

Improving order picking response time at Ankor's warehouse

R. Dekker¹, M.B.M. de Koster², K.J. Roodbergen², H. van Kalleveen¹

¹ Faculty of Economics, Erasmus University Rotterdam, The Netherlands

² Rotterdam School of Management, Erasmus University Rotterdam, The Netherlands

Please refer to this article as:

**Dekker, R., De Koster, M.B.M., Roodbergen, K.J., and Van Kalleveen, H. (2004),
Improving order-picking response time at Ankor's warehouse, *Interfaces* 34(4), 303-313.**

Abstract:

Ankor is a wholesaler of tools and garden equipment, carrying such well-known brands as Skandia. Its warehouse is under continuous pressure to improve its efficiency while it is confronted with several special requirements in its order picking, like the requirement to retrieve heavy products first to prevent damage to other, breakable products. Our research goal was to determine a good combination of policies for storage assignment (assigning products to storage locations) and routing (determining the sequence in which to retrieve products from storage to meet customer demand) for Ankor's situation. We adapted existing solution techniques for this problem, overcoming the special characteristics of Ankor's operations. With these adapted techniques, we cut the average route length in the order-picking operation by 31 percent. As a result of our study, Ankor implemented a new storage and routing strategy. The study also showed further improvement potentials in the picking process, which Ankor adopted as well. All improvements led to a reduction in the number of order pickers of more than 25 percent.

Of the activities conducted in warehouses, many simultaneously affect a large number of units (for example, a full pallet). In order picking however, warehouses often handle single units to satisfy customer demand. It is typically one of the most time-consuming activities and a large contributor to the operational costs in a warehouse (Tompkins et al. 2003). Clearly improving order picking is an important way to save costs. Furthermore, because customers nowadays tend to order late and expect quick delivery, warehouses must improve the efficiency of their order picking.

Factors influencing order-picking efficiency are operating procedures, changing products' demand, the equipment, and the racking and layout of the warehouse. Operating procedures can often improve efficiency without the large capital investments needed to alter racks or equipment. The time required for picking an order includes the times for traveling between product locations, picking products, and such activities as packaging, dropping off products, and acquiring information for the next order. Of these components, travel time is usually the largest (Tompkins et al. 2003) and is therefore the prime candidate for improvement. Travel time is determined largely by two factors: routing policies (for calculating a path that visits all required locations) and storage assignment policies (for assigning storage locations to incoming products).

For a restricted set of layouts, researchers have developed efficient routing policies that find a shortest route (Ratliff and Rosenthal 1983, De Koster and Van der Poort 1998, Roodbergen and De Koster 2001a). For other layouts, one has to resort to methods that give optimal solutions, but not in polynomial time, such as branch-and-bound (see Little et al. 1963) or to heuristics. Heuristics quickly create a picking sequence, which is, however, not guaranteed to be the shortest route possible. Besides of their flexibility for various layouts, heuristics are easy to work with and to implement in warehouse-management software. Heuristics can also easily be adjusted to changes in layout and to predetermined picking priorities (breakable products last). Therefore, warehouses generally prefer heuristics over optimal routing. Petersen (1997), Petersen and Schmenner (1999), and Roodbergen and De Koster (2001b) discuss several heuristics.

Firms use storage-assignment policies to assign products to storage locations in the warehouse. Three common policies are random storage, dedicated storage, and grouped storage (Hausman et al. 1976). Random storage assigns each product to an arbitrarily selected

empty storage location in the warehouse. Dedicated storage assigns each product to a fixed location. Grouped storage assigns products based on product type, physical product characteristics, number of picks, or any combination of these three characteristics. ABC storage is a well-known example of grouped-storage assignment. Firms classify as A, B, or C products based on their pick frequency and assign each product group to a certain area of the warehouse (Malmborg 1995).

The extensive research on these two subjects concerns mainly routing policies or storage policies. Daniels et al. (1998) deal with the combination of routing and storage assignment. They consider a warehouse with multiple locations for each product. Petersen (1999) uses simulation to determine good combinations from a set of predetermined storage-assignment policies and routing methods.

Ankor B.V.

Ankor is a wholesaler in Leidschendam, the Netherlands, that specializes in tools, hardware, and gardening equipment. In its warehouse, it stores over 18,000 products with diverse characteristics. Based on these characteristics it categorizes products in three groups: (1) the 384 fastest-moving products, (2) products whose longest edge is over 80 centimeters, called *non-conveyables* (Nocos), and (3) regular products that do not fall within the other two groups. Its customers are primarily retail store chains, mainly in the do-it-yourself business and garden centers.

XXXXXXXXXXXXXXXXXXXXXXXXXXXX

INSERT FIGURE 1

XXXXXXXXXXXXXXXXXXXXXXXXXXXX

Each product group has its own storage area (Figure 1). We restrict ourselves here to the largest group of regular products, which contains 17,000 of the total of 18,000 products. Storage of this group used to be based on an ABC-classification. Warehouse managers used the head and the tail of each aisle on the ground floor to store A-products (the fastest-moving of the regular products). They used the middle sections of the aisles on the ground to store B-products (medium-moving products). They stored C-products (slow-moving products) on the

mezzanine. The customers, however, want the firm to deliver products separated according to product type, that is, tools, hardware, and gardening equipment. To do this, the warehouse managers developed a storage policy calling for the storage of regular products to be based on product type. Each product type is now stored in its own specified area. Breakable and unbreakable products have separate sections in the picking area. By retrieving the breakable products last, pickers minimize damage to products. Within these categories managers make no further classification based on demand frequency. However, since the warehouse did not move most of the unbreakable products to make this change and since demand for many products has remained fairly stable for years, some products are still stored according to the former ABC-storage method (that is, fast-moving products are often at the heads and tails of aisles, and slow-moving products are in the middle). This effect has diminished though, since managers no longer relocate products if demand changes and they no longer locate new products by demand frequency.

