Waste Management 34 (2014) 1564–1576

Contents lists available at ScienceDirect

Waste Management

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

Inventory routing for dynamic waste collection

Martijn Mes *, Marco Schutten, Arturo Pérez Rivera

Department Industrial Engineering and Business Information Systems, School of Management and Governance, University of Twente, The Netherlands

ARTICLE INFO

Article history:
Received 20 January 2014
Accepted 12 May 2014
Available online 6 June 2014

Keywords:
Inventory routing
Simulation optimization
Optimal learning
Transportation
Waste collection

ABSTRACT

We consider the problem of collecting waste from sensor equipped underground containers. These sensors enable the use of a dynamic collection policy. The problem, which is known as a reverse inventory routing problem, involves decisions regarding routing and container selection. In more dense networks, the latter becomes more important. To cope with uncertainty in deposit volumes and with fluctuations due to daily and seasonal effects, we need an anticipatory policy that balances the workload over time. We propose a relatively simple heuristic consisting of several tunable parameters depending on the day of the week. We tune the parameters of this policy using optimal learning techniques combined with simulation. We illustrate our approach using a real life problem instance of a waste collection company, located in The Netherlands, and perform experiments on several other instances. For our case study, we show that costs savings up to 40% are possible by optimizing the parameters.

1. Introduction

During the last decades, there has been a growing interest in Vendor Managed Inventory (VMI). In VMI, the replenishment decisions are being made by a supplier based on various inventory and supply chain policies (Angulo et al., 2004). The combined decision on when to replenish the customers' inventories, how much product to deliver, and in which way to route the vehicles that execute the delivery, is also known as the Inventory Routing Problem (IRP). Answering all these questions simultaneously is a challenging task, considering that the decisions taken at a certain moment in time for a given planning horizon influence the decisions made later within or beyond this horizon (Baita et al., 1998). By now, various methodologies have been developed to cope with this challenge and to achieve higher service levels for customers, while simultaneously lowering the costs for the suppliers.

Accurate information about current and future customers' inventories and demand is vital for the decisions to be sound during the entire planning horizon (Aghezzaf, 2008). However, the inherent variability in the demand (and thus the inventories) makes it difficult to have a precise prediction, and hence creates an additional layer of complexity to the already difficult IRP. The problem becomes even tougher if we are dealing with companies serving a large numbers of customers. This typically occurs in urban areas, where customers are located closely to each other. Examples include vending machine replenishment (Rusdiansyah and Tsao, 2005), supermarket replenishment (Gaur and Fisher, 2004), and municipal waste collection (Russell and Igo, 1979). The latter is also the topic of this paper.

A particular application of the IRP with a large number of customers, variability in the demand, and a long planning horizon (say several weeks), is the Waste Collection Problem (WCP). In the special case in which the waste collection company plans the emptying of containers dynamically (as opposed to static or periodic scheduling and routing) and bases this planning on the amount of waste inside the containers (which can be known through the uses of sensors in each container), the WCP becomes a special case of the IRP. The difference is related to reverse flows (the purpose of visiting a "customer" is collecting rather than delivering something) and the decision on how much to collect is not relevant since containers will always be fully emptied. Solution methodologies for IRPs also work for WCPs as long as they support decisions for uncertain demand (waste deposits) and a large number of customers (waste containers), which are usual settings for a WCP.

This paper is motivated by a case study at the waste collection company Twente Milieu, located in The Netherlands. Different types of municipal waste containers are used in The Netherlands. The most important types are mini containers (one per household, have to be put along the side of the road on prespecified days) and block containers (shared by multiple households). Since 2009, Twente Milieu also makes use of underground containers, which are also shared by multiple households, but have a number of advantages compared to the mini containers and block containers, see Mes (2012). Initially, these underground containers were mainly placed at apartment buildings and commercial buildings

* Corresponding author. Tel.: +31 534894062; fax: +31 534892159.
E-mail address: m.r.k.mes@utwente.nl (M. Mes).

http://dx.doi.org/10.1016/j.wasman.2014.05.011
0956-053X/

(e.g., at restaurants), but their use is now extended to all sorts of living areas. In this paper we focus on the collection from these underground containers, which are equipped with sensors that provide insight into the fill levels of containers at any point in time. The objective of this study is to use this information to efficiently plan the emptying of underground containers. Given the stochastic nature of waste deposits, it is difficult to design robust plans for all possible demand realizations. To solve the problem for a sufficiently long planning horizon, a way of ''learning'' from historical inventory levels must be incorporated such that better predictions can be done for the future.

