)^2}{V_F}$$

The cost is a function of the sum of the rootflow of each sku assigned to each flowpath. We will call the sum of the rootflows of skus in a given flowgroup the *total rootflow* of the flowgroup, and we will denote the sum of the rootflows of all skus as \mathcal{F} . Similarly, the sum of the picks of all skus assigned to a given flowgroup is the *total picks* of a flowgroup.

Therefore we can simplify the expression for the total cost of restocking the system by letting $\phi_{R,I,F}$ represent the total rootflow assigned to flowpath R, I, F , so that $\phi_{R,I,F} = \sum_{j \in S(R,I,F)} \sqrt{f_j}$. The expression for the restock cost of the inventory system, abbreviated $r(\phi_{R,I,F})$, can now be written

$$r(\phi_{R,I,F}) = \left(c_{RI} \frac{\phi_{R,I,F}^2}{V_I} + \frac{(\sqrt{c_{IF}} \phi_{R,I,F} + \sqrt{c_{RF}} (\mathcal{F} - \phi_{R,I,F}))^2}{V_F} \right).$$

By definition, $\phi_{R,I,F}$ can only assume values that correspond to the total rootflow of a subset of S . Define a variable $\phi_{R,I,F}^c$ that can take on values in the interval $[0, \mathcal{F}]$. Then we can use theory of continuous functions to minimize $r(\phi_{R,I,F}^c)$; the minimizer of $r(\phi_{R,I,F}^c)$ is denoted $\phi_{R,I,F}^{c*}$.

There may not be skus whose rootflows sum exactly to $\phi_{R,I,F}^{c*}$, and thus we may not be able to achieve the minimal restock cost given by $r(\phi_{R,I,F}^{c*})$. However, a simple greedy heuristic will give a feasible value of $\phi_{R,I,F}$ that differs from $\phi_{R,I,F}^{c*}$ by at most the rootflow of a single sku; in a typical application with thousands of skus, the resulting restock cost will deviate very little from the minimal. Below we first describe an algorithm for finding $\phi_{R,I,F}^{c*}$, and then we state the greedy heuristic that finds a feasible value of $\phi_{R,I,F}$ that is close to the value of $\phi_{R,I,F}^{c*}$.

Algorithm to determine $\phi_{R,I,F}^{c*}$

If $c_{RF} < c_{IF}$: $\phi_{R,I,F}^* = 0$.

Else if $\left(\frac{V_F}{V_I}c_{RI} + c_{IF}\right) < \sqrt{c_{RF}c_{IF}}$: $\phi_{R,I,F}^* = \mathcal{F}$.

Else:

$$\phi_{R,I,F}^* = \frac{c_{RF} - \sqrt{c_{RF}c_{IF}}}{(c_{RF} - \sqrt{c_{RF}c_{IF}}) + \left(\left(\frac{V_F}{V_I}c_{RI} + c_{IF}\right) - \sqrt{c_{RF}c_{IF}}\right)} \quad (2.4)$$

Proof Minimizing the continuous restock cost function $r(\phi_{R,I,F}^c)$ subject to the constraint that $0 \leq \phi_{R,I,F}^c \leq \mathcal{F}$ is a convex program. By convexity, the minimizer can be determined by first finding the global minimizer of $r(\phi_{R,I,F})$: if the value falls in the interval $[0, \mathcal{F}]$, then it is an optimal solution to the constrained problem. Otherwise, the global minimizer occurs at the nearest boundary point. Let $\phi_{R,I,F}^G$ be the global minimizer of $r(\phi_{R,I,F})$. Then

$$\phi_{R,I,F}^G = \frac{c_{RF} - \sqrt{c_{RF}c_{IF}}}{(c_{RF} - \sqrt{c_{RF}c_{IF}}) + \left(\left(\frac{V_F}{V_I}c_{RI} + c_{IF}\right) - \sqrt{c_{RF}c_{IF}}\right)} \mathcal{F}$$

Algebra shows that

- $\phi_{R,I,F}^G < 0$ when $c_{RF} < c_{IF}$, and
- $\phi_{R,I,F}^G > 1$ when $\left(\frac{V_F}{V_I}c_{RI} + c_{IF}\right) < \sqrt{c_{RF}c_{IF}}$. \square

The algorithm below takes the continuous minimizer $\phi_{R,I,F}^{c*}$ as input and finds a feasible value of $\phi_{R,I,F}$ that differs from the continuous minimizer by at most the rootflow of a single sku

FindFeasibleRootflow($\phi_{R,I,F}^{c*}$)

1 $TotalRootflow = 0$

2 For $j = 1$ to n

3 $TotalRootflow = TotalRootflow + \sqrt{f_j}$

4 If $TotalRootflow > \phi_{R,I,F}$, exit loop

5 Next j

6 $SkusInModeF = j$

The value of $TotalRootflow$ at the end of the algorithm will be the total rootflow assigned to flowpath R, I, F , where skus 1 through $SkusInModeF$ are assigned to the flowpath.

CHAPTER 3

SLOTTING THE AVON INVENTORY SYSTEM TO MINIMIZE PICKING AND RESTOCKING COSTS

In the next two chapters, we will consider the costs of picking skus as well as the cost of restocking them when slotting a multi-tier inventory system. We begin by examining a simplified version of the Avon inventory system that was described in chapter 1. We will show that the simplified inventory system is equivalent to a two-tier system with two forward modes, meaning that we can then find a minimum cost slotting by using principles developed for multi-mode inventory systems in [2] and [6]. In the next chapter, we will extend these results to a general multi-tier inventory system. To evaluate our model in a real-world situation, we compared the slotting recommended by our model to the actual slotting used in the Avon warehouse; the results are summarized in chapter 5.

As described in chapter 1, the inventory system in place at the Avon warehouse outside Atlanta has seven flowpaths (illustrated in figure 1.10). For ease of computation, we have simplified the model to that shown in figure 3.1. Skus can be picked from the A-frame, denoted mode A , or from a mode of shelving and flowrack that represents an amalgamation of the manual lines and cart pick areas, which we call mode J . All skus that are picked




Figure 3.1: The simplified Avon inventory system.

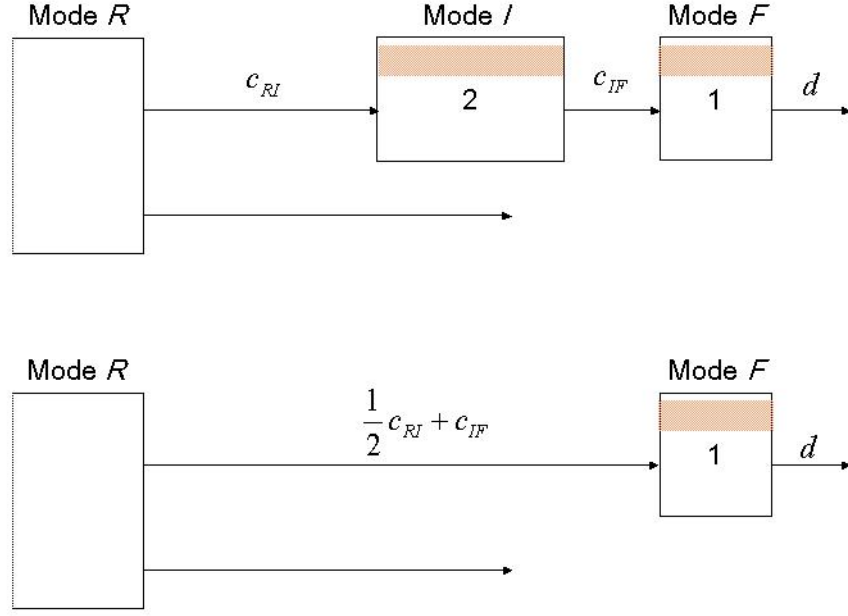


Figure 3.2: Example of equivalent inventory systems

from mode A have a supply in an intermediate mode I , which also comprises several bays of flowrack. The two possible flowpaths in the system are thus R, I, A and R, J . It is cheaper to pick skus from the A-frame than from flowrack, so $d_A < d_J$.

3.1 *Equivalence to a two-tier inventory system*

To show the relationship between the simplified Avon inventory system and a multi-mode system with two forward modes, we will make use of the following definition.

Definition 3.1 *Two inventory systems with the same number of flowpaths are equivalent if for each partition of the skus into flowgroups, the total cost of picking is the same in both systems and the cost of restocking is the same in both systems.*

By way of example, consider the top inventory system in figure 3.2: one flowpath is R, I, F , and the other is R . Mode I has 2 cubic feet of storage space available, and mode F has 1 cubic foot. Assume that skus have been partitioned into flowgroups. By results in [7] (theorem 1.1), sku i in flowgroup $S(R, I, F)$ will be allotted $2 \cdot \sqrt{f_i} / (\sum_{j \in S(R, I, F)} \sqrt{f_j})$ in mode I , and $1 \cdot \sqrt{f_i} / (\sum_{j \in S(R, I, F)} \sqrt{f_j})$ in mode F . Thus every time a sku in mode F

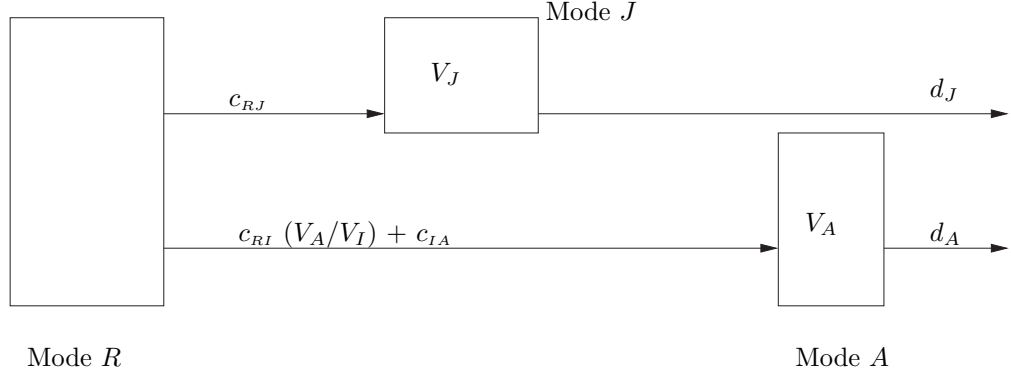


Figure 3.3: Two-tier inventory system equivalent to the simplified Avon inventory system

is restocked, half of the total capacity dedicated to that sku in mode I is depleted. So for every restock of mode F , we are charged for half a restock of mode I . The total cost to the entire flowpath R, I, F each time the forward mode is restocked is then

$$\frac{1}{2}c_{RI} + c_{IF} \quad (3.1)$$

Now consider the bottom inventory system in figure 3.2, with a reserve mode R' and a forward mode F' identical to mode F . Each restock of mode F' from mode R' costs $c_{RI}/2 + c_{IF}$. Thus for any partition of a set of skus into flowgroups, the total cost of picking each flowgroup is the same in both the top and the bottom inventory systems, and the total cost of restocking each flowgroup also is the same. By our definition, the forward-reserve inventory system on the bottom in figure 3.2 is equivalent to the one on the top.

