Impact of Inventory Storage and Retrieval Schemes on Productivity

by

Charlene (Charlie) A. Lieu

B.S. Computer Science (Engineering Track), Texas A&M University (1998)

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

Master of Business Administration
and
Master of Science in Electrical Engineering

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

Signature of Author ______________________________
MIT Sloan School of Management
Department of Electrical Engineering and Computer Science
May 06, 2005

Certified by ______________________________
Alvin W. Drake, Thesis Advisor, Professor Emeritus
Department of Electrical Engineering and Computer Science

Certified by ______________________________
Charles H. Fine, Thesis Advisor, Chrysler Leaders for Manufacturing Professor
MIT Sloan School of Management

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

Accepted by ______________________________
Arthur C. Smith, Chairman, Department Committee on Graduate Theses
Department of Electrical Engineering and Computer Science

THIS PAGE INTENTIONALLY LEFT BLANK

Impact of Inventory Storage and Retrieval Schemes on Productivity

by
Charlene (Charlie) A. Lieu

Submitted to the Sloan School of Management
and the Department of Electrical Engineering and Computer Science
on May 06, 2005 in Partial Fulfillment of the Requirements for the Degrees of

Master of Business Administration
and
Master of Science in Engineering

ABSTRACT

The operational management of high volume, multi-line distribution warehouses is a monumental undertaking, which only a handful of companies in the world have chosen to tackle. Amazon.com is amongst the few, and has further differentiated itself because of its direct to customer method of distribution and complex order mixes. There is no other retailer that carries and directly delivers as many different products (over 4 million different unique items) in as wide range of product categories (from music to cosmetics to electronics to garden hoses) in as high of volume as Amazon.com.

The nature of Amazon's retail model and its organic growth over the past decade has made its fulfillment centers a complex beast to decipher. Decisions on the fulfillment center floor are composed of intricate balances between demand constraints, equipment bottlenecks, storage limitations and labor costs, making the true cost associated with each variable dependent on every other variable.

The goal of this thesis is to document a practical exploration of inventory storage and retrieval schemes and its relationships to productivity (and subsequently cost), as well as identify implementable changes that yields higher throughput, lower lead time for order fulfillment, and ultimately dollar savings.

Of particular interest are ***operationally transparent process changes***, which improve processes in a manner that minimize impact on the fulfillment center floor. This concept will be the central theme of all recommendations resulting from this thesis.

Thesis Advisors
Alvin W. Drake
Professor Emeritus, Department of Electrical Engineering and Computer Science

Charles H. Fine
Chrysler Leaders for Manufacturing Professor, MIT Sloan School of Management

ACKNOWLEDGEMENT

To Chad Nusbaum: thank you for your remarkable guidance and support over the years as I struggle to find my way in the world. I am forever grateful for the countless pearls of wisdom, as well as your unwavering friendship and faith that drive me endlessly forward.

To Al Drake, my engineering advisor: thank you so much for your attention, time tested insights, grandfatherly advice, sanity checks and academic support both as my thesis advisor and as my department curriculum advisor.

To Charlie Fine, my management thesis advisor: thank you for finding time in your busy schedule to support me and for offering your wealth of experience.

To Tim Collins: thanks for your supervision in Reno and for supporting me in my exploration of a wide range of opportunities that made my internship such an amazingly rich experience.

To Jonathan Shakes: thanks so much for taking me on in Seattle so late in the game, and for your support, time, guidance, and conversation.

To the Pick 2PT - thanks for welcoming me into your group, for making me feel like part of the team, and for making the completion of my projects possible.

To the shift 4 pick team and the night shift managers in Reno: thanks for your warmth, acceptance, encouragement and support during my tenor as area manager. My time with you guys was one of the greatest experiences I've had in my professional career.

To David Jackson: thanks for providing valuable data on Lexington and the manual fulfillment centers and for being an unbiased sounding board.

My most heart felt thanks to Amazonians Sarah Marsh, Tracie Murphy, and Elizabeth Cassidy: without whom I would have never survived. I am grateful to each of you for sharing your experience, knowledge, wisdom, patience, and most of all, for keeping me sane.

To Susan: no matter what life brings, you are the one constant. Though fate made us sisters, it is love that made us friends.

To Vivi and Tim: your love, faith and encouragement from around the world have been constant sources of light, inspiration, and motivation.

And finally, to everyone who weathered the past two years with me: my sincerest thanks for your friendship and support – the splendors of life has been made sweeter by your presence.

TABLE OF CONTENTS

LIST OF FIGURES

CHAPTER 1: EXECUTIVE OVERVIEW

1.1 Introduction

Fulfillment center (FC) design is an exercise in balancing tradeoffs. The design of physical inventory storage, inventory processes, as well as the implementation of computational algorithms that drives fulfillment could significantly impact the overall operational cost.

The thesis work presented is the culmination of six months of work at Amazon.com and comprised of two related projects in inventory storage and retrieval. This thesis will be an attempt to aggregate analysis of various alternatives and look at tradeoffs that specifically impact fulfillment productivity. The result will be an integration of data analysis and computer modeling that demonstrate potential wins and losses of various physical designs and algorithmic schemes.

1.2 Project Goals

There are two distinct but related topics this work will address:

1. Investigation of alternative operational schemes – a significant portion of Amazon.com's inventory is currently stored in a random fashion. The outbound fulfillment directed by handheld wireless scanner devices, which are driven by in-house proprietary software algorithms. Hence, the first goal of this thesis work is to investigate alternative schemes that could potentially take advantage of the nature of the products or the storage facilities in a way that is operationally transparent (i.e. without changes to the standard operating procedures). The work will be a practical look at extracting more out of the fulfillment process with proposed changes that could be easily implemented without impacting production.

2. Physical layout analysis – physical design of the fulfillment center has historically not had much emphasis. However, the cost of build, storage and labor could vary significantly depending on the design of the FC, particular shelving layout. The

second goal of this thesis work is an investigation into the effects of physical layouts on product fulfillment cost. The aim is to deliver a simple tradeoff analysis tool that leverages this understanding for future designs.

1.3 Minimizing Operational Impact

An overarching theme in this work is the idea of improving processes in a manner that does not have dramatic operational impact. Of particular interest is to find solutions that fulfill the principle of ***operationally transparency.*** Because major operational changes can often translate to costly procedural implementations, the main criteria for alternative inventory approaches under consideration are:

1. Low implementation cost
2. Low production impact
3. No long term maintenance costs
4. No SOP changes and retraining required on the operational floor

1.4 Outcome Summary

Project 1 resulted in an agreement to implement software changes to reflect the "3 to 5" inventory storage scheme. Initial data analysis suggests that by implementing 3 to 5 picking in the random stow areas, we can achieve some small, but immediate wins. The implementation is slated to be completed in 2005. The details of this project can be found in Chapter 3.

Project 2 resulted in a trade off analysis tool delivered with a number of linked backend models. Initial analysis showed that the newest facility in Japan should be redesigned to follow the "niche" shelving methodology. However, further analysis using the trade off tool showed that although superior according to the model, there were not enough potential benefits to implement given the lack of absolute sensitivity in the model. The details of this project can be found in Chapter 4.

CHAPTER 2: BACKGROUND AND CONTEXT

2.1 State of Online Retail

According to Jupiter research, in 2004, consumers spent $66 billion in online retail purchases. Online retail is one of the fastest growing market segments, and with its increasing prevalence, even bricks and mortar companies are starting to get into the game. For companies like Toys R Us and Target, the solution is to partner with an existing online retail and fulfillment operation like Amazon.com. For others like Walmart, going online is a proposition the company is tackling in house.

In either case, the fulfillment operation of an online retail business differs dramatically from historic fulfillment paradigm where producers deliver to regional distribution center, which then serves distinct geographical areas. The replacement of a physical storefront with an electronic one means that warehouses can store nearly everything in any location and fulfill across the country or around the globe as long as FedEx (or your favorite parcel delivery service) delivers to the address.

2.2 Corporate Overview

As the "world's largest everything store" Amazon.com is leading the charge on online retail and fulfillment. The company boasts an impressive network of international warehouses, a comprehensive sales and customer information site, and net sales of $6.9 billion in 2004[1]. When adjusted for retail sale and compared against total online retail sales, Amazon.com sales contributed to just short 10% of total online retail sales, an impressive number by any measure. Although Amazon.com enjoys a significant lead over its nearest competitor, such as Buy.com, the competition is stiffening. More and more direct competitors are entering the market, making the pressure to maintain competitive advantage ever greater. Fortunately, the company's operational discipline along with its strong software development legacy makes Amazon.com a formidable giant in the online retail industry.

[1] Amazon.com, 2004 Annual Report

Unlike many dot com's of the late 1990's, Amazon.com has survived the collapse of the internet bubble with the company's relentless drive forward. The company is particularly well known for its commitment to customer satisfaction and continuously develops new features to make its website more useful for its customers. However, the company leads the online retail pack for a number of other reasons. For example, one of the company's key competitive advantages is its strong software-driven corporate focus. CEO Jeff Bezos have often been heard boasting that "Amazon.com is a software company, we just happen to be in retail." This *software-centricity* afforded Amazon.com the ability to offers millions of discrete items and to house electronic storefronts to thousands of independent merchants as well as "brick and mortar" retail giants such as Target and Old Navy.

Furthermore, as a low margin operation, the company needed to maintain an edge by focusing on operational excellence to deliver fast, low cost, high throughput, high mix order fulfillment essential for sales, growth and continued customer satisfaction. To this end, the company has committed to "continuous improvement" by establishing a six sigma program. It has also recently began looking at the future of fulfillment by establishing forward-looking cross-discipline teams dedicated to exploring alternative paradigms for fulfillment.

Amongst the various topics under investigation is the physical design of FC and how that impacts the in- and out-bound efficiencies within the warehouses. Additionally, investigation of alternative fulfillment algorithms could potentially yield significant returns in efficiencies.

CHAPTER 3: PARADIGM SHIFT IN INVENTORY SCHEME

To truly understand the context and opportunity in storage and retrieval, it is important that we look at the current theories in academia as well as the approach Amazo.com takes. Hence, the first two sections of this chapter will provide background for the sections to follow.

3.1 Current Academic Thoughts in Inventory Storage

Although there are many approaches to inventory storage and management within fulfillment and distribution centers, academic researchers generally group them into two basic philosophies[2]:

1. *fixed storage* – assigns a specific product to a fixed location. This method allows you to know exactly where your product should be, but can result in empty or wasted storage space.

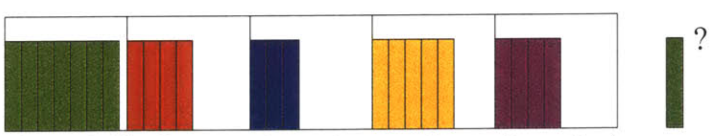

Figure 1: fixed storage has set spaces for each product

2. *random storage* – allows products to be stored in any empty storage location in which the product fits. This method yields higher space utilization, but is difficult to visually deduce whether product is out of place or missing without accessing a central index (most likely a database).

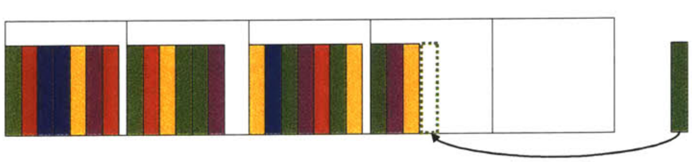

Figure 2: store anything anywhere with random storage

The table below summarizes the major differences between fixed and random storage:

[2] http://www.inventoryops.com/

FIXED LOCATION	RANDOM LOCATION
Location encodes identity	Identity independent of location
Specific location for each item	Any item can be placed in any location at any time
Misplacement is bad, but detectable	Difficult to know if an item is misplaced
Free space limited by identity restrictions	Liberated use of free space
Well Ordered	No Order, location enforced by identification tags
Simple physical count	Very difficult for physical count
Ability to stratify and organize products	Very difficult to organize by product characteristics

Table 1: fixed versus random storage[3]

Current research suggests that random storage is more efficient than fixed location storage in a warehouse with high inbound to outbound packet size ratio and high product mix[4]. The concept is similar to computer memory, where large packets of data are written onto random access memory (RAM), and smaller packets are read from the RAM for faster I/O. Amazon's warehouses are similar to RAM. Inbound have large number of items received in the form of cases or pallets, which have quantities of hundreds or thousands per "bundle." Outbound have small number of items shipped in the form of a order, averaging 2 to 3 items per "bundle." This quality, together with Amazon's focus on space utilization makes the company a good place for random storage.

