Chapter 7
Affinity Based Slotting in Warehouses with Dynamic Order Patterns

Monika Kofler, Andreas Beham, Stefan Wagner, and Michael Affenzeller

Abstract. There has been a wealth of research on warehouse optimization since the 1960s, and in particular on increasing order picking efficiency, which is one of the most labor intensive processes in many logistics centers. In the last ten years, affinity based slotting strategies, which place materials that are frequently ordered/picked together close to each other, have started to emerge. However, the effects of changing customer demand patterns on warehousing efficiency have not been investigated in detail. The aim of this chapter is to extend the classic storage location assignment problem (SLAP) to a multi-period formulation (M-SLAP) and to test and compare how various allocation rules, and in particular an affinity based policy, perform in such dynamic scenarios. A first benchmark instance for the M-SLAP is presented.

7.1 Introduction

The storage location assignment problem (SLAP) was first formulated by Hausman, Schwarz and Graves [8], who had identified a need for research on the design and scheduling of automated warehousing systems. Automated warehousing systems were introduced in the 1950s but became increasingly pervasive in the 1970s. Hausman et al. [8] also developed a widely-used taxonomy for assignment policies, distinguishing between dedicated storage, randomized storage and class-based storage, which are described in detail in Section 7.2.

Frazelle [4] has since shown that the SLAP is NP-Hard. Previous studies mainly focused on finding optimal solutions for instances of this combinatorial optimization problem. However, a lack of commonly shared benchmark instances makes comparisons between research results very challenging. Researchers either rely on randomly generated instances or use proprietary data from a real-world warehouse that can not

Monika Kofler · Andreas Beham · Stefan Wagner · Michael Affenzeller
University of Applied Sciences Upper Austria, School of Informatics,
Communication and Media, Softwarepark 11, 4232 Hagenberg, Austria
e-mail: {mkofler,abeham,swagner,maffenze}@heuristiclab.com

R. Klempous et al. (eds.), *Advanced Methods and Applications in Computational Intelligence*,
Topics in Intelligent Engineering and Informatics 6,
DOI: 10.1007/978-3-319-01436-4_7,

be released to the public. Despite this, the main research efforts in this field are still dominated by the search for best possible solutions. Common SLAP optimization approaches therefore result in assignments that completely change the - usually randomly slotted - initial assignment. If such a solution should be put into practice, it would involve extensive re-arrangements akin to filling a warehouse from scratch. While this is a suitable approach for filling an initially empty warehouse or to determine the potential for improvement, it can not be implemented easily in an already operating warehouse. In a real-world scenario the potential efficiency gains should instead be weighed against the re-arrangement efforts. In this paper we will extend the classic SLAP formulation for multi-period scenarios, thus allowing an evaluation of storage assignment strategies under changing conditions, and use an extended objective function that considers re-arrangement, putaway and order picking efforts.

The chapter is structured as follows: Section 7.2 gives an overview about previous research on slotting, with a focus on random (7.2.1), turn-over based (7.2.2) and affinity based (7.2.3) slotting. Since affinity based slotting strategies are a fairly recent scientific development, a detailed mathematical formulation of the Pick Frequency / Part Affinity (PF/PA) score as one representative of the group is provided in Section 7.2.4. In Section 7.3 we introduce the multi-period storage location problem (M-SLAP) and distinguish between re-warehousing, which involves the re-arrangements of large parts of the warehouse, and healing, which moves only a small number of goods at a time. To test the effectiveness of various slotting strategies on the M-SLAP, a benchmark dataset has been released, which is described in Section 7.4. Test configurations and results are given in Section 7.5, followed by a brief discussion and outlook in Section 7.6.

7.2 Introduction to Slotting

The goal of slotting - also called storage assignment optimization, inventory slotting, or inventory profiling - is to determine the *best* place to store each stock keeping unit (SKU) in a warehouse. Two of the most common incentives for companies to slot their warehouse include the need to squeeze more SKUs into an already overflowing warehouse and the desire to reduce overall handling costs and efforts [2]. In many order picking environments the travel time to retrieve an order has been found to be the largest component of labor, amounting to 50% or more of total order picking time [21]. Therefore this study will focus on the reduction of travel efforts, or more specifically aisle changes, as an approximation of travel time. It should be noted that other factors, such as load balancing accross warehouse zones, work ergonomics (e.g. to reduce bending and reaching activity [19]) or pre-consolidation (to reduce downstream sorting), can also be of importance for a particular scenario.

Generally speaking, slotting is a two-stage process that first assigns a SKU to a product class and afterwards assigns the class to storage locations within the warehouse. Within a class the SKUs are usually arranged via a simple policy such as random or closest location. If there is only a single class, the approach is called random storage, which is discussed in more detail in section 7.2.1. Conversely, if the number of classes equals the number of SKUs, the policy is called dedicated storage. Class-based storage is situated somewhere in between random and dedicated storage. Choosing the right number of classes, product-to-class assignment strategy and storage locations for each class is dependent on the particular warehousing scenario, in particular the layout, material handling equipment, routing strategy and order profile. Detailed reviews of the various storage assignment strategies can be found in [7] and [3]. The primary literature considered in this paper is also summarized in Table 7.1.

Table 7.1. Overview of slotting strategies with an emphasis on the Cube per Order Index (COI) and affinity-based methods such as Order Oriented Slotting (OOS) and Pick Frequency / Part Affinity (PF/PA) Slotting

Type	Year	Citation	Method
Random	1996	C. Malmborg [16]	Random vs. dedicated storage
Turnover-based	1963	Heskett [9]	COI
	1976	Kallina and Lynn [10]	COI
Affinity-based	1989	Frazelle and Sharp [5]	Correlated storage
	2005	Garfinkel [6]	Correlated storage
	2007	Mantel, Schuur, and Heragu [17]	OOS
	2008	Kim and Smith [11]	Improving search
	2009	de Ruijter et al. [20]	OOS, Parameter tuning
	2010	Kofler et al. [13]	PF/PA Slotting
	2010	Wutthisirisart [24]	Minimum Delay Algorithm
	2011	Kofler et al. [14]	PF/PA Slotting, Healing
Review Articles	2007	Gu, Goetschalckx, and McGinnis [7]	Various Approaches
	2007	deKoster, Le-Duc, and Roodbergen [3]	Various Approaches

Most of the existing scientific studies present results obtained by completely re-slotting a warehouse according to some assignment strategy, a process commonly referred to as re-warehousing. Migrating from a randomly slotted warehouse to class-based storage, changing the number of classes or the class characteristics in a live operating warehouse is more problematic. Warehouse managers might wish to implement a target assignment gradually during normal operations, rather than interrupt the normal workflow to move hundreds of items. In addition, an optimal storage assignment created by re-warehousing might become outdated before long due to fluctuations in product demand caused by seasonal variations and product life cycle characteristics. In this study, we therefore also distinguish between construction and improvement approaches. Construction algorithms build a feasible slotting from scratch, assuming that the warehouse is initially empty. On the other hand, improvement algorithms try to enhance an existing feasible solution. We will compare the results of both construction and various improvement strategies, distinguishing further between greedy and heuristic search methods. The approaches are summarized in Table 7.1 and described in more detail in subsection 7.2.1 to 7.2.3.