Theorem 3.1 *The simplified Avon inventory system has an equivalent two-tier inventory system with two forward modes.*

Proof Assume that skus have been assigned to flowpaths in the simplified Avon inventory system and that space is allotted to the skus in each mode to minimize the cost of restocking the mode. By results in [7] (theorem 1.1), if sku i in flowpath R, I, A occupies volume v_{iA} in mode A , it will occupy $(V_I/V_A) \cdot v_{iA}$ in mode I . It follows that each restock of the sku in mode A contributes to V_A/V_I restocks of mode I . The cost to the entire flowpath of one restock of mode A is thus $c_{RI}(V_A/V_I) + c_{IA}$.

If mode I were removed from the inventory system and sku i were restocked in mode A




Figure 3.4: Avon system where all flowrack is considered one mode.

directly from the reserve mode at a cost per restock of $c_{IA} + c_{RI}(V_A/V_I)$, the total cost of restocking sku i in tier A would be unchanged. Thus for any partition of skus into flowgroups, the cost of restocking each flowgroup would be the same as in the simplified Avon inventory system. Since the cost of picking is unchanged from the simplified Avon inventory system, the two-tier inventory system formed by removing mode I and setting the cost of restocking mode A from mode I equal to $c_{IA} + c_{RI}(V_A/V_I)$ is equivalent to the simplified Avon inventory system. This system is pictured in figure 3.3. \square

Since we can effectively think of the Avon system as a two-tier inventory system with two forward modes, we can apply results concerning slotting multi-mode inventory systems from [2] (theorem 1.2) to find a near-minimum cost slotting of the simplified Avon inventory system by considering only the $n + 1$ well-ranked slottings.

3.2 Optimal allocation of storage resources

When visiting the Avon Products warehouse, we observed that the total supply of flowrack in the inventory system has been divided between mode I and mode J . We then wondered if each mode had received the proportion of flowrack that would minimize the cost of restocking the system.

To answer this question, we modeled the simplified Avon inventory system as if the combined flowrack storage were one mode, mode $I \cup J$, of volume $V_{I \cup J}$. We let the variable α represent the proportion of mode $I \cup J$ dedicated to skus in flowgroup $S(R, I, A)$, as shown in Figure 3.4. The optimal proportion of flowrack to be used to support the A-frame




Figure 3.5: A two-mode system equivalent to the Avon system with one tier of flowrack.

is the value of α in the minimum cost slotting, which we denote α^* .

The cost of restocking each flowpath in mode $I \cup J$ remains as in the simplified Avon inventory system: the cost of restocking mode A from mode $I \cup J$ is c_{IA} , the cost per restock of mode $I \cup J$ for a sku in flowpath R, I, A is c_{RI} , and the cost per restock of mode $I \cup J$ for a sku in flowpath R, J is c_{RJ} . This assumption is reasonable if the optimal proportion of flowrack to dedicate to skus in each flowgroup is fairly close to the current proportion, since this means that restock costs in each area will remain approximately the same.

Finding a near-minimum cost slotting We observe that to find a near-minimum cost slotting of this inventory system, we still need only consider well-ranked assignments of skus to flowgroups. To see this, assume that we know the value of α^* . Then the combined-flowrack Avon inventory system is equivalent to the multi-mode system shown in figure 3.5, and as in the previous section, we use results in [2] to find a near-minimum cost slotting by considering only well-ranked inventory systems.

To determine the best well-ranked inventory system, we will use the fact that the well-ranked assignment of skus to flowgroups with the lowest total picking and restocking costs must set α such that the cost of restocking the flowgroups in the inventory system is minimized. Therefore we need to compare each well-ranked assignment of skus to flowgroups where α is chosen for each assignment to minimize the cost of restocking the assigned flowgroups in the inventory system.

For any assignment of skus to flowgroups in the combined-flowrack Avon inventory

system, the cost of picking and restocking the inventory system is

$$\left(d_A \sum_{i \in S(R,I,A)} p_i + d_J \sum_{i \in S(R,J)} p_i \right) + c_{RI} \frac{\phi_{R,I,A}^2}{\alpha V_{I \cup J}} + c_{RJ} \frac{\phi_{R,J}^2}{(1-\alpha) V_{I \cup J}} + c_{IA} \frac{\phi_{R,I,A}^2}{V_A} \quad (3.2)$$

Theorem 2.2 states that

$$\alpha^* = \frac{\sqrt{c_{RI}} \cdot \phi_{R,I,A}}{\sqrt{c_{RI}} \cdot \phi_{R,I,A} + \sqrt{c_{RJ}} \cdot \phi_{R,J}}. \quad (3.3)$$

Substituting this value of α^* into the cost function in expression 3.2, the cost of picking and restocking the system for a given assignment of skus to flowgroups can be written

$$\left(d_A \sum_{i \in S(R,I,A)} p_i + d_J \sum_{i \in S(R,J)} p_i \right) + \frac{(\sqrt{c_{RI}} \phi_{R,I,A} + \sqrt{c_{RJ}} \phi_{R,J})^2}{V_{I \cup J}} + c_{IA} \frac{\phi_{R,I,A}^2}{V_A} \quad (3.4)$$

We thus find a near-minimum cost slotting by evaluating expression 3.4 for the $n + 1$ well-ranked assignments of skus to flowgroups. Once this slotting is identified, we substitute the corresponding values of $\phi_{R,I,A}$ and $\phi_{R,J}$ into expression 3.3 above to find the value of α^* , the optimal proportion of flowrack that should be dedicated to supporting the A-frame.

CHAPTER 4

MINIMIZING PICKING AND RESTOCKING COSTS IN GENERAL MULTI-TIER INVENTORY SYSTEMS

In this chapter we show how to slot an arbitrary multi-tier inventory system to minimize total picking and restocking costs. We first show that any multi-tier inventory system has the same fundamental structure as a two-tier inventory system (a multi-mode inventory system) with the same number of flowpaths. Consequently, using results from [2], we need only consider well-ranked assignments of skus to flowpaths to find a near-optimal slotting.

For each well-ranked assignment of skus to flowpaths, we show how to allot space to each flowgroup in each mode to minimize the total picking and restocking costs. To find a near-optimal slotting of the multi-tier inventory system, we select the well-ranked assignment of skus to flowpaths with the lowest total picking and restocking costs.

4.1 A multi-tier inventory system is equivalent to a multi-mode inventory system

We first show that we can obtain a near-optimal slotting of a multi-tier inventory system by considering only well-ranked assignments of skus to flowpaths.

Theorem 4.1 *A multi-tier inventory system with P possible flowpaths is equivalent to a multi-mode system with P forward modes.*

Given an arbitrary multi-tier inventory system, assume that skus have been partitioned into flowgroups, and assume that we know how much space that each flowgroup will occupy in each mode in the optimal solution. Consider a flowpath in the inventory system with Q tiers, numbered $1, \dots, Q$, where skus assigned to the flowpath are picked from tier Q . (See figure 4.1.) Let the space allotted to the flowgroup in tier q be denoted V_q , and let each restock of tier q from its predecessor tier cost c_q . By results in [7] (theorem 1.1), if sku i

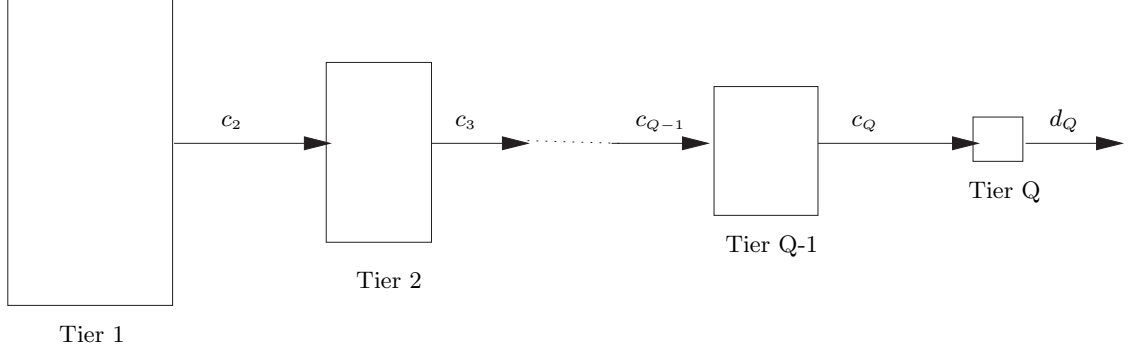


Figure 4.1: Flowpath with Q tiers.




Figure 4.2: Multi-tier inventory system where space in each tier has been allotted to flowgroups.

in flowpath occupies volume v_{iQ} in tier Q , it will occupy $(V_q/V_Q) \cdot v_{iQ}$ in tier q . Thus each restock of a sku in tier Q contributes to (V_Q/V_q) restocks of the sku in tier q . The total cost to the flowpath for each restock of the flowgroup in tier Q is then

$$\sum_{q=1}^Q c_q \frac{V_Q}{V_q} \quad (4.1)$$

The total cost per restock of the flowgroup in tier Q is the same as if we restock a sku in the flowgroup in tier Q directly from the reserve mode with a cost per restock as given in expression 4.1. Since this holds for all P flowpaths, the multi-tier inventory system is equivalent to a multi-mode inventory system with P forward modes. \square

As a result of this equivalence, we know by results in [2] that we need consider only well-ranked assignments of skus to flowpaths to find a near-optimal solution.

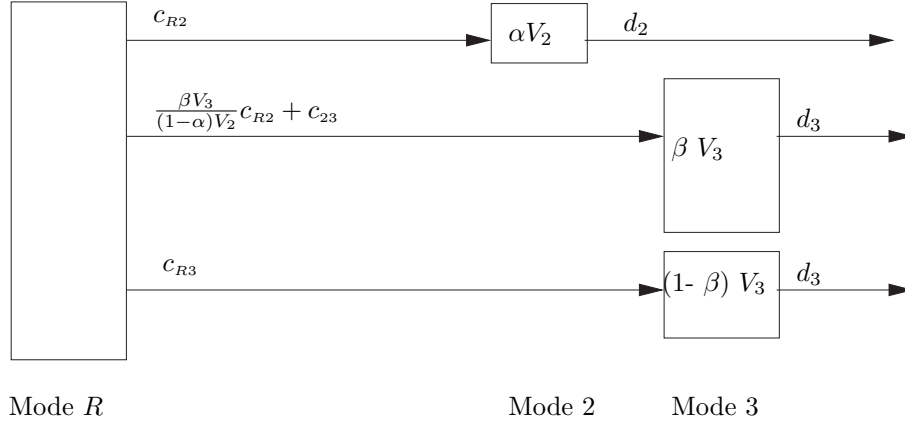


Figure 4.3: Multi-mode inventory system equivalent to system in Figure 4.2.

Example In the multi-tier inventory system in figure 4.2, flowgroup $S(R, 2)$ receives αV_2 cubic feet in mode 2. Flowgroup $S(R, 2, 3)$ is allotted $(1 - \alpha)V_2$ cubic feet in mode 2, and βV_3 cubic feet in mode 3. Flowgroup $S(R, 3)$ is allotted $(1 - \beta)V_3$ cubic feet in mode 3. Consider a sku assigned to flowpath $R, 2, 3$, for example. We know that each restock of the sku in mode 3 will contribute to $\beta V_3 / (1 - \alpha)V_2$ restocks in mode 2. Thus the cost of restocking the skus on flowpath $R, 2, 3$ is the same as restocking them directly from the reserve mode to a forward mode of volume $(1 - \beta)V_3$ at a cost of $(\beta V_3 / (1 - \alpha)V_2)c_{R2} + c_{23}$. Applying this same logic to the other flowpaths, we see that the inventory system in figure 4.2 is equivalent to the inventory system in 4.3.