For each product type within the regular products category, we determined the total number of products, the number of breakable products, the percentage of picks (out of a total of 956,967 picks per year), and the percentage of storage locations needed (Table 1).

Regular products	Total number of products	Breakable products	Picks percentage	Percentage of storage locations
Tools	8538	338	68.3	38.0
Hardware	8150	529	24.2	47.1
Gardening	859	325	7.5	14.9
Total	17547	1192	100	100

Table 1: Subdivision of the regular products in tools, hardware, and gardening equipment.

Most of the picks are for tools. (All of the fast-moving products are also tools.) Although the hardware category contains fewer products than the tools category, it requires more storage locations because hardware products are generally stocked in greater numbers and are larger. Ankor introduced the gardening equipment in 1999; therefore it generates only 14.9 percent of the total picks. The breakable products in each product type must be stored together so employees can retrieve them at the end of a route.

Our research objective

Early 2001, Ankor's management realized that order picking is a major contributor to its operational costs, and travel time for order picking time depends largely on operational procedures. Ankor asked us to look into the order picking process (the routing), hoping for efficiency improvement. We decided to investigate possibilities for improving the operational procedures, in particular combinations of routing and storage policies.

Although the literature includes research on routing and storage assignment, the situation at Ankor differs from the warehouse layouts and general assumptions described in the literature. Ankor's layout includes multiple cross aisles, two floors, noncoinciding start and end points (on the mezzanine), and some dead-end aisles (Figure 1). For such situations, we found no standard operating policies in the literature. Furthermore, Ankor ships different product types that may not be mixed in the packing process (leading to restrictions on the routing during order picking). Each product type also includes breakable products that must be picked last in the route.

In fact, such restrictions and nonstandard layouts are fairly common in practice, but we found no publications on combinations of storage policy and routing policy for these situations.

In the remainder of this paper we will first discuss alternative routing policies. Based on characteristics of these routing policies we define and evaluate storage policies. We will end this paper with an overview of the implementation.

Routing heuristics

We created a simulation model for the warehouse area containing regular products in the programming language Visual-C++. This model generates picking locations from an order and calculates the length of the route created by a number of routing policies. The input is based on real orders and real picking frequencies for the products in weeks 6 and 7 of 2001. The data set consists of 1,098 orders with a total of 27,790 picks. Orders in the data set vary from 1 to 325 lines with an average order size of 25.31 order lines. The data set is large enough to represent the order pattern. Furthermore, seasonal aspects are hardly present for the regular products. Actually, the warehouse sees seasonal variation mainly for the Nocos group (most of the gardening items), which we did not analyze.

Each replication of the simulation chooses one random order from this set with equal probability. For each line of the chosen order, the program generates a pick location. To generate a location, we look at two factors: (1) the product type (tools, hardware, or

gardening) and (2) the pick frequency of products of this product type. The location generated must be in the warehouse area designated for the product type and the probability selecting a particular location depends only on the pick frequency of the product stored in that location. We based pick frequencies on sales data from the second half of 1999 and the first half of 2000. Because of the huge number (17,000) of locations in the storage area for regular products, we made a simplification in the simulation model. Instead of using actual locations, we used 376 clustered locations, each representing all the locations contained in two meters of an aisle. Thus, the pick frequency for one clustered location equals the sum of all pick frequencies of the products stored in those two meters of an aisle. The number of actual locations contained in a clustered location ranges from 24 to 80 and the average is 45. This variation is due to differences in locations' sizes.

To estimate the average time needed to retrieve an order, we collected information about working speed. We found that order pickers are productive for about seven hours and ten minutes of the working day of eight hours. Their average walking speed is one meter per second (not including time for products and picking them). This speed is fairly high compared to similar operations because employees do much of their walking without a cart. The measured start-up time equals 130 seconds per order and picking time equals 47 seconds per order line. In these 47 seconds, an order picker reads the screen of his or her terminal, identifies the location, retrieves the product from its location, and scans the product, packs the product and scans the box in which the product is packed.

An order-picking route basically consists of traversing all of the cross aisles entirely and entering aisles to pick products when needed (Figure 1). The order pickers leave their carts in the cross aisles when they enter aisles to retrieve the products, generally from the cross aisle closest to the product. In the mezzanine, however, most aisles can be entered from only one end (Figures 1, 2, and 3). On the ground floor, pickers enter aisles on the two sides alternately, while on the mezzanine, they work on one side of the cross aisle at a time.

We considered heuristic routing methods (Figure 2): the midpoint heuristic (Hall, 1993], the largest gap heuristic, the S-shaped (or transversal) heuristic, and the combined heuristic. The midpoint heuristic differs from the traditionally used routing policy in that the picker occasionally travels through an entire aisle to go either to the next cross aisle or to the end point of the route. Thus, for the ground floor, the midpoint heuristic policy is equivalent to the traditional routing policy plus some shortcuts. For the mezzanine, it allows the picker to

retrieve products from both sides of the cross aisle simultaneously.

