



MIT Open Access Articles

Stowage decisions in multi-zone storage systems

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation	Yuan, Rong et al. "Stowage Decisions in Multi-Zone Storage Systems." International Journal of Production Research (November 2017): 1–11 © 2017 Informa UK Limited, trading as Taylor & Francis Group
As Published	https://doi.org/10.1080/00207543.2017.1398428
Publisher	Taylor & Francis
Version	Author's final manuscript
Citable link	http://hdl.handle.net/1721.1/113206
Terms of Use	Creative Commons Attribution-Noncommercial-Share Alike
Detailed Terms	http://creativecommons.org/licenses/by-nc-sa/4.0/

Stowage Decisions in Multi-Zone Storage Systems

Rong Yuan[†] • Tolga Cezik[‡] • Stephen C. Graves[†]

[†]*A. P. Sloan School of Management, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139-4307, rongyuan@mit.edu, sgraves@mit.edu*

[‡]*: Amazon.com, Seattle, Washington 98109, cezikm@amazon.com*

June 2017, revised September 2017

ABSTRACT

The stowage decision determines how arriving products are distributed in a storage system or warehouse. In particular we consider the zone-stowage decision for large warehouses that are organized into distinct storage zones. An example would be a multi-floor warehouse where each floor is a storage zone. Each storage zone has limited picking capacity; we want to stow the product inventory across the storage zones so as to be able to meet uncertain demand requirements with the limited picking capacity in each zone. Determining how to spread the inventory across the storage zones is the zone-stowage decision that we consider in this paper. With a simulation study, we identify two zone-stowage policies that are effective at balancing the picking workload across different storage zones. The first zone-stowage policy achieves a chaining-inspired allocation by splitting the received quantity for each product across two storage zones; the second zone-stowage policy explicitly tracks the expected workload for each storage zone, termed the velocity of the zone, and then stows arriving products to the storage zone with the smallest velocity.

Key Words: Stowage Decision, Storage Systems, Flexibility, Warehousing Systems

1. Introduction

The intent of this research is to examine a key operational decision, namely the zone-stowage decision, as it arises in multi-zone storage systems. This stowage decision determines how arriving products are distributed across multiple storage zones. We use storage system to denote a facility that receives and holds inventory that is then used to fulfill orders, e.g., a warehouse or order fulfillment center. In a multi-zone storage system, the storage space is physically segmented into distinct, parallel storage zones. Each storage zone has its own dedicated storage area and operational resources. Product that is received by the storage system can be stowed in any zone; and the inventory in each storage zone can be picked to fulfill any order. Items from an order can be picked concurrently from different storage zones and then sent downstream to a sortation system at which the items are “assembled” into an order and prepped for shipping (Figure 1). This type of storage system is typical in large order-picking warehouses, such as the fulfillment centers for online retailers. These fulfillment centers can often cover one million square feet of floor space, spread over several floors, with each floor being one or more storage zones.

For the large fulfillment centers that we consider, products are received primarily in cases or eaches (i.e., individual units), with only high-volume items being received in pallets. However, the product is stowed into the storage system as eaches. The motivating context is a robotic mobile fulfillment system in which inventory is stored on mobile pods, and picking and stowing occurs at stationary stations on the boundary of the storage field. See D’Andrea and Wurman (2008), Enright and Wurman (2011), Lamballais et al. (2017), Zou et al. (2017) and Yuan (2017) for details on these semi-automated storage systems.




Figure 1: Material Flow from a Receive Station to the Multi-Zone Storage System

In these storage systems, the products first arrive at the receiving stations. One then needs to decide for each product, how many units to stow to each of the storage zones. Spreading the inventory of a product across multiple zones provides some level of operational flexibility in that there are multiple options for meeting a demand. However, there is an additional cost to do this. For instance, if product were received in a case, then at the receiving dock someone needs to open the case, separate the units into sub-batches, and then place each sub-batch into a tote that can then be conveyed to the different storage zones. Alternatively, if the product were to be stowed in a single zone, then the case could be directly conveyed there, without additional touches.

In light of these trade-offs, one simple policy is to stow all units of a product to a single, dedicated storage zone. The advantage of this policy is that it is simple, as all demand for an item must be picked from a single storage zone. However, the disadvantage of this policy is that it provides little flexibility for picking as all of the product's inventory is in one zone; if the picking workload assigned to the zone exceeds its picking capacity, then some picks will be delayed which can lead to orders missing their shipping deadlines. Another common stowage policy is the random stowage policy where the inventory of an item is stowed to a random storage zone upon receiving each replenishment. Whether this policy can provide any flexibility for picking depends on the batch size and the replenishment frequency of the item. If the inventory replenishment occurs infrequently in very large batches, then it is likely that most inventory is still stored in a single storage zone under the random stowage policy; however, if the replenishment occurs frequently with small batches, then a random stowage policy could result in spreading the inventory across multiple storage zones, creating flexibility. Other policies might split each replenishment batch into sub-batches and then stow each sub-batch in a different storage zone. Such policies will break up the inventory of a product across multiple zones, but will incur some overhead or additional effort to create and process the sub-batches. For example, to stow a case that contains 24 units of an item to two storage zones, one needs to open the case at the receiving station, separate the units into two sub-batches with (say) 12 units in each sub-batch, and then place each sub-batch onto a cart or tote, which can be sent to the designated storage zone.

The zone-stowage decision is important as it determines the inventory profile for each storage zone, which then determines what items can be picked from what zones. An unbalanced inventory

distribution across the storage zones may easily result in unbalanced picking workloads, especially when demand is volatile; that is, some storage zones have a greater workload relative to their capacity whereas others have a lesser relative workload. This situation can be problematic in storage systems for which there is limited ability to move picking capacity from one zone to another, e.g., when each zone has a fixed number of picking stations. When the workload exceeds the picking capacity for a zone, orders can miss their due dates and/or expediting expenses are incurred. When a zone has excessive picking capacity, then operating costs increase due to underutilized resources.