4.2 Finding the best well-ranked slotting

To determine the well-ranked slotting with the lowest total picking and restocking costs, we will use the fact that the well-ranked assignment of skus to flowgroups with the lowest total picking and restocking costs must allocate space to each flowgroup to minimize the cost of restocking the multi-tier inventory system. (For a given assignment of skus to flowgroups, pick costs are fixed, but the cost of restocking the flowgroups can be changed by adjusting the amount of space allotted to each flowgroup in each mode.)

For each well-ranked assignment of skus to flowgroups, we use theorem 2.2 to determine the space in each mode to allot to each flowgroup to minimize the cost of restocking the assigned flowgroups in the inventory system. Using this space allocation, we then compute

the total picking and restocking cost for the given well-ranked assignment of skus to flowgroups. Then to find a near-minimum cost slotting, we select the well-ranked slotting with the lowest total picking and restocking cost.

Example Assume that skus have been assigned to flowpaths in the multi-tier inventory system in figure 4.2, where ϕ_g is the total rootflow in flowgroup $S(g)$. By theorem 2.2, the cost of picking and restocking the inventory system for a given assignment of skus to flowgroups is

$$\begin{aligned} \sum_{j \in S(R,2)} d_2 p_j + \sum_{j \in S(R,2,3)} d_3 p_j + \sum_{j \in S(R,3)} d_3 p_j \\ + \frac{(\sqrt{c_{R2}}\phi_{R,2} + \sqrt{c_{R2}}\phi_{R,2,3})^2}{V_2} + \frac{(\sqrt{c_{23}}\phi_{R,2,3} + \sqrt{c_{R3}}\phi_{R,3})^2}{V_3}. \end{aligned} \quad (4.2)$$

To evaluate which well-ranked assignment of skus to flowpaths gives the lowest total cost of picking and restocking the multitier inventory system in figure 4.2, for each well-ranked assignment we evaluate expression 4.2. The assignment with the lowest total cost corresponds to a near-optimal slotting of the multi-tier inventory system.

Future research The fact that we must consider only well-ranked assignments of skus to flowpaths reduces the number of slottings we need to consider when solving the problem by enumeration, but with several modes and thousands of skus, finding a near-optimal solution by enumeration can be computationally difficult. Bartholdi and Hackman present a method for finding the near-optimal slotting of a multi-mode inventory system in $O((\log n)^M)$ time in [2]. We can use this result to efficiently find the near-optimal slotting of a multi-tier inventory system if no flowgroups share a mode in the system: we need only derive the equivalent multi-mode inventory system and apply the method.

If the multi-tier inventory system has a mode that is shared by flowgroups, however, then the cost to restock a flowgroup in the equivalent two-tier inventory system depends on the amount of space that the flowgroup is allotted in each shared mode. As shown in theorem 2.2, this amount of space depends on the total amount of rootflow assigned to each flowgroup, meaning that the restock costs for the equivalent two-tier system will be

different for each slotting that we consider. Because of this, we cannot directly apply the faster search method derived in [2]. A goal for future research is to identify properties of multi-tier inventory systems with shared modes that allow us to either adapt the search method in [2] or to develop similar search methods to reduce the number of potential slottings we must consider to find a near-optimal solution.

CHAPTER 5

APPLICATION: SLOTTING THE AVON INVENTORY SYSTEM

We used our model to construct a slotting for the set of skus stocked in the Avon inventory system during a ten-day period in February, 2002. (The system is described in detail in chapter 1.) Every ten business days represents a new sales campaign for Avon; we analyzed the inventory system during campaign 4. During this campaign, the warehouse processed orders for 6,848 skus, with total sales of over 11 million pieces, as seen in table 5.1.

Table 5.1: Profile of campaign 4 in the Avon warehouse

Profile of campaign 4	
number of skus available	6,848
orders expected	154,262
picks expected	8,993,024
piece sales forecast	11,317,916

The highest-selling skus during campaign 4 were much more popular than the lowest-selling skus, as shown by the graph in figure 5.1. The 100 skus with the highest expected piece sales accounted for 44% of the total piece forecast of the campaign, while the slowest-selling 3000 accounted for only 0.2%.

We focused on activity in the section of the warehouse where orders are picked and sent to shipping, known as the order fulfillment area.

5.1 Estimating model parameters

The model used to analyze the Avon inventory system is shown in Figure 5.2. To apply this model to minimize picking and restocking costs in the order fulfillment area, we needed to know

- the volume of each storage mode




Figure 5.1: Piece forecast of the top 500 skus, in descending order of sales.




Figure 5.2: Configuration of Avon's warehouse in Suwanee, GA.

- the costs of picking from each mode and the costs of restocking each mode
- the picks and flow for each sku during the campaign
- which skus could be dispensed from the A-frame

We were given a file with information about the piece sales forecast available to planners when they slotted the inventory system for the campaign, as well as a record of the actual slotting used during the campaign.

5.1.1 Computing the volume of each storage mode

In the Avon warehouse, each sku assigned to a given storage mode is allotted approximately the same amount of storage area in that mode. We estimated these values for each storage mode based on measurements taken during warehouse visits. To estimate the total volume available for storage in each mode, we summed the volume allotted to each sku stored there in campaign 4.

5.1.2 Computing picking and restocking costs

We used this model to show the minimum cost of labor required to pick and restock skus in the warehouse, since this is the major variable cost in the warehouse. All other costs that relate to picking and restocking activities, such as the cost of equipment, management and support staff, and utilities, are assumed to be fixed and are not incorporated into the model.

To estimate the cost per restock for each mode, we summed the costs of the labor dedicated to restocking each area over the course of the campaign, and we divided the sum by the estimated number of restocks that were made in that area, assuming that skus were allotted volume as computed above. Each manually-picked sku in the pick-to-light section of flowrack was stored in six separate storage locations, each requiring a separate restock.

Similarly, to estimate the cost per pick in each storage area, we summed the costs of labor that was dedicated to picking in that area, and divided the sum by the number of picks of skus assigned to that area by Avon. Since no labor is used to pick skus from the A-frame, skus stored there have a pick cost of zero.

Table 5.2: Actual and optimal solutions for campaign 4

Comparison of Actual vs. Optimal Solution for Campaign 4		
	actual solution	ideal solution
campaign cost	\$290,880	\$230,819
skus in A-frame	1,590	1,208
Restocks		
of A-frame	57,028	31,608
of intermediate storage area	3,439	1,167
of manual storage area	49,252	6,381
Picks		
from A-frame	3,858,131	3,743,806
from manual storage area	5,134,893	5,249,218

5.1.3 Computing picks and flow

We estimated picks and flow from the piece sales forecast available the week before the campaign, since it is the information the warehouse managers use when planning the slotting.

5.2 Results

The actual slotting used in campaign 4 assigned 1,590 skus to the A-frame and 5,258 to the manual storage area. Under the assumptions of our model, it had a total picking and restocking cost over the ten-day campaign of \$290,880. The near-minimum cost slotting found using the model assigned 1,208 skus to the A-frame and the remaining 5,640 to the manual storage area, with a total cost over the ten-day campaign of \$230,819, a savings of over 20%.

A summary of the differences between the actual slotting and the ideal slotting proposed by the model can be found in Table 5.2. A breakdown of the costs in each solution is given in Figure 5.3. The values for the ideal solution reflect adjustments we made to ensure that each sku in the A-frame is assigned an integral number of channels. This increased the total number of restocks of the A-frame from 31,608 in the optimal solution of the fluid model to 32,167, increasing the total restock cost only slightly.

5.3 Recommendations for slotting the order fulfillment area

Based on the results of our analysis, we recommended the following changes to the current slotting procedure to reduce the total picking and restocking costs.

1. Store fewer skus in the A-frame.
2. Store the correct amount of each sku in its assigned storage mode.
3. Make more skus dispensable from the A-frame.
4. Reduce the cost of picking skus from the manual storage area.
5. Restock skus in the A-frame less often.

The first two recommendations would require only slight adjustments to the current operations in the warehouse; the last three would be more complicated to implement, since they may involve changes in product packaging, the layout of the storage area, and the warehouse management software.

Store fewer skus in the A-frame To the management of the warehouse, the A-frame is the most desirable picking location for fast-moving skus. Their current slotting policy places as many of the highest-selling dispensable skus in the A-frame as possible, while ensuring that no channel stocks out so quickly that the replenishers cannot keep up with the machine. Our results indicated that this procedure places too many skus in the A-frame and gives too little space to the fastest-movers, thus driving up the cost of restocking the A-frame. According to our model, almost 25% fewer skus should be picked from the A-frame. Picking these skus from the manual storage area instead cuts the cost of restocking the intermediate storage area and A-frame by \$25,000 each campaign, while increasing the pick cost of skus in the manual storage area by only \$5,000.

Store the correct amount of each sku in its storage area The actual slotting of the warehouse in campaign 4 required an estimated 109,719 restocks, while the optimal solution




Figure 5.3: A cost breakdown of the actual and optimal solutions.

requires only 39,156. The total restock cost drops from \$95,040 to \$30,619, a savings of over \$64,000.

The actual solution required more restocks than the optimal for two reasons. First, in campaign 4, over half the skus in the warehouse had been allotted space that was more than 50% above or below the optimal amount. One possible explanation is that management may decide not to move a sku from a large location to a smaller one, even when sales are forecasted to decline, because they believe that the act of moving a sku has a high penalty cost. From our observations of the sku-moving process, however, the cost of physically moving a sku from one location to another is fairly low. The dislike for moving skus is exacerbated by the fact that skus are moved in preparation for the next campaign while orders are being picked for the previous campaign, meaning that restockers have two jobs during this time.

Second, the layout of the the pick-to-light section of the manual storage area requires that each sku be stocked in six storage locations. No location is easily accessible from another. The optimal solution assumes that the entire supply of each sku can be stored

together in the manual storage area, reducing the number of restocks needed. Making this change would require changes in both the layout of the area and the warehouse management system, however, neither of which is an easy undertaking.

Make more skus dispensable from the A-frame Some skus cannot be dispensed from the A-frame due to flimsy packaging or size or an odd shape. In campaign 4, 58% of skus were non-dispensable, including one-quarter of the 20% with the highest viscosities. Non-dispensable skus are picked manually and are subject to the higher pick costs in that area. If all skus were dispensable, however, the total cost of picking and restocking would be reduced to 64% of the current optimal cost. Fewer skus would be placed in the A-frame than in the current optimal solution, but they would be the skus with the highest viscosity, many of which are now picked manually. Even if just the 100 skus with highest viscosities were all dispensable, and the dispensability of all other skus remained unchanged, we could reduce the current cost of picking and restocking by \$47,000. Figure 5.4 compares the total cost of the current solution with the total costs of solutions where more skus are dispensable.

Making skus dispensable would require a joint effort with the marketing and packaging departments, which have historically designed products without considering the effect of their design on the cost of their distribution.

Reduce the cost of picking in the manual storage area In the optimal solution, almost 60% of the picks in campaign 4 were made from the manual storage area—the cost of picking from the manual storage area made up 87% of the total cost of the campaign in the optimal solution. A small reduction in the cost per pick in that storage area would therefore lead to a large decrease in total cost: if the cost per pick were reduced by 12.5% (equivalent to removing one picker per line), the total cost of picking and restocking would drop from \$230,819 to \$213,090, a savings of over \$17,000. Compare this to a more expensive solution: if the size of the A-frame and intermediate storage area were doubled and skus were stored in optimal amounts, the total cost of picking and restocking would be \$216,076, meaning that establishing another A-frame would save Avon only about \$14,000.

