Order fulfillment: warehouse and inventory models

Dijkstra, Arjan Stijn

Chapter 1

Introduction

In the past decade, adoption of the internet and mobile devices have made it increasingly common for consumers to order products online. For example, in the Netherlands the share of products bought online has risen from 10% in 2012 to 15% in 2017 (Blauw Research, 2012; Thuiswinkel, 2017). Worldwide, this share is projected to rise from 7.4% in 2015 to 15.5% in 2021 (Statista, 2018a). In line with this trend, many online retailers that sell products to consumers online have emerged. Online retailers can operate only online, but a significant part of online retailers are multi-channel retailers that employ brick-and-mortar stores as well.

The cost of logistics for online retailers is substantial. For example, in 2017 online fashion retailer Zalando's order-fulfillment costs were 23.3% of revenue, compared to 10.3% for marketing (Zalando, 2018). Balancing these costs with customer service levels is the central problem in online retail logistics (Fernie & Sparks, 2009). Facing high order-fulfillment costs and a rapidly evolving market place, many online retailers are reconsidering their complex logistical processes in order to stay competitive.

A number of characteristics of customers buying online present challenges to the logistics of online retailers. Examples of such characteristics include the role of consumer trust and large demand peaks. First, many online retailers promise short deadlines for delivery of online orders, facing rising expectations from customers. Failure to meet delivery promises decreases consumer satisfaction (Koufteros et al., 2014). Furthermore, consumers facing failure in delivery are less likely to make future purchases (Rao et al., 2011). Hence, the logistics of an online retailer are a matter of trust (Xu et al., 2009; Kim et al., 2009). Secondly, customer demands are irregular in time, with large demand peaks being common both day-to-day and throughout the year. The holiday season in December sees a surge in demand in many product categories. For example, in the UK the total value of online retail sales in December exceeds September of the same year by approximately 60% (Statista, 2018b).

Order fulfillment is one of the most challenging aspects of the logistics of an online retailer (Agatz et al., 2008). Order fulfillment is the process of supplying customers with the ordered products. When ordering, consumers select their preferred delivery method, such as in-store pick up or home delivery. Before customers can be provided with their products, the products need to be retrieved from inventory in a warehouse. Retrieving a set of products from inventory locations in a warehouse in response to a customer request is called order picking (De Koster et al., 2007). After delivery of the order, customers may return products in many situations. For example, purchases of most products sold online can be canceled within 14 days without cause by law in the EU. As a consequence, online retailers see a return flow of products to their warehouses and stores, potentially interfering with the rest of their logistics. These products may end up in locations which already have sufficient inventory, whereas other locations may have limited inventories. Careful management of return flows can help to prevent such situations by allocating returned products to locations with small inventories.

Furthermore, many online retailers choose to display inventory information of products in their online store (Aydinliyim et al., 2017). As a consequence, stock outs usually result in lost sales, due to either the impossibility to order or customers' reluctance to wait for out-of-stock products. In case inventory is not displayed online, eliminating stock outs may increase long-term customer equity by as much as 56% (Jing & Lewis, 2011).

Managing an order fulfillment process to meet customer expectations efficiently is challenging. Two main aspects of managing the order fulfillment process are warehouse management and inventory management (Hübner et al., 2015; Ishfaq et al., 2016). In this thesis, we study both warehouse management and inventory management related to online retail.

Warehouses facilitate order picking and the receipt, storage and shipment of goods (Gu et al., 2007). Traditionally, order picking is the most costly activity in warehouses, due to the large amount of labor and capital involved in the process (Tompkins et al., 2010; Marchet et al., 2015). In online retail, order picking is even more critical, as the order-throughput achieved by an order picking system is an important driver in meeting delivery promises (Gong & De Koster, 2008).

Ability to adapt to large demand peaks is an important factor in the design of order-picking systems for online fulfillment (Bozer & Aldarondo, 2018). This ability is often achieved through employing order picking systems with human order pickers, as these scale more easily in throughput capacity (Dallari et al., 2009). Order-picking systems in which humans collect products for customer orders are called picker-to-parts order-picking systems. The location at which products are stored has a large impact on throughput capacity in picker-to-parts order-picking systems (Van der Gaast, 2016; Petersen & Aase, 2004). By storing products in a way that often demanded products are more easily retrieved, a substantial increase in system performance may be achieved.

Inventory management in online retail aims to balance product availability against the cost of keeping and moving inventory and the cost of surplus inventory. The cost of keeping inventory includes depreciation, insurance, storage and cost of capital. Having products available is a prerequisite for order fulfillment. Inventory control studies the problem of determining the quantity of products to keep in inventory at inventory holding locations. While the inventory control of a single location is far from trivial (Zipkin, 2008; Bijvank & Vis, 2011), inventory control for multiple locations is even more complex.

For example, challenges arise when multi-channel retailers want to take advantage of their offline stores. Offline stores could be used for order fulfillment by incentivizing customers to collect their orders in a store nearby (Bretthauer et al., 2010). However, doing so requires careful consideration of the multiple inventory levels of different offline stores, as stores with critically low inventory levels should be shielded from online demand (Mahar et al., 2012). Online customers returning products in offline stores complicates the situation even further (Mahar & Wright, 2017).

