Order fulfillment: warehouse and inventory models

Dijkstra, Arjan Stijn

Chapter 6

Summary and conclusions

The past decade, the share of products bought online has been rising, and it is projected to rise even further. The process of supplying products bought online to customers is called order fulfillment. Order fulfillment is an important and costly aspect of running an online business. Therefore, streamlining order fulfillment processes has been a focus of many online retailers, which are companies that sell products to consumers online. In this thesis, we have considered the order fulfillment process from both a warehouse management and an inventory control perspective.

6.1. Summary warehouse management

The efficiency of warehouses in order fulfillment processes is mainly determined by the order-picking process. Order picking is the process of retrieving a set of products from storage in response to customer requests. A number of operational decisions impact the performance of the order-picking process, measured as the expected time to retrieve an order, including order-picker routing and product storage assignment.

The routing of order pickers is central to many operational decisions in order-picking systems, as the length of order-picking tours are an important factor in determining the time to pick per order. The route an order picker travels to pick an order is called the order-picking tour. Many heuristics exist to find good order-picking tours, while the optimal order-picking tour can be found using a dynamic program. In the Storage Location Assignment Problem (SLAP), an assignment of products to storage locations is determined such that the expected length of the order-picking tour is minimized. Therefore, the performance of a given assignment is influenced heavily by the routing method employed.

In Chapter 2, we have considered the SLAP for a single-block warehouse. Formulas for the expected length of the order-picking tour have been derived for any storage assignment for the return, S-shape, largest-gap and midpoint routing heuristics. The formulas were used to derive a number of properties of the optimal storage assignment. For example, we have proven that the optimal storage assignment under return routing sorts products in each aisle such that the most frequently demanded products are closest to the front cross-aisle. Using the optimality properties and the formula for the expected length of the order-picking tour, a dynamic program was constructed that finds optimal class-based storage assignments for return routing, and very good storage assignments for the other routing heuristics. In experiments, we have found shapes of storage assignments that are different from those previously found in literature.

In Chapter 3, we have developed a method to determine the expected length of the optimal order-picking tour. We revisited the dynamic program (DP) proposed by Ratliff & Rosenthal (1983) and developed a DP with a single stage corresponding to each aisle, instead of the original two. Based on the DP, we have constructed a decomposition of the length of the optimal order-picking tour into a sum of lengths per aisle. The decomposition served as the basis of a stochastic DP that determines the expectation of these lengths for each aisle, which could be used to determine the expected length of the optimal order-picking tour. Furthermore, we have shown that the state space of the stochastic DP is bounded, which implies the stochastic DP can determine the expected length of the optimal order-picking tour in polynomial time. In numerical experiments we have shown the stochastic DP can compute the expected length of the optimal order-picking tour for instances of realistic size. Previous methods to approximate the expected length of the order-picking tour are shown to have deviations from the actual expectation of up to 19%.

6.2. Summary inventory control

Inventory control in online retail aims to balance product availability against the cost of keeping and moving inventory and the cost of surplus inventory. Inventory control in online retail is complicated by lost customer demands in case of stockouts, and product returns that raise inventory levels at locations they are returned at. Pooling inventory can be a useful tool in achieving higher product availability against the same cost. Many different ways of pooling inventory exist, including lateral transshipments and keeping inventory for multiple locations facing customer demand in a central warehouse.

In Chapter 4, we have studied inventory pooling in a supply chain network consisting of a single warehouse and multiple retailers (or stores) under periodic review and deterministic lead times where excess demand at the retailers is lost. The retailers face external stochastic customer demand and are replenished by the warehouse. The warehouse is replenished by an external source with infinite supply. The objective was to find replenishment quantities for each location in the network, such that the long-run average cost per period was minimized, where the total cost consisted of unit holding and penalty cost. We have formulated a Markov Decision Process to find optimal replenishment policies for all locations in the network. Additionally, we have studied the class of echelon base-stock policies for the network. The Markov Decision Process was adapted to evaluate echelon base-stock policies and has been used to find the best base-stock policies for a number of instances. However, setting the base-stock levels with such an approach is computationally expensive. Furthermore, using the base-stock policy resulting from an equivalent inventory system with backordered demand (i.e., ignoring the lost-sales characteristic) resulted in an average increase of the average total cost by 4.5 to 8.5% compared to the best base-stock policy. Therefore, we have constructed an approximation procedure for the long-run average costs for a given echelon base-stock policy. Then, echelon base-stock policies were found using the approximation procedure. In numerical experiments, the policies found were shown to be typically within 1% of the best base-stock policy.