A key idea to avoid unbalanced inventory distribution is to create options or flexibility for the picking activity. That is, we create more options for picking by having the inventory of each item stowed in multiple storage zones; in effect, the item can be picked as long as one of the storage zones has adequate picking capacity. The concept of flexibility has been widely adopted in the manufacturing and service industries. In this paper, we explore how to accomplish flexibility in a multi-zone storage system, as used by online retailers for which demand is highly variable. In this instance, the major cost of creating flexibility for picking is the additional operating costs to divide a replenishment batch into sub-batches and then prepare each sub-batch to be conveyed to its storage zone.

The rest of the paper is organized as follows. We first review the related literature in section 2. We describe the simulation model and the eight stowage policies we consider in section 3. We show the numerical results in section 4 and finally conclude our work in section 5.

2. Literature Review

The literature of the operational decisions within warehouses or fulfillment centers generally falls into two categories, namely storage decisions and order-picking decisions.

The storage decisions determine where and how to store the inventory in the warehouse; the order-picking decisions determine how to efficiently pick items from the storage system to fulfill a set of customer orders. The major research topics related to the storage decision are the forward-reserve allocation problem, the zoning problem, and the storage location assignment problem. The

forward-reserve problem considers how to allocate items between the reserve area and the picking area within a warehouse; the zoning problem studies how to create multiple zones for picking and how to spread the inventory for each item over the multiple zones; the storage location assignment problem decides where to store the inventory within a storage zone.

The major research topics related to the order picking decisions comprise batching, routing, and sorting problems. The batching problem studies how to group the orders together for a single picking tour; the routing problem considers how to sequence the tasks in a picking tour to construct the most efficient picking route; the sorting problem determines how to efficiently sort and assemble the picked items into customer orders.

The focus of this paper is on a zoning problem as it arises as part of the storage decisions. We will limit our literature review to this topic area. We refer the reader to the review papers De Koster et al. (2007), Gu et al. (2007) and Staudt et al. (2015) for more extensive discussions of the overall literature on warehouse operational decisions.

Compared to other planning issues, the zoning problem has received less attention despite its important impact on the performance of order-picking systems (De Koster et al. 2007). There are a couple of issues that have been addressed in zoning literature. One issue is the layout design of the zones; for this issue, the research literature has primarily considered how to configure the zones, in terms of size and shape, in order to minimize the operational costs. Some examples are given by Gray et al. (1992), Petersen (2002) and Le-Duc and De Koster (2005).

Another issue is how to assign items over the multiple zones, which is directly related to the zone-stowage decisions we consider in this paper. Malmberg (1995) studies the assignment policy of the items based on the Cube-per-Order Index (COI) for multi-zone storage systems. Jewkes et al. (2004) consider the optimal zoning assignment decision for a specific sequential zone picking system where pickers work at home bases within their zones and must return to their home bases after each picking tour. Jane and Laih (2005) propose a heuristic algorithm for a multi-zone system that assigns each item to a single zone accounting for item affinity, namely the likelihood that two items will appear in the same order. There are a couple of papers that do consider the zoning

problem with the objective of workload balancing. Gray et al. (1992) discuss the benefits of allocating items uniformly across storage zones according to their demand class in order to achieve a balanced workload. Jane (2000) considers primarily the objective of workload balance across multiple storage zones in a progressive-zoning order-picking system. The paper proposes a heuristic algorithm that assigns each product to a single zone based on the demand history of the items. Onal et al. (2017) recently discuss an explosive storage policy by which the items are stored at many storage locations in large online retailing fulfillment centers. Although the paper does not focus on multi-zone storage systems, their conclusion on the benefits of spreading the inventory of an item over many storage locations is consistent with our study on using zone-stowage policies to create flexibility for picking.

Our work differs from the previous work in that instead of defining the optimal storage profile in a multi-zone storage system, we focus explicitly on the stowage policies. That is, our research examines the operational decision when a product is received, namely to which storage zones to stow the product. The storage profile at each storage zone is then a consequence of the stowage policy. Our intent is to find stowage policies that provide flexibility in order picking so as to achieve workload balancing across the storage zones.

The key decision of the zone-stowage problem is what quantity of each item to stow in each storage zone. The main objective of our work is to balance the workload across the different storage zones in order to avoid the situation of not having enough picking capacity to satisfy uncertain demand by some given shipping deadlines. The key concept we consider is the picking flexibility, which depends on the distribution of the inventory across the storage zones. For each item, the more storage zones the item is stowed in, the more picking options there will be. In this sense, there is an analogy with process flexibility as discussed by Jordan and Graves (1995). Process flexibility is "...being able to build different types of products in the same manufacturing plant or on the same production line at the same time." Stowing an item in multiple storage zones allows for each of the zones to pick ("build") the item. The work of Jordan and Graves (1995) has been extended to many contexts including manufacturing, supply chain, and service sectors. For instance, Chou et al. (2010) and Simchi-Levi and Wei (2012) further develop the underlying theory for process flexibility for asymptotically large systems and finite systems respectively. We refer the reader to

Buzacott and Mandelbaum (2008) for a broad overview on the research on flexibility in manufacturing and service systems, and to Chou et al. (2008) and Graves (2008) for more focused reviews on process flexibility.