The remainder of this chapter is organized as follows. First, we introduce and outline our contributions to warehousing in online retail. Secondly, we introduce and outline our contributions to inventory control in online retail. Finally, we discuss the organization of the thesis.

1.1. Warehousing in online retail

Many different decisions impact warehouse operations, including warehouse layout design, product storage and the design of the order-picking system (De Koster et al., 2007). We focus on the location in which products are stored in the warehouse, which is an important factor in order-throughput (Petersen & Aase, 2004). The storage location assignment problem (SLAP) considers the assignment of products to storage locations such that order picking is efficient (Gu et al., 2007). Typically, storage location assignments are class-based. Under class-based storage, products are grouped together in classes with similar demand frequencies. Classes are assigned storage locations in the warehouse and products within a product class are stored randomly in these locations. Alternative approaches to storage assignment are random storage and full turnover storage, in which products are assigned to dedicated storage locations. Class-based storage offers efficiency gains over random storage, while allowing more flexibility than full turnover storage. The flexibility accommodates both fluctuating product inventory and easy inclusion of new products. Furthermore, the number of products that need to be stored is typically large, especially in online retail. Using class-based storage facilitates the construction of storage location assignments.

The efficiency of a given storage location assignment is affected by the operational execution of the order-picking process (Van Gils et al., 2018). The manual order-picking process can be divided in three operational aspects: routing, batching, and sorting (Gu et al., 2007). Routing considers the route an order-picker is traveling to target locations through the warehouse. Batching is commonly approached by the combination of multiple customer orders into a single order-picker route, saving average travel distance per order (Gademann & Van de Velde, 2005; Matusiak et al., 2017). Sorting considers how and when multiple customer orders, which were picked together, can be sorted into individual packages per customer ready for shipping (Johnson & Meller, 2002; Meller, 1997).

Determining short routes for the order pickers to travel in the warehouse is crucial for maximizing the number of orders that can be picked per hour at a given capacity. For some layouts, optimal routes can be found with a polynomial-time dynamic program (Ratliff & Rosenthal, 1983; Roodbergen & De Koster, 2001). For larger warehouses, both local-search and meta-heuristics have been proposed to find short routes (Theys et al., 2010; De Santis et al., 2018). Additionally, some aspects can make the routing challenging. Examples include order-picker interaction (Pan & Shih, 2008; Parikh & Meller, 2010), and product returns (Schrotenboer et al., 2017). In practice, however, often simple routing heuristics are used to route order pickers (Dekker et al., 2004). Travel-time models that approximate the average order picking time have been proposed for many different routing methods (Hall, 1993; Le-Duc & De Koster, 2005; Rao & Adil, 2013a). The routing method used has a large impact on the average order picking time of a given storage location assignment (Petersen & Aase, 2004).

Storage location assignments are usually evaluated by the average order picking time (Van Gils et al., 2018). Studies on storage location assignment have approximated the average order picking time for intuitive storage location assignments (Petersen & Schmenner, 1999; Petersen & Aase, 2004; Caron et al., 1998). Attempts to improve on the intuitive storage assignments include local-search from these storage location assignments (Dekker et al., 2004; Le-Duc & De Koster, 2005), meta-heuristics (Pan et al., 2015) and integer linear programming (Muppani & Adil, 2008; Ene & Öztürk, 2012). Above mentioned papers all use approximate travel-time models. Exact average order picking times have been determined for some settings only, such as a warehouse with only a single aisle (Eisenstein, 2008) or simplified routing methods (Jarvis & McDowell, 1991).

In Chapter 2, we study the storage location assignment problem in a warehouse with two cross aisles. Eisenstein (2008) and Jarvis & McDowell (1991) have shown in simplified settings that a demand model assuming independent product demands leads to tractable travel-time formulas. We extend their approach to different routing methods in a multi-aisle warehouse. Specifically, we provide exact formulas to determine the average time to pick an order for commonly used routing methods. We use these methods to prove properties of the optimal solution of the storage location assignment problem for these routing methods. These properties aid in the construction of a dynamic program (DP) that provides optimal solutions for the return routing method in polynomial time for a fixed number of product classes, and good performance for the other routing methods.

In Chapter 3, we present an analytic method to determine the exact expected time to pick an order when using optimal routing in a warehouse with two cross aisles. We describe a stochastic DP, building upon the deterministic DP proposed by Ratliff & Rosenthal (1983). We prove that our stochastic DP has a polynomial running time, based on properties of the state space of our DP. Numerical experiments demonstrate the existing approximate method from literature has deviations of up to 19% with our exact expectation, highlighting the need for more accurate methods.

1.2. Inventory control in online retail

Inventory control has been studied in many different settings (Axsäter, 2003). We focus on a number of topics in the literature on inventory control that are relevant for online retailers: lost sales, returns and inventory pooling.