The largest gap heuristic is the same as the midpoint heuristic, except that the picker goes up an aisle as far as the largest gap instead of to the middle. The largest-gap is the largest part of the aisle without pick locations for this route. For example, in an aisle with four locations on either side, with two items to be retrieved from the two central locations on one side, the picker would enter the aisle from both ends with the midpoint heuristic, but would enter from one end only with the largest-gap heuristic (Figure 3: the second aisle in the middle block from the right on the mezzanine; the largest gap is not between the two adjacent pick locations but at the head of the aisle). The largest-gap heuristic always performs better than the midpoint heuristic (Hall 1993).

The S-shape (or transversal) heuristic (Hall 1993) is quite different from the midpoint and largest-gap heuristics. The order picker travels through a cross aisle to an aisle with a pick location, travels through that aisle to another cross aisle, continues along this cross aisle to the next aisle on the same side with a pick location and travels through this aisle to return to the original cross aisle. This picker repeats this process to gather all the items from one block. The picker then repeats the process for the other blocks. This heuristic is probably the simplest and most frequently used heuristic in practice.

To construct a route with the combined heuristic, one uses dynamic programming. This heuristic is identical to the S-shape heuristic in the sequence in which the picker visits aisles. The combined heuristic, however, decides for each aisle whether the picker should traverse it entirely or return to the cross aisle from which he or she entered the aisle. By using this method, one can look one aisle ahead. For each aisle, the combined heuristic chooses the option (traverse entirely or return) that gives the shortest combination with the next aisle (Figure 4). Roodbergen and De Koster (2001b) give a full description of this algorithm.

One can also calculate optimal routes using a procedure for determining traveling salesman routes (the order-picking problem is a special case of the traveling salesman problem) developed by Jonker and Volgenant (Volgenant and Jonker 1982, Jonker and Volgenant 1984). Pickers can traverse aisles with pick locations entirely, enter and leave from one cross aisle, or enter and leave twice from two cross aisles. The picker can sequence picks in various ways by switching between blocks and aisles to reduce travel distance.

Each of these routing heuristics is an adaptation of an existing heuristic. None of the authors who describe these heuristics discusses sequencing restrictions. However, the Ankor

warehouse has these restrictions: picked products must be sorted by type, and breakable products must be picked last. We can use two approaches to deal with these restrictions. First, we can modify the routing policies so that pickers bypass a product if picking it would violate a restriction. However, we must then incorporate a mechanism to allow the picker to return to that product when picking it no longer violates a restriction. Though possible, this approach would necessitate complex routing policies and it would likely produce only minor gains. Why would one visit a location twice, if one could position the breakable items near the end of the route? The second approach would be to keep the existing routing policies and, based on their behavior, position products in the warehouse to meet the restrictions automatically. For example, we can position breakable products in the right most aisles of each block on the ground floor; order pickers always pass these products at the end of their routes. The optimal algorithm disregards restrictions in calculating the shortest route and therefore serves only as a benchmark with which to compare the viable alternatives.

XXXXXXXXXXXXXXXXXXXXXXXXXXXX

INSERT FIGURE 2

INSERT FIGURE 3

INSERT FIGURE 4

XXXXXXXXXXXXXXXXXXXXXXXXXXXX

To obtain average route lengths, we used our simulation model to generate 10,000 orders. This number of replications leads to small standard deviations in travel distance, implying narrow confidence intervals, so that the results are useful in making comparisons. The calculation times are very modest. Each of the heuristics takes only a fraction of a second to calculate an order-picking route. The optimal algorithm is much slower and calculation times vary; however, it calculates most routes within a few seconds (Table 2).

Routing policy	Average route length on the ground floor (m)	Average route length on the mezzanine (m)	Average total route length per order (m)	Standard deviation of average total length (m)	Difference with traditional method (%)
Traditional	157.05	37.24	194.29	1.53	-
Midpoint	129.41	34.85	164.26	1.34	-15.5
Largest gap	128.59	34.85	163.43	1.33	-15.9
S-shape	147.01	36.92	183.93	1.49	-5.3
Combined	134.64	35.16	169.80	1.35	-12.6
Optimal	119.36	31.18	150.54	1.47	-22.5

Table 2: The average route length on the ground floor and the mezzanine, the total route length and standard deviation of the routing policies, all in meters.

If Ankor were to use the optimal routing, it could cut the average length of a picking tour from 194.29 meters to 150.54 meters. The four heuristics produce various results. The S-shape heuristic performs only slightly better than the firm's traditional routing. By using one of the other heuristics, it could save between 12 and 15 percent, only a little less than it would save using the optimal routing. The midpoint heuristic is especially remarkable because it greatly resembles Ankor's traditional routing policy but achieves great savings. The remains of the storage policy based on an ABC-classification tailored to the firm's traditional routing policy may favor the midpoint heuristic, which resembles the traditional policy. Since the combined heuristic follows a very different concept to construct routes, it is outperformed by the simpler heuristics that resemble Ankor's traditional routing policy.

Storage policies

To improve the performance of the routing policies, we defined several ABC-storage policies, intuitively tailoring each policy to one of the routing heuristics. In doing so we took several constraints into account:

- (1) Product types (tools, hardware, gardening) must be retrieved separately.
- (2) Breakable products must be retrieved after all unbreakable products.
- (3) Some locations must be empty in all zones (ABC zones plus breakables) in which, for example, to store new products.