Jordan and Graves (1995) find for systems with parallel servers or facilities that limited flexibility can perform nearly as well as a fully flexible structure, if properly configured. The key idea is to deploy flexibility in a way that chains together the products and facilities. In this paper, we explore how these ideas might translate and apply to the multi-zone storage system. In particular, we seek to understand the effectiveness of stowage policies that stow each item to a limited number of zones, relative to a stowage policy in which all items are stowed in all zones (i.e., a fully flexible configuration).

3. Evaluation Model for Comparing Stowage Policies

4. Model Framework

To evaluate the effectiveness of different stowage policies, we use a single-period, idealized model. The inputs of the model are the stochastic item demand, the fixed capacity limits for each storage zone, the inventory available to stow, and the stowage policy under consideration. In particular, we first assume that we start with an empty multi-zone storage system and an inventory of products to be stowed. We then stow the products according to a specified stowage policy. We then realize a single period of demand on the storage system, and determine how best to pick the demand from the inventory stored across the multiple storage zones. For given capacity constraints on how many units can be picked from each zone, the objective is to meet as much of the demand as possible within the single period. The intent of this modeling framework is to provide a basis for getting insights into the relative performance of different stowage policies.

Obviously, this framework is a simplification of reality as one would seldom if ever have an empty storage system into which to stow the inventory. But by starting with an empty system, this would seem to provide the ideal conditions for the implementation of each policy. As such, this should allow for the cleanest comparison between stowage policies.

We also think a single-period model is sufficient for our purposes where the length of the period would depend on how stowage and resource planning are done. Indeed, the period should roughly correspond to the frequency with which stowage is done. For instance, if stowage occurs in each shift, then the period could be one shift as there would be an opportunity to adjust the picking assignments and staffing for each zone in each period. If stowage occurs primarily, say, during the night or early morning hours, then the period might be one day, reflecting the fact that you can only make adjustments once a day to the inventory profiles.

We use simulation to evaluate each policy, where we take the demand on the storage system as being randomly generated. The major steps in the simulation include:

- Step 1: We generate the inventory to be stowed to the storage system for each item (refer to assumption A4 for the generation procedure), assuming that all storage zones are empty initially.
- Step 2: According to the given stowage policy, we assign the inventory of each item to the storage zones, subject to any given storage capacity limits.
- Step 3: We generate a single-period demand realization for each of the items stored in the storage system (refer to assumption A2 and A3 for the distribution of the demand).
- Step 4: We solve an optimization problem to fulfill the demand (from Step 3) as much as possible subject to the fixed picking capacity limits and to the inventory availability in each zone (from Step 2).

Formally, in the simulation framework above, Step 1 determines the inventory, namely the number of units of each item to be stowed in the system. Step 2 assigns these items to storage zones to create an inventory profile. That is, we determine parameter a_{ij} , the number of units of item i to be stowed to zone j , for each item and each storage zone; this will be done for each stowage policy under consideration. Step 3 generates parameter d_i , the demand realization of item i to be fulfilled. And the last step solves the following max-flow problem:

$$\begin{aligned}
& \max \sum_{i=1}^I \sum_{j=1}^J x_{ij} \\
s.t. \quad & \sum_{j=1}^J x_{ij} \leq d_i \quad \forall i \\
& \sum_{i=1}^I x_{ij} \leq c_j \quad \forall j \\
& 0 \leq x_{ij} \leq a_{ij} \quad \forall i, j
\end{aligned} \tag{1}$$

where x_{ij} is the decision variable, denoting the units used from zone j to satisfy the demand of item i ; c_j is the single-period picking capacity of zone j and is assumed to be fixed and given for each storage zone; I is the number of items and J is the number of storage zones. The objective is to pick as many units as possible subject to the constraints. We solve this optimization with Gurobi optimization solver.

For each inventory profile (as given by a_{ij}), we repeat steps 3 and 4 for a large number of demand realizations, so as to obtain statistically significant estimates of the performance of each stowage policy. As there is randomness in each stowage policy, we also simulate a large number of inventory profiles for each stowage policy.

5. Model Assumptions

We have designed an experiment to test and compare various stowage policies. The following assumptions provide the specifications for this experiment.

- A1. We assume there are 6 identical storage zones ($J = 6$), each with a fixed picking capacity. We assume that there is no space constraint.
- A2. We assume there are 100 items ($I = 100$) stored in the system, and the demand rate (expected single-period demand) for item i is given by the exponential function $\mu_i = \beta e^{is}$ where i is the index of item ($i = 0, 1 \dots 99$) and s and β are the shape and scale parameters respectively. (We refer the reader Bender (1981) for more discussion on modeling Pareto's law in the context of inventory systems). For this exponential function, the demand represented by the top 20% of the items only depends on the parameter s . We then create four demand patterns by adjusting the parameter s to allow the top 20% of the items to account for 80%, 70%, 60%, and 50% of the total expected demand respectively. We name those four demand patterns as Ultra-high

Skewness, High Skewness, Medium Skewness and Low Skewness. Furthermore, we set the parameter β accordingly so that the four cases have the same total expected demand of 2000 units per time period. We illustrate these four demand patterns in Figure 2 and summarize the parameter settings in Table 1. We note that the number of items in a typical storage system might easily be on the order of tens to hundreds of thousands; we limit our study to much fewer just for computational ease.