7.2.1 Random Slotting

In the random storage paradigm incoming SKUs are assigned randomly to suitable, available storage locations. The advantages of random storage are ease of implementation and balanced picker traffic across the warehouse. On the downside this may result in longer travel times. Random storage is frequently used in practice and as a performance baseline in the scientific literature.

7.2.2 Slotting by Turnover Based Metrics

Early attempts to optimize slotting in dedicated storage warehouses were based on the idea that fast-moving items should be located in easily accessible pick areas. Heskett extended this simple policy and proposed the cube per order index (COI) rule [9], which ensures that heavy or fast-moving SKUs are stored in more desirable locations close to ground level. Kallina and Lynn discussed the implementation of the COI policy in practice and proved that the COI rule is optimal under certain conditions [10]. One such condition is that there is no dependency between picked items in the same tour, which is unfortunately not the case for most order picking scenarios [7]. Modifications of the COI rule have since been published, which also consider inventory costs, zoning constraints or work ergonomics [16].

7.2.3 Slotting by Affinity

In order picking environments a picker usually retrieves multiple items per order, processing the individual order lines according to some routing strategy. Items that are frequently ordered together are said to be correlated or affine [6]. The idea behind slotting by affinity is that storing affine items close to each other will reduce the total travel time. Unfortunately this is not universally true for all warehouse scenarios and depends on the specific warehouse layout, material handling equipment, picker routing strategy and order profile. Even if two items are ordered together the picker will not necessarily pick them in sequence [22]. Moreover, narrow aisles that do not allow reverse back out or large orders might require a full traversal of the warehouse anyway. In such scenarios storing by affinity could even have a negative effect on picking efficiency by causing congestion in aisles were many fast-moving items are stored together.

Slotting by affinity was first introduced by Frazelle et al., who implemented a class-based storage strategy called *correlated storage* [5]. The algorithm starts with the most popular product and subsequently adds affine items to the class until a capacity constraint is reached. The generated classes are then placed within the warehouse according to their total popularity.

Garfinkel developed a local search improvement heuristic based on 2-exchange and cyclic exchange with the goal to minimize multi-zone orders. The moves are evaluated with different correlation measures and the algorithm is benchmarked against various construction approaches [6]. Kim and Smith introduced a similar two-phase dynamic slotting heuristic procedure that generates an initial assignment using a pick frequency measure such as the COI rule and afterwards uses pairwise interchanges to move affine parts closer to each other [11].

Mantel, Schuur and Heragu created the order oriented slotting (OOS) problem [17]. They present integer linear programming (ILP) models for two small warehouse scenarios but also two heuristics for the optimization of larger, real-world instances. Their so-called *interaction frequency based quadratic assignment heuristic* (IFH-QAP) forces items that frequently occur in the same order to be close together, while at the same time ensuring that fast movers are not allocated too far from the I/O-point. The objective function is based on two measures that are multiplied with the routing specific distances for each SKU: The *popularity* of an SKU denotes how often it was ordered. The *interaction frequency* of two SKUs equals the number of orders that contain them both. The impact of each component on the target function is tuned via a weight parameter, which can be adjusted empirically or automatically as shown in [20]. Kofler et al. devised the pick frequency / part affinity (PF/PA) objective function [13], which extends the IFH-QAP heuristic such that one SKU may be stored in multiple locations throughout the warehouse. Moreover relative values of interaction frequency and popularity are used to get comparable results, independent of the time window that is considered for a given order profile.

Another interesting approach that considers SKU affinity is the Minimum Delay Algorithm (MDA) [24], which was inspired by linear placement algorithms. In this construction heuristic approach SKUs are placed in a fashion that reduces the delay (= additional traveling distance) for the other orders.

Due to the similarities between [6], [11], [17] and [13] the PF/PA score is used exemplarily in this study as a representative of improvement heuristics for affinity-based slotting. A detailed algorithm description is given in the following section.

7.2.4 Pick Frequency / Part Affinity Score

The pick frequency / part affinity (PF/PA) score was first introduced in [13] and slightly revised in [14]. The approach combines affinity based storage and storage by pick frequency. To increase comprehensibility and reproducibility of the results presented in this study, we provide a full mathematical formulation of the score, as previously published in [14]. Let's assume:

$$\begin{gathered} S = \text{set of all storage locations } s_k; \\ 0 < k <= m \text{ where } m = |S| \end{gathered} \tag{7.1}$$

$$dist(s_k, s_l) = \text{routing specific distance between storage location } s_k \text{ and } s_l \tag{7.2}$$

In addition to the storage locations that can hold products, at least one input/output location *origin* must be defined that denotes the shipping dock. The travel times / distances between locations are routing-specific.

Usually, the storage locations and the distance matrix need to be determined only once for a given warehouse and can later be re-used for different problem instances. Conversely, the following parameters are likely to change over time and need to be retrieved from the enterprise resource planning or warehouse management system. Most important, the set P lists all products that are present in a particular assignment.

$$P = \text{set of all products } p_i; \quad 0 < i <= n \text{ where } n = |P| \tag{7.3}$$

For each product p_i we need to know the number of picking orders in which the product occurs. We use the relative number of orders $orderRatio(p_i)$ to get comparable scores as results, independent of the time window that is considered from a given order profile. Similarly, the affinity matrix stores the ratio of all orders in which a particular product pair occurs together. Finally, the current warehouse assignment defines how many products p_i are stored at location s_k. The set of locations $L(p_i)$ stores all locations of a particular product.