The first constraint mandates separate areas for the three product types. Each product type includes unbreakable and breakable products. To retrieve the two separately, within each product-type area, the warehouse should reserve a section for storing breakable items. Order

pickers must visit this section at the end of their routes. Changes in the demand will change the structure of the different ABC classes. Therefore, the warehouse must maintain a few buffer locations (empty locations that can be used in the future). These buffer locations take up space and thus increase travel times. If the warehouse had no buffer locations, it might have to put a (new) fast-moving product in a C location, which would greatly increase travel times. For maximum future efficiency we situated some buffer locations at the borders of the different classes, so we could use them for classes on either side.

To make an ABC-storage classification, we examined the regular, unbreakable products (table 1). For the unbreakable products, we generated typical Pareto pick-frequency curves, which we used to subdivide each product type into ABC classes. We subdivided each product type into three classes. We found no hard rules in literature for making such subdivisions. Hausman et al. (1976) and Graves et al. (1977) suggest that the A zone should be very small. However, their studies concern automated storage and retrieval machines that return to the depot after every pickup. Because the machine moves constantly from and to the depot, item placement is skewed accordingly. Situations with multiple aisles and multiple picks per route, such as Ankor's, may do better with large A zones in which the order picker can spend time without returning to the depot or moving to the next zone. Petersen and Schmenner (1999) support this idea, using an A zone of 20 percent of the available picking area and performing experiments with 40, 60 and 80 percent of the picks from the A zone. We decided on the following proportions for the three zones:

- Class A contains up to 70 percent of the total picks and at most 15 percent of the products.
- Classes A and B together contain up to 90 percent of the picks and at most 50 percent of the products.
- Class C contains the remaining products.

We fill class A is filled with products first. Repeatedly we add to class A the product with the highest number of picks that is not yet assigned to a class, until the number of picks in class A reaches 70 percent of all picks or until 15 percent of the products are in class A, whichever comes first. We fill classes B and C with products consecutively in the same fashion.

To calculate the storage space needed for each class, we calculated how many locations are occupied by the products assigned to each class. Based on how much space the products traditionally occupy, we can determine the storage space needed for the class (Table 3).

Class	Tools			Hardware			Gardening		
	percent of products	percent of picks	number of locations	percent of products	percent of picks	number of locations	percent of products	percent of picks	number of locations
A	10	70	23	15	65	31	15	62	7
B	17.5	20	25	26	25	43	27	28	9
C	72.5	10	82	59	10	81	58	10	18
Breakable	N/A	N/A	7	N/A	N/A	13	N/A	N/A	16
Buffer	N/A	N/A	6	N/A	N/A	9	N/A	N/A	6
Total	100	100	143	100	100	177	100	100	56

Table 3: This table shows the subdivision of the product types (tools, hardware, gardening) into class A, class B, class C, breakable products, and buffer locations, based on our storage policy.

We can group the routing heuristics according to the storage policies they require. First are the S-shape and the combined heuristics, which never require the picker to enter an aisle twice. Second are the midpoint and largest-gap heuristics which may require a picker to enter an aisle twice. Third is the optimal algorithm. Based on the three groups we defined two new storage policies. We did not construct a storage policy tailored to the optimal algorithm, because we use it only as a benchmark. We also specified locations for the A, B, and C classes.

For the storage policy covering the S-shape and combined heuristics, which we call the combined storage policy, we ensure that each aisle contains products of only one product type. The reason for this is that the picker enters an aisle only once and retrieves product types separately.

Next we can assign locations to each of the five classes, class A, class B, class C, breakable products, and buffer locations. The routing heuristics specify that the order picker travels through the first aisle of each block at the end of each route. Therefore, we use these aisles with shelf racks to store breakable products. We store products from the A class, the fastest moving products, in aisles closest to the starting point, products in the B class next, and products in the C class in aisles farthest from the starting point. We create proportional buffer locations at the borders between the A, B, and C classes (Figure 5).

XXXXXXXXXXXXXXXXXXXXXXXXXXXX

INSERT FIGURE 5

XXXXXXXXXXXXXXXXXXXXXXXXXXXX

For the storage policy for the midpoint heuristic and the largest gap heuristic, which we call the midpoint storage policy, we defined the areas for the product types by half aisles, because order pickers seldom pass the midpoints of aisles. Ankor’s traditional storage policy defined the areas in the same way.

We assigned breakable products and buffer locations to the same locations we used for the combined storage policy. The preferred locations are the aisles closest to the starting points, and locations at the heads and tails of aisles. The reason is that these routing heuristics call for the order picker to return to the cross aisle from which he or she came, making it time consuming to visit the middles of aisles. Thus we make a trade-off between locations at the heads and tails of aisles and locations in the middle of aisles close to the depot (Figure 6).

For the mezzanine, the two storage policies are identical for the two outer blocks, since the pickers can enter the aisles in those blocks from one side only. For the middle block, they differ since the first group of heuristics has pickers entering each aisle only once, whereas the midpoint and largest-gap heuristics may allow pickers to enter each aisle twice. These storage policies are not optimal, but they respect the properties of the routing heuristics and ensure that A-products in particular are well located. We also conducted an experiment to further improve the location assignment by applying 2-opt.

XXXXXXXXXXXXXXXXXXXXXXXXXXXX

INSERT FIGURE 6

XXXXXXXXXXXXXXXXXXXXXXXXXXXX

We executed the same 10,000 orders with both storage policies, determining the average route length for each routing policy (Table 4).

Routing policy	Traditional storage policy	Combined storage policy	Midpoint storage policy
Traditional	194.29	179.06	169.10
Midpoint	164.26	138.02	138.20
Largest gap	163.43	136.25	137.55
S-shape	183.93	140.36	153.96
Combined	169.80	134.73	141.42
Optimal	150.54	122.40	125.43

Table 4: This table shows the average route length (in meters) of six routing policies when executed in combination with Ankor’s traditional storage policy and the new storage policies.