We would need to conduct a thorough analysis of the manual storage area in order to




Figure 5.4: Total cost of picking and restocking if more skus were dispensable.

definitively identify ways to reduce the cost per pick there, but one option that should be explored is changing the configuration of the pick-to-light flowrack. This section has six picking lines, each with an identical product mix. Each row is formed from eight U-shaped pods, and the warehouse management system sends empty totes down a conveyor belt that passes by each pod. One picker is assigned to pick skus from each pod; if those skus happen to be slow movers, the picker will finish picking an order quickly and will have nothing to do until the next tote appears, while pickers in busy pods will cause bottlenecks in the system.

Picking by bucket brigades, as described in [3], would reduce the amount of downtime that some pickers currently experience and possibly mean that fewer pickers would be needed. Each picking line in the pick-to-light section of the manual storage area would need to be converted from the current series of eight U-shaped cells to a row of flowrack so that pickers could walk up and down the entire line.

Restock skus in the A-frame less often The operating system of the A-frame sounds an alarm when the supply of a sku in one channel is only half-empty. This causes the A-frame replenishers to restock skus that occupy many channels when the supply is only

slightly depleted, meaning that Avon restocks the A-frame more often than necessary. If a signal to restock were triggered when the *total supply* of a sku is nearly depleted, the number of restocks of the A-frame could be cut *almost in half*, meaning that the amount of labor dedicated to restocking the A-frame could also be reduced by almost half, reducing the labor costs by over \$50,000 each campaign.

Overcoming this limitation of the A-frame will require changes on the part of the manufacturer, however: the A-frame is designed so that each sku has one primary pick channel that is depleted each time a sku appears on an order. Only if an order requires multiple skus is the supply in the other channels depleted.

5.4 Allocation of storage resources in the Avon warehouse

As mentioned in Section 3.3.2, Avon has divided the total supply of flowrack in the order fulfillment area into two sections: one part serves as a supply tier for the A-frame, and the rest is used in the manual storage area.

In practice, however, the A-frame is filled only about 6 feet deep. Our analysis did not factor in the cost of moving leftover supply to another storage mode, but if the penalty is high enough to justify the practice of only partially filling the flowrack, the unused flowrack would be put to better use in the manual storage area.

To see if picking and restock costs are indeed minimized if less flowrack is used to support the A-frame, we modeled the Avon inventory system as if the flowrack in the inventory system were a single mode (see Figure 3.3.4) and defined a variable to represent the fraction of this mode that should be used to support the A-frame. When this model was solved to optimality, calculations showed that the flowrack that supports the A-frame should have a depth of 10 feet, and each restock should completely fill the space available. This indicates that labor currently spends more time restocking the flowrack than is optimal.

CHAPTER 6

SLOTTING A FORWARD-RESERVE SYSTEM OVER MULTIPLE PERIODS

The Avon warehouse is reslotted every two weeks to accommodate new skus and to reflect new demand forecast information for skus already in the warehouse. Since the warehouse stocks many seasonal and promotional skus, the demand for a sku can change quite dramatically from one two-week period to the next, making the assigned storage mode and storage amount inappropriate in the period being slotted. Between 200 and 300 new skus are typically added to the warehouse every two weeks, and about the same number are removed from the warehouse.

Reslotting the Avon warehouse requires a significant number of labor hours. Restockers must be diverted from their primary jobs to box up the remaining quantities of any skus being assigned to a new storage location. Engineers may be called to make mechanical adjustments to the A-frame. One employee works half-time solely to oversee the transfer of skus into and out of the A-frame. Other employees must update the warehouse management system to reflect changes in storage locations.

When planners assign a sku to a storage location for a given period, they consider the demand forecast for not only that period, but for upcoming periods as well. For example, assume that a sku is introduced into the warehouse during a period in which it has low viscosity. If the planners considered only the demand forecast in that period, the sku would be assigned to the shelving area. If the sku will have a significantly higher viscosity in the next period, however, they may assign the sku to the A-frame initially to avoid the cost of reassigning the sku to the A-frame the next period.

In this chapter we will develop several heuristics to slot a forward-reserve inventory system to minimize the total cost of picking, restocking, and reslotting skus over multiple time periods. We call this the *multi-period forward-reserve problem*. At the end of the




Figure 6.1: Forward-reserve inventory system in the multi-period forward-reserve problem.

chapter, we describe an application of this research to the Avon inventory system.

6.1 Model

Inventory system We will consider the forward-reserve system pictured in figure 6.1. Skus can be picked from either the forward mode or the reserve mode, where the forward mode has space V_F and the reserve mode is assumed to have unlimited space. All skus have a supply in the reserve mode.

Time horizon The time horizon over which we slot the forward mode depends on whether we receive information on a rolling or stationary horizon. A stationary horizon is a natural choice when the entire set of skus has the same fixed lifetime in the inventory system. For example, a set of skus may be stocked in the inventory system for a limited number of weeks; then all skus will be removed and a completely new set of skus will be stocked there. Since information is not relevant for any other periods, in this case we slot the forward mode in periods $1, \dots, T$ to minimize costs over these periods.

In other industries, information about picks and flow in any given period may be known for only a few periods in the future. We will assume that when forecasts have such a rolling horizon, we have picks and flow information for T periods in a given period: information for that period and $T - 1$ future periods. In this case we may initially slot the forward mode to minimize costs over periods $1, \dots, T$, but only the slotting in period 1 is fixed. In period 2 we will reslot the forward mode in periods $2, \dots, T + 1$ to minimize total costs over those periods, taking into account the new information received.

Sku information The projected number of picks of sku i in period q is denoted p_{iq} , and the flow is f_{iq} . We will refer to the quantity $\sqrt{f_{iq}}$ as the rootflow of sku i in period q . The sum of the rootflows of all skus in the forward mode in period q is called the *total rootflow* in the forward mode in period q .

Reslotting costs There is a cost to change the storage mode to which a sku is assigned from one period to the next, which we refer to as a *reassignment cost*. Let w be the cost of assigning a sku to the forward mode from the reserve mode plus the cost of later assigning it to the reserve mode again. The entire cost is charged when a sku is assigned from the reserve mode to the forward mode. The cost charged is independent of the amount of stock that must be transferred from one mode to another; this is in keeping with our assumption that the cost of restocks is independent of the amount restocked. This also reflects warehouses where administrative costs associated with reassigning a sku are much greater than the cost of labor to transfer the supply of the sku.

The reassignment cost is the only cost in this model that comes from reslotting. We assume that there is no cost to change the amount of space a sku occupies in a mode from one period to the next, as long as the sku remains in the same mode. This assumption reflects the fact that inventory levels in a storage mode can often be adjusted incrementally as restocking takes place. Since the cost of a restock is assumed to be independent of the amount restocked, a new storage amount can be obtained by adjusting the size of the last restock in a given time period.

Picking and restocking costs The cost per pick from the forward mode is d_F , and the cost per pick from the reserve mode is d_R , where $d_F < d_R$. The savings per pick achieved by assigning a sku to the forward mode is thus $d_R - d_F = s$. If a sku is assigned to the forward mode, there is a cost of c each time the supply of a sku is restocked there. This cost is assumed to be independent of the amount being restocked, and skus are restocked instantaneously when the supply in the forward mode is depleted. Picking and restocking costs remain constant from one period to the next.

To remain consistent with earlier research on forward-reserve inventory systems, we

develop heuristics to slot the warehouse to maximize net total pick savings over T periods, defined as the total pick savings achieved over T periods minus total restock costs and move costs.

6.2 *Heuristics from the single-period forward-reserve problem*

The simplest approach to solving the multi-period forward-reserve problem is to use heuristics that were developed for the single-period forward-reserve problem. One option is to ignore move costs entirely when planning the slotting: we will call this the *one-period heuristic*. If this heuristic is used to solve the multi-period forward-reserve problem, the cost of slotting the forward-reserve system over T periods is then the sum of the minimal picking and restocking cost for each period, plus the cost of reassigning any skus that are in a different storage mode from one period to the next. This is an optimal solution if $w = 0$. If reassignment costs are a very small factor in the total cost of slotting the system over multiple periods, then a warehouse manager may be satisfied with ignoring move costs when planning a slotting.

One improvement we can make to the one-period heuristic is to consider move costs period by period. If sku i is assigned to the forward mode in period 1, but it was in the reserve mode in the previous period, then its net pick savings in period 1 is $sp_{i1} - w$. To simplify notation, we will let I_{iq} be an indicator variable with value 1 if sku i is in the reserve mode in period q , so that the pick savings of sku i in period q can be expressed as $sp_{iq} - wI_{i(q-1)}$.

It is proved in [7] (with explicit development in [1]) that to minimize picking, restocking, and moving costs in period q in the single-period forward-reserve problem, the skus most suitable for placement in the forward mode are those with the highest values of

$$\frac{sp_{iq} - wI_{i(q-1)}}{\sqrt{f_{iq}}}, \quad (6.1)$$

One natural heuristic to solve the multi-period forward-reserve problem, then, is to calculate the value in expression 6.1 for each sku in period 1 and then assign skus to the forward mode in descending order of this value, stopping when the marginal restock cost of the next

sku outweighs its marginal net pick savings. The process is repeated in each subsequent period, calculating the net pick savings of each sku based on the slotting in the previous period. When this heuristic is used to solve the multi-period forward-reserve problem, we will refer to it as the *one-period-with-setup-costs heuristic*.

One drawback of this heuristic is that no future trends are considered. For instance, if two skus have the same viscosity in period q , but the first has higher viscosities in subsequent periods than the second, then the first sku should have a stronger claim to the forward mode in period q when looking at a multi-period horizon.

Another problem with the one-period-with-setup-costs heuristic is its lack of flexibility: the cost of moving a sku must be deducted from the pick savings *in a single period*. If the cost of reassigning a sku is high, it may be desirable to amortize the move cost over several periods, on the expectation that the sku will remain in the forward mode for those periods.

To address these problems, heuristics should consider multi-period assignments when evaluating the strength of the claim of a sku to the forward mode.

6.3 *How to evaluate a multi-period assignment*

If a sku will be assigned to the forward mode in period 1 for t periods, then the strength of its claim to the forward mode is based on the tradeoff of total costs and total savings over a t -period horizon. Since the claim depends on both the sku and on the sequence of periods for which the sku will be in the forward mode, we must evaluate the strength of the claim of *sequences* to the forward mode. The sequence where sku i is assigned to the forward mode for periods p, \dots, q is represented $(i, p:q)$. We refer to this as assigning sequences to the forward mode. We say that a sequence is made up of one or more consecutive *sku-periods*.

In this section we will show that the strength of the claim of sequence $(i, p:q)$ to the forward mode is

$$\frac{\sum_{t=p}^q sp_{it} - wI_{i(t-1)}}{\sum_{t=p}^q \sqrt{f_{it}}}. \quad (6.2)$$

We refer to the quantity in expression 6.2 as the *sequence viscosity* for sequence $(i, p:q)$.