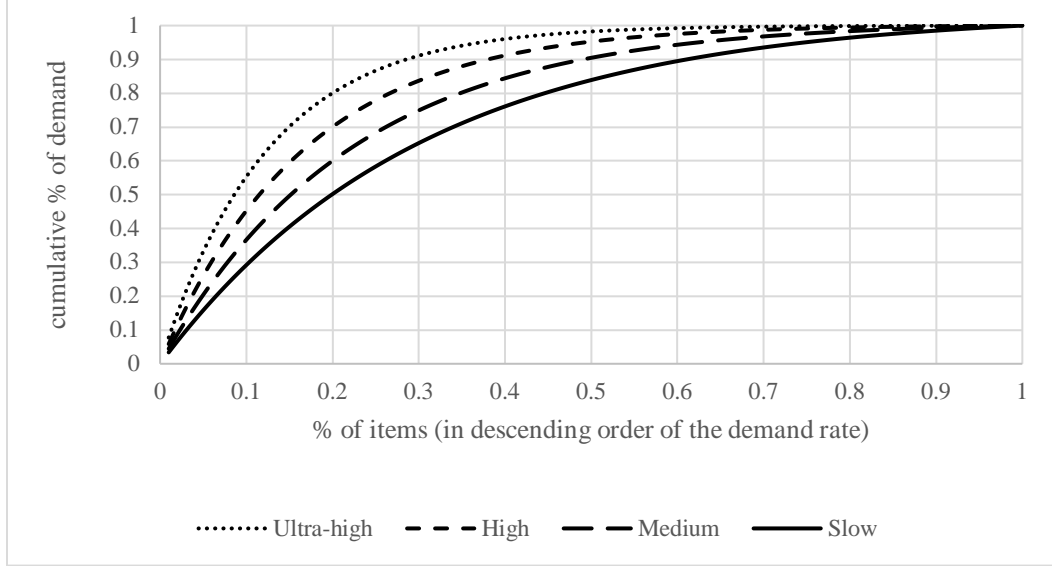


Figure 2: Demand Patterns: % Demand to % Items

	Ultra-High	High	Medium	Low
s	0.0805	0.0600	0.0450	0.0330
β	0.0535	0.3074	1.0345	2.5700

Table 1: Parameter Settings for Demand Patterns

- A3. We assume the demand of each item i follows a Poisson distribution with rate μ_i .
- A4. We assume that for each item the amount of inventory to be stowed is equal to four times its mean demand, rounded up to the nearest integer. Thus, for each item the inventory cover is four periods. As a consequence, for this setting, there is a very low probability that the assigned demand will exceed the available inventory. In particular, the stock-out probability of item i is

$$\Pr(d_i > [4\mu_i]) = 1 - \Pr(d_i \leq [4\mu_i]) = 1 - e^{-\mu_i} \sum_{t=0}^{[4\mu_i]} \frac{\mu_i^t}{t!}.$$

This function has local maximums at $\mu_i = 0.25, 0.5, 0.75, 1$. The stock-out probability is 0.026

for $\mu_i = 0.25$ ($\lceil 4\mu_i \rceil = 1$), 0.014 for $\mu_i = 0.5$ ($\lceil 4\mu_i \rceil = 2$), 0.007 for $\mu_i = 0.75$ ($\lceil 4\mu_i \rceil = 3$), and 0.004 for $\mu_i = 1$ ($\lceil 4\mu_i \rceil = 4$). For both the medium and low skewness cases, the demand rate for each item is greater than one, and hence there is a very low probability of stock-out for all items. For the other two cases, there are more items in the tail: for the high skewness case, 20% of the items have demand rates less than one, whereas for the ultra-high skewness case, 20% of the items have demand rates less than 0.25. Hence in these cases, our simulation setup will induce higher rates of stock-outs for these items with very low demand rates.

6. Stowage Policies

We test a set of plausible zone-stowage policies in the simulation. We use dedicated stowage as the base-case policy (Policy 1). We then structure a variety of policies that are inspired by the chaining strategy from Jordan and Graves (1995); each of these policies will split the inventory of each item across multiple storage zones (Policy 2-5). We also evaluate the effectiveness of a mixed strategy where some items are managed with a dedicated stowage policy while others use a chaining strategy (Policy 6-7). The final policy is a dedicated stowage policy that tries to balance the workload among storage zones by explicitly stowing items to the storage zone with the least aggregate velocity (Policy 8). We explain below the specification and settings of each stowage policy.

- *Policy 1 (no chain)*: For each item a storage zone is randomly selected and the entire inventory of the item is stowed to the storage zone. This is a base-case policy where the storage zones do not share common inventory items.
- *Policy 2 (3 short chains)*: Each item has a likelihood of 1/3 to be stowed to Zone1&2, 1/3 to Zone3&4 and 1/3 to Zone5&6; then the inventory of the item is divided in half, with half going to each of the two selected zones. With this policy, each pair of storage zones (Zone1&2, Zone3&4 or Zone5&6) forms a two-zone chain which allows for the picking capacity to be shared between the two storage zones within the chain.
- *Policy 3 (2 short chains)*: Each item has a likelihood of 1/2 to be stowed within Zone1&2&3 and 1/2 to be stowed within Zone4&5&6; then the inventory of the item is divided in half and

is stowed in two of the three selected zones, randomly chosen from the triplet. With this policy, each triplet of storage zones (Zone1&2&3 or Zone4&5&6) forms a three-zone chain.

- *Policy 4 (full chain)*: Each item has a likelihood of $1/6$ to be stowed to Zone1&2, Zone2&3, Zone3&4, Zone4&5, Zone5&6, and Zone6&1 respectively; the inventory of the item is divided in half with half going to each of the two selected zones. With this policy, all the storage zones are chained in the sense that the picking capacity of any storage zone can be shared with any other storage zone through the chaining structure.
- *Policy 5 (random pairs)*: Each item is stowed in a pair of zones that are randomly chosen; that is, each pair is equally likely to be chosen. The inventory of the item is divided in half with half going to each of the two selected zones. As there are many items in the system, this policy further connects the storage zones through the common items stored in two different zones.