In Chapter 5, we studied policies to determine transshipments of cross-channel returned products, i.e. products that are bought online and returned to an offline store. We considered a finite periodic sales season. At the end of each period of the sales season, products are returned with a certain probability. At that moment, cross-channel returned products may be transshipped to the online store. A Markov Decision Process (MDP) was formulated to determine optimal transshipment policies. However, due to the curse of dimensionality the MDP cannot solve instances with a large number of stores. Therefore, we introduced a heuristic based on the expected costs during the sales season. The heuristic compares the expected costs of keeping a product in the offline store with the costs of shipping it to and keeping it at the fulfillment center of the online store. Whenever the expected costs corresponding to the offline store exceed the expected costs corresponding to the online store, a transshipment takes place. The costs of the heuristic policy resulting from this approach were shown to have a maximum deviation of 1.59% from the costs corresponding to the optimal transshipment policy in experiments. Furthermore, two static heuristic policies were introduced in which products are either always or never shipped back to the online store. We have observed that our heuristic transshipment policy is more effective than static policies in dealing with imbalances resulting from cross-channel returned products.

6.3. Discussion and future research

In this thesis we have studied a number of models related to online order fulfillment. In all these models, stochasticity played an important role. In Chapter 2 and Chapter 3, uncertain order contents lead to uncertain order picker routes. Uncertain customer demands are a key ingredient to many inventory control models; Chapter 4 and Chapter 5 are no exception. In Chapter 4, the additional stochasticity found in the uncertain mean of lead time demands was central to the approximation of long-run costs. In Chapter 5, the dependence of returns on actual observed demand added a layer of uncertainty uncommon in inventory control models with returns.

Our approach in the chapters on warehouse management was to formulate the length of the order-picking tour as a dynamic program. In Chapter 2, resulting dynamic programs proved to be trivial for all routing heuristics, with the exception of the largest-gap heuristic. Therefore, the expected route length could be expressed by closed-form formulas, which in itself contained a decomposition of the expected route length per aisle. This decomposition aided the construction of a dynamic program for the storage location assignment problem, which considered aisles one by one.

In Chapter 3, a reformulation of the original dynamic program to find the optimal order-picking tour by Ratliff & Rosenthal (1983) provided a number of opportunities for further analysis. While the length of the optimal order-picking tour per aisle can be determined, it proved useful to consider the additional increment from the minimum length partial-tour in any class in the previous aisle, to the minimum length partial tour in any class and to the minimum length finished tour in the aisle, respectively. These increments could be determined numerically in a stochastic dynamic program, which has a polynomial running time in case distances in the warehouse are discrete.

In the warehouse management chapters, an important assumption in making the transition from the deterministic to the stochastic case was the assumption of independent product demands. In warehouse literature, it is common to assume product demands are independent. However, there are two ways of interpreting this assumption. In this thesis, we have modeled the demands as independent Bernouilli trials. An alternative is to fix the total number of products on an order and use independence sampling without replacement to determine the products of an order. An interesting question for future research is how results from the two demand models are related and if models from the former can be extended to the latter. A possible approach would be to extend the dynamic programming methods with an extra state variable representing the total number of products on the order up to the aisle considered.

Other interesting future research directions involve the SLAP. While we have obtained optimal storage location assignments for return routing and shown that the storage location assignment DP provides near-optimal solutions for S-shape routing in Chapter 2, algorithms for determining optimal storage assignments for midpoint, largest-gap and optimal routing have not been determined.

Our approach in the chapters on inventory control has been to formulate a Markov Decision Process to determine optimal policies. The curse of dimensionality implied optimal policies could only be determined for relatively small instances. Therefore, our aim has been to determine heuristic policies that scale well to larger problem instances. These policies could then be compared with the optimal policies found by the MDP. In Chapter 4, an approach based on decomposition of costs led to near-optimal heuristic policies. The costs were decomposed in costs per retailer dependent on warehouse echelon inventory level, which were then combined through the distribution of this warehouse echelon inventory level. As in similar papers on the OWMR-system, linear rationing facilitated decomposition. Future research may investigate how other rationing rules can be used in the approximation, for example by using appropriate expectations of inventory positions of the retailers. Furthermore, a model could be studied in which more than two echelons are considered.

In Chapter 5, the costs were approximated for two different situations, assuming no future transshipments take place: one in which the product under consideration was transshipped and one in which it was not. The heuristic policy then determined transshipments based on the comparison of the costs of the two different situations. An interesting avenue for future research would be how transshipments of multiple products should be determined in case there is a fixed constant transshipment cost when shipping multiple products. In this case, considering products one by one would probably not yield good policies, as the decision considers multiple products at once. Furthermore, in our model we assumed the initial inventory levels to be fixed. Taking future transshipment decisions into account when determining initial stock levels is an interesting problem.