However, there has not been much published research on hybrid approaches that are intermediates between random and fixed storage. One such method is indexed storage, in which products are not always in a fixed location, but are also not randomly stowed. This storage method stratifies products and storage locations and directs stows based on predetermined characteristics such as size or product types.

[3] Ho, Stephen (MIT). Auto-ID Center Lecture: "Investigating Intentially Fragmented Media Search"
[4] Saenz, Norman (Carter & Burgess). IDII Whitepaper: "Don't Waste Your Space!!!"

3.2 Amazon.com – Operational Process Overview

3.2.1 Inventory Storage

Currently, Amazon.com employs a mix of random and indexed storage. Indexed storage is called directed stow at Amazon and is used for pickable pallet or case locations. The majority of inventory are located in random stow bins, and there are no rules that enforces a storage scheme in these bins, only the way in which items are stowed[5].

The only guidance related to stow location are "best practice" guidelines such as:

- "larger books on the top shelf" (due to limited size of bins)
- "heavier items at the bottom shelf" (for safety)
- "no storage of similar items in immediately above, below, to the right or left" (for error minimization).

Inventory can be allocated anywhere there is room, hence living up to the name of "random stow". One could imagine under this practice, the randomness of products in the bin (herein referred to as *bin entropy*) is very high.

3.2.2 Pick Assignment

According to the software team in charge of picking, the "pick assignment," or linkage of physical inventory to a customer order, follows a number of complex rules. Most noteworthy is the "use to exhaust" rule, which dictates that when an identical item is located across multiple locations, the location closest to being empty is the location from which the item will be picked from.

As we will see in later sections, although practical in terms of creating continuous open storage space making stowing easier, this inventory linkage scheme does not necessarily constitutes the most efficient labor usage in terms of inventory

[5] SOP#096 for random stow, April 2004

extraction, nor does it constitute a significant departure from a purely random inventory storage scheme.

3.2.3 Fulfillment

After a particular item in the warehouse is linked to an order, it is fulfilled through a process called "picking" where someone manually go to the location specified and extract the inventory. The pick process begins with a batch of items in need of extraction. A path generation algorithm then creates a path that is downloaded to a hand held computer scanner, which directs a human being (known as a picker) to follow the path to find the inventory in question. The picker will then place the inventory into a bin (or "pick tote") for routing to the shipping department. Often, the pick totes are carried on roller carts (called "pick carts") from location to location.

3.3 Project Summary

There are two primary parts to project 1:

1. an academic exploration of inventory storage schemes and its relationships to picking and stowing productivities
2. recommendation for alternative retrieval that will yield higher throughput, lower lead time for order fulfillment, and ultimately dollar savings.

3.3.1 Opportunity

Consider the following:

- Last year, the warehouse in Reno (RNO1) alone spent approximately $4M in labor for picking. On average, picking is anywhere between 10 – 15% of the total fulfillment labor expenditure at Amazon.com[6].

	Q1	Q2	Q3	Q4	FY2003
Picking Hrs	55,069	46,463	49,207	120,912	271,651
Total Picking Cost	$925K	$781K	$827K	$2,031K	$4,564K

[6] From 2003 financial report, and given total cost of $16.80 per labor hour.

- Currently at RNO1, 70 – 75% of the products are located in some 227,000+ random stow bins, and anywhere between 60 – 85% of the total yearly picking labor is expended in picking out of random stow bins.

	Q1	Q2	Q3	Q4	FY2003
Estimated % time spent in random stow	85%	85%	85%	60%	
Picking cost in random stow	$786K	$663K	$703K	$1,219K	$3,371K

This means that improvements applicable specifically to the random stow scheme could potentially have non-negligible dollar impact.

One of Amazon.com's claims to fame is how fast inventory move through its warehouse floors, meaning the bins "decay" (ie inventory depletes from the location) and replenish on a regular basis. Inventory completely decays and replenishes approximately18-19 times per year. This roughly translates to a complete inventory turnover once a month, with inventory moving much faster at year end due to the seasonal spike caused by holiday shopping. The speed at which products move through the warehouse is clearly something that could leverage to extract additional productivity, hence alternatives will be explored to see how the product scheme could be optimized.

3.3.2 Project Deliverables

1. Overview of alternative inventory storage strategy
2. Analysis of benefits and impact
3. Recommendation for implementation

3.3.3 Project Approach

1. Understand current operational paradigm of "random" storage and retrieval
2. Explore alternative schemes
3. Investigate possibility of implementing alternative scheme

3.4 Survey of Alternative Paradigms

There are endless variations of alternatives for inventory storage and retrieval paradigms. For the purpose of this discussion, they will be selected and divided broadly into two categories: "natural decay" schemes and "non-decay" schemes.

Natural decay schemes are inventory paradigms that depend on the "decay" (or fulfillment extraction) of the inventory in various locations to create virtual zones with faster inventory movement. This class of solutions changed the product storage distribution in a transparent way. Namely, no labor is expended to categorize inventory and direct its storage, and instead areas naturally stratify by virtue of product movement itself. These virtual zones are bounded by decay rules set forth in the inventory extraction software, and ultimately contribute to higher overall productivity.

Not to limit our boundary of solutions, natural decay methods are only one class of ideas being investigated. There are other storage and picking schemes that may also yield similar benefits without affecting product storage distribution. These "non-decay" schemes may work in lieu of or in conjunction with the natural decay algorithms.

3.4.1 Natural Decay Schemes

Essentially, we can leverage the computer driven nature of the Amazon fulfillment and change picking in such a way that we extract additional benefits from the fast inventory movement, namely by using the decay of products to change the configuration of our product storage that constantly evolve to match product sales velocity. Because a random storage warehouse have multiples of the same item stored in multiple locations (upwards of thousands of distinct locations), preferential assignments of inventory based on location could create virtual zones.

To do so, pick assignments (i.e. association of customer order to physical inventory) could be redesigned such that priority for inventory is given based on some preset criteria. Hence, items in a preferred location would be extracted first. The result would be heavier picking in computationally defined virtual hot zones without losing the appearance of "random" stow.

Some examples of preferential schemes which would create faster moving zones are as follows:

- ***3 to 5 focus*** – 3 to 5 is a concept similar to the "strike zone" idea at McMaster-Carr, which refers to shelves between eye and knee level. These locations have been determined to be most conducive to picking safety, but we postulate that there are also labor savings because these locations are easier to extract inventory from. "3 to 5" shelves, as the name suggests, are shelves between 3 feet to 5 feet.

  These shelves have the highest picking efficiency because this zone allows picker to pick without having to tip toe or bend down to reach and see items. Under this scheme, inventory which resides within this 3 to 5 zone will be favored, creating faster inventory movement in those locations as shown in the figure below. This scheme may yield shorter pick time in addition to safety gains.

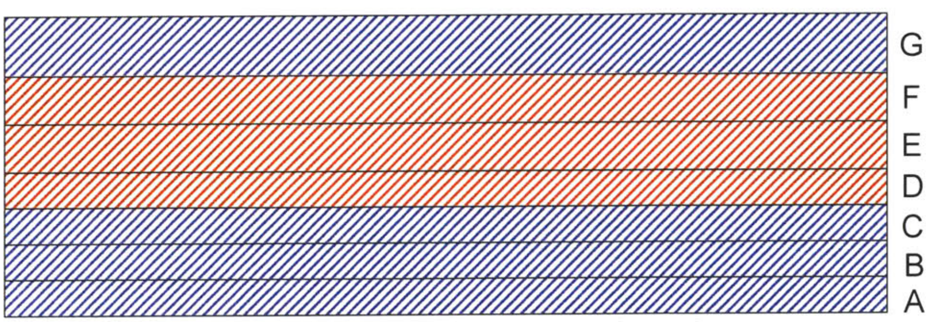

Figure 3: shelf face zone shown in red will have faster inventory movement

  Currently, the storage of items in random stow are distributed such that fastest moving items are evenly spread in the shelving space. Hence, if you are looking a storage shelf (or "bin face") you would expect to see fast and slow moving inventory interspersed fairly evenly as in the left of Figure 4 below.

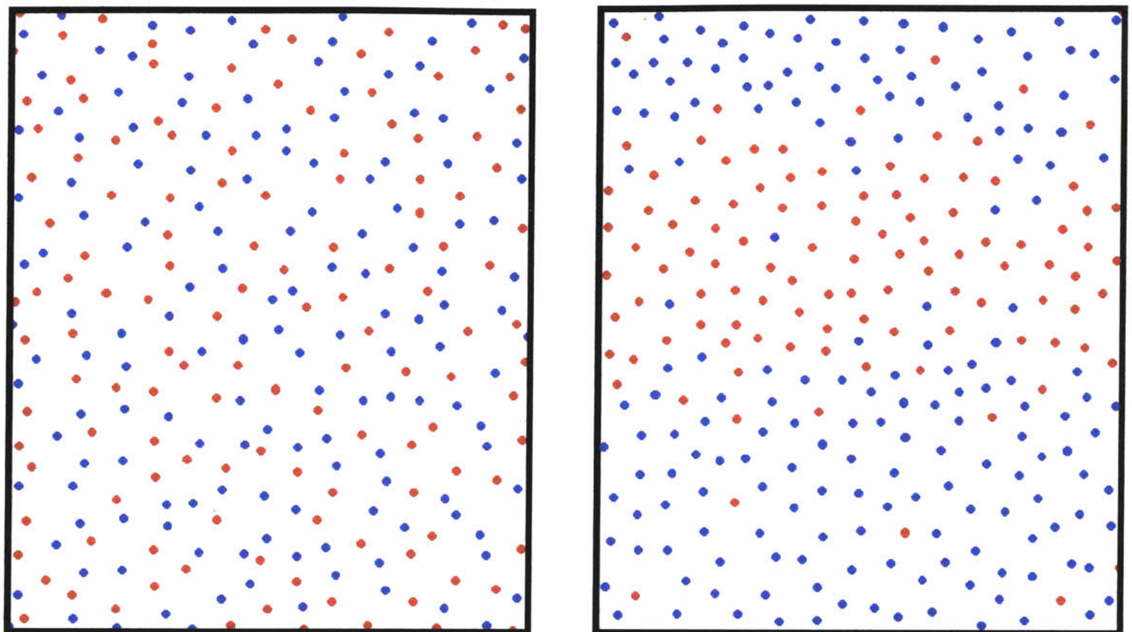

Figure 4: fast moving items in random scheme (left) vs "3 to 5" scheme (right)

Over time, with inventory extraction favoring the 3 to 5 shelves, one would expect the movement to speed up in the 3 to 5 feet range and the slower moving items to stay in the exterior of 3 to 5 as shown in the right of Figure 4.

- ***Zoning*** – This concept gives extraction preference to inventory closest to the conveyor, by associating priority based on an item's distance to a conveyor. Proximity to conveyors are desirable because it means less travel distance to drop off pick totes. *Distance to conveyor* also happen to equate to *distance to pick totes*, as pick carts are generally "parked" in the major aisles next to the conveyors because there is no room to take the carts down minor aisles.