$$orderRatio(p_i) = \text{relative number of orders in which } p_i \text{ occurs} \tag{7.4}$$

$$affinity(p_i, p_j) = \text{relative number of orders in which products } p_i \text{ and } p_j \text{ occur together} \tag{7.5}$$

$$quantity(p_i, s_k) = \text{number of packing units of product } p_i \text{ stored at location } s_k \tag{7.6}$$

$$L(p_i) = \text{set of all locations where } quantity(p_i, s) > 0 \tag{7.7}$$

The entities defined in 7.3-7.7 can be calculated from order picking histories or demand forecasts and the current warehouse assignment. We can now define the objective functions in Equation 7.8 and 7.9.

$$PF = \sum_{i=0}^{n} \frac{orderRatio(p_i)}{|L(p_i)|} \cdot \sum_{s \in L(p_i)} dist(s, origin) \tag{7.8}$$

The *total pick frequency score* (PF), as defined in Equation 7.8, ensures that frequently picked products are placed in more favorable storage locations near the I/O point. For each product it detects all current storage locations $L(p_i)$, calculates their distance to the origin and weights each distance with the expected number of picks given the number of previous orders. The picks are uniformly distributed on all storage locations, independent of the actual stored quantities in the different locations.

$$\begin{aligned} PA = & \sum_{i=0}^{n} \sum_{j=0}^{n} \frac{affinity(p_i, p_j)}{|L(p_i)| \cdot |L(p_j)|} \\ & \cdot \sum_{s_k \in L(p_i)} \sum_{s_l \in L(p_j)} dist(s_k, s_l) \end{aligned} \tag{7.9}$$

The *total part affinity score* (PA) takes all pairs of products p_i and p_j, and retrieves all respective storage locations $L(p_i)$ and $L(p_j)$ from the current assignment. The distance between each resulting storage location pair is calculated and weighted with the part affinity divided by the number of location pairs $|L(p_i)| \cdot |L(p_j)|$. The term reduces to zero for products with no part affinity, therefore the calculation can be sped up by only looking at products p_j that have an affinity greater than zero with a given product p_i.

The resulting multi-objective evaluation function for assignments is computed as weighted sum of the two objective functions given in Equation 7.8 and 7.9 such that

$$\text{PF/PA score} = \alpha \cdot PF + \beta \cdot PA. \tag{7.10}$$

PF and PA have different ranges. The parameters α and β can be adjusted automatically ensure that both factors contribute equally to the objective as proposed in [20]. We prefer tuning them manually, usually assigning an intentionally higher weight to α as discussed in [14].

7.3 Multi-period Warehouse Slotting

Both turn-over and affinity based slotting strategies rely on historical SKU order profiles and/or demand forecasts for decision-making. However, slotting is usually not a one-off event since SKU demands are subject to change over time. An optimal storage assignment might become outdated due to demand fluctuations, modifications to the picking line, infrastructure changes, or variations in the order mix etc. We therefore propose an extension of the classical SLAP, called the multi-period storage location assignment problem (M-SLAP).

The M-SLAP problem formulation was inspired by the field of facility layout problems (FLP), where researchers face a similar challenge: Changing material flows between departments during the planning horizon might make re-arrangements of the departments necessary, in order to keep material handling costs low. In the so-called dynamic FLP formulation the planning horizon is split into multiple consecutive periods. Within each period, the estimated flow data is assumed to be constant. The optimization generates a separate layout for each of these periods, which can be treated as classic, static FLPs. However, re-arrangement costs caused by layout adjustments between periods of different material flows are considered in the evaluation of the dynamic layout. An overview of dynamic layout algorithms can be found in [1].

Likewise, the M-SLAP problem consists of a sequence of warehouse assignments and their associated SKU demand profiles. To evaluate the quality of an M-SLAP solution the picking effort for a set of picking orders is simulated on each generated assignment. The objective is to minimize the total travel distance of all pickers, caused by order picking and (re-)slotting. The aim of this new formulation is to evaluate mid- and long-term stability of warehouse assignments generated by various slotting approaches. Most important, it should yield a better cost-benefit analysis of re-slotting effort (= cost) vs. order picking savings (= benefit).

7.3.1 Re-warehousing

Most previous studies focused on achieving optimality in the allocation of SKUs and often re-ordered the existing layout to a very large degree. This process is commonly referred to as re-warehousing. The effort for re-warehousing a real-world logistics center is considerable and requires the movement of hundreds or thousands of items, thus blocking personnel and material handling equipment for an extensive time period. This is why it is conducted infrequently (quarterly, biannually, during holidays) in practice.

The only study known to us that considers re-warehousing over time was conducted by Neuhäuser and Wehking [18]. In this publication they developed a metric to determine suitable re-warehousing intervals for the food retail sector, which is characterized by strong seasonal demand and stock fluctuations. The authors assume that the optimal target storage zone or class is known for each stock keeping unit at each point in time, for instance based on COI, and calculate the (weighted) sum of displaced SKUs over all SKUs as a metric of warehouse entropy. By comparing the cost of operation in warehouses with high and low entropy and also considering the re-warehousing costs, they were able to determine suitable re-warehousing intervals for their scenario in a simulation study. Unfortunately, data and simulation model are not publicly available, which makes a comparison with the other approaches or replication of the results difficult.

7.3.2 Healing

From a warehouse manager's perspective, achieving 'optimality' might not be as important as finding a 'good enough' slotting that is both robust and requires only few stock transfers to implement. For instance, experiments in [14] showed that moving only a limited number of SKUs can significantly reduce the total travel distances. In this previous study picker travel distances could be reduced by 23% by moving only 60 pallets. The total distance optimization potential for the problem instance was over 60%, however this required the movement of 1400 pallets. For many real-world scenarios it might be more prudent to move only a few SKUs per day, but iteratively over a longer time period. We refer to this process as *healing* [14].

7.3.3 M-SLAP: Optimization and Evaluation

The typical M-SLAP optimization steps are depicted in Figure 7.1 for a single period. The result of this chain is a new warehouse assignment, which is used as input for the next period, plus cost estimates for re-arrangements, order picking and putaway effort. The total costs of an M-SLAP is the sum over all costs over all periods.