The best combination was the combined routing heuristic with the combined storage policy having an average route length of 134.73 meters, which is an improvement of 31 percent compared to the average route length of 194.29 meters of Ankor’s traditional methods. The best solution (using optimal routing) suggests a potential for an additional six percent improvement only. Therefore, we have little room for inventing better routing methods.

We questioned whether we could further improve these results by modifying the storage policy. After all, we constructed layouts intuitively by predicting the behavior of the routing heuristics. We designed a 2-opt procedure (Lin and Kernighan 1973) that can interchange zone assignments between storage locations. The improvement possible with this procedure appeared to be less than one percent. A 2-opt procedure does not guarantee optimality, but we could obtain only a marginal improvement with this procedure, which suggests that any possible additional improvements with other storage policies are likely to be minor.

Reducing travel time can increase productivity. Under the traditional circumstances it takes an order picker 1,514 seconds to retrieve an average order of 25.31 lines. This time includes picking time (25.31×47), start-up time (130) and travel time (194.29). Using the best possible heuristic alternative (combined routing together with the combined storage policy), we reduced this time to 1,454 seconds, a decrease of four percent. An important reason why total picking time decreases by only four percent is that the item retrieval time (47 seconds per line) and the order start-up time form a high proportion of total time. This study prompted Ankor to find improvements in these areas as well.

Implementation

As a result of our study, by the end of 2001, Ankor had changed its routing policy and its storage policy. Because of the restriction, that order pickers must retrieve product types separately, it cannot use some routing heuristics at all or cannot use them in combination with a certain storage policy. These infeasible options are the optimal routing method and the largest-gap heuristic, the midpoint heuristic in combination with the combined storage policy, and the S-shape and combined heuristics combined with the midpoint storage policy. These exclusions left the company with three options: the midpoint heuristic with the midpoint storage policy, the S-shape heuristic with the combined storage policy, and the combined heuristic with the combined storage policy. The results of the three remaining options vary between 134.73 and 140.36 walking meters per order. Because the methods differed very little, the company decided to stay as close to its existing methods as possible and therefore to implement the midpoint heuristic with the midpoint storage policy, which should reduce the average route length to 137.60 meters, a reduction of 28.9 percent.

To put the new policy into practice, Ankor had to move many products, a laborious project. By spring 2002, it had finished the first phase. The storage policy it implemented is similar to the midpoint storage policy (Figure 6), although it spread the A-articles on the ground floor out over the heads of a few more aisles. This option reduced congestion in the first aisles. Ankor bought new picking carts that are easier to maneuver within the aisles.

Besides this implementation, Ankor has made several further improvements to reduce the picking time. With Ankor's traditional methods, picking and administration time consumed 76 percent of the total time, which consists of travel, picking, and administration time. Reducing travel time makes picking time a higher percentage of the total time, which asks for attempts to reduce picking time as well. Ankor has reduced picking time by adopting scanners with a shorter response time than the old scanners, which could take several seconds to confirm picks, by picking individual units rather than packages containing several units that had to be unpacked at receiving, and by calculating box sizes accurately, saving on packing time by the pickers. Since our study, Ankor has increased the number of products it carries from 18,000 to 19,000 and the percentage of picks in the tools area from 68 to 80 percent (tools take somewhat less pick time than gardening equipment).

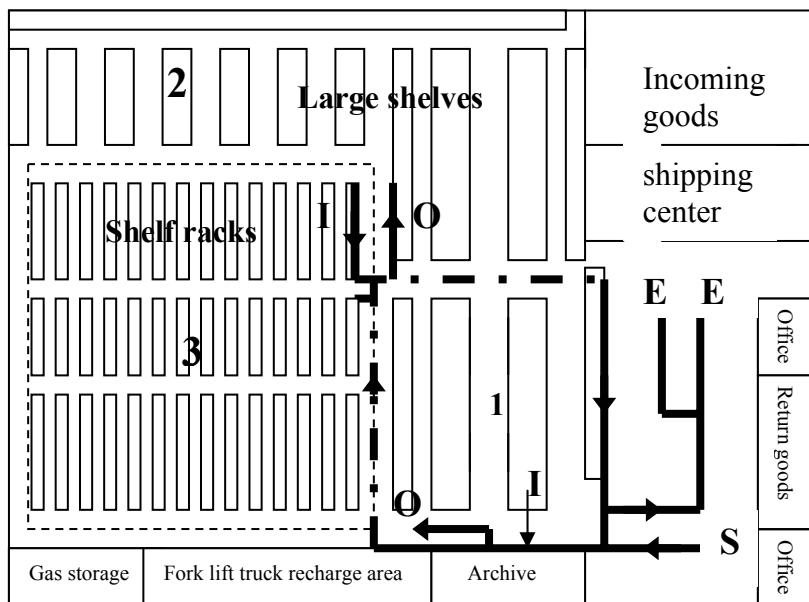
Because of all these changes, Ankor has reduced the number of pickers it employs from 20 in

2001 to 12 to 15 (depending on demand peaks) in 2002, which saves the company about € 140,000. It is difficult to determine how much of this reduction to attribute to the new storage and picking policies. They played an important role, but the order-picking and administration times also improved as a side effect of our research. Often in practice, thoroughly examining a process highlights opportunities for improvement.