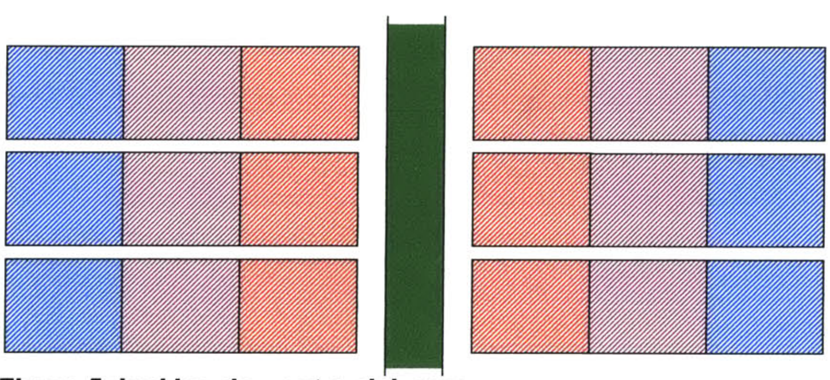

Figure 5: looking down at a pick area

This strategy will lower the picker travel times per item in the pallet and case picks because the paths will now be more tightly grouped around the closest location according to the time map. Furthermore, totes will enter the conveyor at a more optimal location, lowering tote travel time.

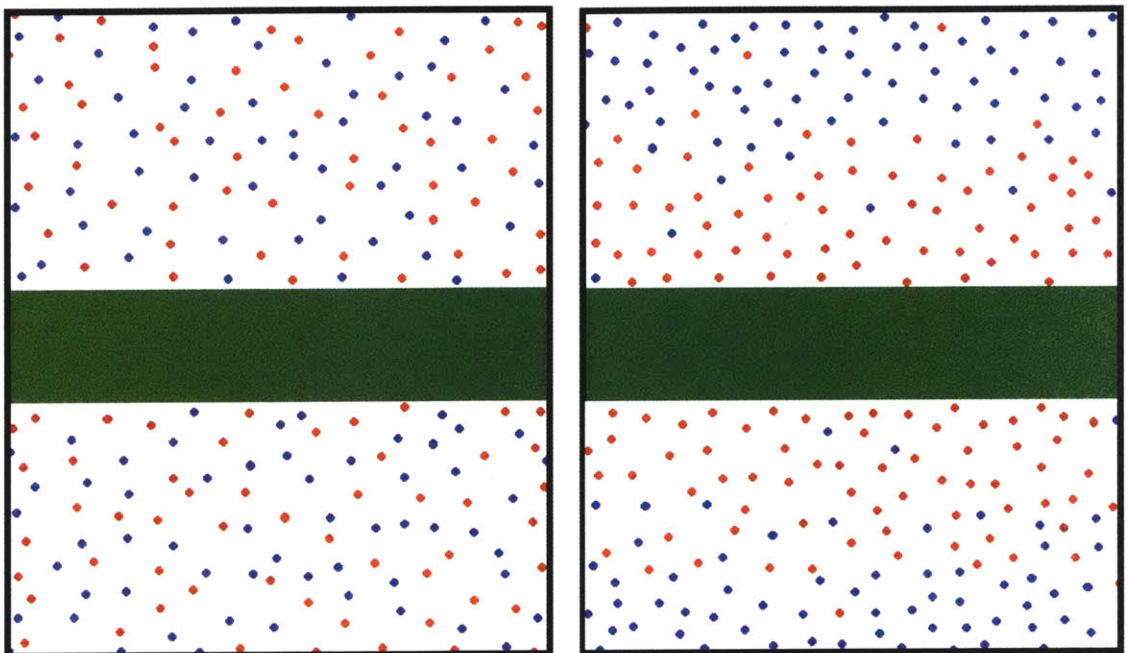

Figure 6: fast moving items in random scheme (left) vs "zoned" scheme (right)

Similar to the random stow bins in the 3 to 5 range, one could imagine that at instantiation, ***FC entropy***, or the randomness of products on the warehouse floor, is very high because the floor is completely randomized. However, over time, there will be a propensity for high velocity products to be stowed nearer to conveyor than slow moving products (as shown above).

The beauty of "natural decay" is that there's no additional cost incurred as compared to most process optimizations. The computer already directs people to locations, hence any change is transparent to the operational floor, and retraining is not necessary.

3.4.2 Non-Natural-Decay Solutions

There are a number of proposals currently aimed to extract additional fulfillment productivity that do not depend on the decay of the inventory, a few of these ideas are as follows:

- ***Random directed stow*** – ASINS with the highest sales velocity would be placed in "prime" locations closest to the conveyor. This scheme is similar to the zoned picking, but the implementation will focus on stow. The resulting storage locations will yield shorter travel time as the items with highest sales velocity will be extracted most often, and thus by placing it close to the conveyor, the travel time between the item location and the conveyor is reduced for that item over its fulfillment lifetime. With lower tote travel and picker travel time, the cycle time for the picking operation will shrink and ultimately, lead times from customer order to shipment will also become shorter.

  Unfortunately, this solution is suboptimal for a number of reasons. First, demand forecasting as a whole is at best only partially accurate, hence projection of sales velocity is rarely dependable. Furthermore, product placed into a prime location would occupy that location until it is fully exhausted, which means that when that product slows in its sale, it would still occupy the same favored location unless labor is expended to move it to a "less prime" location further from the conveyor.

- ***Directional picking*** – after observing and participating in the outbound processes over a two week period, there's strong evidence that the current pick path is vertical in orientation (as shown in the left of the figure below). The orientation of the pick path was later confirmed by a database dump of the current baseline pick path routing.

  This makes for more bending and reaching, and probably adds to search time and potential repetitive injury hazards. Something to consider is to change the orientation of pick direction to consider horizontal movement. This will save picker time in terms of not having to bend down or reach up (during which the picker could not do anything else). However, the horizontal picking will add to the travel time (during which a picker could theoretically parallel process and be looking for a product).

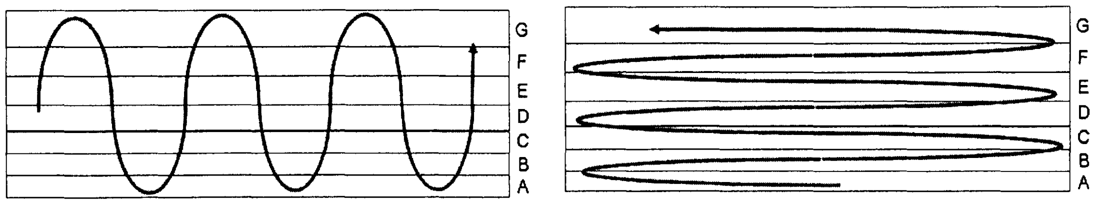

Figure 7: possible pick orientations – vertical or horizontal – that is the question

- ***Fast pick area*** – this solution would take advantage of the fastest moving ASINs (Amazon Specific Identification Number, the company's own version of SKU) by creating a "fast pick" area, where pallets of "best sellers" will have assigned space in a prime location near the receiving dock and conveyors. This scheme would have tangible benefits for the fasted moving inventory, but the cost trade offs are not completely clear. Obviously there will be additional cost incurred in building and maintaining such an area, but there may also be wider systemic effects such as the dilution of the fastest moving products from distributing throughout the floor, making the pick densities thinner in areas outside of "fast pick". Additionally, Sarah Marsh, an HBS intern at the Reno site, produced analysis of the fastest moving ASINS shows that there will be high turnover (nearly 40% per week amongst the top 500 ASINS) of products within the fast pick area, which will translate to high maintenance cost and potential errors due to the movement of inventory. That analysis is discussed in more detail in the last chapter of this thesis.

3.5 In Depth Analysis of the "3 to 5" Approach

After some initial discussions with various project stakeholders, it was collectively decided that the 3 to 5 concept warrant further exploration – primarily because the concept of "3 to 5" is simple, as its implementation. The idea is to create algorithmic rules that will create "virtual zones" to which faster moving inventory will naturally gravitate.

Before we talk about how this might happen, we will examine the current state, so to create a baseline from which we derive assumptions and to which we can compare results.

3.5.1 Baseline State

One might speculate that given the relative ease of stowing in the 3 to 5 zone, that there will already be some existing "favoring" in that range (ie stowers will already stow to the 3 to 5 locations more than other locations). However, early indicators points to no such favoritism. Spot interviews conducted on the production floor indicated that people rather stow to the entire bin face than to walk further with their inventory carts.

If the warehouse is consistently near full capacity, then inventory will have to be placed anywhere space is available, and the initial sense is that the warehouse is indeed always near capacity, hence the distribution of inventory is fairly even through the warehouse. However, to test the even-ness of current pick distribution, a test was performed to see the percentage of products that currently comes out of the 3 to 5 zone. To this end, I wrote a script that parses the production data daily for four weeks in order to obtain a percentage of random stow products currently coming out of 3 to 5. A snapshot of the first two weeks is displayed below.

day	toal captured picks	total random picks	total 3 to 5 random picks	% random stow from total	%3 to 5 from random stow	note
7/28/2004	388	85	44	21.91%	51.76%	test case
8/2/2004	38,605	29,145	14,094	75.50%	48.36%	
8/3/2004	59,312	43,576	21,402	73.47%	49.11%	
8/4/2004	46,218	33,461	16,042	72.40%	47.94%	
8/5/2004	53,553	39,747	19,056	74.22%	47.94%	
8/6/2004	43,254	33,108	15,886	76.54%	47.98%	
8/7/2004	40,619	29,526	13,969	72.69%	47.31%	
8/8/2004	37,454	28,028	13,377	74.83%	47.73%	
8/9/2004	48,556	34,451	16,155	70.95%	46.89%	
8/10/2004	59,655	44,566	20,704	74.71%	46.46%	
8/11/2004	79,694	56,879	26,893	71.37%	47.28%	
8/12/2004	86,663	61,837	29,671	71.35%	47.98%	
8/13/2004	50,958	36,119	16,898	70.88%	46.78%	
8/14/2004	31,993	24,148	11,454	75.48%	47.43%	
8/15/2004	40,066	30,663	14,566	76.53%	47.50%	
Total	**716,988**	**525,339**	**250,211**	**73.27%**	**47.63%**	

Monday
Tuesday
Wednesday
Thursday
Friday
Saturday
Sunday

Table 2: percent of random stow products extracted from 3 to 5 during production

Looking over the entire 30 day period, the average percentage of inventory coming out of 3 to 5 turned out be 47.21%, which is a little less than half of picks. This is roughly what one would expect given that 3 to 5 shelves constitutes a little less than half of the total bins. Hence, we could conclude that there is currently no favoring of 3 to 5, and a new inventory binding rule that favors those shelves would push the 3 to 5 percentage up.

3.5.2 Implementation Details

The algorithm we wish to implement is as follows:

When multiple items exist in random stow bins, we would assign the item in 3 to 5, and then progressively out ward.

"3 to 5" comprised of the following set of bins from locations below:

- Shelves D, E, and F for random stow library locations
- Shelves B, C, D for random stow library deeps
- Shelves C, D, E for flow racks.

When similar items are located across multiple 3 to 5 locations, the "use to exhaust" rule would then apply, meaning the bin with the least amount of existing inventory will be the location from which to pick from.

Pick association of this sort will force faster product movement in the bin, resulting in a shorter ***bin half life***[7] in the 3 to 5 shelves by natural decay. There will be a propensity for fast moving products to gravitate to a 3 to 5 location and slow moving products to decay into the outer regions. The beauty of this scheme is that the changes are transparent to the floor associates.

Given the product depletion speed given earlier (full replen 18-19 times a year), we assume that the bin half life is shorter than one month, which theoretically, means

[7] True to standard definition, half life refers to the time it takes for something (in this context, bins) to undergo a particular process (in this context, product depletion or decay).

that within as little as three weeks we can begin to see the formation of virtual 3 to 5 zones.

A point of consideration in terms of implementation is that to fully leverage the new stow zones, it may be important to reevaluate the pick path algorithms to eliminate potential new inefficiencies and extract new benefits. For example, dynamic path planning, which is discussed in the last chapter of this thesis, may be a complementary strategy that will help extract the most out of a new storage scheme.

3.5.3 Concerns

There is concern that the density of picks for the same distance traveled will become diluted because focusing on 3 to 5 will extend the area (and hence distance) that a picker will need to cover in order get the same amount of material. However, if one considers the picks globally and think about how the warehouse is always near capacity, then there really isn't room for inventory to dilute itself. The pick path may actually become tighter because over a bin face, the density of materials will remain the same, just that the picks will concentrate around the 3 to 5 shelves. If anything, the pick path should shorten. However, the definitive proof requires additional experimentation or testing via limited implementation.

