Transportation Research Part B 117 (2018) 455–472

Contents lists available at ScienceDirect

Transportation Research Part B

journal homepage: www.elsevier.com/locate/trb

The delivery problem: optimizing hit rates in e-commerce deliveries

Alexandre M. Florio[a,*], Dominique Feillet[b], Richard F. Hartl[a]

[a] University of Vienna, Oskar-Morgenstern-Platz 1, Vienna 1090, Austria
[b] École des Mines de Saint-Étienne and UMR CNRS 6158 LIMOS, CMP Georges Charpak, Gardanne F-13541, France

ARTICLE INFO

Article history:
Received 20 April 2018
Revised 6 August 2018
Accepted 18 September 2018

2010 MSC:
90B06
90C15

Keywords:
Small package delivery
e-commerce
Last mile logistics
Column generation
Branch-and-price

ABSTRACT

Unsuccessful delivery attempts, or failed hits, are still a recurring problem in the fulfillment of e-commerce orders to private customers. In this paper, we consider a parcel delivery company interested in optimizing the rate of successful deliveries. By doing so, the company is able to offer a differentiated service, increasing customer satisfaction, and reducing the costs related to failed delivery attempts. In order to achieve this, routes must be designed in a way that visiting times are convenient for the customers. Revisits to some customers may also be planned, so that the expected number of successful deliveries increases. We propose availability profiles to represent the availability of customers during the delivery period. Using these profiles, we are able to compute the expected number of successful hits in a given route. We model the delivery problem as a set-partitioning problem, and solve it with a branch-and-price algorithm. The corresponding pricing problem is solved with a labeling procedure, in which reduced cost bounds are employed to discard unpromising partial routes. We show that the reduced cost of route extensions is bounded by the optimal solution to an orienteering problem, and efficiently compute bounds for that problem within the labeling procedure. Computational experiments demonstrate the effectiveness of the approach for solving instances with up to 100 customers. A tradeoff analysis suggests that significant hit rate improvement can be achieved at the expense of small additional transportation cost. The results also indicate that flexibility regarding maximum route duration translates into an improved hit rate, and that planning revisits may reduce expected unsuccessful deliveries by more than 10%.

1. Introduction

The increasing popularity of e-commerce is unfolding challenges and opportunities in the parcel delivery market. On the one hand, the high costs associated with last mile logistics incentivize parcel delivery companies, or *providers*, to implement efficient delivery operations. On the other hand, customers are becoming increasingly more demanding in regard to delivery services (e.g., next- or same-day delivery, delivery to alternative locations and collection points). In this context, providers are becoming more and more motivated to develop new delivery practices to achieve a good compromise between efficiency and customer convenience (Savelsbergh and Van Woensel, 2016).

* Corresponding author.
E-mail addresses: alexandre.de.macedo.florio@univie.ac.at (A.M. Florio), feillet@emse.fr (D. Feillet), richard.hartl@univie.ac.at (R.F. Hartl).

https://doi.org/10.1016/j.trb.2018.09.011
0191-2615/

As in any other competitive and mature market, opportunities for service differentiation also exist in the parcel delivery market. Premium service offerings (e.g., express delivery) attempt to cater the needs of customers requiring shorter shipping times. These expedited services are already offered by most providers. In addition to delivery *at the right day*, increased customer satisfaction can also be achieved by delivery *at the right time* (Goebel et al., 2012). A delivery at the right time, or *successful delivery*, or simply *hit*, happens when the customer is available for receiving the parcel at the moment he is visited by the provider. Successful deliveries save the customers from visiting the nearest parcel shop to collect the items, and decrease the chances of product returns (Rao et al., 2014).

From an operational perspective, successful deliveries are also much desirable since they avoid costs related to new delivery attempts, further package handling, or contracting of third-party parcel shops for storing the parcels. However, ensuring successful hits is a task far from trivial, since many customers are not available at their home addresses during the entire delivery period. Moreover, except for premium delivery services, time windows are not common in the delivery of small parcels (Wong, 2008), and in many cases are not negotiated when orders are placed (Vanelslander et al., 2013). Fixing time windows for all orders leads to severe constraints during route execution, and ultimately results in very high operational costs (Boyer et al., 2009).