Literature

- Daniels, R.L., J.L. Rummel, R. Schantz. 1998. A model for warehouse order picking. *European Journal of Operations Research* 105(1) 1-17.
- De Koster, R., E.S. Van der Poort. 1998. Routing orderpickers in a warehouse: A comparison between optimal and heuristic solutions. *IIE Transactions* 30(5) 469-480.
- Graves, S.C., W.H. Hausman, L.B. Schwarz. 1977. Storage-retrieval interleaving in automatic warehousing systems. *Management Science* 23(9) 935-945.
- Hall, R.W. 1993. Distance approximations for routing manual pickers in a warehouse. *IIE Transactions* 25(4) 76-87.
- Hausman, W.H., L.B. Schwarz, S.C. Graves. 1976. Optimal storage assignment in automatic warehousing systems. *Management Science* 22(6) 629-638.
- Jonker, R., A. Volgenant. 1984. Nonoptimal edges for the symmetric traveling salesman problem. *Operations Research* 32(4) 837-846.
- Lin, S., B.W. Kernighan. 1973. An effective heuristic algorithm for the traveling salesman problem. *Operations Research* 21(2) 498-516.
- Little, J.D.C., K.G. Murty, D.W. Sweeney, C. Karel. 1963. An algorithm for the traveling salesman problem. *Operations Research* 11(6) 972-989.
- Malmberg, C.J. 1995. Optimization of cube-per-order index warehouse lay-out with zoning constraints. *International Journal of Production Research* 33(2) 465-482.
- Petersen, C.G. 1997. An evaluation of order picking routing policies. *International Journal of Operations and Production Management* 17(11) 1098-1111.
- Petersen, C.G. 1999. The impact of routing and storage policies on warehouse efficiency. *International Journal of Operations and Production Management* 19(9-10) 1053-1064.
- Petersen, C.G., R.W. Schmenner. 1999. An evaluation of routing and volume-based storage policies in an order picking operation. *Decision Sciences* 30(2) 481-501.
- Ratliff, H.D., A.S. Rosenthal. 1983. Orderpicking in a rectangular warehouse: A solvable case

- of the traveling salesman problem. *Operations Research* 31(3) 507-521.
- Roodbergen, K.J., R. De Koster. 2001a. Routing order pickers in a warehouse with a middle aisle. *European Journal of Operational Research* 133(1) 32-43.
- Roodbergen, K.J., R. De Koster. 2001b. Routing methods for warehouses with multiple cross aisles. *International Journal of Production Research* 39(9) 1865-1883.
- Tompkins, J.A., J.A. White, Y.A. Bozer, J.M.A. Tanchoco. 2003. *Facilities Planning*, John Wiley and Sons, New York.
- Volgenant, A., R. Jonker. 1982. A branch and bound algorithm for the symmetric traveling salesman problem based on the 1-tree relaxation. *European Journal of Operational Research* 9(1) 83-89.




Main floor




Mezzanine

*

Figure 1: This layout shows the storage areas for fast-moving products (marked 1), for Nocos (marked 2) and for regular products (marked 3), both for the main floor and the mezzanine. The dashed line in the left image indicates the mezzanine's location. The aisles in both outer storage blocks with racks on the mezzanine can only be entered from one side (the middle of the floor), as the other side of the aisles is fenced. The bold line is a conveyor used in order picking. Pick orders are assigned to pick bins that start on the conveyor at position S. A picker collects the bins at the conveyor output station assigned to his or her storage area (marked O) and drops them, after picking, at a conveyor input station (marked I). The conveyor brings the completed pick bins to the shipping area (E).




Main floor




Mezzanine

Figure 2: In this example of the traditional routing policy for Ankor’s warehouse’s main floor and mezzanine, pick locations are shown as black rectangles. The starting points of routes are labeled S.



Main floor



Mezzanine

Figure 3: In this example of the midpoint routing heuristic on the Ankor warehouse main floor and mezzanine, we depict pick locations as black rectangles and we indicate the starting points of routes with S. Largest gap differs only slightly from this heuristic: pick-up items distributed around an aisle middle will usually be picked up starting from one of the cross aisles only.