3.5.4 Benefits

By favoring the "3 to 5" bins, the movement of inventory will favor the middle of the bin face. This will have three major benefits:

1. more picks will come from bins between 3 feet to 5 feet range, and each of those picks will on average yield 2 to 4 seconds faster pick times
2. furthermore, more picks mean more space will open up in the 3 to 5 zone, which benefits stowing operations
3. slower moving products will naturally decay out to the extremities because if a slow moving item is shelved in both zones, the ones in 3 to 5 will be extracted first, leaving additional copies in the slower moving zone. Hence the outer zones will tend to have slower moving items, creating natural stratification of fast movers in the fast moving zones

Although this scheme will benefit both pickers and stowers in terms of time and safety, the benefit calculation will focus around picker time. From initial rough estimates a picker could spend anywhere from 3 – 10 seconds more if the item being picked is out of the 3 to 5 zone.

Given current random distribution, 50% of picks will be outside of the 3 to 5 zone (the baseline analysis above seems to suggest more). Assume under the new scheme, only 25% of picks will be outside of the zone, and conservatively assuming the low end average of 3 additional seconds per item picked outside of the zone, the estimated time savings can be calculated as follows:

A slightly below average picker picks 150 items an hour in random stow[8], so 1200 items a day. Of those 600 are outside of the zone currently. With the new scheme only 300 out of the 3 to 5 in the new scheme, which means we have 300 additional "outside of zone" picks right now.

(300 items /day) * (3 sec / item) * (5 days / week) * (52 weeks / year) =
65 hours saved per picker each year.

65 labor hours * $16.80 / labor hour = $1092 per picker per year

Additionally, working in "3 to 5" would likely have safety benefits that are difficult to fully quantify. According Mike Strakal, a physical therapist and president of The Center for Physical Therapy and Occupational Rehabilitation

Researchers at San Francisco State University found that reducing the arm extension for can minimize muscle tension as well as stress in neck of workers[9].

---

[8] This is a very conservative estimate as pickers often pick in the 200's in random stow bins.
[9] Peper, Erik, et. al. "Repetitive Strain Injury- Assessment and Training Protocol." Electromyography: The Biofeedback Foundation of Europe, 1997 (http://www.bfe.org/protocol/pro09eng.htm)

Companies such as Barber Foods (a food producer), Creation Windows (a manufacturer of windows[10]) 3M, and Goldkist (the second largest chicken processor in the U.S.) have all implemented programs that redesigned processes to minimize repetitive stress injuries, and have saved the respective companies as much as 40% in repetitive stress injury claims[11].

[10] Sheley, Elizabeth. "Preventing Repetitive Motion Injuries." HR Magazine, October 1995.
[11] Grossman, Robert. "Make Ergonomics Go." HR Magazine, April 2000.

CHAPTER 4: PHYSICAL STORAGE DESIGN

4.1 Project Summary

The design of physical inventory space impacts two major metrics important to a fulfillment operation: productivity and space utilization. However, there is an inverse relationship between labor productivity and space utilization because the more densely packed a specific location is, the more difficult it is for stowers to find space to put additional product and for pickers to identify the correct item for extraction. Very roughly, the relationship between productivity and space utilization looks something like the following:

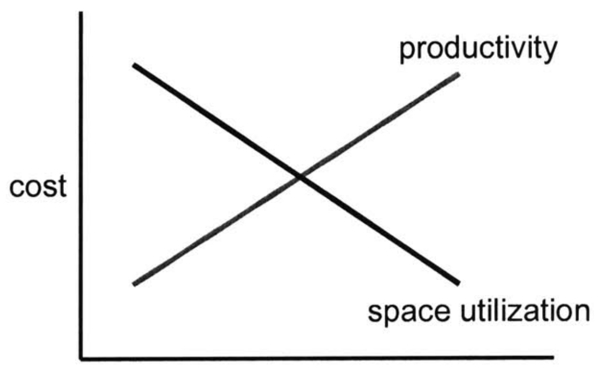

Figure 8: the relationship between productivity and space utilization

The goal of the project is to understand the optimal design paradigm which will yield the best space utilization and highest pick productivity.

More broadly, in the grand scheme of warehouse design, the final configuration of various characteristics would significantly impact the overall cost of fulfillment. The work presented here is an attempt to aggregate analysis of various tradeoffs that specifically impact productivity and space utilization.

4.1.1 Opportunity

Although there are many factors that contribute in different ways to productivity, as outlined in sections to follow, this chapter will look closely at some major physical design elements of a warehouse – namely the arrangement of shelving.

The inspiration came from the fact that construction of a new warehouse in Glasgow, Scotland (GLA1) was recently completed with a newly formulated "pick niche." This new design is more expensive to construct both in terms of material and labor, but the design was expected to yield gains in product density (more inventory per ft^2) as well as faster picking when compared with the traditional floor layout.

Each design results in trade-offs between three major cost components:
Capital Investment + Storage Cost + Labor Cost

This chapter will explored how a variety of shelving design impact these cost components and will show data that suggests the niche design can have potential cost wins over traditional straight aisle design.

4.1.2 Project Deliverables

This project will have two major deliverables:

1. Models of general shelf configurations for the floor
2. Trade off analysis

The result is wrapped up in the form of a simple excel tool. The tool shows the tradeoffs between selected probable designs, and leveraged data from computer modeling in order to determine the potential wins and losses of each design.

4.2 Impact of Physical Storage on Productivity and Product Density

There are a number of factors that impacts both productivity and density:

- Shelving layout – niche vs. straight line designs directly impact how many shelves could fit into a set amount of square footage. Furthermore, the design would directly impact the pick path generated and the distance necessary to travel to fulfill an order batch.
- Aisle length – the length of the aisle impacts how many shelves could fit into a warehouse area. Also impacts pick paths, hence productivity.

- Aisle breaks – creating breaks in long shelving would allow intra-aisle travel but also take away shelves that would otherwise house products.

There are some additional factors that impacts productivity specifically:

- Batch size / pick density – the size of pick batch should impact the pick density and therefore, change the travel path.
- FC size – the size of the area from which product will be fulfilled from would change the travel path; the smaller the area, the shorter the path traveled to cover the area.
- Carrying capacity – the volume of items a picker could carry could be increased through shoulder baskets, which mitigates the need to return to the pick cart, increasing productivity.
- Lighting – the higher the product density and the closer the shelving are to each other, the better the lighting has to be in order to properly locate the correct item. In other words, the worse the lighting, the longer it will take to fulfill an item.

A variety of preliminary analyses was performed through computational modeling on four of the above factors: shelving layout, FC size, aisle depth and pick density. Whenever possible, the analyses were conducted with representative production data. However, as some factors can not be modeled in production, pure in-silico modeling (or computer simulation) was employed.

These analyses provided a sense of how some of the factors could impact design and three of the factors were then built into a computer simulation that provided the backend data for the trade off analysis tool.

4.2.1 Effects of Fulfillment Center Size

Using production data from Amazon's Philadelphia warehouse (PHL1), the effects of varying warehouse size (FC size) was simulated by creating filters that selected picks that transition only in the areas in question.

For each size (measured in number of total bins) in question, the warehouse was divided into equal size regions and performance results from all regions averaged to capture the different product types and locations.

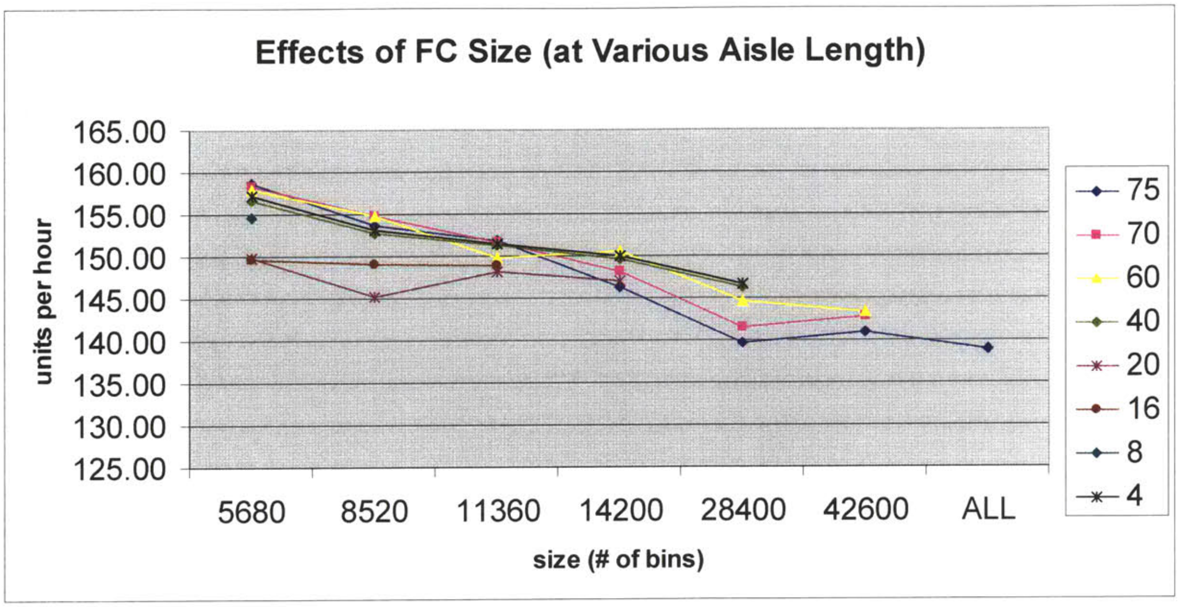

Figure 9: The effects of fulfillment center size on productivity

Generally speaking, the smaller the FC, the higher the units per hour (UPH) because the small the area you have to cover, the shorter the path you travel.
In practice, one could simulate the effects of a smaller FC by creating pick areas in which the entire batch of orders will be picked from, as in the case of RNO1.

4.2.2 Effects of Aisle Length

Similarly, using production data from PHL1, the effects of various aisle lengths were simulated under different warehouse (FC) sizes. We started with the full data from the non-pallet and non-case locations of the fulfillment center. This gave us a total aisle length of 75 feet, and gradually dividing the full FC into areas of varying aisle lengths, which we (once again) averaged.

The result of this analysis is represented in the graph below:

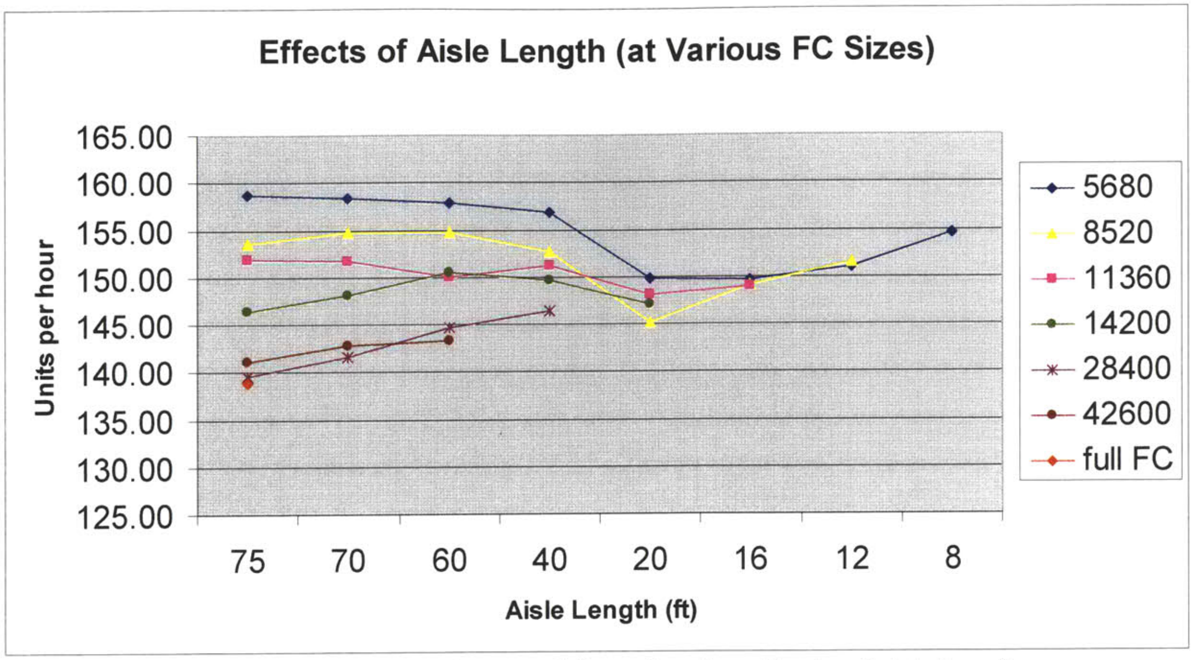

Figure 10: results from production data modeling showing effects of aisle length

As a follow up, a computer simulation of a theoretical fulfillment center with no conveyors, no aisle breaks, and no separate pick areas was built to model the effect of aisle length.