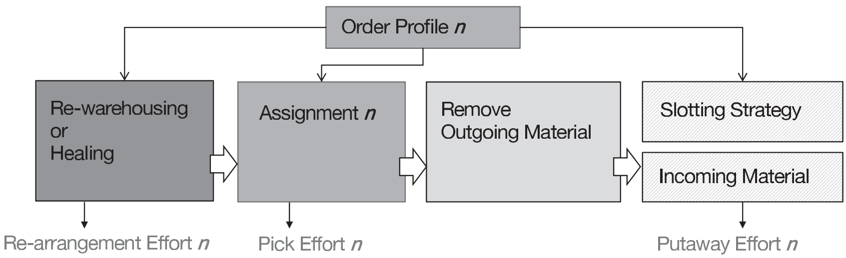

Fig. 7.1. The M-SLAP optimization process for a single period

We explain the individual steps in more detail, starting with the first period:

1. Re-warehousing or healing: This step is optional. In the first period, a warehouse assignment can either be retrieved from a real-world warehouse management system or generated via a heuristic (=re-warehousing). In each subsequent period, either re-warehousing, healing or no action can be conducted. If re-warehousing/healing is performed the re-arrangement effort counts towards the total costs.
2. Assignment 1: The assignment is then evaluated with the test order profile of period 1, resulting in a cost estimate for the pick process in this period.

3. Remove outgoing material: The test order profile should be representative for a period but does not have to include all pick operations that take place within the period. It is therefore assumed that all SKUs that are present in period 1 but not in stock in period 2 will be removed in one step. This is one of the abstractions of the M-SLAP. However, if the temporal spacing of the periods is very small, such as a day or an hour, this error becomes negligible. The removal has no associated costs since they are already considered in the picking process.
4. Likewise, all incoming SKUs that are newly available in period 2 will be added in one step. A slotting strategy is used to assign the new SKUs to locations, which can use the order profile as input to calculate intra-period turn-over rates or PF/PA scores. The result of this step is a new assignment, which is fed into period 2. The cycle starts over at step 1.

7.4 M-SLAP Benchmark Data

A benchmark instance for the M-SLAP was generated using anonymized data from the logistics center of an Austrian company in the automotive sector. The data was taken from the high-rack pallet warehouse, which is operated with man-to-goods order picking. In order to comply to confidentiality agreements and to make the dataset more suitable as a benchmark, we simplified the data and environmental constraints in the following way:

- We assume that the warehouse is rectangular with 12 parallel, two-sided aisles, which amounts to 24 identical racks. Each rack is 27 pallets deep and 9 levels high. In total the warehouse can therefore hold 5832 pallets.
- All storage locations are equal, meaning that SKUs can be stored in any location without additional costs. We assume that only one container size (a euro pallet) is used throughout the warehouse.
- The SKUs are retrieved from the warehouse with fork lifts. Aisle changes are the most time-consuming step for this particular scenario, therefore the target cost value is the number of aisle changes required to pick the orders, slot incoming material or perform re-arrangement operations. We assume that all products within one aisle are equally easy to reach.
- Each SKU is stored in only one storage location, independent of the stored amount. There is no mixed storage, meaning that SKUs cannot share a storage location.
- Storage locations can be empty. Due to inventory changes not all SKUs are in stock in all periods. However, it is assumed, that all SKUs that are requested in the test orders of are particular period are in stock in sufficient quantities.
- There is an interim storage area, where incoming/outgoing SKUs are placed, which is located at the head of the first aisle. The pickers always start from and return to the interim storage area.

- One order is picked by one picker (no zoning) in one go (no order splitting or batching).
- Pickers process the individual positions of an order ascending by aisle index.

The M-SLAP benchmark instance consists of five periods, each of which provides an initial (random) storage assignment and a 1-month demand forecast. The entire data set spans one year from December 2010 until December 2011. Over the year, almost 12,000 different SKUs are stored in the warehouse, however in any given period only 4,300 to 5,200 pallets are in stock. A detailed overview of the stock movements is given in Table 7.2 and Figure 7.2 also illustrates the fluctuating inventory levels.

Table 7.2. Inventory levels and incoming/outgoing SKUs between the periods

	Inventory	Outgoing	Incoming
DEC10	4708	-1316	+1696
MAR11	5088	-1199	+1417
JUN11	5306	-1171	+1080
AUG11	5215	-1061	+1199
DEC11	5353		

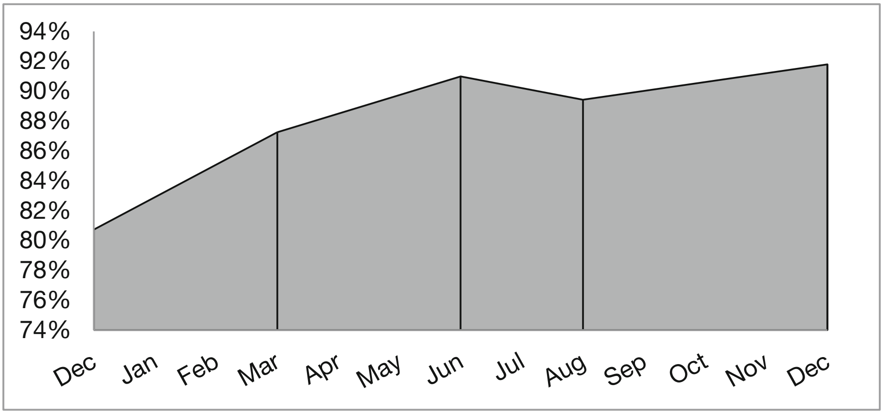

Fig. 7.2. Illustration of the inventory levels per period in the published M-SLAP benchmark instance

The M-SLAP benchmark instance and best known solutions will be published at http://dev.heuristiclab.com/trac/hl/core/wiki/AdditionalMaterial.

7.5 Experimental Setup and Results

7.5.1 Algorithms

All experiments were conducted using HeuristicLab [23], an open source framework for heuristic optimization. We implemented the random and cube per order index (COI) construction methods. In addition, we employed first improvement local search and simulated annealing [12] to optimize a randomly generated initial warehouse configuration subject to the PF/PA objective.

Simulated annealing (SA) is a metaheuristic that was modeled after the annealing process in metallurgy. One of the advantages of simulated annealing is that it offers a strategy to escape from local optima by employing a temperature parameter to guide the search. Contrary to greedy search techniques, which only accept moves that improve the fitness of a solution, SA accepts *uphill* moves with a certain probability. As the algorithm proceeds, a step-wise reduction of the temperature according to a pre-defined cooling scheme reduces the likelihood that bad moves get accepted. The simulated annealing algorithm runs were all configured in the following way: Exponential annealing scheme, start temperature 100, end temperature 1E-06, 4 million iterations.

The same random initial solution was used for all improvement algorithms to make the results comparable. The algorithms could be configured with two different move generators:

- Random Swap Move Generator: Generates one 2-swap move that switches the content of two random locations in the warehouse.
- Sampling Swap Move Generator: Generates n 2-swap moves, sorts them by quality and returns the best move.