In lost-sales models, demand is lost when it cannot be met from stock. In contrast, when excess demand is backordered, it is fulfilled later. In retail settings fulfilling demand later is often not possible and lost sales are a more common assumption (Kapalka et al., 1999; Ehrenthal et al., 2014). For some models with backorders, the optimal policy can be expressed with a single or two parameters (Porteus, 2002). However, the optimal policy for the corresponding lost-sales models is complex (Zipkin, 2008; Bijvank &

Vis, 2011).

Return rates may be substantial for online retailers. Returned products increase the inventory at the location they are returned to. The literature on inventory control with resalable returns mainly considers single location models (Kelle & Silver, 1989; Buchanan & Abad, 1998; Mostard & Teunter, 2006). In these models, initial inventory levels are determined while taking return rates into account. In contrast, literature on models with returns at multiple locations are scarce.

Inventories of the offline stores and online store can be pooled. By pooling inventories, the risk of demand fluctuations is shared by multiple locations. Typically, this results in less lost sales when inventories remain the same. Pooling can be achieved by pooled replenishments and lateral transshipments. Demand allocation can be seen as a form of inventory pooling as well. If a product is available at multiple locations, online retailers can decide which location serves the demand (Mahar et al., 2009; Xu et al., 2009; Acimovic & Graves, 2015).

Pooled replenishments are a classical approach to inventory pooling (Schwarz, 1981). With pooled replenishments, a warehouse orders inventory from an outside supplier. This warehouse replenishes the stores. The main difference in pooled replenishment literature is whether or not the warehouse is keeping inventory. If the warehouse is not allowed to hold inventory, individual replenishments for the stores are determined when shipments from the outside supplier arrive. This way, uncertain demands during the lead-time from the outside supplier are shared by all stores (Acimovic & Graves, 2017). When the warehouse is allowed to keep inventory, replenishments to the stores happen independently from the shipments of the outside supplier. In situations where the warehouse is only supplying a single store and demand is backordered, optimal policies can be characterized completely (Clark & Scarf, 1960). Under some assumptions, this characterization also holds for optimal policies for systems with multiple stores (Clark & Scarf, 1960; Eppen & Schrage, 1981; Federgruen & Zipkin, 1984). Finding the optimal policy, however, is not simple. Diks & De Kok (1998) propose an algorithm to determine a good heuristic policy for the case with backorders. Literature on models with lost sales either assume no lead time between the warehouse and the stores (Nahmias & Smith, 1994; Paul & Rajendran, 2011), or continuous replenishment (Hill et al., 2007; Alvarez & van der Heijden, 2014). Both assumptions may not be realistic in an online retailing setting.

Lateral transshipments are movements of inventory between stores, instead of between the warehouse and a store. Lateral transshipment can either be proactive or reactive (Paterson et al., 2011). Where reactive lateral transshipments take place in response to a customer demand, proactive lateral transshipments seek to balance inventory before stockouts happen. Reactive transshipments have been studied in the context of online retailing (Ramakrishna et al., 2015; Zhao et al., 2016). Returned products offer an opportunity for proactive lateral transshipments, as handling is required for returns regardless where they end up.

In Chapter 4, we study base-stock policies for a one-warehouse-multiple-retailers (OWMR) system with lost sales. We find the optimal replenishment policy using a Markov Decision Process (MDP). We numerically determine the best base-stock policy and show that it performs close to optimal. Furthermore, we develop a cost approximation of base-stock policies. Numerical experiments illustrate that base-stock policies found using this cost approximation are close to optimal; costs are typically within 1% of the costs corresponding to the best base-stock policy.

In Chapter 5, we study the proactive transshipment of resalable returned products. We study a finite selling season, consisting of a discrete number of periods. Products sold during a period can return during that period with a given probability. Products sold online can be returned to offline stores. Returned products are allowed to be transshipped to the online order fulfillment location or kept on hand in the store. We find the optimal transshipment policy, using a Markov Decision Process. Furthermore, we propose a transshipment heuristic based on an approximation of the costs in the remainder of the selling season that can be attributed to a product. Experiments indicate that the transshipment heuristic performs close to the optimal policy. Furthermore, the heuristic outperforms static policies in which the transshipment decision for all products is the same during the entire selling season.

1.3. Publications

This thesis is based on the following papers:

Chapter 2
Dijkstra, A.S. & Roodbergen, K.J. (2017). Exact route-length formulas and a storage location assignment heuristic for picker-to-parts warehouses. *Transportation Research Part E: Logistics and Transportation Review, 102*, 38–59.

Chapter 3
Dijkstra, A.S., Van der Heide, G., & Roodbergen, K.J. (2019). *The expected length of the optimal order-picking tour in a rectangular warehouse.* Manuscript in preparation.

Chapter 4
Dijkstra, A.S. & Bijvank, M. (2019). *Base-stock policies for two-echelon retail inventory systems with lost sales.* Manuscript in preparation.

Chapter 5
Dijkstra, A.S., Van der Heide, G., & Roodbergen, K.J. (2017). Transshipments of cross-channel returned products. *International Journal of Production Economics.* Advance online publication. doi:10.1016/j.ijpe.2017.09.001.