This modeling was done completely "in silico," which meant all data was generated by the computer model given the parameters of a set configuration. This model was approached in a way that the total size of the fulfillment center (x) is conserved, but the aisle length and the number of rows differ in a relationship that met the following: *aisle length * total number of rows = x.*

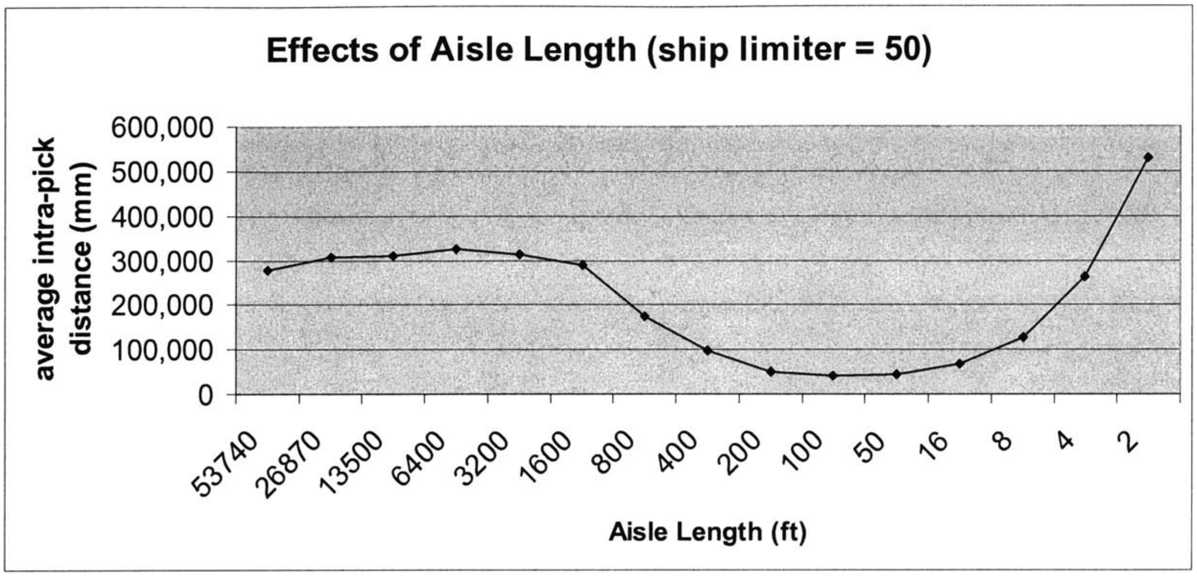

Figure 11: results from computer modeling showing effects of aisle length

The results matched very closely to the pattern seen in earlier analysis using production data from PHL1. In this instance, the "ship limiter" variable is a proxy for pick density, and "intra-pick" distance (distance between two consecutive picks) is a proxy for productivity (represented by units per hour in the previous analysis). The "ship limiter" is set to 50 because that is the batch sizing at PHL1.

The result of this model served two purposes:

1. to show the undiluted effects of the aisle length on pick distance, which is our proxy for pick efficiency
2. to compare the relevance of a pure in-silico computational simulation to data obtained in reality

4.2.3 Effects of Pick Density

Armed with some assurance that a fully in-silico model more or less matched reality, another simulation was built for analyzing pick density. This was done because there was no way to perform this modeling with real data as we can not ask the warehouses to vary pick density in production.

Similar to the previous model, pick density is proxied by "ship limiter" (as the ship limiter or batch size increase, so does the pick density) and an "intra-pick" distance is calculated to proxy for productivity.

A number of other factors were varied to create some representative designs, for which models were built and analyzed. Specifically, this exercise compared the long aisles and a few reasonable niche designs under different pick density levels.

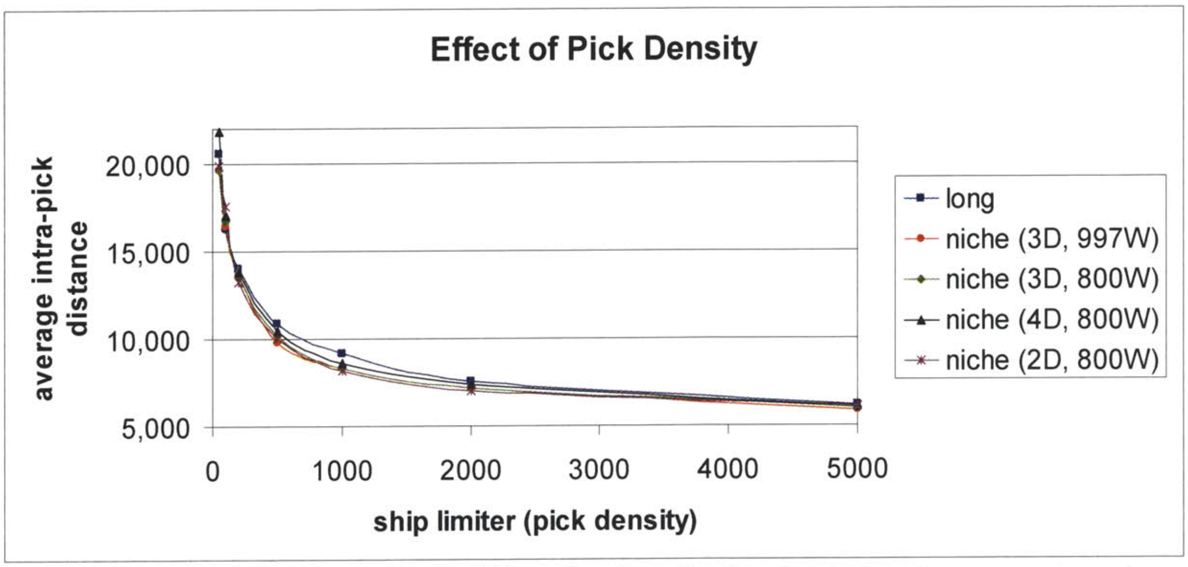

Figure 12: results from computer modeling showing effects of pick density

4.3 Shelving Design Experiment and Analysis

4.3.1 Shelving Design Approach

There are two major design concepts under consideration:

1) Straight aisle shelving: our current standard FC layout.

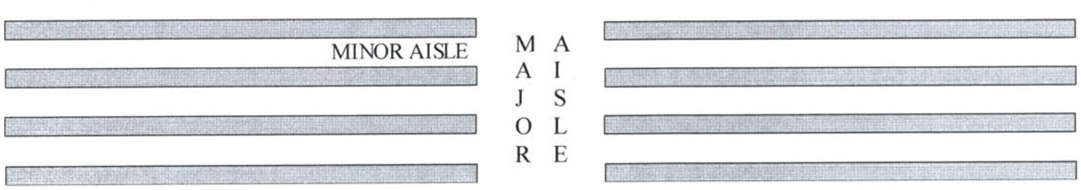

2) Niche shelving: previously postulated could yield higher pack densities and potentially higher pick rates.

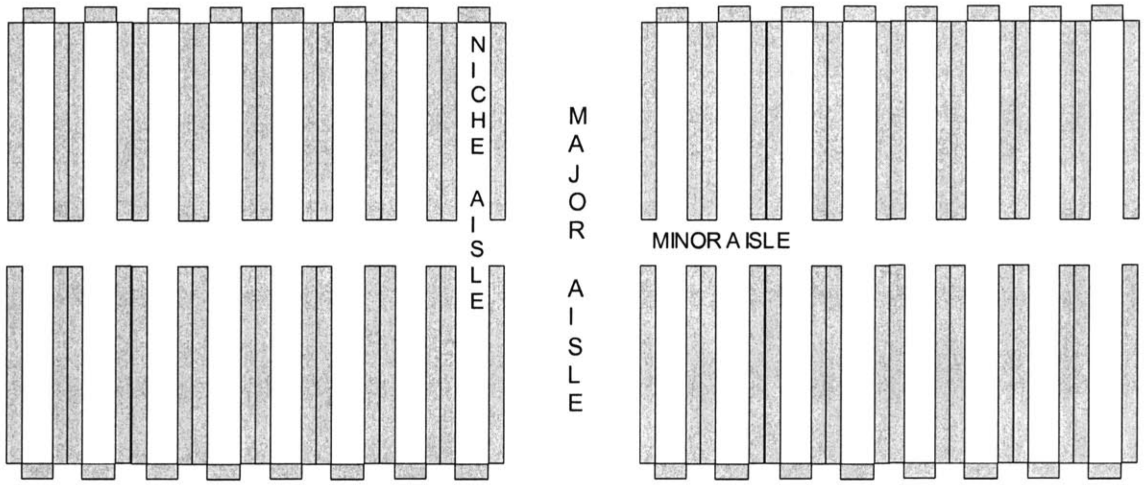

Design Constraints and Additional Considerations:

- For egress regulation reasons, the minimum width of aisles is set to 800 mm and niche aisle length maximum is set at 4 meters. Although, according to Kevin Scharetg, there are no rules governing minimum aisle width and depth.
- As aisles become narrower, it will become more difficult to light the product in a way that does not compromise the quality of fulfillment.

4.3.2 Project Approach

1. Design layout in roughly 60 sq meter space with center conveyor
2. Model the following configurations:
    - Straight line shelving, 1000 mm minor aisles, intra-aisle break every 7000 mm
    - 2000 mm (2 shelves) deep, 800 mm wide niche design
    - 3000 mm (3 shelves) deep, 800 mm wide niche
    - 3000 mm (3 shelves) deep, 994 mm wide niche
    - 4000 mm (4 shelves) deep, 800 mm wide niche

    To simplify the problem, I chose to ignore supporting structures such as posts as well as exits under the premise that those structures will more or less impact every design similarly.
3. Gather data on build cost:

Use existing quotes for each of the existing builds: PHL1 for long aisle shelving and GLA1 for niche shelving. Shelving costs were normalized to materials used in GLA1 because shelving there was chosen to support potential multi-level structures. New facilities built will all use these stronger shelves, and the PHL1 quotes did not reflect that pricing.

4. Understand pick labor cost in term of time to pick:
   a) Model travel time
   b) Derive average time at bin as well as relationship between travel distance and travel time from existing FC data
   c) Account for cart return behavior through FC experiment

4.3.3 Cost Component Analysis

As previously mentioned, there are three major components of cost under investigation in this analysis: storage, build, and labor. Each component will be examined in more detail below.

4.3.3.1 Storage

Cost of storage has two primary components:

- Yearly rent cost
- Maximum storage density

Maximum pack density is directly proportional to amount of shelving, so we can use number of shelves as a proxy for storage density.

4.3.3.2 Build Cost

Cost of building shelving has two primary components:

- Material: more costly for niche because additional material (bins sides, brackets, and braces) necessary to construct niches.
- Labor: more costly for niche because of the complexity of creating niches.

Niche bins require greater capital investment expenditure, but this is meant to be off set by improvements in pick productivity and thus a lower cost of pick labor per pick.

4.3.3.3 Pick Labor

Picking labor is the most significant contributor to the overall cost of a pick. It could be as much as 90% of the total cost of the pick.