We also varied the α and β parameters in the PF/PA objective function. Setting $\alpha = 1$ and $\beta = 0$ results in a turnover-based objective function similar to COI. Setting $\alpha = 0$ and $\beta = 1$ results in an optimization of part affinity only, which reduces the within-order distances but does not consider SKU turnover rates. Finding a good trade-off between placement by part affinity and placement by retrieval frequency is unfortunately not trivial. The parameters α and β can be adjusted automatically to ensure that both factors contribute equally to the objective as proposed in [20]. However, we found that empirically sampling the parameters produced better results as described in detail in [13]. We fixed $\beta = 1$ and conducted tests for $\alpha \in \{1, 2, 3, ... 20, 30, 40, 50, 60\ 70, 80, 90, 100\}$. For the investigated scenario a setting of $\alpha = 10$ and $\beta = 1$ was found to be most effective. All optimization runs were conducted in a high performance computing environment on an 8-core machine with 2x Intel Xeon CPU, 2.5 Ghz and 32GB memory.

7.5.2 Results

7.5.2.1 Re-warehousing

First we employed re-warehousing on each period separately to assess the maximum optimization potential. As shown in Table 7.3 and Figure 7.3, optimizing by part affinity alone (PA) did not produce very good results, only reducing the total travel distance by 11% to 21%, depending on the observed period, compared to an initial random slotting. Although part affinity slotting optimizes the number of aisle changes within an order, the placement relative to the shipping dock is not optimized. Conversely, the optimization potential for turn-over based or mixed approaches is quite large and reduced the number of aisle changes in picking by up to 97%.

Table 7.3. Re-warehousing with different slotting strategies can reduce the pick effort (given in aisle changes) significantly

	Random	PA	COI	PF	PF/PA
Dec 10	9,878	8,746	218	248	214
Mar 11	17,024	13,386	452	474	440
Jun 11	12,672	10,500	410	454	416
Aug 11	15,138	12,080	372	402	352
Dec 11	16,596	14,030	602	618	572

The COI construction heuristic and simulated annealing with pick frequency (PF) as target performed equally well, which is not surprising, since no weight information was available for the data set. In this case, slotting by COI and slotting by pick frequency reduce to the same objective function. A properly parameterized simulated annealing run with the pick frequency objective therefore converged towards the quality obtained with the COI construction heuristic. The combined PF/PA slotting performed marginally better than turnover based slotting (PF, COI) in all periods except one.

Although the improvements seem vast, all of the above approaches require a complete re-warehousing, meaning that no SKU is on its initial place in the target assignment. Table 7.4 lists the re-warehousing efforts to move all SKUs from their initial positions in the random slotting scenario to the respective target assignments, as generated by the different approaches. The re-warehousing effort assessment was conducted under idealized conditions, assuming that the two forklifts work in tandem and buffer the pallets in the intermediate storage area. Each forklift retrieves a pallet that needs to be moved to a different aisle, brings it to the intermediate storage area, picks up a waiting pallet and moves it to its target aisle, where - it is assumed - the next pallet is already waiting for transport to the intermediate storage area. Thus the fork lifts never run empty.

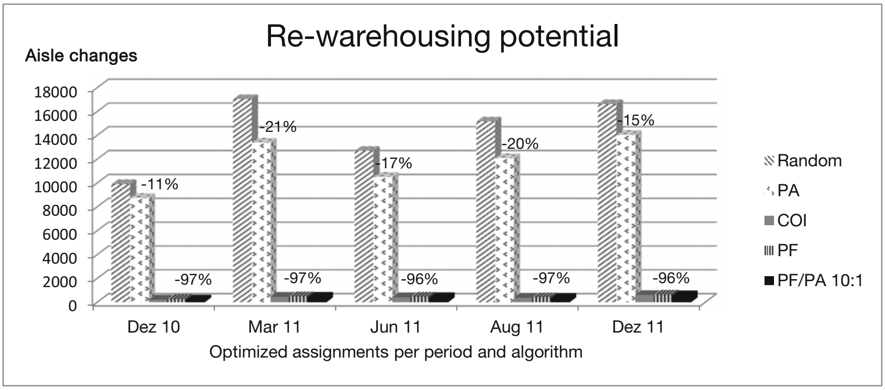

Fig. 7.3. Re-warehousing potential in each period: The number of picker aisle changes can be reduced by up to 97% compared to the initial random slotting, if a complete re-arrangement is conducted. (Data from Table 7.3)

Table 7.4. Number of aisle changes required to re-arrange the assignments from table 7.3

	PA	COI	PF	PF/PA
Dec 10	75,506	57,786	75,436	75,734
Mar 11	81,612	63,594	81,944	81,716
Jun 11	85,078	67,646	86,030	85,558
Aug 11	83,326	65,298	83,930	84,158
Dec 11	86,022	68,166	85,618	86,262

Table 7.4 shows that re-warehousing with PF, PA and PF/PA require approximately the same amount of effort in terms of aisle changes to realize the new assignment. Realizing a COI slotting is ’cheaper’ but nevertheless the involved effort cannot be redeemed quickly. In each period we utilize an order picking preview of roughly one month. If we interpolate this data, it would take roughly five months for a complete COI re-warehousing to pay off - that is, if the order profiles do not change in the meantime. In the Section 7.5.2.2 we will therefore discuss a more efficient way to increase order picking efficiency by simply switching to a different putaway strategy for incoming SKUs.

7.5.2.2 Putaway

Since a complete re-warehousing is very time-consuming, the second set of experiments focused on the effect that different putaway strategies for incoming goods have on an existing layout. We used the assignments generated with random, COI or PF/PA slotting for the first period in Section 7.5.2.1 as starting point and observed how the warehouse performance develops when newly arriving SKUs are slotted randomly or according to COI. In these experiments we do not only consider the order picking efforts but also the putaway efforts for newly arriving SKUs between two periods. Re-warehousing efforts, on the other hand, are not considered, because they have already been investigated in the previous section.