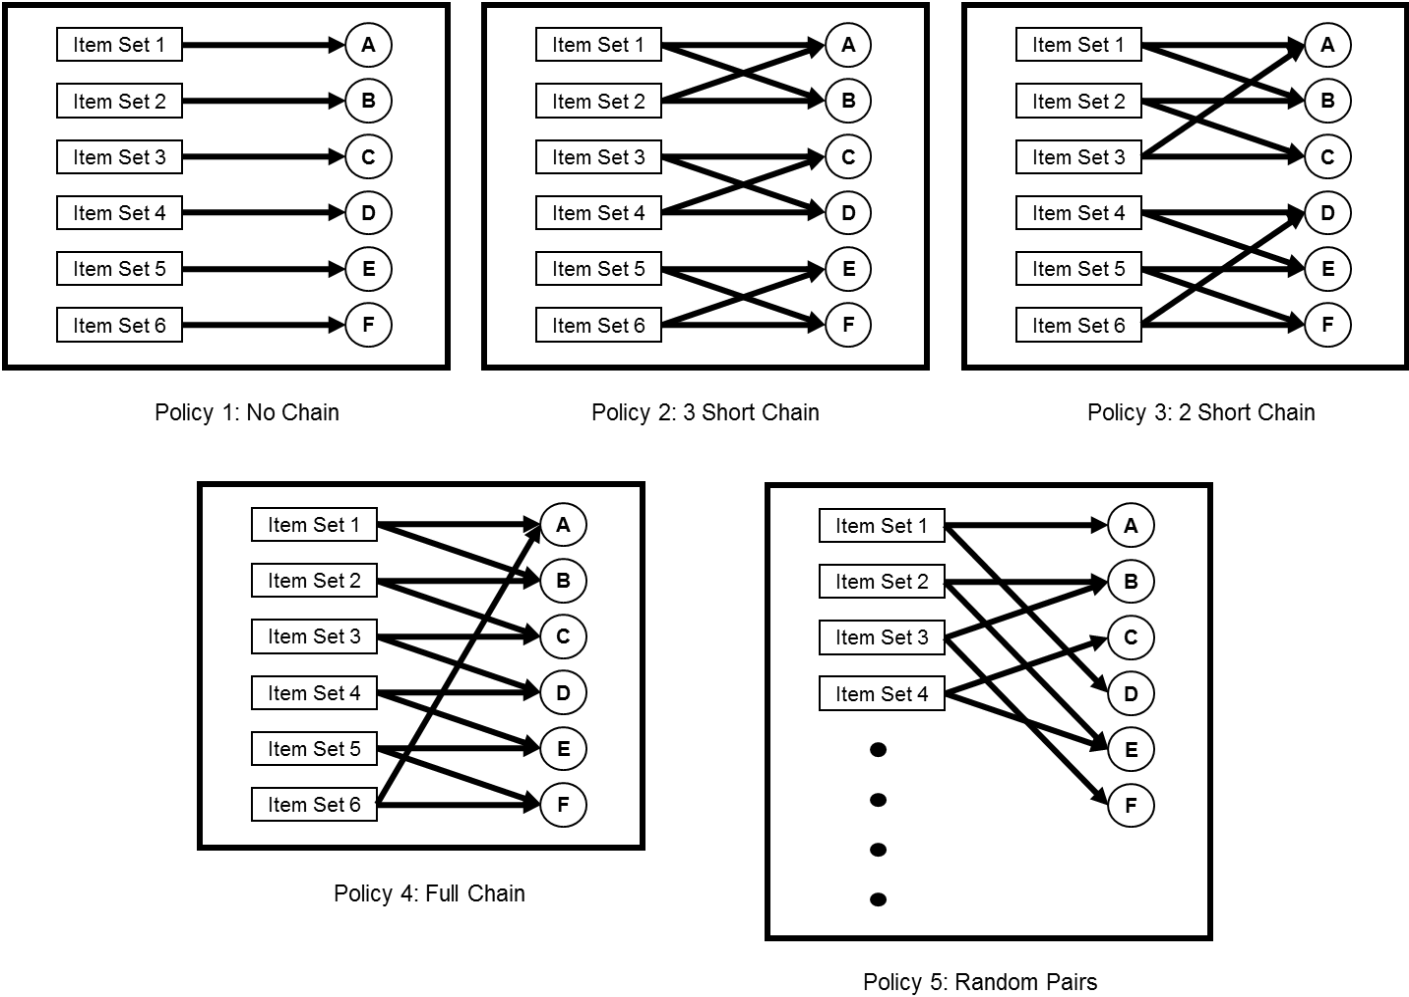


Figure 3: Storage policies. In each case an item is first assigned randomly to an “item set.” Policies 1, 2, 3, 4 each have 6 item sets as shown, with their assignments to the zones (A, B, C, D, E, F). For policy 5 there are $\binom{6}{2}=15$ items sets, of which only four are shown in figure.

In Figure 3 we depict the policies 1 to 5 for a 6-zone storage system. For each policy an item is first assigned randomly to an “item set.” For policies 1, 2, 3, 4 there are six item sets as shown, where each item set has an assignment to the zones (A, B, C, D, E, F). For policy 5 there are $\binom{6}{2} = 15$ items sets, corresponding to the number of zone pairs; we depict only four of these in the figure.

- *Policy 6 (mixed)*: The items are divided into two categories randomly; the first category has 80% of the items and each of the items is stowed according to policy 1: a storage zone is randomly selected and the entire inventory of the item is stowed to the storage zone. The

second category has 20% of the items and each item is stowed according to policy 5: the item is stowed in a pair of zones that are randomly chosen. The inventory of the item is divided in half with half going to each zone. This is a mixed stowage policy with less overhead than policies 2, 3, 4 or 5 as only 20% of the items have their inventory split.