Ship Limiter	Long	Niche.3D.997W	Niche.2D.800W	Niche.3D.800W	Niche.4D.800W
50	20,512	19,579	19,810	19,602	21,884
100	16,193	16,387	17,535	16,662	17,042
200	13,981	13,586	13,207	13,515	13,811
500	10,827	9,769	10,037	10,137	10,468
1000	9,183	8,283	8,141	8,324	8,608
2000	7,542	7,147	7,020	7,155	7,370
5000	6,118	5,857	6,053	5,998	6,089

Table 3: pick performance of each design under consideration (at different pick densities)

Cost of picking labor is directly proportional to pick time and can be obtained by multiplying time spent by the hourly labor rate.

Pick time is made up of three key parts:

- Overhead time – includes cart prep, start time, tote collection, etc. The largest overhead is cart returns, which increase with number of picks
- Time at bin -- largest proportion of total pick time per item (at high pick density)
- Travel distance / time -- largest portion of total pick time per item (at low pick density)

<u>Cart Return</u>

Since cart return consists of a non-negligible part of the total time expenditure, an experiment to model cart return behavior was conducted in collaboration with staff at PHL1. The idea is that if we can find some way to predict when pickers return to cart and how long they stay at the cart, then we can model the effect of cart return. After some brainstorming, it is postulated that a picker would return to cart when the distance to the next item is significantly further than the distance to cart, so it makes sense to return to cart before moving on to next item.

An experiment was designed so that the scanners were used to track when pickers returned to the tote cart to unload inventory as well as when pickers move the carts. Exhibit 1 shows the experiment summary and instructions given to collaborators at PHL1. The experiment ran for three days and began with 12 pickers, but ended with all of picking on night shift.

The data obtained by the scanner were written into log files identified by user names, and the content showed when the picker scan an item, when the picker return to cart, as well as when the picker move the cart, and after parsing, files which follows the following format were obtained:

```
Item1
[move cart]
item2
item3
item4
[cart return]
item5
item6
....
```

The number of *items in hand* was obtained by summing the number of items between each cart movement or cart return, and the *delta distance* at the time of cart return was calculated. The delta distance is simply the difference between the distance to the next item and the distance to cart.

For each unique *items in hand*, the *delta distances* were averaged. These numbers were then plotted against each other to yield the graph below.

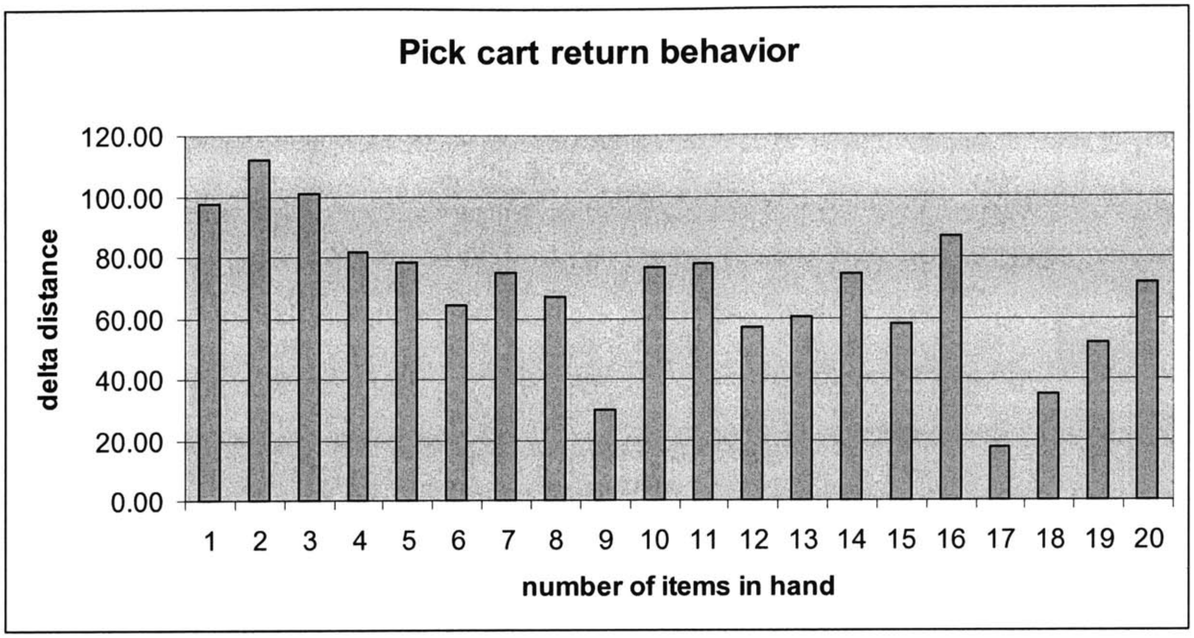

Figure 13: pick cart behavior

After looking at the information obtained, only *items in hand* under 20 included in the analysis. This is due to two reasons:

1. As we get higher in the number of items, there were progressively less replicates, which make the measurement inaccurate.
2. It is highly unlikely that a picker can hold many more than 20 *items in hand*, even with the help of a shoulder basket. It is more likely that any result showing more than 20 *items in hand* is the result of the picker forgetting to scan the cart when they returned to their cart, making the measurement inaccurate.

A pivot analysis of the standard deviation was performed, and looking at Figure 14, it seemed clear that given the high standard deviation that the data itself is highly volatile. This means that given the same number of items in hand, there are no clear indication of whether a picker would return to cart or not based on the delta of distance from next bin location and the distance to cart.

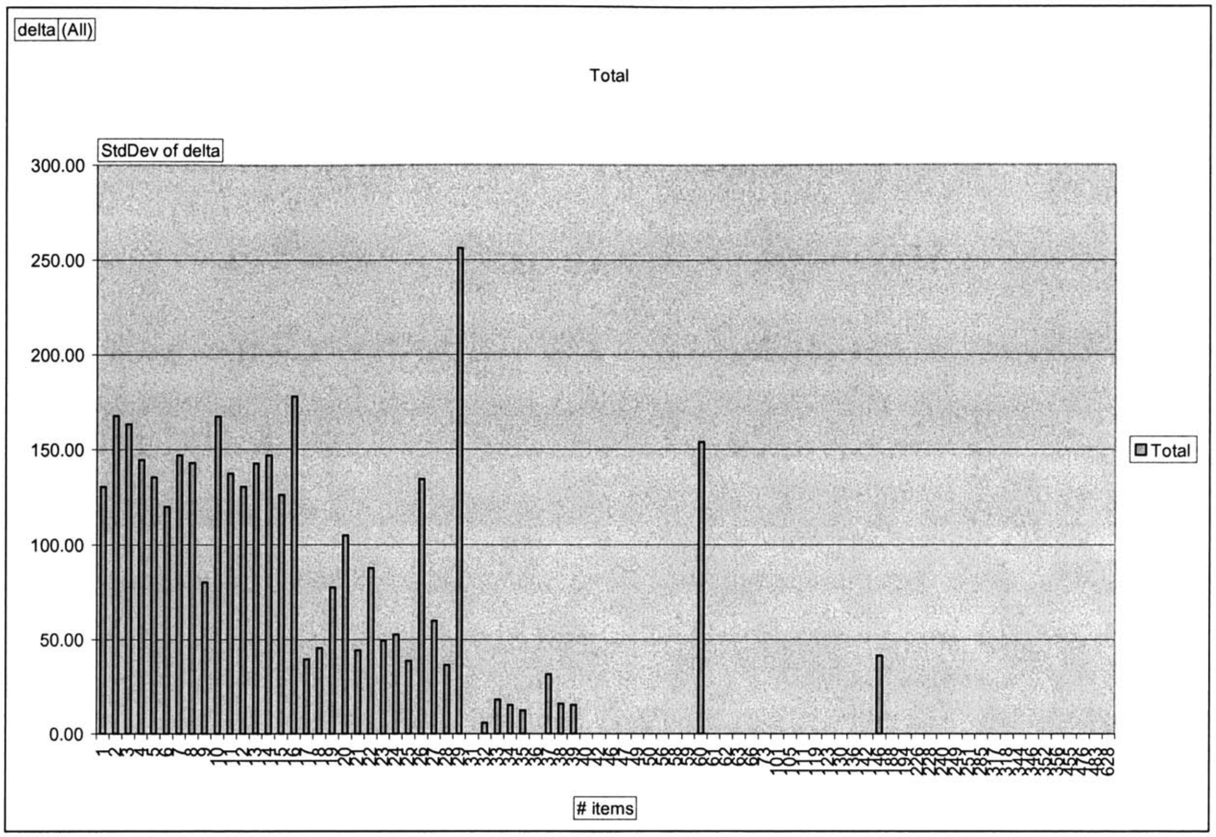

Figure 14: standard deviations for each *item in hand*

Upon further analysis, through floor observations and self experience, there are a number of other factors that would contribute to cart return:

- Is the direction of the pick path taking the picker away from the cart?
- Are there too many items currently in hand to be able to pick another item, necessitating a cart return?
- Do people even think strategically about when they should return to cart or is it something done on a whim?

Cart return, it seem is a much more complex phenomenon to model than first expected. Unfortunately, there does not appear to be a set of rules that can be derived to properly model cart return behavior. It may be possible to model the process given more time, but since cart return is only one component of a much larger analysis, a representative number was selected to represent cart return.

Disregarding *items in hand* numbers greater than 30, the average items picked between cart return is 7 items. Hence, the effects of cart return was accounted for by adding in cart return travel time for every 7 picks.

Travel Distance and Time

Taking data from warehouse, we can see a clear relationship between average time and distance between consecutive picks (herein referred to as "intra-pick distance") in the graphs below. This relationship between distance and time is important because it allows for the conversion of distances to time and ultimately, to labor cost.

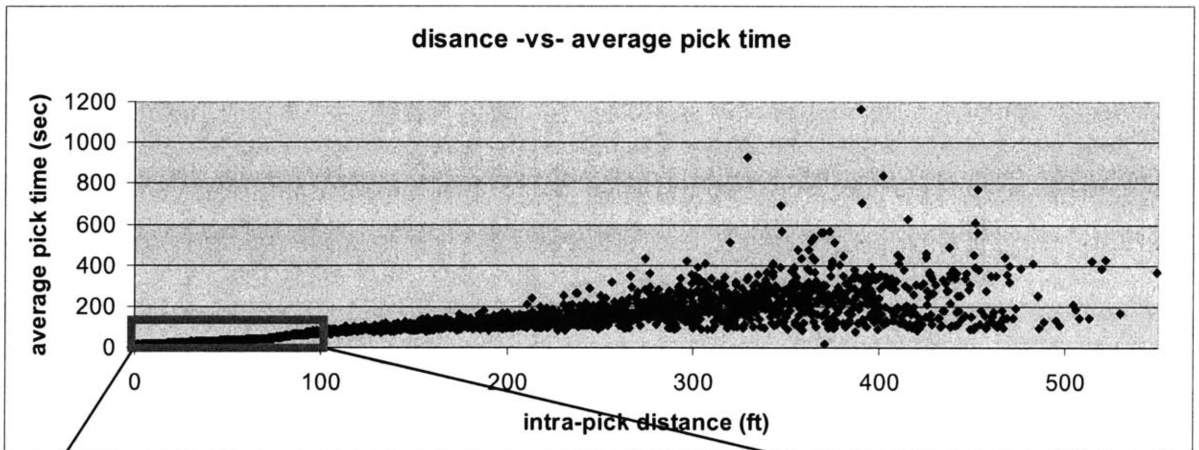

Figure 15: production data of distance traveled as it relates to time

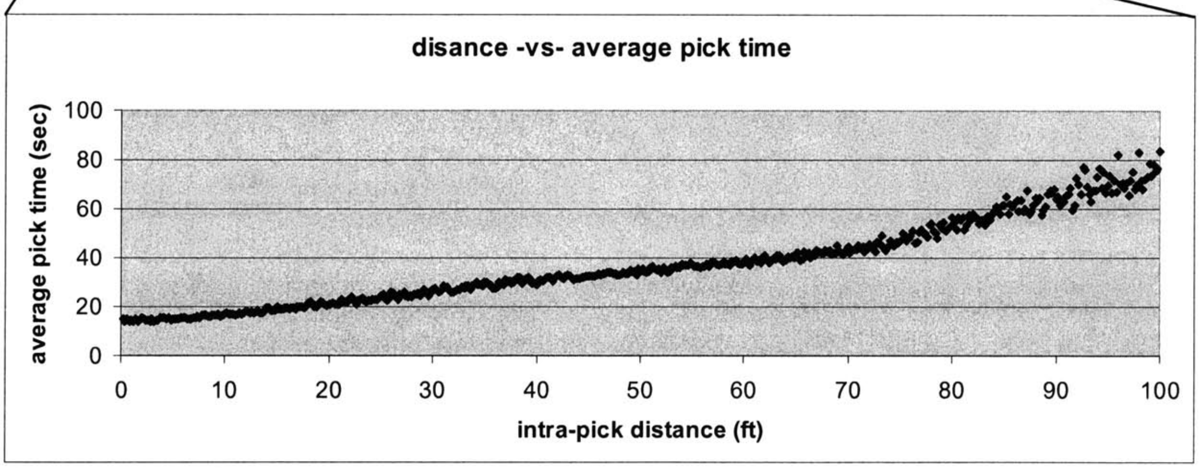

Figure 16: magnified section showing direct correlation between time and distance

Each datapoint represented in the above figures is an average of travel times for each intra-pick distance. Intra-pick distance is calculated as the physical distance between consecutive pick locations.