The sequence viscosity represents an average net pick savings per unit of rootflow that the sku would realize if assigned to the forward mode for the corresponding periods. The

Table 6.1: Example sku to show how sequence viscosity is computed

reassignment cost = 3 pick savings = 1		Period			
sku	mode in period 0		1	2	3
a	Reserve	picks	3	12	12
		rootflow	2	3	1

total net pick savings is distributed over each period of the sequence so that the net pick savings achieved in a single period is proportional to the rootflow in that period. Thus the viscosity of each sku-period of the sequence then equals the sequence viscosity. *In the rest of this chapter, when we refer to the net pick savings of a sku in a period, we are referring to the distributed pick savings of the sequence of which the period is a component.*

As an example, consider sequence $(a, 1:3)$, where picks and flow information for each period is given in table 6.1. The net pick savings of the sequence is 24, while the total rootflow of the sequence is 6: the sequence viscosity of $(a, 1:3)$ is thus $24/6 = 4$. Assume the sequence is assigned to the forward mode. Then we say that in period 1 the sku realizes a net pick savings of $2 \cdot (24/6) = 8$, even though the actual pick savings in the period is less. In period 2 the sku realizes a pick savings of $3 \cdot (24/6) = 12$, and in period 3 the sku realizes a savings of $1 \cdot (24/6) = 4$. The viscosity in each period of the sequence is then 4.

6.3.1 Intuitive justification of ranking by sequence viscosity

The space available in the forward mode is the limiting resource, and in the optimal solution of the multi-period forward-reserve problem, skus will be assigned to the forward mode in each period to maximize total pick savings over T periods within the limits of this resource. Sequences with the highest ratios of pick savings to space consumed are the strongest candidates for placement in the forward mode.

As an example, consider over three periods an arbitrary sku i that was not in the forward mode in period 0. We are considering assigning sku i to the forward mode for some number of periods. For each period that sku i is assigned to the forward mode, assume that we know the space it will be allotted there and the pick savings it will achieve, as listed in

Table 6.2: Example sku for intuitive justification of sequence viscosity

reassignment cost = 3 pick savings = 1		Period			
sku	mode in period 0		1	2	3
i	Reserve	picks	3	12	12
		space if in forward mode	1	3	2

table 6.2.

If sequence $(i, 1:3)$ is assigned to the forward mode, it occupies a total of 6 period-units of space, as shown in diagram A of figure 6.2, and has a total net pick savings of 24, since sku i was not in the forward mode in period 0. The ratio of savings to period-units of space is thus $24/6 = 4$.

If instead sequence $(i, 1:2)$ is placed in the forward mode, then the sku occupies 4 period-units of space in the forward mode, as shown by diagram B of figure 6.2. The sku achieves a total net pick savings of 12, giving a ratio of savings to space of $12/4 = 3$.

Alternatively, if sequence $(i, 2:3)$ is placed in the forward mode, then the sku occupies 5 period-units of space in the forward mode, as shown by diagram C of figure 6.2. The sku achieves a total net pick savings of $12 + 12 - 3 = 21$, giving a ratio of savings to space of $21/5$.

To assign sku i to the forward mode for a set of nonconsecutive periods, consider them as two independent sequences and consider assigning each to the forward mode separately.

In the above derivation, we have assumed that we know in advance the amount of space that sku i will occupy in the forward mode each period that it is assigned there. In reality, however, this quantity can not be determined in advance because it depends on which other products are assigned to the forward mode. To see this, observe that because there is no cost to change the *amount* of space a sku receives in the forward mode each period, the skus assigned to the forward mode in a given period will be allotted space to minimize the cost of restocking them there. Results in [7] show that this quantity depends on the other skus assigned to the forward mode. Letting ϕ_q represent the total rootflow of the skus in the forward mode in period q and letting v_{iq} represent the volume occupied by sku i if it is

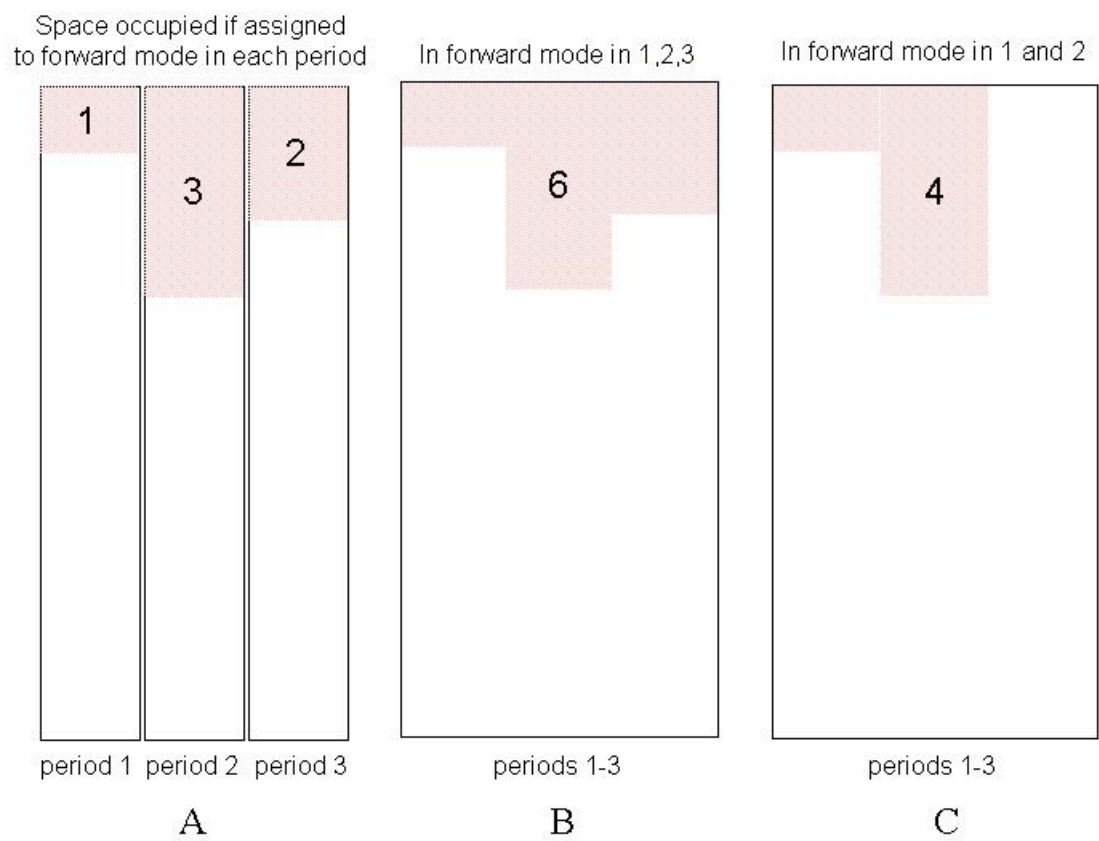


Figure 6.2: Space occupied by sku i in the forward mode over three periods.

assigned to the forward mode in period q , then by [7],

$$v_{iq} = \frac{\sqrt{f_{iq}}}{\phi_q} V_F \quad \text{ft}^3$$

Thus if sequence $(i, 1:3)$ in the above example is assigned to the forward mode, sku i will occupy

$$\frac{\sqrt{f_{i1}}}{\phi_1} V_F + \frac{\sqrt{f_{i2}}}{\phi_2} V_F + \frac{\sqrt{f_{i3}}}{\phi_3} V_F$$

cubic feet over the three periods.

If we make the simplifying assumption that the value of ϕ_q is the same for all periods q , then the total volume occupied by sku i in periods 1,2, and 3 is proportional to

$$\sqrt{f_{i1}} + \sqrt{f_{i2}} + \sqrt{f_{i3}}.$$

By assuming that the total rootflow in the forward mode is the same in each period, we assume that the forward mode is restocked approximately the same number of times in each period. One period may see twice as many skus in the forward mode as another, but we assume that in both periods, the total number of restocks is the same. The model is most appropriate for those warehouses where the demand for restocks of the forward mode is consistent over time, or where changes in demand occur gradually. The model may not adequately capture the activity in warehouses with occasional surges in which the forward mode must be restocked much more often than normal.

Sku i thus has a strong claim to be placed in the forward mode for periods 1, 2, and 3 if it has a high value of

$$\frac{\sum_{t=1}^3 sp_{it} - wI_{i0}}{\sum_{t=1}^3 \sqrt{f_{it}}}.$$

6.3.2 Theoretical justification of ranking by sequence viscosity

To more rigorously develop the criterion to evaluate the strength of the claim of a sequence to the forward mode, we present a mathematical programming approach. Let x_{iq} be a binary variable with value 1 if sku i is in the forward mode in period q , and let v_{iq} be the amount of space sku i is allotted in the forward mode in period q . Define a function that expresses the cost of restocking sku i in period q as $r(v_{iq}) = cf_{iq}/v_{iq}$ if $x_{iq} > 0$ and 0 otherwise. Then

a mathematical program that corresponds to the multi-period forward-reserve problem over periods $1, \dots, T$ is

$$MP(x_{iq}, v_{iq}) : \quad \max \sum_{q=1}^T \sum_{i=1}^n [(sp_{iq} - wI_{iq-1}) x_{iq} - r(v_{iq})] \quad (6.3)$$

$$s.t. \sum_{i=1}^n v_{iq} x_{iq} \leq V_F \quad \forall q \quad (6.4)$$

$$v_{iq} \geq 0 \quad \forall i, q \quad (6.5)$$

$$x_{iq} \in \{0, 1\} \quad \forall i, q \quad (6.6)$$

$$I_{iq-1} = (1 - x_{iq-1}) \quad \forall i, q \quad (6.7)$$

To obtain a single knapsack-type constraint, we sum constraints 6.4 over all periods q to arrive at the following relaxation:

$$R(x_{iq}, v_{iq}) : \quad \max \sum_{q=1}^T \sum_{i=1}^n [(sp_{iq} - wI_{iq-1}) x_{iq} - r(v_{iq})] \quad (6.8)$$

$$s.t. \sum_{q=1}^T \sum_{i=1}^n v_{iq} x_{iq} \leq TV_F \quad (6.9)$$

$$v_{iq} \geq 0 \quad \forall i, q \quad (6.10)$$

$$x_{iq} \in \{0, 1\} \quad \forall i, q \quad (6.11)$$

$$I_{iq-1} = (1 - x_{iq-1}) \quad \forall i, q \quad (6.12)$$

Since skus assigned to the forward mode in period q will be allotted space to minimize the number of restocks of the forward mode in period q , we substitute the values of v_{iq} derived by Hackman and Rosenblatt in [7] into the above math program. These values are a function of the total rootflow assigned to the forward mode in period q , which we denote

ϕ_q . Then $R(x_{iq}, v_{iq})$ is transformed into

$$R(x_{iq}, \phi_q) : \quad \max \sum_{q=1}^T \left(\sum_{i=1}^n (sp_{iq} - wI_{iq-1}) x_{iq} \right) - c \frac{(\phi_q)^2}{V_F} \quad (6.13)$$

$$s.t. \sum_{q=1}^T \sum_{i=1}^n \frac{\sqrt{f_{iq}}}{\phi_q} x_{iq} \leq T \quad (6.14)$$

$$x_{iq} \in \{0, 1\} \quad \forall i, q \quad (6.15)$$

$$I_{iq-1} = (1 - x_{iq-1}) \quad \forall i, q \quad (6.16)$$

If we assume that the value of ϕ_q is approximately the same in each period, then $R(x_{iq}, \phi_q)$ becomes

$$R(x_{iq}, \phi) : \quad \max \sum_{q=1}^T \left(\sum_{i=1}^n (sp_{iq} - wI_{iq-1}) x_{iq} \right) - c \frac{(\phi)^2}{V_F} \quad (6.17)$$

$$s.t. \sum_{q=1}^T \sum_{i=1}^n \sqrt{f_{iq}} x_{iq} \leq T\phi \quad (6.18)$$

$$x_{iq} \in \{0, 1\} \quad \forall i, q \quad (6.19)$$

$$I_{iq-1} = (1 - x_{iq-1}) \quad \forall i, q \quad (6.20)$$

The sequences with the strongest claim to the forward mode in $R(x_{iq}, \phi)$ are those with the highest sequence viscosities. It is hoped that the optimal solution to the relaxed problem is a good solution for the original problem.