- *Policy 7 (mixed smart):* The items are first ranked according to their demand rates. The items are then divided into two categories. The 80% of the items with the lowest demand rates are stowed by policy 1 and the remaining 20% of the items with the highest demand rates are stowed by policy 5. This policy is similar to Policy 6 but where we segment the items based on their demand velocity.
- *Policy 8 (dynamic balance):* Items are first randomly sequenced and then stowed sequentially according to this sequence; for each item we select the storage zone that has the smallest aggregate demand rate and stow the entire inventory of the item to this zone. We note that with this policy, we require the knowledge of the aggregate velocity of each storage zone to make the stowage decision. We compute the velocity for each zone by $\sum_{i \in A_j} \mu_i$ where A_j is an index set for the items that have already been assigned to storage zone j .

We note that in policy 2 through 7, when the received quantity R is an odd number, we assign $\frac{R-1}{2}$ units to each of the two chosen zones, and then assign randomly the “extra unit” to one of these zones. Finally, we note that only policy 8 depends on the sequence with which the stowage decisions are made. Both policy 7 and 8 require some knowledge of the item demand rates; however, policy 7 only needs to be able to differentiate the items into low versus high-velocity classes, whereas policy 8 requires knowledge of the exact demand rate for each item.

7. Numerical Results

8. Comparison of Different Stowage Policies

We compare the stowage policies listed in section 3.3 with the simulation setup described in 3.1. We define an inventory profile as a realization of the assignment of the inventory units for each item under a certain stowage policy. We created n inventory profiles for each stowage policy. We

use $A_k^p = \{a_{ij}\}$ to denote the k^{th} inventory profile for the stowage policy p . We generate m demand realizations for all items according to assumptions A2 and A3. We use D_l to denote the l^{th} demand realization.

For the base case, we assume a medium skewness demand pattern ($s = 0.045, \beta = 1.0345$). The expected total demand of the items is 2000 units per time period. We set the picking capacity equal to 370 units per time period for each storage zone; thus, the total picking capacity is 2220 units per period and the average resource utilization is around 90%. We let $n = 100$ and $m = 100$; that is, for each stowage policy we generate 100 inventory profiles, each of which is simulated against 100 (common) demand realizations for a total of 10,000 simulations for each policy.

We solve the max-flow problem in section 3.1 to obtain the single-period unfulfilled demand U_{kl}^p for the inventory profile A_k^p with the demand realization D_l for each stowage policy p . We note that we adjust this measure to account for any system inventory shortages, so as to reflect the true unfulfilled demand that is caused by an imbalanced inventory. That is, we set

$$U_{kl}^p \leftarrow U_{kl}^p - \sum_i \max(0, d_{l,i} - a_{k,i}^p)$$

where $d_{l,i}$, $a_{k,i}^p$ are the demand realization and total inventory of item i for the test case under consideration. We report for each stowage policy the average unfulfilled demand

$$\mu^p = \frac{\sum_k \sum_l U_{kl}^p}{nm}. \quad (2)$$

We also record the standard deviation of the unfulfilled demand under each inventory profile k for each stowage policy p

$$\sigma_k^p = \sqrt{\frac{\sum_l \left(U_{kl}^p - \frac{\sum_l U_{kl}^p}{m} \right)^2}{m}}. \quad (3)$$

To analyze the variability of the unfulfilled demand produced by the stowage policies, we report the average standard deviation for each stowage policy p

$$v^p = \frac{\sum_k \sigma_k^p}{n}. \quad (4)$$

We re-express μ^p and v^p as a percentage of the expected total demand in Table 2.

	μ^p	v^p
Policy 1 (no chain)	15.61%	2.73%
Policy 2 (3 short chains)	7.66%	2.13%
Policy 3 (2 short chains)	3.46%	1.37%
Policy 4 (full chain)	0.01%	0.03%
Policy 5 (random pairs)	0.00%	0.01%
Policy 6 (mixed)	7.82%	2.17%
Policy 7 (mixed smart)	0.18%	0.21%
Policy 8 (dynamic balance)	0.63%	0.87%

Table 2: μ^p and v^p Represented as Percentage of the Expected Total Demand

From this analysis, we observe that

- The performance of the dedicated stowage policy (policy 1) is quite poor, and will often result in overloaded storage zones that cannot keep up with the demand.
- There are two strategies that can dramatically improve the performance. One is to spread out the inventory of each item across the zones so as to create chains (policies 2, 3, 4, 5, 6, 7). The other is to assign inventory to zones in a way that balances the expected workload (policy 8).
- The performance of the system gets better when “the chain is longer” (policies 2, 3, 4) -- indicating that the chain structure provides additional flexibility for picking. For example, the full chain structure (policy 4) always outperforms two shorter chains (policy 3), which outperforms three even shorter chains (policy 2). Essentially, with a longer chain there is more flexibility to “move” the picking capacity across the storage zones to satisfy the demand.
- Furthermore, the random pairs policy does slightly better than the full chain. This is because the random pairs policy generates higher connectivity among the storage zones. This effect can be captured by the so-called expansion index. We refer to Chou et al. (2008) for more detailed discussion on this.
- If we were to stow 1/6 of the inventory in each zone and ignored integrality requirements, then the solution of the max flow problem (1) is just the system inventory shortage; thus, the adjusted unfulfilled demand is zero. This is clearly the best possible stowage policy and is the analog to a system with full flexibility in this context. Hence, we observe from Table 2 that the

full chain and random pairs policies effectively achieve the performance of the fully-flexible system that spreads the inventory of each item evenly across the storage zones. This suggests that there may be little value to split the inventory of an item to more than two storage zones. As in Jordan and Graves (1995), we observe that limited flexibility, deployed in the right way (i.e., policies 4 and 5), achieves the performance of full flexibility.