Table 7.5. How the quality (= number of aisle changes) of an initial random, COI or PF/PA slotted warehouse develops over time when random or COI slotting strategies are used for incoming material

	Rand-Rand	Rand-COI	COI-Rand	COI-COI	PF/PA-Rand	PF/PA-COI
DEC 10	9,878	9,878	218	218	214	214
PUT1	18,716	12,214	24,766	18,612	20,920	14,890
MAR 11	17,868	13,642	15,058	6,578	12,782	6,996
PUT2	14,978	14,284	17,344	16,182	17,412	16,944
JUN 11	13,408	10,308	11,404	5,582	10,502	5,592
PUT3	11,860	8,998	12,420	9,690	13,166	10,820
AUG 11	14,868	10,654	12,076	5,696	11,422	5,802
PUT4	11,774	11,504	12,750	12,638	14,626	14,580
DEC 11	15,370	11,826	13,648	8,208	14,208	8,018
Total	128,720	103,308	119,684	83,404	115,256	83,860

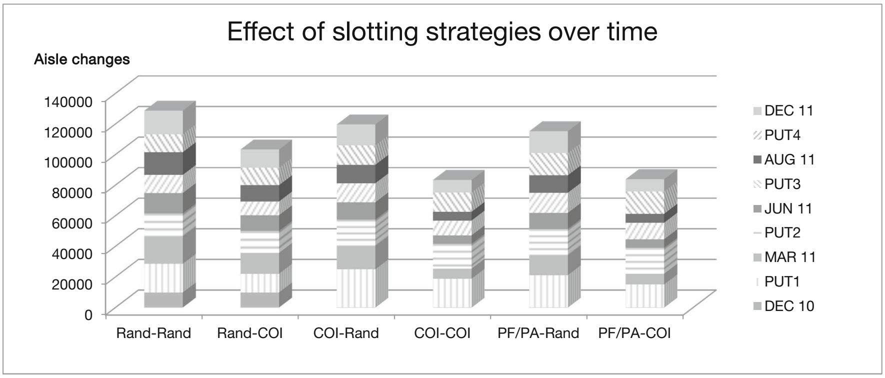

Fig. 7.4. Effects of putaway slotting strategies on an initial assignment (Illustration of Table 7.5)

Table 7.5 and Figure 7.4 summarize the results and illustrate how random slotting can defragment a warehouse after a while. For instance, in December 2010 almost 10,000 aisle changes are required to fulfill the demand forecast on a randomly slotted warehouse but less than 220 on a recently PF/PA or COI slotted warehouse. However after one year of random putaway for incoming SKUs the effect of re-warehousing has almost evaporated for the COI-Rand and PF/PA-Rand scenarios. The total annual number of aisle changes for these scenarios are 115,256 and 119,684, which amounts to a relative improvement of 10.5% and 7% compared to Rand-Rand. Most importantly, the initial re-warehousing efforts (cf. Section 7.5.2.1, Table 7.4) are not yet included in these numbers. If these are considered as well, COI-Rand and PF/PA-Rand actually perform worse than Rand-Rand. The conclusion is, that re-warehousing once per year does not pay off for this M-SLAP benchmark if random slotting is employed for putaway during the year.

Before the second set of results with COI as putaway stategy is discussed, we would like to point out a singularity of the test data set. We claim that order picking is the most effort-intensive process in the warehouse, yet the putaway effort in Table 7.5 is frequently larger than the picking effort. This can be easily explained: As discussed in Section 7.4 the available benchmark data only provides a 1-month demand forecast for each period and not all picking orders between the snapshots. The intervals between periods range between two and four months, therefore the pick efforts are roughly twice to four times as large in reality.

Figure 7.4 illustrates the same results as a graph, showing pick and putaway efforts in turns. One interesting result that the graph shows very well is that the putaway efforts between period 1 and period 2 (PUT1) are much larger for well slotted warehouses than for a randomly slotted warehouse. This can be explained in the following way: Both the COI construction heuristic and the PF/PA slotting algorithm will result in assignments where all the best locations are occupied with SKUs. An illustration of this effect can be seen in Figure 7.5 for a warehouse where the effort is calculated as distance (in meters) from the shipping dock.

We consider this overfitting, because all SKUs drift towards the good storage locations, and the less desirable storage locations are left empty. For incoming material, this means that hardly any storage locations are available, even if the newly arriving SKU has a very high pick frequency. Therefore, putaway efforts are very large directly after COI or PF/PA re-warehousing.

Figure 7.6 also summarizes the results of Table 7.5 but aggregates the putaway and pick efforts. Once again, it should be noted that the total order picking efforts over the entire year are larger since only forecasts for five months were available. It can be seen that an initial random slotting benefits most from a switch to COI slotting. Also, COI-COI and PF/PA-COI yield the overall best results, but one would have to consider if the initial re-warehousing efforts required would be worth the additional efforts. In Section 7.5.2.3 we therefore use healing strategies as another measure to improve the annual warehouse efficiency without the prohibitive costs of re-warehousing.

7.5.2.3 Healing

In the last set of experiments we investigated how conducting a small number of healing moves at the beginning of a period can impact the total warehouse performance. The goal of this additional step is to address changing demand patterns and re-slot SKUs already present in the warehouse according to the demand forecast

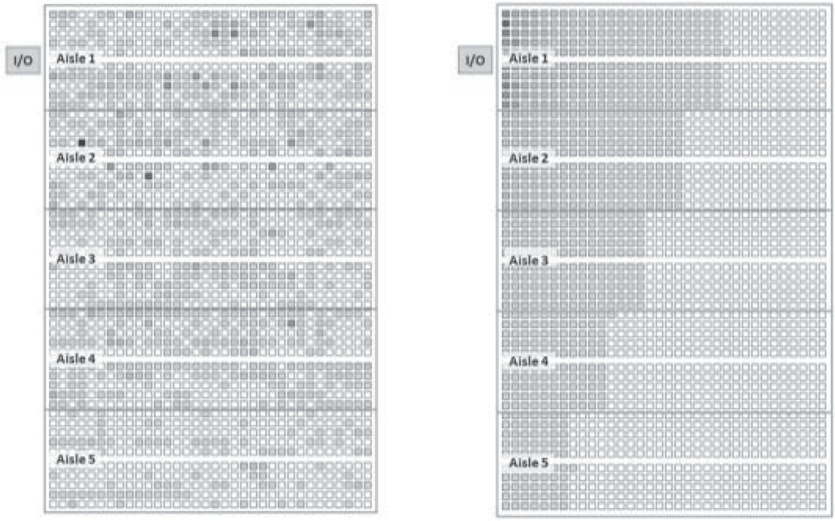

Fig. 7.5. In this illustration, the storage locations in a bird's eye view of a warehouse are colored according to the pick frequency of the assigned SKUs. The white locations indicate empty storage locations or locations with SKUs that are never ordered. An initial random slotting (left) is compared to an (turnover) optimized slotting (right). In the optimized assignment, all the closest locations to the shipping dock in the upper left are occupied.