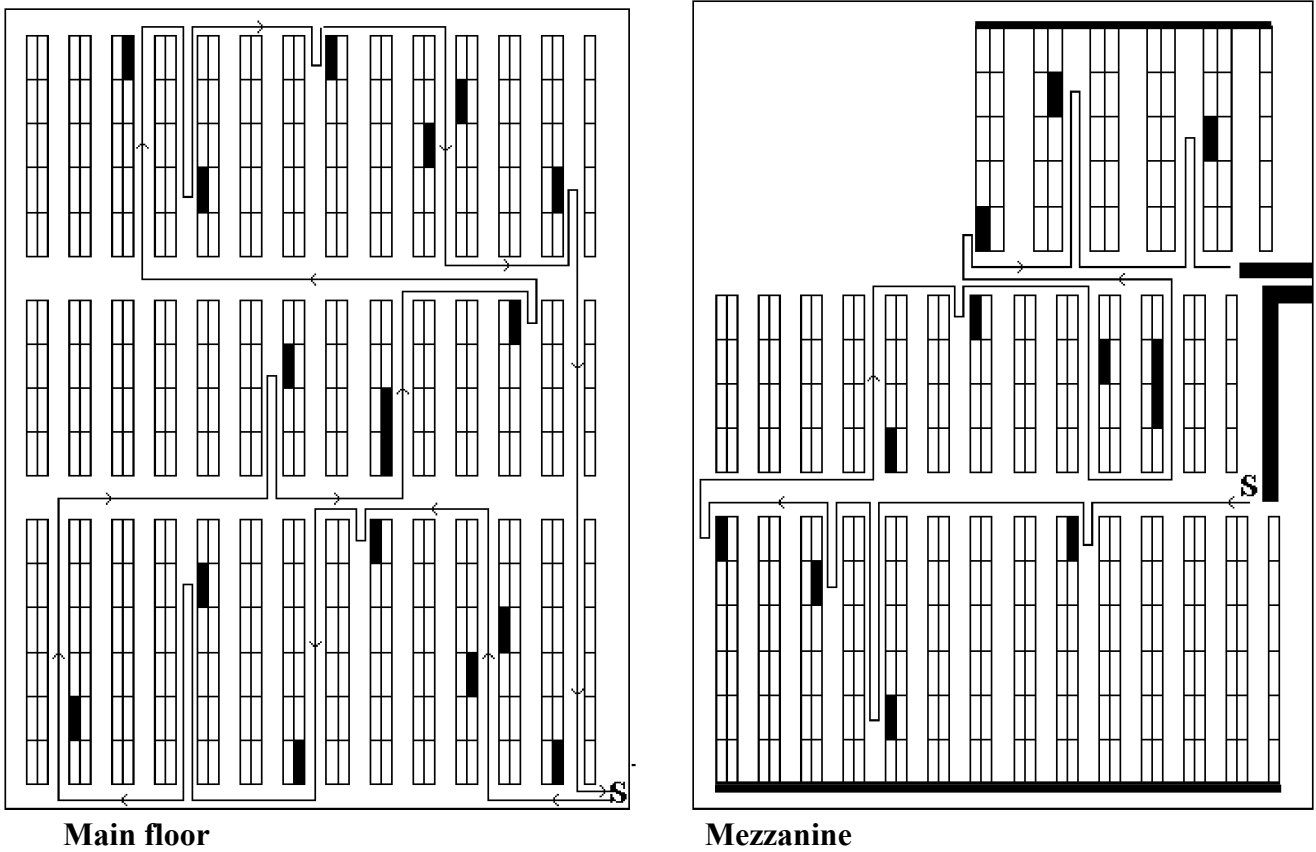


Figure 4: In this example of the combined routing heuristic for Ankor’s warehouse main floor and mezzanine, we show pick locations as black rectangles and starting points of routes as S.