- From the mixed and mixed-smart policies, we see that we might only need to spread a fraction (20%) of the items in order to get the benefits from chaining. In particular the mixed-smart policy shows that only the high-volume items need to be split. This is quite important from a practical perspective, as there are operational costs to splitting the inventory as discussed earlier. The evaluation of the mixed-smart policy shows that by only splitting the high-volume items we can get quite near to the performance of the best stowage policies, namely the full chain and random pairs policies.
- The dynamic balance policy can also be quite effective, although it lags relative to the full chain, random pairs and mixed-smart policies. If it were not possible or too costly to split an item's inventory across storage zones, then the dynamic balance policy would seem most reasonable as it works to assure that the inventory is “evenly” distributed across the storage zones. However, this policy requires knowledge of the demand rates for all items. In comparison, the other policies, with the exception of the mixed-smart policy, do not require knowledge of the item demand rate; and for the mixed-smart policy, we only need to be able to separate the items into high versus low demand.

9. Sensitivity Analysis

We perform sensitivity analysis with respect to the resource utilization, the demand skewness, and the inventory level in this section. The intent is to get a deeper understanding of the relative performance of the different zone-stowage policies, as reported in the prior section. Of particular interest is to test the effectiveness of the stowage policies discussed in 3.3 under the unfavorable scenarios such as high resource utilization level, high demand skewness, or heterogeneous inventory coverage.

Sensitivity Analysis towards Resource Utilization Level

In this test, we analyze the zone-stowage policies under four resource utilization levels. We set the

demand pattern parameters as $s = 0.0465, \beta = 1$ (medium skewness) in this test. We report the average unfulfilled demand in Table 3 for each of the four resource utilization levels where the base case is 90%. As expected, we observe that the performance of the system gets worse when the resource utilization increases, but the chaining and random-pairs stowage policies continue to work very well even when capacity is very tight. Indeed, the relative improvement from a full chain or random-pairs policy grows as the utilization increases. These two policies, which create the most connectivity across the storage zones, exhibit the least sensitivity to the impact from increased utilization.

Policy	85%	90%	95%	99%
Policy 1 (no chain)	12.61%	15.61%	21.37%	24.87%
Policy 2 (3 short chains)	5.00%	7.66%	11.99%	15.20%
Policy 3 (2 short chains)	2.78%	3.46%	7.21%	9.92%
Policy 4 (full chain)	0.00%	0.01%	0.31%	1.37%
Policy 5 (random pairs)	0.00%	0.00%	0.01%	1.09%
Policy 6 (mixed)	5.22%	7.82%	11.27%	15.52%
Policy 7 (mixed smart)	0.02%	0.18%	1.19%	2.90%
Policy 8 (dynamic balance)	0.15%	0.63%	2.33%	4.93%

Table 3: μ^P for Different Resource Utilization Levels

Sensitivity Analysis towards Demand Skewness

In this test, we analyze the stowage policies under low, medium, high and ultra-high demand skewness. Figure 2 depicts the demand patterns. We assume a resource utilization of 90% in this test. We report the average unfulfilled demand in Table 4 for each demand pattern we consider. As expected, we observe that the performance of the system gets worse when there is more skewness in the demand. Again the full chain and random pairs policies are quite insensitive to the increase in demand skewness and continue to be effectively equivalent in performance to a full flexibility policy. It is noteworthy that policy 7 (mixed smart) works very well even for the ultra-high skewness demand pattern as the high demand items are always split into two storage zones.

Policy	Low Skewness	Medium Skewness	High Skewness	Ultra-high Skewness
Policy 1 (no chain)	14.46%	15.61%	17.84%	24.46%
Policy 2 (3 short chains)	6.18%	7.66%	9.59%	12.66%
Policy 3 (2 short chains)	3.16%	3.46%	5.58%	5.35%

Policy 4 (full chain)	0.03%	0.01%	0.24%	0.17%
Policy 5 (random pairs)	0.00%	0.00%	0.03%	0.06%
Policy 6 (mixed)	6.46%	7.82%	9.17%	14.86%
Policy 7 (mixed smart)	0.08%	0.18%	0.36%	0.69%
Policy 8 (dynamic balance)	0.39%	0.63%	1.18%	2.28%

Table 4: μ^p for Different Demand Skewness

Sensitivity Analysis towards Inventory Level

In this test, we relax assumption A4 that set the inventory level for each item. Now, we assume that each item has an inventory equal to x times its demand mean where x is assumed to be uniformly distributed in the range of $[u_1, u_2]$. We assume Medium Skewness for the demand, and an average resource utilization of 90% in this test. We compare the results for the average unfulfilled demand for the inventory level in the range of [3,5], [2,6] and [1,7] in Table 5. (The base case has $x = 4$ for all items.) We observe that the performance of each of the stowage policies is in general insensitive to variability in the inventory level, with the exception of the dynamic balance policy; this suggests that stowing items according to the aggregate demand rates may be less effective when the relative amount of the inventory varies across the items.