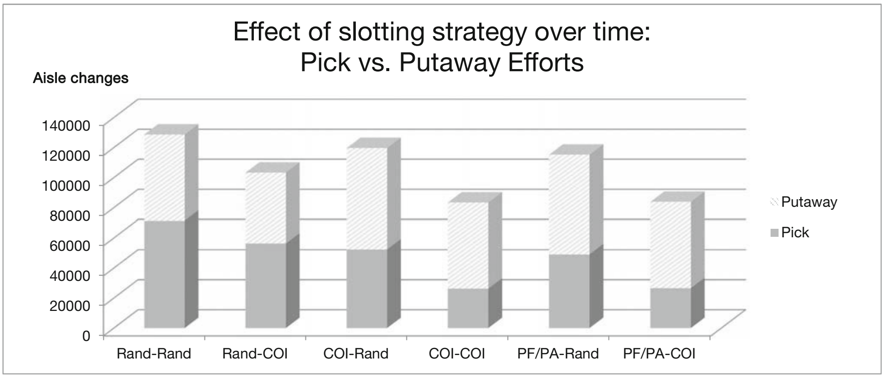

Fig. 7.6. Effects of putaway slotting strategies on an initial assignment: Putaway vs. pick efforts (Illustration of Table 7.5)

for the upcoming period. These experiments therefore implement the full M-SLAP process as described in Section 7.3.3 and illustrated in Figure 7.1.

We used two of the scenarios presented in the previous section as a starting point:

- Rand-Rand: The first period was slotted randomly and a random slotting strategy was used for putaway between subsequent periods.
- Rand-COI: Once again, the first period was slotted randomly, but putaway was conducted with the COI strategy between periods.

A local search algorithm was started at the beginning of period 2-5 to find 50 good 2-swap moves. In each of the 50 iterations the best of 10,000 randomly generated moves was applied. The re-arrangement effort was once again considered in the total efficiency assessment.

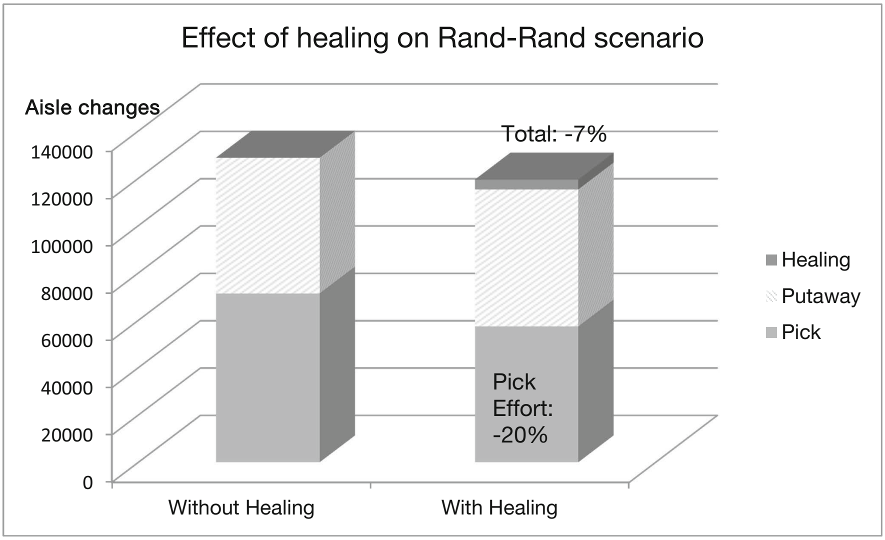

Fig. 7.7. Warehousing efforts for the Rand-Rand scenario from Section 7.5.2.2 without (left) and with healing (right)

Figure 7.7 shows the results for the Rand-Rand scenario with and without healing. It can be seen that healing reduced the number of aisle changes in order picking by 20%. However, the total improvement is only 7% compared to a scenario without healing, because the cleanup swaps require some effort themselves and the putaway effort rises slightly. This can be attributed to the overfitting phenomenon described in the previous section: Due to healing fewer ’good’ storage locations will be available for putaway.

This effect is even more pronounced in the Rand-COI scenario as depicted in Figure 7.8. Once again, the order picking effort could be significantly reduced (-19%) compared to the same scenario without healing. However, simultaneously the putaway effort increases by 13% and when the healing effort is considered the total gain amounts to less than 1%.

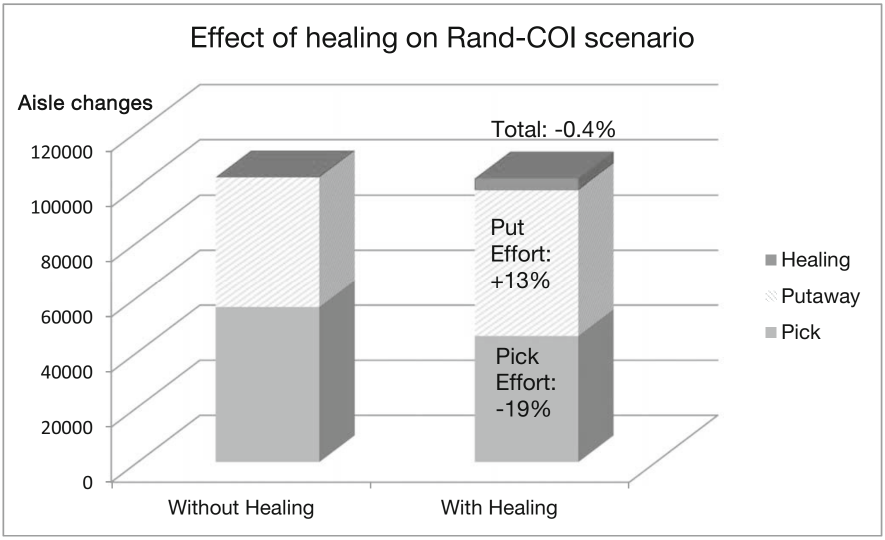

Fig. 7.8. Warehousing efforts for the Rand-COI scenario from Section 7.5.2.2 without (left) and with healing (right)

The presented experiments are not extensive enough to offer guidance for the 'best' M-SLAP strategy, not even for the given scenario. However, they illustrate clearly that focusing solely on pick efficiency improvements offers a very biased warehouse efficiency assessment, especially if real-world warehouses should be optimized during on-going operations.