6.4 Multiperiod heuristics

When deciding which skus should be assigned to the forward mode in a given period, the multi-period heuristics consider the picks and flow forecasts of a sku not only in that period, but also in subsequent periods. The periods from which a heuristic considers the picks and flow information is called the *information horizon* of a heuristic.

The information horizon available depends on whether we receive information on a rolling or stationary horizon. With a stationary horizon, the information horizon in period q is q, \dots, T . With a rolling horizon, the information horizon in period q is $q, \dots, q + T - 1$.

When the single-period heuristics are used to solve the multi-period forward-reserve problem as in section 6.2, they slot the forward mode entirely in one period before slotting the forward mode in the next period. The first two multi-period heuristics we consider use the same iterative approach when solving the multi-period forward-reserve problem. For example,

- In a problem with a stationary horizon, skus are assigned to the forward mode in period 1 using information about their picks and flow in periods $1, \dots, T$. Then skus are assigned to the forward mode in period 2 using information from periods $2, \dots, T$, and so on.
- In a problem with a rolling horizon, the forward mode is slotted in period 1 using information from periods $1, \dots, T$. Then the forward mode is slotted in period 2 using information about each sku in periods $2, \dots, T + 1$, and so on.

In both cases, some sequences that were assigned to the forward mode in a given period may be terminated early when assignments are made for subsequent periods. For example, sku i may be assigned to the forward mode in period 1 on the strength of the sequence $(i, 1:3)$, but when the forward mode is slotted in period 2, removing sku i from the forward mode may allow higher cost savings.

Table 6.3 lists the characteristics of each heuristic we develop.

6.4.1 The T -period heuristic

The T -period heuristic is based on the idea that since we are slotting the forward mode for periods $1, \dots, T$ periods, we want skus that are good candidates during all those periods. In period 1, for each sku i we calculate the sequence viscosity for the sequence $(i, 1:T)$. We then assign skus to the forward mode in period 1 in descending order of this value until the marginal restock cost of adding the next sku to the forward mode outweighs the marginal net pick savings in period 1. We formally state the algorithm as if slotting the inventory system in period 1 given a rolling horizon.

Table 6.3: Characteristics of heuristics for the multi-period forward-reserve problem

Heuristic	Sequential assignment by period or global assignment	Re-assignment costs considered	Horizon Type	Information horizon when slotting period q	Sequences considered within the information horizon	
					start of sequence	end of sequence
Ignore move costs	sequential	no	Stationary	q	q	q
			Rolling	q	q	q
One-period-with-reassignment-costs	sequential	yes	Stationary	q	q	q
			Rolling	q	q	q
T-period	sequential	yes	Stationary	q, \dots, T	q	T
			Rolling	$q, \dots, q+T-1$	q	$q+T-1$
Best viscosity	sequential	yes	Stationary	q, \dots, T	q	$q, q+1, \dots, T$
			Rolling	$q, \dots, q+T-1$	q	$q, q+1, \dots, q+T-1$
Best sequence	global	yes	Stationary	q, \dots, T	q, \dots, T	$q, q+1, \dots, T$
			Rolling	$q, \dots, q+T-1$	$q, \dots, q+T-1$	$q, \dots, q+T-1$

Table 6.4: Example skus for the T -period heuristic and best-viscosity heuristic

reassignment cost = 3 pick savings = 1		Period					
sku	mode in period 0		1	2	3	4	T -period viscosity
a	Reserve	picks	3	12	12	10	4.25
		rootflow	2	3	1	2	
b	Forward	picks	16	5	2	2	4.167
		rootflow	2	1	1	2	

The T -period heuristic

For each sku i

 Calculate the sequence viscosity of the sequence $(i, 1:T)$

Sort skus in descending order of this value.

Do until the marginal cost of the next sku outweighs the marginal net pick savings in period 1

 Assign skus to the forward mode in period 1 in rank order.

The T -period heuristic has a worst-case running time of $O(n \log n)$.

One drawback of the T -period heuristic is that even though it has an information horizon of periods $1, \dots, T$, it considers only one sequence for each sku within that horizon, namely the sequence $(i, 1:T)$ for each sku i .

Example Consider the skus in table 6.4: The sequence $(a, 1:4)$ has a higher sequence viscosity than sequence $(b, 1:4)$, but sku b has a higher viscosity in period 1 than sku a . This means that if the T -period heuristic assigned sku a to the forward mode and sku b to the reserve mode in period 1, then we could increase the savings of the heuristic solution by keeping sku b in the forward mode for period 1, and leaving sku a in the reserve mode until period 2.

The T -period heuristic devalues the claim of sku b to the forward mode because it is a sku with declining viscosity: in period 1, it has a high viscosity relative to sku a , but its viscosity begins to decrease in the next period. When the T -period heuristic computes the viscosity of the T -period sequence $(b, 1:4)$ for sku b , it takes into account the low-viscosity

later periods, even though they may be of secondary importance when we are trying to decide where the sku should be placed in period 1.

The next heuristics we present will consider more sequences within the information horizon $1, \dots, T$.

6.4.2 The best-viscosity heuristic

Like the T -period heuristic, the best-viscosity heuristic finds a slotting of the forward mode over T periods by considering each period in sequence. Instead of judging each sku on the sequence viscosity of a T -period sequence, however, the best-viscosity heuristic judges a sku on the sequence for which the sku achieves the highest sequence viscosity. When slotting the forward mode in period 1, for each sku i the heuristic evaluates the sequence viscosity of sequences $(i, 1:1), (i, 1:2), \dots, (i, 1:T)$. The highest value is the *best viscosity* for sku i ; skus are assigned to the forward mode in period 1 in descending order of their best viscosity until the marginal restock cost of adding the next sku outweighs the marginal net pick savings in period 1. We formally state the algorithm as if slotting the inventory system in period 1 given a rolling horizon.

The best-viscosity heuristic

For each sku i

Calculate the best viscosity over periods $1, \dots, T$

Sort skus in descending order of the best viscosity

Do until the marginal cost of the next sku outweighs the marginal net pick savings in period 1

Assign skus to the forward mode in period p in rank order.

The best-viscosity heuristic has a worst-case running time of $O(Tn + n \log n)$. Since the value of T is typically several orders of magnitude less than the value of n in a typical warehouse, the worst-case running time is $O(n \log n)$.

Example Table 6.5 lists the sequence viscosities in period 1 for the skus in table 6.4. The best viscosity of sku a is 4.25 from either the sequence $(a, 1:3)$ or $(a, 1:4)$, and

Table 6.5: Sequence viscosities of skus in table 6.4

	Sequence viscosities				
sku	$(\cdot, 1:1)$	$(\cdot, 1:2)$	$(\cdot, 1:3)$	$(\cdot, 1:4)$	best viscosity
a	0	2.4	4	4.25	4.25
b	8	7	5.75	4.17	8

the best viscosity of sku b is 8 from the single-period sequence $(b, 1:1)$. According to the best-viscosity heuristic, sku b has a stronger claim to the forward mode in period 1 than sku a ; the low viscosities of sku b in future periods are not considered by the best-viscosity heuristic when slotting the forward mode in period 1.

The best-viscosity heuristic is more flexible than the best-viscosity heuristic, since it considers more than one sequence when ranking skus. It may, however, weight future periods too heavily. Skus with low viscosities in period 1 and very high viscosities in later periods may have a sufficiently high best viscosity to be assigned to the forward mode in period 1. These skus may displace skus already in the forward mode that are better candidates for the forward mode in period 1—in effect, the best-viscosity heuristic can move a sku to the forward mode too soon.

A better approach to slotting the forward mode would be to explicitly consider the periods during which a sku is a good candidate for the forward mode. We present two methods for doing this: in the next section we discuss a heuristic that finds local improvements to the solution given by the best-viscosity heuristic, and in the following section we present a heuristic that slots the forward mode in all periods simultaneously.

6.4.3 The best-viscosity-improvement heuristic

The best-viscosity heuristic assigns sequences to the forward mode that have the highest sequence viscosities given that the sku is moved to the forward mode *in period 1*. In some cases, however, postponing the assignment of a sku to the forward mode to a later period may improve the overall solution.

Example Sku c in table 6.6 has a best viscosity of 6.125 with sequence $(c, 1:4)$, while sku d has a best viscosity of 5 with sequence $(d, 1:1)$. The best-viscosity heuristic may assign

Table 6.6: Example skus for the best-viscosity-improvement heuristic

reassignment cost = 2 pick savings = 1						
sku	mode in period 0		period 1	period 2	period 3	period 4
c	Reserve	picks	9	12	10	20
		rootflow	2	3	1	2
d	Forward	picks	15	5	2	2
		rootflow	2	1	1	2

sku c to the forward mode and sku d to the reserve mode in period 1, even though sku d has a higher viscosity (7.5) in period 1. We can improve the solution to the multi-period forward-reserve problem by leaving sku d in the forward mode for period 1 and moving sku c to the forward mode in period 2.

If sku c is not assigned to the forward mode until period 3, the sequence viscosity of sequence $(c, 3 : 4)$ increases to 9.33, contributing a greater savings per period-unit of rootflow to the solution of the multi-period forward-reserve problem. We assume that in period 3 we will assign sequence $(c, 3 : 4)$ to the forward mode; we then evaluate the marginal contribution of sequence $(c, 1 : 2)$ separately. If there is a sequence with higher best viscosity over these periods than sequence $(c, 1 : 2)$, then the sku is a better candidate for the forward mode in period 1 than sku c .

The marginal contribution of sequence $(c, 1 : 1)$ is $(9 - 2)/2 = 3.5$. The marginal contribution of sequence $(c, 1 : 2)$ is $(9 + 12)/(2 + 3) = 4.2$. In the second case, we do *not* need to include the reassignment cost when computing the sequence viscosity. We have assumed that sequence $(c, 3 : 4)$ will be assigned to the forward mode and accounted for the cost of one reassignment of sku c to the forward mode in the viscosity of that sequence. No additional reassignments are required if sku c is also assigned to the forward mode in periods 1 and 2. Sequence $(d, 1 : 1)$ has a viscosity of 5; since this is larger than the best viscosity of sku c , sku d is a better candidate than sku c for the forward mode in period 1.

The best-viscosity-improvement heuristic is applied after each iteration of the best-viscosity heuristic; it re-evaluates the claims to the forward mode of skus assigned there by the best-viscosity heuristic. Let the skus assigned to the forward mode in period 1 by the

best-viscosity heuristic be the *assigned* skus in period 1. We formally state the algorithm as if slotting the inventory system in period 1 given a rolling horizon.