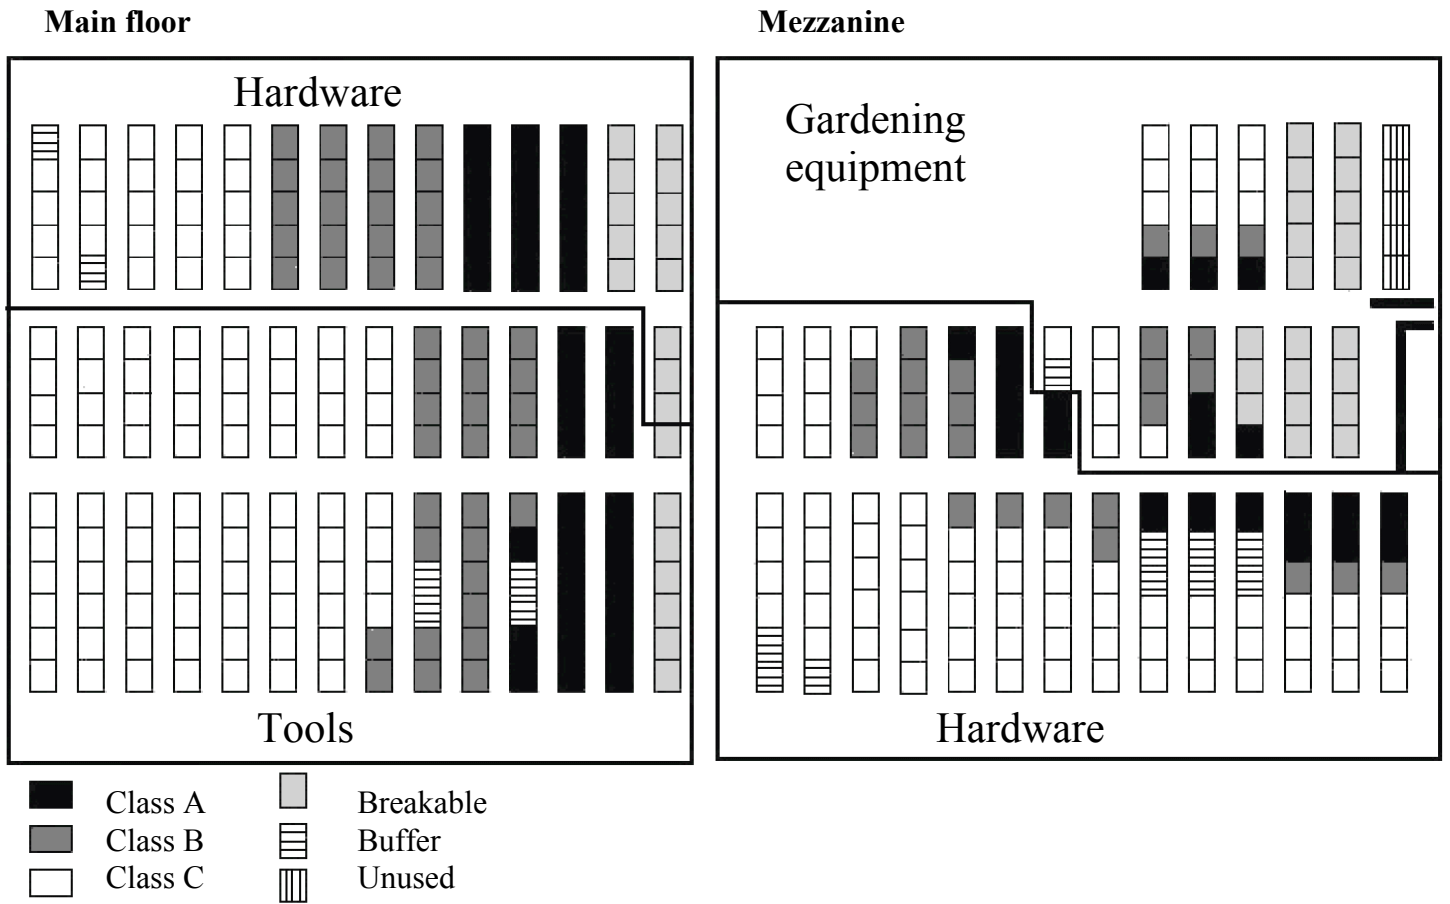


Figure 5: This figure shows the combined storage policy

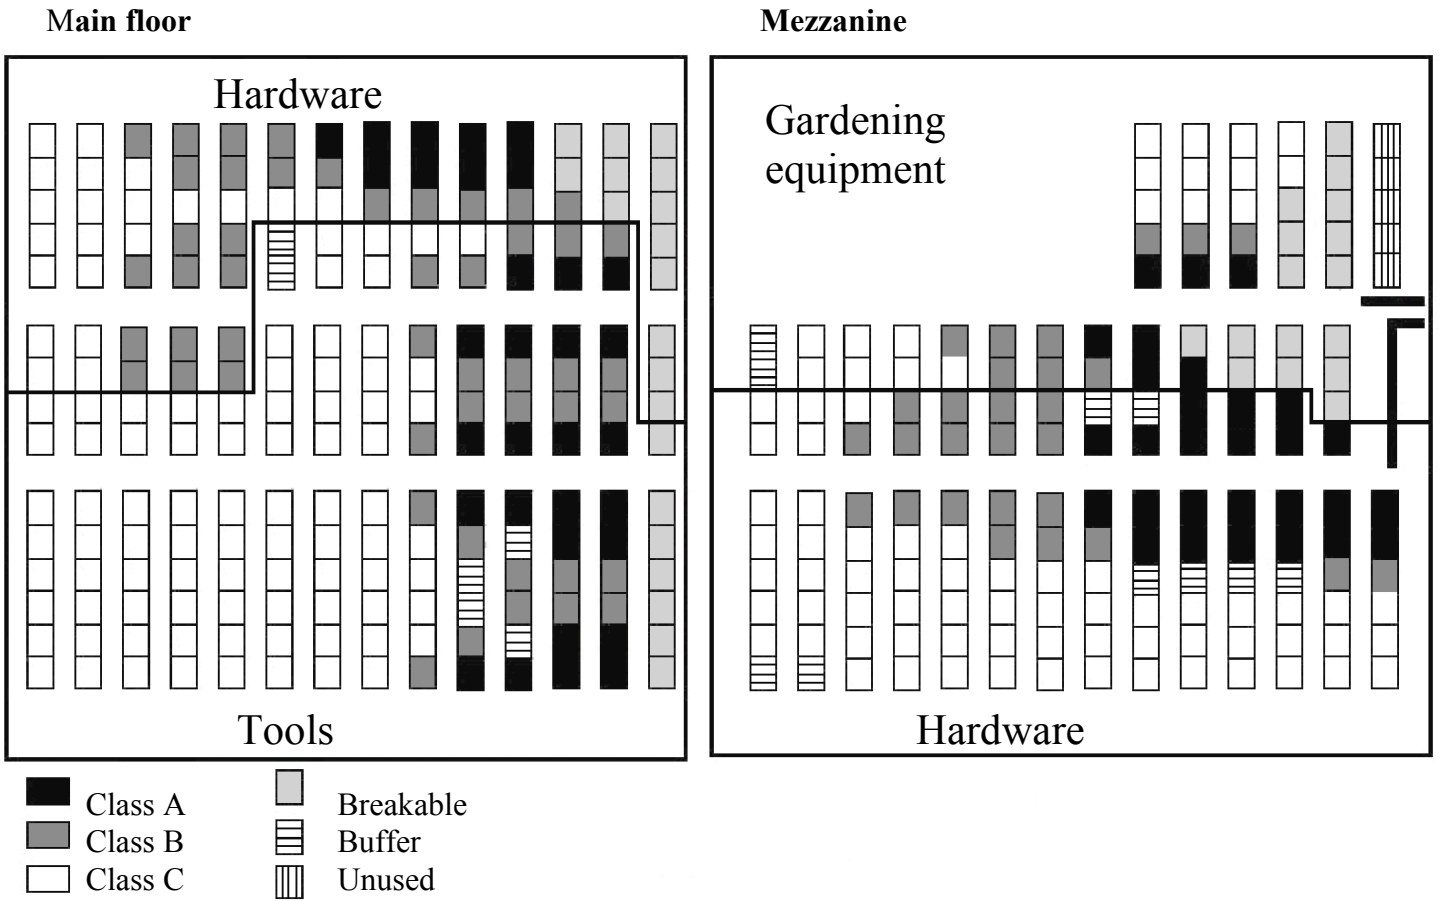


Figure 6: This figure shows the midpoint storage policy