7.6 Conclusion and Outlook

In this paper we introduced a multi-period formulation of the classic storage location assignment, called M-SLAP, and released a first benchmark instance for the problem. The M-SLAP consists of multiple planning periods with different inventory levels and SKU demand patterns. A new objective function does not only assess order picking efforts but also considers re-arrangement and putaway efforts over all periods. Some preliminary results for the benchmark instance were given, illustrating how different initial slottings (random, COI and PF/PA) develop over time given a random or COI putaway strategy. Moreover, the impact of conducting a small number of healing moves in each period was investigated. First results indicate that considering only order picking efforts is a very limited view on the problem. Moreover, the dynamic change of demand patterns over time makes it even more crucial to consider re-arrangement efforts and putaway efforts. In dynamic warehousing scenarios an 'optimal' slotting can degenerate very quickly, therefore the effort invested into re-arrangements must be weighed against both robustness and efficiency gains.

We plan to conduct comprehensive tests on the presented M-SLAP instance to find algorithms and reference settings that perform well on such a dynamic scenario. We also hope to encourage other researchers to model and optimize M-SLAP instances, since we believe that the dynamic multi-stage scenario offers many new algorithmic challenges and also facilitates transfer of research results into practice. Finally, we are currently conducting M-SLAP tests in the real-world logistics center of a project partner, combining putaway, re-warehousing and healing strategies to find a balanced strategy that is both practical as well as 'optimal'.

Acknowledgements. This paper is an updated and extended version of [15]. The work described in this chapter was done within the Josef Ressel-Centre HEUREKA! for Heuristic Optimization sponsored by the Austrian Research Promotion Agency (FFG). HeuristicLab is developed by the Heuristic and Evolutionary Algorithm Laboratory[1] and can be downloaded from the official HeuristicLab homepage[2].

References

1. Balakrishnan, J., Cheng, C.: Dynamic layout algorithms: a state-of-the-art survey. Omega 26(4), 507–521 (1998)
2. Bartholdi, J.J., Hackman, S.T.: Warehouse and distribution science (2010), Textbook available at `http://www.warehouse-science.com` (accessed September 22, 2012)
3. de Koster, R., Le-Duc, T., Roodbergen, K.J.: Design and control of warehouse order picking: A literature review. Eur. J. Oper. Res. 182, 481–501 (2007)
4. Frazelle, E.: Stock location assignment and order picking productivity. PhD thesis, Georgia Institue of Technology (1990)
5. Frazelle, E., Sharp, G.: Correlated assignment strategy can improve order-picking operation. Ind. Eng. 4, 33–37 (1989)
6. Garfinkel, M.: Minimizing multi-zone orders in the correlated storage assignment problem. PhD thesis, School of Industrial and Systems Engineering, Georgia Institute of Technology (2005)
7. Gu, J., Goetschalckx, M., McGinnis, L.F.: Research on warehouse operation: A comprehensive review. Eur. J. Oper. Res. 177(1), 1–21 (2007)
8. Hausman, W., Schwarz, L., Graves, S.: Optimal storage assignment in automatic warehousing systems. Manage Sci. 22(6), 629–638 (1976)
9. Heskett, J.: Cube-per-order index - a key to warehouse stock location. Transport Distrib. Manage 1963, 27–31 (1963)
10. Kallina, C., Lynn, J.: Application of the cube-per-order index rule for stock location in a distribution warehouse. Interfaces 7(1), 37–46 (1976)

[1] `http://heal.heuristiclab.com/`

[2] `http://dev.heuristiclab.com/`

11. Kim, B., Smith, J.: Dynamic slotting for zone-based distribution center picking operation. In: 10th International Material Handling Research Colloquium, Dortmund, Germany, pp. 577–599 (2008)
12. Kirkpatrick, S., Gelatt, C.D., Vecchi, M.P.: Optimization by simulated annealing. Science 220, 671–680 (1983)
13. Kofler, M., Beham, A., Wagner, S., Affenzeller, M., Reitinger, C.: Rassigning storage locations in a warehouse to optimize the order picking process. In: Proceedings of the 22th European Modeling and Simulation Symposium (EMSS 2010), Fez, Morocco (2010)
14. Kofler, M., Beham, A., Wagner, S., Affenzeller, M., Achleitner, W.: Re-warehousing vs. healing: Strategies for warehouse storage location assignment. In: Proceedings of the IEEE 3rd International Symposium on Logistics and Industrial Informatics (Lindi 2011), Budapest, Hungary, pp. 77–82 (2011)
15. Kofler, M., Beham, A., Wagner, S., Affenzeller, M., Achleitner, W.: The multi-period storage location assignment problem. In: Proceedings of IEEE APCAST 2012 Conference, Sydney, Australia, pp. 38–41 (2012)
16. Malmborg, C.: Storage assignment policy tradeoffs. Int. J. Prod. Res. 33, 989–1002 (1996)
17. Mantel, R., Schuur, P., Heragu, S.: Order oriented slotting: a new assignment strategy for warehouses. Eur. J. Ind. Eng. 1(3), 301–316 (2007)
18. Neuhäuser, D., Wehking, K.H.: Der Lagerorganisationsgrad als Steuerungsgröße für optimale Reorganisationszyklen in Kommissioniersystemen. Logist J. Proc. 7, 1–11 (2011), doi:10.2195/LJ_proc_neuhaeuser_de_201108_01
19. Petersen, C.G., Siu, C., Heiser, D.R.: Improving order picking performance utilizing slotting and golden zone storage. Int. J. of Oper. Prod. Manage 25(10), 997–1012 (2005)
20. de Ruijter, H., Schuur, P.C., Mantel, R.J., Heragu, S.S.: Order oriented slotting and the effect of order batching for the practical case of a book distribution center. In: Proceedings of the 2009 International Conference on Value Chain Sustainability, Louisville, Kentucky (2009)
21. Tompkins, J., White, J., Bozer, Y., Frazelle, E., Tanchoco, J., Trevino, J.: Facilities planning. Wiley, New York (1996)
22. Waescher, G.: Supply chain management and reverse logistics. In: Order Picking: A Survey of Planning Problems and Methods, pp. 323–347. Springer, Heidelberg (2004)
23. Wagner, S.: Heuristic optimization software systems - modeling of heuristic optimization algorithms in the heuristiclab software environment. PhD thesis, Institute for Formal Models and Verification, Johannes Kepler University, Linz, Austria (2009)
24. Wutthisirisart, P.: Relation-based item slotting. Master's thesis, University of Missouri (2010)