In this paper we introduce the *delivery problem* (DP), which is the problem a provider faces when trying to maximize the rate of successful deliveries. Instead of adopting only time windows, we introduce a more general concept called *availability profile*. Every customer is associated with an availability profile. These profiles map their availabilities for receiving deliveries throughout the delivery period, and can be estimated from historical data (Van Duin et al., 2016), or used to model specific situations such as time windows, or even scenarios where there is no information about preferred delivery times. The objective in the DP is to design a set of routes, each not exceeding some given maximum route duration, such that the expected number of successful hits is maximized. At least one delivery must be attempted to each customer, but multiple visits to the same customer are allowed. Planned revisits increase the chances of a successful delivery to the revisited customers, and consequently contribute towards the objective we try to maximize.

We model the DP as a set-partitioning problem, and solve it with a branch-and-price (Barnhart et al., 1998) algorithm. The corresponding pricing problem resembles an orienteering problem (Golden et al., 1987) with time-dependent rewards, and is solved by a labeling procedure. The time-dependency of the availability profiles means that strong label dominance rules cannot be proposed. In order to control the combinatorial growth of labels, we make extensive use of reduced cost (or completion) bounds. We show that the reduced costs of all possible route extensions from a particular label are bounded from below by the optimal solution to an orienteering problem. Within the labeling procedure, we efficiently compute orienteering problem bounds at every new label that is generated, to determine whether the label should be discarded or not. One of the bounds employed, the *longest walk to the depot* bound, is a new bound to the orienteering problem introduced in this paper.

The proposed algorithm is exact when revisits are forbidden, and instances with up to 100 customers could be solved to optimality. When revisits are allowed we employ a simple recourse action during route execution, in which customers that already received their parcels are skipped. In this case, the branch-and-price approach works as a heuristic solution method. In the computational experiments we solve different versions of the DP, which are obtained by allowing revisits at any moment, by allowing revisits only after all customers in the route have been visited, or by forbidding revisits entirely. We also introduce and solve a variant of the DP where the total transportation cost is constrained. The results allow us to measure by how much the hit rate improves, on average, when availability profiles are explicitly considered in the design of routes. We further estimate the additional improvements in the hit rate when revisits are allowed, and perform an analysis of the tradeoff between hit rates and transportation costs.

The remainder of this paper is organized as follows: in Section 2 we review the relevant literature; In Section 3, we formally introduce customer availability profiles, and present a mathematical formulation of the DP; In Section 4, we present the branch-and-price algorithm for solving the DP; In Section 5, we introduce a variant of the DP where the total transportation cost is constrained. In Section 6, we present and discuss the computational results; Finally, in Section 7 we discuss some of the implications of our work, and indicate future research directions.

2. Related Literature

The framework of time windows has been traditionally considered in the vehicle routing literature to deal with scenarios where deliveries must occur within some specified time frame. The vehicle routing problem with time windows (VRPTW, see Desaulniers et al., 2014) has been used as a base model for solving these types of routing problems. Applications of the VRPTW for modeling pick-up and delivery courier problems (Sungur et al., 2010), for optimizing the on-time delivery probability to customers (Zhang et al., 2016), and for territory-based routing (Schneider et al., 2014) have been proposed. These models assume that time windows are available, i.e., they have either been informed by the customers, or have been committed to by the logistics provider (self-imposed time windows, see Jabali et al., 2015).

Even though most e-commerce customers would appreciate having the option of selecting a convenient delivery time for their orders (Xu et al., 2008), time windows, in most cases, are not offered at the moment an order is placed (especially when the delivery is made by a parcel delivery company, see Vanelslander et al., 2013). The main problem is that committing to time windows for all customers, even if rather large ones (e.g., 3 hours), could cause a significant impact in the delivery costs (Boyer et al., 2009; Punakivi and Saranen, 2001). This impact would be proportionally larger in the delivery of low-