Policy	4	[3,5]	[2,6]	[1,7]
Policy 1 (no chain)	15.61%	17.36%	16.46%	16.78%
Policy 2 (3 short chains)	7.66%	8.09%	7.87%	8.58%
Policy 3 (2 short chains)	3.46%	3.66%	4.33%	3.64%
Policy 4 (full chain)	0.01%	0.24%	0.08%	0.02%
Policy 5 (random pairs)	0.00%	0.00%	0.00%	0.00%
Policy 6 (mixed)	7.82%	8.23%	7.71%	8.15%
Policy 7 (mixed smart)	0.18%	0.04%	0.28%	0.11%
Policy 8 (dynamic balance)	0.63%	1.05%	2.29%	4.59%

Table 5: μ^p for Inventory Level Uniformly Distributed Over the Demand Mean

10. Conclusion

In this research, we analyze zone-stowage policies by means of a simulation study. We identify two very effective stowage strategies that could potentially improve the picking performance by preserving the balance of the inventory distribution. The first strategy is inspired by chaining as developed for achieving process flexibility. The strategy suggests splitting the arriving inventory

for each item in half and stowing each half in a different storage zone. Our tests show that the best assignment is to randomly choose the pair of storage zones to which to send the two sub-lots of inventory. Alternatively, we find that stowing the inventory according to a full chain is nearly as effective. In addition, we find that an effective policy only needs to apply the random-pairs strategy to the high velocity items, even under high resource utilization and high demand variability. This is of practical importance as there may be additional material handling cost or complexity associated with splitting the items.

The second strategy, as embedded in the dynamic balance policy, stows the arriving items based on the aggregate demand load in each zone. The intent is to stow the inventory in a way that maintains workload balance across the storage zones. But to implement this strategy will require knowledge of the aggregate velocity of the items stored in the storage zones. Furthermore, our implementation of this policy assumes that the inventory of each item is dedicated to a storage zone, which makes the determination of the aggregate velocity relatively straight-forward; if an item's inventory is spread over multiple zones, then this determination becomes more challenging.

Future research might extend this work under less stringent assumptions. For instance, we ignored space capacity constraints and we limited our analysis to a single period setting. It would be interesting to explore how these zone-stowage policies perform in a dynamic multi-period setting with constraints on both picking and space.

Reference

Bender, Paul S. (1981). Mathematical modeling of the 20/80 rule: theory and practice. *Journal of Business Logistics* 2.2, 139-157.

Buzacott, J. A., & Mandelbaum, M. (2008). Flexibility in manufacturing and services: achievements, insights and challenges. *Flexible Services and Manufacturing Journal*, 20(1-2), 13.

Chou, M. C., Teo, C. P., & Zheng, H. (2008). Process flexibility: design, evaluation, and applications. *Flexible Services and Manufacturing Journal*, 20(1-2), 59-94.

Chou, M. C., Chua, G. A., Teo, C. P., & Zheng, H. (2010). Design for process flexibility: Efficiency of the long chain and sparse structure. *Operations Research*, 58(1), 43-58.

- De Koster, R., Le-Duc, T., & Roodbergen, K. J. (2007). Design and control of warehouse order picking: A literature review. *European Journal of Operational Research*, 182(2), 481-501.
- De Koster, R. B., Le-Duc, T., & Zaerpour, N. (2012). Determining the number of zones in a pick-and-sort order picking system. *International Journal of Production Research*, 50(3), 757-771.
- Enright, J., & Wurman, P. R. (2011). Optimization and Coordinated Autonomy in Mobile Fulfillment Systems. In *Automated action planning for autonomous mobile robots*, pp. 33-38.
- Graves, S. C. (2008). Flexibility principles. In *Building Intuition* (pp. 33-49). Springer US.
- Gray, A. E., Karmarkar, U. S., & Seidmann, A. (1992). Design and operation of an order-consolidation warehouse: Models and application. *European Journal of Operational Research*, 58(1), 14-36.
- Gu, J., Goetschalckx, M., & McGinnis, L. F. (2007). Research on warehouse operation: A comprehensive review. *European Journal of Operational Research*, 177(1), 1-21.
- Jane, C. C. (2000). Storage location assignment in a distribution center. *International Journal of Physical Distribution & Logistics Management*, 30(1), 55-71.
- Jane, C. C., & Laih, Y. W. (2005). A clustering algorithm for item assignment in a synchronized zone order picking system. *European Journal of Operational Research*, 166(2), 489-496.
- Jewkes, E., Lee, C., & Vickson, R. (2004). Product location, allocation and server home base location for an order picking line with multiple servers. *Computers & Operations Research*, 31(4), 623-636.
- Jordan, W. C., & Graves, S. C. (1995). Principles on the benefits of manufacturing process flexibility. *Management Science*, 41(4), 577-594.
- Lamballais, T., Roy, D., & De Koster, M. B. M. (2017). Estimating performance in a robotic mobile fulfillment system. *European Journal of Operational Research*, 256(3), 976-990.
- Malmberg, C. J. (1995). Optimization of cube-per-order index warehouse layouts with zoning constraints. *International Journal of Production Research*, 33(2), 465-482.
- Onal, S., Zhang, J., & Das, S. (2017). Modelling and performance evaluation of explosive storage policies in internet fulfillment warehouses. *International Journal of Production Research*, 55(20)

5902-5915.

Petersen, C. G. (2002). Considerations in order picking zone configuration. *International Journal of Operations & Production Management*, 22(7), 793-805.

Wurman, P. R., D'Andrea, R., & Mountz, M. (2008). Coordinating hundreds of cooperative, autonomous vehicles in warehouses. *AI magazine*, 29(1), 9.

Simchi-Levi, D., & Wei, Y. (2012). Understanding the performance of the long chain and sparse designs in process flexibility. *Operations Research*, 60(5), 1125-1141.

Staudt, F. H., Alpan, G., Di Mascolo, M., & Rodriguez, C. M. T. (2015). Warehouse performance measurement: a literature review. *International Journal of Production Research*, 53(18), 5524-5544.

Yuan, Rong. (2016) Velocity-based Storage and Stowage Decisions in a Semi-automated Fulfillment System. Ph.D. Dissertation, Massachusetts Institute of Technology.

Zou, B., Gong, Y., Xu, X., & Yuan, Z. (2017). Assignment rules in robotic mobile fulfillment systems for online retailers. *International Journal of Production Research*, 55(20), 6175-6192.