PHL1 as once again used because it is the warehouse that has the most linear layout scheme, which will give us the cleanest time to travel relationship. The data reflects a time period between mid-July to late-October, which was chosen to capturing some off peak and some peak conditions.

There are two major considerations when looking at this data. First, the spread becomes worse as "intra-pick distance" increases, because the longer the distance, the more variety and/or distractions can be introduced during travel. The number of replicate data points for each distance becomes smaller because pick paths tries to minimize pick distances, hence there are fewer number of picks, and there's an upward spike for time points after approximately 70 feet because pickers tend to not return to carts if the next pick is relatively close.

The relationship between pick distance in feet (x) and pick time (y) shown in the above figure yields the following equation: $y = 0.4x + 15$ seconds.

The y-intercept for this equation should be the *time at bin*, or the time spent picking an item if your travel distance is 0. The y-intercept is 15, which is consistent with previous calculations[12] produced by the pick team.

4.4 Tradeoff Analysis Tool

The data of the computer model experiments and trade off analysis were aggregated into an excel analysis tool that helps lay out the effects of the tradeoffs for a number of designs. The tool holds reference tables of raw modeling experiment data and calculates total cost per unit picked as a sum of storage cost, capital investment cost

[12] Janert, Philip, et. al. "Picking Cost Model." Amazon.com report, November 2003.

and cost per pick. Users can then select from a pull down menu of geographical locations for which there is current data for, and then enter numbers for a variety of variables that best represent the (potential) fulfillment center layout under consideration.

	LONG	NICHE-3D-997W	±%	NICHE-2D-800W	±%	NICHE-3D-800W	±%	NICHE-4D-800W	±%
area modeled (sq meters)	6,080	6,156		6,156		6,156		6,400	
shelves in area modeled	5,800	5,676		6,600		6,725		7,200	
units per square meter	343.42	331.93		385.96		393.27		405.00	
sq meter needed for total unit stored	2,912	3,013		2,591		2,543		2,469	
total rent	$ 382,301.32	$ 395,536.38		$ 340,161.29		$ 333,838.59		$ 324,173.64	
rent per unit	$ 0.03186	$ 0.03296	3.46%	$ 0.02835	-11.02%	$ 0.02782	-12.68%	$ 0.02701	-15.20%
cost of lighting	Assume to be the same for now, but want to allocate line item for future modification								
material and equipment cost (per shelf)	$ 133.20	$ 154.34		$ 162.80		$ 154.34		$ 149.64	
installation (labor) cost (per shelf)	$ 20.45	$ 33.30		$ 36.60		$ 33.30		$ 30.00	
total cost (per shelf)	$ 153.65	$ 187.64		$ 199.40		$ 187.64		$ 179.64	
total cost of shelving in area modeled	$ 891,153.11	$ 1,065,060.86		$ 1,316,040.00		$ 1,261,898.21		$ 1,293,440.00	
cost of shelving (per unit stored)	$ 0.00508	$ 0.00621	22.13%	$ 0.00659	29.78%	$ 0.00621	22.13%	$ 0.00594	16.92%
average travel per pick (in mm)	9,182.86	8,282.72		8,140.61		8,324.03		8,607.75	
average time per pick (in sec)	28.19	26.90		26.70		26.96		27.37	
average cost per pick	$ 0.09398	$ 0.08967	-4.59%	$ 0.08899	-5.31%	$ 0.08987	-4.38%	$ 0.09122	-2.93%
total cost per unit	$ **0.13092**	$ **0.12883**	-1.59%	$ **0.12393**	-5.34%	$ **0.12389**	-5.37%	$ **0.12418**	-5.15%
total cost per 10,000 unit	$ **1,309.19**	$ **1,288.35**		$ **1,239.28**		$ **1,238.91**		$ **1,241.80**	

facility location	**Japan**	please use drop menu to select values
units per shelf	Campbellsville	
lifetime of shelving in years	Coffeyville	
total units stored	Delaware Fernley	
inventory turns (per year)	Japan	
total units picked (per year)	Lexington	this value is automatically calculated, please do not adjust
shipment limiter (in units)	**1,000**	<-- please use drop menu to select values
rent per sq meter (per year)	$ 131.29	

Figure 17: snapshot of sample configuration in trade off analysis tool

4.5 Results

Using the tradeoff analysis tool, the cost of pick under a number of scenarios with different configurations and pick densities were examined for a number of countries. Geographically speaking, the location of particular interest is Japan, where a new warehouse is in the final process of design for build.

4.5.1 Numerical Comparisons

A table summarizing cost differences in Japan is shown below. Using the "long aisle" configuration as baseline, a percent change in cost for each configuration was also calculated at different pick densities.

SHIP LIMITER	LONG	NICHE-3D-997W	±%	NICHE-2D-800W	±%	NICHE-3D-800W	±%	NICHE-4D-800W	±%
50	$0.1852	$0.1817	-1.87%	$0.1798	-2.90%	$0.1779	-3.93%	$0.1878	1.40%
100	$0.1645	$0.1664	1.17%	$0.1689	2.69%	$0.1638	-0.41%	$0.1646	0.05%
200	$0.1539	$0.1530	-0.58%	$0.1482	-3.71%	$0.1488	-3.34%	$0.1491	-3.12%
500	$0.1388	$0.1347	-2.94%	$0.1330	-4.17%	$0.1326	-4.48%	$0.1331	-4.11%
1000	$0.1309	$0.1276	-2.53%	$0.1239	-5.34%	$0.1239	-5.37%	$0.1242	-5.15%
2000	$0.1231	$0.1222	-0.72%	$0.1186	-3.65%	$0.1183	-3.87%	$0.1183	-3.91%
5000	$0.1162	$0.1160	-0.22%	$0.1139	-1.99%	$0.1127	-3.00%	$0.1121	-3.55%

Table 4: summary of cost for each design (separated by batch size or "ship limiter")

In every configuration, having bigger batch sizes decreases the total cost.

Although batch size impacts only the labor cost component, that component has significant effect on the final cost because labor cost is the largest contributor of all the components. Looking back at numbers in Figure 17, we see that labor cost is approximately an order of magnitude greater than the cost of build, and labor is 3 to 4 times the storage cost.

Each design impacts the three cost components in different ways, and although some factors (ie aisle width or depth) might impact a particular cost component dramatically, it may not be a component that significantly contributes to the final cost. Additionally, a particular factor may not impact the final cost because it improves one component while negatively impacting another.

4.5.2 Conclusions

From the results, there appears to be a potential 5% in savings by going with one of the niche designs. However, the general consensus is that the analysis as it currently stands is not sensitive enough (nor the savings significant enough) to warrant converting the Japanese FC designs from long aisles to niches. The reason being that niche designs would add additional complexity not only to build, but to both the software and operational analysis.

It is also interesting to note that there may be potential hidden costs to "complicating" the shelving design. The additional complexity of the niche design will force the software to consider one extra level of "dimension" and this may translate to some additional non-negligible software development engineer and/or operational excellence black belt time each time the software has to be updated or each time efficiencies are in need of analysis.

Increase to capital expenditure will be absolute if we go to niche, but the returns in pick efficiencies can not be guaranteed. In the case of Glasgow (GLA1), where niches have been built, the facility does not appear (at this time) to be extracting cost benefits. Part of the reason is that the shelf design was modified during build such that niche aisles were increased from 800 mm to 997 mm. My suspicion is that the change was done to simplify shelving component purchase (997 mm is the standard shelving size that's being used), but the change translated to a loss in the benefit of the niche design. According to the excel model, when batches are set between 50 to 200, which is generally the operational range, it is when the 994 mm wide niche design performs the poorest.

The conclusion is that although niches appear to be a good idea, it should probably not be implemented presently as we can not guarantee its cost returns. There are many other factors in the FC design that could currently be understood better and improved upon. Hence, it may be useful to invest some additional work to understanding the full effects of these design factors and create future fulfillment centers based on those results.

CHAPTER 5: THE FUTURE

5.1 General Observations

Globally speaking, minimization of warehouse entropy (randomness) is highly recommended at Amazon.com because the computational driven nature of warehousing could allow fine tuning of inventory schemes to extract additional productivity from the warehouses. By optimizing around prime pick locations, either around conveyors or "3 to 5", the fulfillment operation could retain the major benefits of random storage (ie high pack density) while deriving higher efficiencies.

5.2 Action Items and Future Projects

There are a number of items that would constitute future work either as action items or as potential projects that are logical extensions of the work presented in this document.

There is one key action item, namely, the implementation of 3 to 5 strategy. Additionally, there are four other potential projects:

- Total Cost Model for Directing Inventory
- Tuning of Physical Build Tradeoff Model
- Dynamic Path Planning
- Directed Stowing and Stocking

5.2.1 Implementation of 3 to 5 Inventory Strategy

After the analysis of various inventory strategies, it was agreed that 3 to 5 should be implemented on a small scale. The Pick 2PT had made commitment to do so in 2004, but schedule slip prevented the completion of this implementation. Hence, the lingering action item is to implement the scheme and perform analysis of

The following pieces of data should be collected:

- [BEFORE] and [AFTER] data on performance: namely units per hour (UPH). Note that it's important to collect this data on a weekly basis as the virtual zones will not form immediately, but rather, will appear over a period of time as bins decay.

- [AFTER] data on 3 to 5 fulfillment percentages: namely, the amount of inventory that is being fulfilled from 3 to 5. The [BEFORE] data had already been collected, which can be found in section 3.5.1 Baseline State.

5.2.2 Total Cost Model for Directing Inventory

The per item cost calculated for each physical space design scenario is a starting point for a number of other analyses. A similar approach could be applied to derive an extensive total cost model for fulfillment. Knowing how the storage location of an item could impact its ultimate fulfillment cost is an important understanding for directing inbound inventory.

For each product, there are a variety of factors that contributes to the total cost of fulfillment. One could imagine that there may be wide variation in final cost depending on where the inventory is placed and as well as the mix within the orders.

The ultimate goal would be to define a total cost model that includes the cost of stow labor, storage and pick labor. The model would give us a general picture of the variables that impact these components and give us estimates such that the sum of would give us a sense of our total fulfillment cost:

Total cost (all items in an order) = $\sum_{x=1}^{i} StowLabor_x + \sum_{y=1}^{j} Storage_y + \sum_{z=1}^{k} PickLabor_z$

Using this estimate we could direct storage of inventory in a way that begins to minimize the cost of fulfillment.

As we have the most extensive amount of data in picking, a starting point for a future project would focused around the third component of this equation – pick labor.