The best-viscosity-improvement heuristic

Run the best-viscosity heuristic in period 1 with a rolling information horizon of periods $1, \dots, T$.

For periods $q = T$ to 2 (in reverse order)

Find all assigned skus in period 1 whose values of best viscosity increase if the sku is not assigned to the forward mode until period q .

Consider these skus part of the unassigned skus in period 1.

For each unassigned sku i

Calculate the best viscosity over periods $1, \dots, q - 1$

Sort the unassigned skus in period 1 in descending order of this value.

While marginal net pick savings of the next sku outweighs the marginal restock cost in period 1

Assign unassigned skus to the forward mode in period 1 in rank order

The best-viscosity-improvement heuristic considers postponing skus to period T first because skus that may replace it in the forward mode achieve the highest possible values of best viscosity when evaluated over the longest possible time horizon. Assuming that T is significantly smaller than n , the best-viscosity-improvement heuristic has a worst-case running time of $O(Tn \log n)$.

The best-viscosity-improvement heuristic was effective at finding low-cost multi-period slottings of the Avon warehouse. More information is provided in Section 6.6.

6.4.4 The best sequence heuristic

Like the previous multi-period heuristics, the best sequence heuristic assigns sequences to the forward mode in descending order of sequence viscosity until a stopping point is reached.

The best sequence heuristic differs from the T -period heuristic and the best-viscosity heuristic in two important aspects, however: first, the best sequence heuristic considers *every* sequence within the information horizon $1, \dots, T$, and not just those beginning with period 1. Second, the forward mode is not slotted completely in one period before it is slotted in the next. For example, the best sequence heuristic may assign the sequence $(i, 4:5)$ to the forward mode in one iteration and sequence $(j, 2:3)$ in the next.

We formally state the algorithm as if slotting the inventory system in period 1 given a rolling horizon.

The best sequence heuristic

For each sku i

Determine the sequence for which the sku achieves the highest viscosity

Rank all sequences in descending order of sequence viscosity

Do until stopping point

Assign to the forward mode the sequence with the highest viscosity over all sequences

For the sku corresponding to this sequence

Compute the sequence with the next highest viscosity

Insert the new sequence into the list ranked by sequence viscosity

The best sequence heuristic has a worst case running time of $O(Tn \log n)$

6.4.4.1 Computing the sequence viscosity of a sequence

The sequence viscosity of a sequence in iteration k is the increase in savings over the solution after iteration $k - 1$ that would be realized by assigning the sequence to the forward mode in iteration k , divided by the rootflow of the sequence (in period-units).

Example Table 6.7 lists the picks and flow of sku a over periods 1–5, and table 6.8 gives the initial sequence viscosities for all possible sequences of sku a over these periods. Initial sequence viscosities are computed assuming that the sku has not yet been assigned to the forward mode in any period 1 – 5.

Table 6.7: Example sku for the best sequence heuristic

move cost = 2 pick savings = 1							
sku	initial mode		period 1	period 2	period 3	period 4	period 5
a	Forward	picks	3	12	12	10	11
		rootflow	2	3	1	2	2

Table 6.8: Sequences for sku a : Iteration 1

Sequence viscosities: Iteration 1						
	ending period of sequence					
starting period of sequence	1	2	3	4	5	
1	1.5	3	4.5	4.63	4.8	
2		3.33	5.50	5.33	5.38	
3			10.00	6.67	6.20	
4				4.00	4.75	
5					4.50	

The sequence with the highest sequence viscosity is $(a, 3:3)$. If this sequence is assigned to the forward mode, the sequence of sku a with the next highest viscosity is computed with this knowledge. Table 6.9 lists the sequence viscosities for the sequences remaining after $(a, 3:3)$ has been assigned to the forward mode.

Two important points arise from the second computation of the sequence viscosities. First, if sequence $(a, 4:4)$ or $(a, 4:5)$ is assigned to the forward mode in the second iteration, no additional reassignments of sku a are required, so no reassignment costs need to be accounted for in the sequence viscosities. The sequence viscosity of sequence $(a, 4:4)$

Table 6.9: Sequences for sku a : Iteration 2

Sequence viscosities: Iteration 2						
		ending period of sequence				
starting period of sequence		1	2	3	4	5
	1	1.5	3.4	—	—	—
	2		4	—	—	—
	3			—	—	—
	4				5.00	5.25
	5					4.50

increases from $(10 - 2)/2 = 4$ to $10/2 = 5$, and the sequence viscosity of sequence $(a, 4:5)$ increases from $(10 + 11 - 2)/(2 + 2) = 4.75$ to $(10 + 11)/(2 + 2) = 5.25$.

Second, if sequence $(a, 1:2)$ is assigned to the forward mode in iteration 2, the number of reassignments of sku a *decreases* from the previous iteration. Sku a is in the forward mode in period 0, and therefore the sequence viscosity of sequence $(a, 1:2)$ is computed assuming a net pick savings of $3 + 12 = 15$ plus a reassignment cost *credit* of 2, bringing the sequence viscosity to $(15 + 2)/(2 + 3) = 3.4$

Observation 6.1 *For a given sku, the sequence with highest sequence viscosity generated after iteration k has no higher a sequence viscosity than the one generated after iteration $k - 1$.*

Simple algebra shows that if the sequence with highest sequence viscosity in iteration k has a higher sequence viscosity than the sequence generated in iteration $k - 1$, then combining the two sequences in iteration $k - 1$ would give a sequence with higher viscosity than the sequence chosen in iteration $k - 1$, meaning that it was not the sequence of highest viscosity in that iteration.

6.4.4.2 Stopping conditions of the best sequence heuristic

As a default, the best sequence heuristic stops when adding the next sequence to the forward mode would decrease the net pick savings in at least one period. This is a conservative stopping point, since the total savings may increase if skus are added to the forward mode in other periods. As an alternative, when adding a sku-period to a certain period no longer increases total pick savings in that period, consider only those sequences that do not contain that period. The heuristic stops when there is no period for which savings increases if any remaining sku-period is added to the forward mode. We will note when the alternative stopping condition is used.

6.5 Bounding the solution of the multi-period forward-reserve problem

An intuitive upper bound on the total net pick savings of the multi-period forward-reserve problem is calculated by slotting the forward mode in each period to maximize the pick savings in each period—ignoring the costs of any reassignments—then summing the resulting savings over all periods. We will call this the no-move-cost-upper-bound.

The drawback of the no-move-cost-upper-bound is that the difference between this value and the optimal solution increases as the reassignment cost increases. We are currently working to develop bounds that are computed with respect to the reassignment cost. We believe that we can use the process of sequentially finding the best sequences of a sku to compute an upper bound on the pick savings that each sku-period would achieve if assigned to the forward mode. With these values and the rootflow of each sku-period, we then find an upper bound on the total pick savings that would be achieved in each period by maximizing the bounded pick savings minus restock costs in each period.

6.6 Case Study

We compared the performance of the heuristics for the multi-period forward-reserve problem by applying each to data for the Avon warehouse over multiple sales campaigns. Avon reslots the warehouse before each campaign, so each campaign represents a period.

Avon forecasts sales in future periods by releasing items for sale to a small portion of their customers two sales periods ahead of most of their customer base. The demand for items among the trial customers is used to forecast the demand of those products among the larger customer base, meaning that at the start of any period, the planners have forecast information for that period as well as the two following periods. This is especially useful for new skus that are introduced with no historical sales data.

Data was available for each sku on the piece sales and the storage location during eight periods. The information for skus in period 1 served to establish an initial location for each sku in period 2. Using data for periods 2–8, slottings were determined for the five periods 2–6.

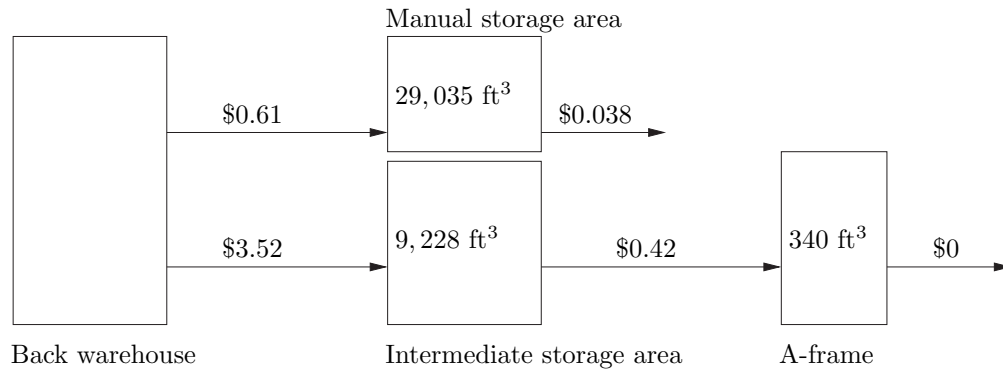


Figure 6.3: Configuration of Avon’s warehouse in Suwanee, GA.

Since Avon does not collect information about the composition of the orders it fills each day, the picks for each sku were set equal to piece sales. This inflated the viscosity of the skus to varying degrees. To get a sense of how this decision affected the viscosities of the skus, we used data we had gathered for the case study in Chapter 5 to compare the picks of each sku to the total quantity sold during one day of Campaign 4 of the previous year. Approximately 20% of the skus in that campaign had total piece sales forecasts that were more than 25% greater than their picks; these skus were fairly evenly distributed among the top two-thirds of the skus ranked by viscosity. We believe that the heuristic solutions will therefore assign slightly more skus to the forward mode than is optimal, but the assigned skus will be the best candidates for the forward mode. The pick savings achieved by the heuristic solutions will be slightly inflated.

The Avon inventory system was modeled as in Chapter 3, shown in figure 6.3. Skus can be picked from the A-frame or from an area of shelving and flowrack called the manual storage area. All skus that are picked from the A-frame have a supply in the intermediate storage area, which comprises several bays of flowrack. The two possible flowpaths in the system are thus *reserve mode, intermediate storage area, A-frame* and *reserve mode, manual storage area*. No labor is required to pick skus from the A-frame, meaning the pick cost is effectively 0, and thus the pick savings from assigning a sku to the A-frame is the cost per pick in the manual storage area, or \$.038.

Since there is no time or labor force dedicated specifically to moving a sku—it is done as picking and restocking is taking place—Avon has no official estimates on the cost to

reassign a sku. After observing the move process, we estimate that the act of physically moving the supply of a sku to another storage location requires time on the order of several minutes. Using the Avon hourly labor rate of \$17, the cost of physically moving a sku can be estimated to be \$1–\$3. The activities associated with updating the warehouse management system and re-engineering storage locations in the A-frame contribute additional costs. We thus ran the heuristics using five different values of reassignment cost: \$2, \$10, \$20, \$100, and \$1,000 per reassignment. The larger move cost values reflect different penalty costs for the activity of moving skus, as well as opportunity costs for labor used to reassign skus.

6.6.1 Profile of the Avon warehouse

The eight periods of data corresponded to the first four months of 2003. There were about 6,800 skus in the warehouse each period (see figure 6.4); the total pieces sold during each period varied from 7.8 million pieces in period 1 to 12.8 million pieces in period 6, as shown in figure 6.5. Figure 6.6 shows the distribution of skus by piece sales in a typical campaign.

On average, 300 new skus were introduced in each period, and approximately the same number left the warehouse each period. About 150 skus were new to the A-frame each period, including both skus new to the warehouse and skus that were reassigned from another storage mode.