Data requirement:

- List of potential factors that may affect the pick productivity:
  - ASIN – the pick labor component is actually ASIN agnostic… this is needed only as a mean to look up product size and weight
  - Product size and weight – will affect storage location
  - Number of units – needed to calculate a total cost, but not needed as a cost calculation variable component
  - Total cubic volume – pick labor is agnostic to this; a constraint, not a variable
  - Bin size
  - Bin geographical location
  - Bin type
  - Bin height from floor
  - Bin physical density
  - Bin ASIN density
  - Forecast velocity
  - Inbound packet relationship
    - Probability of subsequent near-term receives of the same ASIN
    - Probability of NEXT receive of the same ASIN
    - Number of ASINS in other (pickable and reserve) locations
  - Outbound packet relationship
- Pick path and pick time data: extract from FC time map

Experimental approach:

- statistical analysis of all factors, eliminate low contribution / noise factors
- formalize key factors and formulate function for calculation cost (within acceptable tolerance).
- create algorithm to generate the picking cost to feed to larger cost program

Project deliverable:

- data supporting the cost of pick labor based on location data
- algorithm

5.2.3 Tuning of Physical Build Trade-off Model

The work presented in the trade off model described in the previous chapter is by no means comprehensive. It is meant as a starting point for understanding the effects of FC design given the way we currently operate.

For greater model sensitivity, there are a number of non-negligible factors one might want to include in the model, most of which was listed in section 4.2 Impact of Physical Storage on Productivity and Product Densityof the previous chapter.

There are also a number of financial and risk tradeoffs one might want to consider:

- Cost of capital expenditure – how current interest rate would impact your overall expenditure. For example, if interest rate is low and expected to stay low, then you might want to consider spending the money to get the productivity gains for the long run.
- Renting versus owning – currently, the cost model is calculated on rental cost, in $/sq ft. However, if the space is owned, the capital cost will be calculated based on footprint, and there could potentially be significant differences.

5.2.4 Towards Dynamic Path Planning

Another potential future project could center around redesign of pickpath algorithm. Currently, the pick path does not take advantage of dynamic path planning. The pick path is created in batches and does not dynamically take into consideration the last location of the picker when assigning new picks. Hence, the software could bring you from one area to anther via crossing of multiple conveyors and then bring you back to your original location.

Given the picker's last known pick, the software should optimize around the picker's expected location. Otherwise the efficiencies gained from a hot zone or a 3 to 5 scheme would be lost. Specifically, because of the new density around conveyors, it is now doubly important to manage distribution of pickers. This is where global dynamic re-optimization is vital. With each pick wave drop, the picker should get items in his/her row, without the introduction of new people in the same location.

5.2.5 Directed Stow Optimization

A natural extension of the work presented here would be to take a more theoretical look at global storage schemes across the entire FC and assess potential impact of various approaches on productivity. Members of the Picking, Spaceman and Revo teams believe looking at the global storage schemes could potentially be highly beneficial. Future projects should look at different ways of implementing changes for more intelligent directed stow.

In 2004, directed stow inventory is about 27% of Amazon.com's outbound volume. Inventory slated for directed storage is "profiled" (very roughly segmented by types) and placed by cubic volume in the first available space. The space maintenance team as well as the picking team would like to see a more detailed scheme for storage with the following factors taken into consideration:

- Space utilization
- Storage/ stowing throughput
- Picking throughput
- Cycle time for order fulfillment

Since it is difficult to balance all the factors listed above, it may make sense to implement some sort of artificial intelligence algorithm either in the form of a vector machine (SVM) or perhaps neural network such that the system will learn to optimize storage on its own. At the very least, the algorithm could extract profiles of product classes or correlate new product to existing product for which we know the sales velocity decay.

One of the ways to direct product is according to week to week demand by ASIN... placing fastest moving items in virtual "hot zones" for higher density and easier access.

For example, in the graph below[13], we can see that of the top 15 books and top 10 DVD's, most have some similarities in profiles. The goal is to investigate how we might leverage a computational method to recognize the profile and use this and other data to direct storage accordingly.

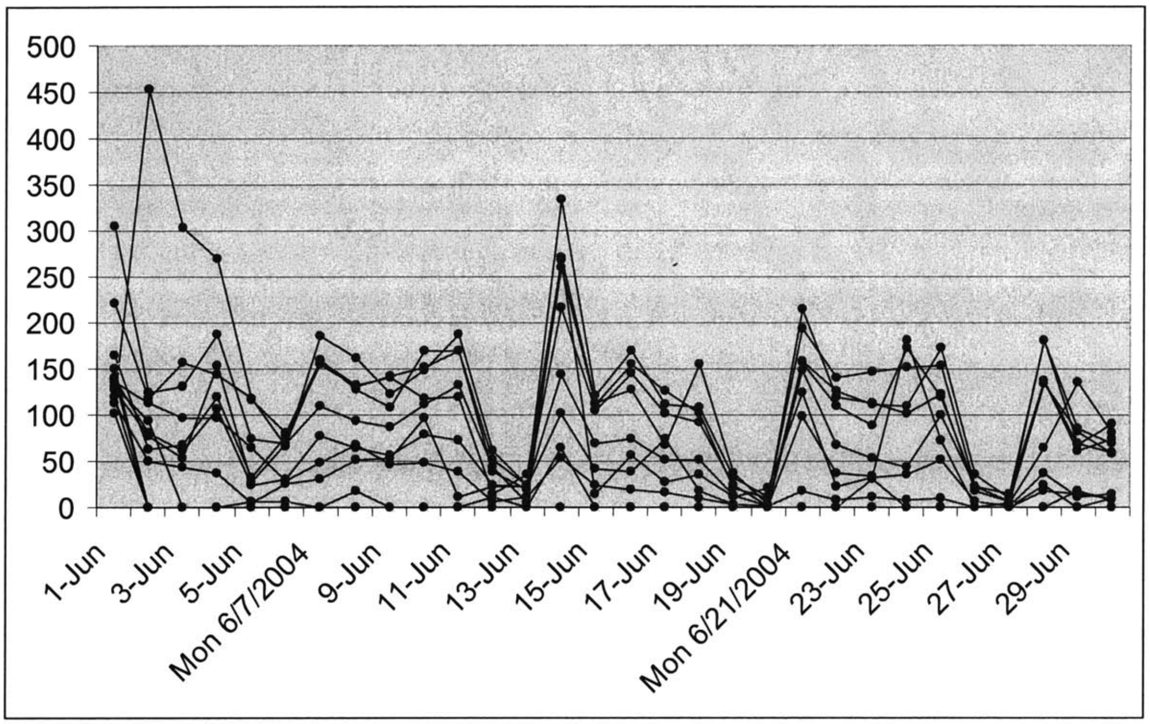

Figure 18: June profile of top 15 book titles

[13] Marsh, Sarah. "Fast Pick Area in RNO1." Amazon.com Report, July 2004.

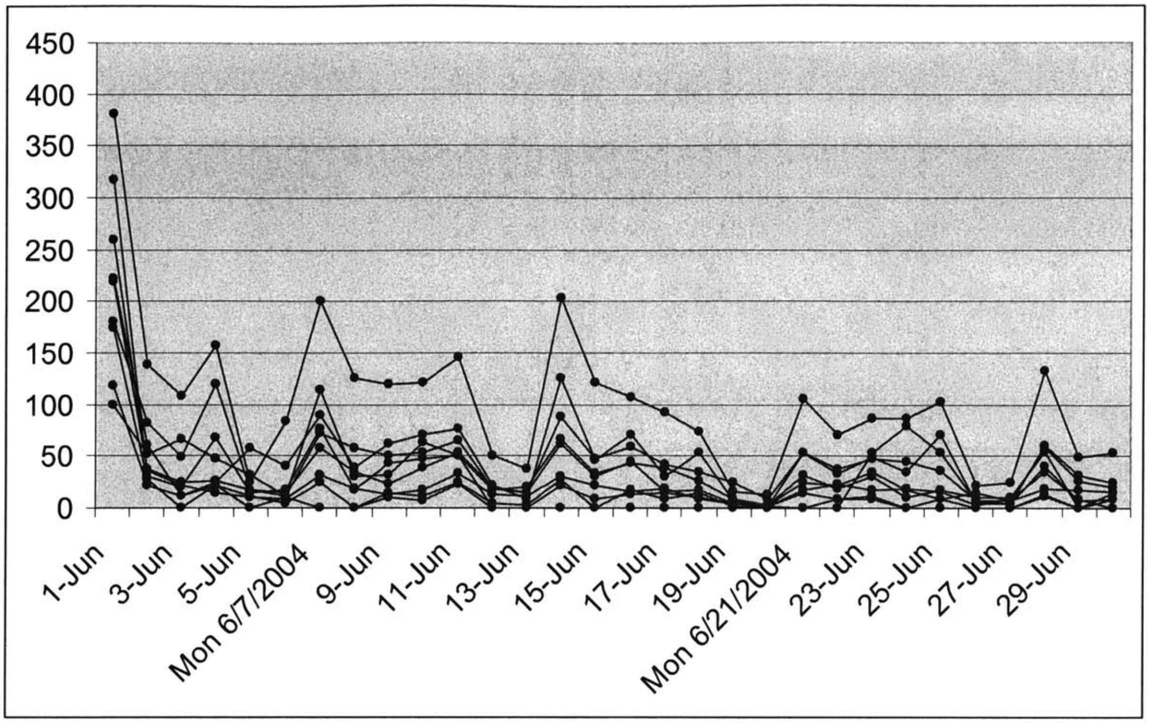

Figure 19: June profile of top 10 DVD titles

Similar to the random stow locations, the same natural decay idea could have applications in directed bins such as flow racks and pallet locations.

Future work should attempt to investigate the idea of applying a similar transparent approach to non-random pick locations as well as investigate the theoretical impact of schemes, some of which opposing, such as:

- ASIN scatter – the scatter of the highest demand products across the FC floor in an attempt to ensure there will always be a pick next to another pick during fulfillment
- Fast Pick – the consolidation of the highest demand products to a central location

Currently, products are already directed to zones based on how fast the product is expected to move. Research could go into the following area:

1. how many zones and product strata would be optimal
2. consolidation strategies
3. prediction algorithms for demand by ASIN.

As mentioned previously, one way to predict demand might be to have the system "watch" for patterns and learn through fuzzy logic. Based on the current expected incoming inventory and current available space, the system will assign items to the bins through some direct correlation of "sales velocity" and "distance from conveyor." Under a good prediction and stowing scheme, the optimized FC floor would have natural "fast pick" areas.

5.3 Closing Remarks

The work presented here only represents some smalls wins that could potentially yield savings, as well as some global understanding of tradeoffs. However, one can see from the small list of potential future projects that this is only the beginning. Amazon.com's fulfillment center, to quote Professor Al Drake of MIT's Electrical Engineering department, is like "a big operational playground full of toys and opportunity."

It is our hope that through this work, we would have set some foundations upon which we can build an array of projects that leads to a better, faster fulfillment center.

APPENDIX A

Exhibit 1

Experiment: data collection for "cart return behavior" modeling

Goal: To gather data to understand picker behavior in terms of "cart returns" (ie the decision between [going to the next pick] and [going back to the cart]).
Understanding picker behavior will--

1) Help us understand the trade off between upfront construction cost and long term operating labor cost, and design new FC's accordingly.
2) Generate potential low cost redesign of existing FC's to gain higher productivity.

Number of pickers required: 10 – 12

- Would like to get as many pickers as possible because the more pickers we get, the more accurate our data will be. More than 12 would be wonderful, but not necessary.
- Would like to get as representative sample as possible (ie really good pickers, average pickers and slower pickers). HOWEVER, it is more important to get pickers who can be depended on to scan their special barcode every time, even if that means the 12 people we get aren't completely representative of the general population.

Number of days: 3

- PHL1, please run experiment for one day only to begin with
- Charlie will process data to ensure integrity and determine if sample size is appropriate
- If data looks good, we will go ahead with day 2 & 3.

Impact: 40 – 50 scans per picker per day

- Expected time loss 5 - 8 minutes per picker per day.

<u>Instructions for Pickers</u>

1. Every time you return to your cart, scan special barcode "cartreturn"
2. You will get an error "invalid barcode" (hit return according to instruction on screen)
3. If you are prompted to scan tote after a pick, scan tote first, then special barcode
4. Continue picking