6.6.2 Comparison of heuristics

The best-viscosity-improvement heuristic and best sequence heuristic (using the alternative stopping condition) consistently found slottings of the warehouse with the highest overall cost savings for all values of reassignment cost except $w = 2$, where the one-period-with-setup-costs heuristic did slightly better. In future research we will study the performance of other stopping conditions for the best sequence heuristic as well as post-heuristic improvement algorithms.

Tables 6.10 and 6.11 compare the performance of the best sequence heuristic with heuristics designed for the single-period forward-reserve problem. A comparison of the savings given by these heuristics is shown in figure 6.7. The importance of using a heuristic designed to solve the multi-period forward-reserve problem becomes more significant as the




Figure 6.4: Number of skus in the Avon warehouse by period




Figure 6.5: Pieces sold from the Avon warehouse by period




Figure 6.6: Distribution of skus by pieces sales in a typical campaign




Figure 6.7: Cost savings per campaign from several heuristics

reassignment cost increases. Figure 6.8 presents a more detailed look at the performance of the three heuristics designed specifically for the multi-period forward-reserve problem.

The upper bound on savings over five periods given by the no-move-cost-upper-bound is \$885,855, and table 6.12 shows the distance between the solution of the best sequence heuristic and the bound (expressed as a percentage of the solution value) for various move costs. For reassignment costs \$2, \$10, and \$20, the heuristic provides a solution that is within 2% of optimal. The gap between the bound and the solution may be unacceptably large for higher move costs, motivating the development of an upper bound that is sensitive to move costs.


If we assumed a rolling information horizon of the slotting period plus *three* future periods, the best sequence heuristic found the lowest-cost multi-period slotting for all values of reassignment cost. The best-viscosity-improvement heuristic did not perform as well with a longer information horizon because it is based on the best-viscosity heuristic and can thus weight future periods too heavily, despite the fact that it was designed to correct this tendency.

Table 6.10: Pick savings in an average period for heuristics

	Reassignment Cost				
	2	10	20	100	1000
one-period	\$176,669	\$176,669	\$172,155	\$152,091	\$-73,629
one-period-with-reassignment costs	\$176,682	\$174,928	\$172,984	\$162,831	\$132,673
best-sequence	\$176,673	\$175,123	\$173,645	\$166,387	\$143,121

Table 6.11: Skus moved to the A-frame in an average period

	Reassignment Cost				
	2	10	20	100	1000
one-period	251	251	251	251	251
one-period-with-reassignment costs	234	190	159	68	6
best-sequence	232	165	134	72	13

**Figure 6.8:** Cost savings per campaign from multi-period heuristics**Table 6.12:** Maximum error of the solution of the best sequence heuristic

	Move cost				
	\$2	\$10	\$20	\$100	\$1000
percent difference from no-move-cost-upper-bound	0.3%	1.2%	2%	6.5%	23.8%

The best-viscosity heuristic was less effective than algorithms designed for single-period use when reassignment costs were low: the myopic one-period-with-setup-costs heuristic realized higher cost savings for all move costs \$20 and under, even with a four-period forecast horizon. The main reason for this is that if the reassignment cost is low compared to pick savings, we are primarily concerned with optimizing total pick savings in the forward mode in each period, and future forecasts are less important.

We also investigated the effects on the Avon inventory system of using information about future periods, but reslotting less frequently. To simulate reslotting the warehouse every third campaign, we aggregated the total picks and flow for periods 1 – 3 into a single 6-week period, and did the same for periods 1 – 4. We then applied the one-period-with-setup-costs heuristic to the two 6-week periods. This approach outperformed heuristics designed for the single-period forward-reserve problem given reassignment costs \$10 or higher, and the average cost savings per campaign differed from that of the best sequence heuristic by less than \$1000 for all reassignment costs except \$1000.

The approach did not perform as well when we slotted the warehouse for the single 6-week period comprising periods 1 – 3 and maintained that slotting for periods 4 and 5: the total net pick savings decreased by about \$40,000 over periods 4 and 5 given reassignment costs of \$2, \$10, or \$20, and by more for reassignment costs \$100 and \$1000.

6.6.3 Implications for the Avon warehouse

Accounting for move costs when slotting the inventory system Our calculations show that the policy that Avon uses to reassign skus in the Avon warehouse is consistent with the near-optimal policy given a reassignment cost slightly less than \$20. With a reassignment cost value of \$20, the best sequence heuristic moved an average of 134 skus into the A-frame each period; Avon moved an average of 153.

If the cost of reassigning skus in the Avon warehouse is mainly a function of the cost of physically transferring skus from one storage mode to another, then a more realistic estimate of reassignment costs is just a few dollars. In that case, reassignment costs can be ignored when slotting the Avon order fulfillment area with no significant increase in

Table 6.13: Cost savings per campaign if future forecasts are known

Demand known for	Move cost				
	\$2	\$10	\$20	\$100	\$1000
1 period	\$176,682	174,928	172,984	62,831	132,673
2 periods, additional savings from 1 period	5	128	455	2,603	8,890
3 periods, additional savings from 1 period	—	—	—	460	1,757

costs. On the other hand, if administrative and engineering costs make this cost more than \$10, then heuristics that take reassignment costs into account should be used. If the actual cost of moving a sku into and out of the A-frame is \$20, a slotting in which this cost is ignored will require 60 more reassignments than the slotting provided by the best sequence heuristic, adding an extra \$1,200 to total operating costs each period.

Value of future forecasts As expected, the lower the reassignment cost, the less important it is to know forecasts for future periods. As seen in table 6.13, knowing the forecast for at least one period beyond the slotting period helped increase the cost savings per campaign. The additional difference in total cost savings from knowing forecasts for two periods beyond the slotting period was negligible if the cost per move was \$20 or less. The importance of future forecasts depends to a large extent on how often new skus are introduced into the warehouse, as well as how varied their sales patterns are over many campaigns.

6.7 *Future Research*

6.7.1 Incorporating different types of reassignment costs

In the problem just described, we modeled the cost of changing the slotting of the inventory system from one period to another as a cost per sku moved to the forward mode. This cost structure is an appropriate reflection of costs that occur with every sku moved, such as the cost to physically move the supply of the sku from one storage location to another, and the cost to update the warehouse management system. But reslotting the warehouse often involves other types of costs. For example, in the Avon warehouse, the mechanical limitations of the A-frame contribute to another type of cost: the width of the dispenser

channels that hold the supply of each sku in the A-frame must be adjusted if the new skus assigned to the location have different dimensions than the previous ones. For technical reasons, the engineering department must do this, and so there is a labor cost for adjusting the width of a channel.

In the previous model, we assumed that the cost to change the volume of a sku is negligible, but this may not be a good assumption in some applications. For example, in an automatic picker such as an A-frame, the control software may be programmed to deplete the supply of a sku in a given channel, and the restockers can fill it with another sku as if they are making a replenishment. The procedure may be more complicated for skus stored in shelving or flow rack, since restockers may have to move the skus to a storage location with a different depth.

6.7.2 Considering multi-mode inventory systems over multiple periods

A goal for future research is to build on our results for forward-reserve systems to develop methods for slotting multi-mode inventory systems over multiple time periods.

CHAPTER 7

CONCLUSIONS

In this research we have developed and analyzed a model for multi-tier inventory systems, a more general class of inventory systems than a forward-reserve or multi-mode inventory system. By showing that a multi-tier inventory system has an equivalent multi-mode inventory system, we can find a slotting for a multi-tier inventory system where picking and restocking costs are close to minimal.

Multi-tier inventory systems can be commonly found in warehouses; those warehouses that use an A-frame to pick items may have a storage mode located nearby that serves to provide quick replenishments of the A-frame. Our algorithm for finding a near-minimum cost slotting of such systems can provide guidance for warehouse personnel who slot such systems. Based on the results of our case study at Avon Products, Inc., our research can help identify ways to reduce total picking and restocking costs on the order of 10-20% by appropriately managing which skus are assigned to each storage mode, and the quantities in which they are stored there.

If storage modes are shared among flowgroups in a given multi-tier inventory system, enumeration is the only provably effective method of finding an near-minimum cost slotting. For a multi-tier inventory system with P flowgroups, the worst-case running time is on the order of $O(n^P)$, which is computationally impractical for inventory systems with more than a few flowpaths. The most important area we have identified for future research is therefore to develop faster methods for finding good slottings of general multi-tier inventory systems. We believe that an effective approach may be similar to that taken in [2] where it was proved that a near-minimum cost slotting can be found for a multi-mode inventory system with M modes in a worst-case running time of $O((\log n)^M)$.

In all multi-tier inventory systems that we analyzed, we have assumed that the cost of picking a sku from a given storage mode is the same for all skus picked there. Because

of this, the picking modes can be ranked so that each sku achieves the lowest pick cost in the top-ranked picking mode, the next lowest pick cost in the next highest-ranked picking mode, and so on. We believe it will be valuable to study multi-tier inventory systems where the cost of picking and restocking a sku in a given storage mode is a function of the sku. This will allow us to model inventory systems where, for example, some skus cannot be picked from certain picking modes.

In this research we have also considered how to design good slottings for inventory systems that are reslotted periodically. There is often a cost incurred each time a sku is moved from one storage mode to another in a reslotting, and so any procedure to find a good slotting for such inventory systems must consider sku demand patterns over more than one time period. Solving such problems to optimality is computationally intensive and impractical for a normal sized warehouse. For this reason, we developed heuristics to find good slottings of a forward-reserve inventory system with respect to total picking, restocking, and move costs over multiple periods. We tested the performance of all heuristics on data for multiple sales periods from the Avon warehouse and compared the performance of each heuristic with an upper bound on the total net pick savings of the optimal slotting of this warehouse. All the heuristics are easily codable on software that is commonly available in most warehouses. In future work, tighter bounds will be determined and heuristics will be developed to find good multi-period slottings for inventory systems with multiple forward modes.

REFERENCES

- [1] BARTHOLDI, J. J. and HACKMAN, S. T., “Warehouse science.” Not yet published.
- [2] BARTHOLDI, III, J. J. and HACKMAN, S. T., “Simple ranking properties for the multi-mode forward-reserve problem.” Working paper.
- [3] EISENSTEIN, D. D., *Self-Organizing Production Systems*. PhD thesis, Georgia Institute of Technology, August 1992.
- [4] FRAZELLE, E. H., *World-Class Warehousing*. Atlanta: Logistics Resources International, Inc., 1996.
- [5] FRAZELLE, E., HACKMAN, S., PASSY, U., and PLATZMAN, L., “The forward-reserve problem,” in *Optimization in Industry 2* (T.A. CIRIANI, R. L., ed.), pp. 43–61, John Wiley & Sons Ltd., 1994.
- [6] HACKMAN, S. T. and PLATZMAN, L. K., “Near-optimal solution of generalized resource allocation problems with large capacities,” *Operations Research*, vol. 38, pp. 902–910, September-October 1990.
- [7] HACKMAN, S. T. and ROSENBLATT, M. J., “Allocating items to an automated storage and retrieval system,” *IIE Transactions*, vol. 22, pp. 7–14, March 1988.
- [8] SADIQ, M., LANDERS, T. L., and TAYLOR, G. D., “An assignment algorithm for dynamic picking systems,” *IIE Transactions*, vol. 28, pp. 607–616, 1996.