Taking into account the size of the container network, the stochastic nature of waste deposits, and the need to use a long planning horizon, on top of the interrelatedness of the multiple decisions, it is clear that not all solution approaches for the IRP are suitable. Modelling the planning decisions and solving the model for real-life problem settings and instances are challenging tasks. Exact solutions, such as mathematical programming, are not suitable to solve larger problem instances (McLeod and Cherrett, 2008). Additionally, mathematical programming models usually assume deterministic demands. Stochastic modelling approaches, such as Stochastic Dynamic Programming and Markov Decision Processes, also become computationally intractable due to large state spaces and high-dimensional value functions that cannot be solved analytically (Kleywegt et al., 2004). For these reasons, different types of heuristic approaches have been proposed in the literature. In their review of various heuristics for the IRP, Abdelmaguid et al. (2009) show that these heuristics involve parameters or settings that influence their performance. Even if ways of determining the parameters are given, they usually do not incorporate any form of coping with uncertainty; variability in demand realizations may thus diminish their performance.

In this paper, our main goal is to develop a fast and parameterized heuristic for solving the IRP for waste collection, together with a methodology to determine the best parameter settings for our heuristic. Since the performance and the quality of a particular heuristic heavily depend on choosing the right values of its parameters, we propose the use of techniques from optimal learning (Powell and Ryzhov, 2012). This paper makes the following contributions: (i) we propose a practical and simple heuristic for solving the IRP with many customers, (ii) we show how simulation optimization can be used for tuning the parameters of our heuristic in the best way for a given problem setting, and (iii) we provide insight into the dependency of the parameters of our heuristic with respect to several network characteristics (e.g., density of the container network, fluctuation in waste deposits, etc.). We illustrate our approach using the case study at Twente Milieu. This company has implemented the dynamic collection policy, where underground waste containers are scheduled to be emptied based on sensor information and by using the heuristic as presented in this paper.

The paper is organized as follows. In Section 2, we briefly present the key points addressed in the scientific literature about the problem under consideration. In Section 3, we describe our model and present the assumptions of the IRP for waste collection. Following this, we explain our parameterized heuristic approach for solving the problem in Section 4. In Section 5, we describe the way an optimal learning algorithm can be applied to this problem and specifically to our heuristic. We present the experimental design and the insights of this study in Section 6. We end with conclusions in Section 7.

2. Literature

In the Inventory Routing Problem (IRP), three questions have to be answered: (i) when to visit a customer, (ii) how much product to deliver during the visit, and (iii) how to route the vehicles (Campbell and Savelsbergh, 2004). The IRP combines two problems classes: the Vehicle Routing Problem (VRP) and Vendor Managed Inventory (VMI). Some IRPs are considered as extensions of the VRP (Moin and Salhi, 2007). In a VRP, a company limits itself to receiving customer orders and finding the best way to satisfy and deliver them. On the other hand, in an IRP, the customer orders are determined by the company, usually guided by some service level agreement. This case of customer stocks being replenished without an explicit customer order is known as VMI. VMI decisions focus on determining the size and time of replenishment. The combination of VMI and VRP decisions makes the IRP a challenging problem.

To cope with uncertainty in customer inventories, IRPs are usually solved repeatedly for a multi-period planning horizon (i.e., using rolling horizon procedures). These frequent decisions have the effect that previous decisions influence current and future ones, as explained by Baita et al. (1998). Therefore, the length of the planning horizon has an impact on the way the problem should be conceptualized and tackled. For a thorough categorization of the characteristics of dynamic routing and inventory problems, we refer to Baita et al. (1998) and Kleywegt et al. (2002). Here, we elaborate on only two of these characteristics: the planning horizon and the uncertainties in demand.

The planning period of IRP studies vary from a single period to an infinite horizon. Nevertheless, most researchers agree that the interrelatedness of decisions through time has an impact on the long-term planning objective. Since early studies of IRPs, authors have developed ways of measuring the long-term effect when using single period models. For example, Dror and Ball (1987) solve a series of single period problems, model the long-term effect through the use of penalties, incentives, and expected changes in costs, and optimize the output of the single period problems in accordance to the long-term objectives. Chien et al. (1989) also tackle the IRP long-term decision effects with single period problems, with the difference that they pass inventory and cost information from one period to the next one, and therefore make decisions taking into account information from other periods. The problem has also been studied the other way around: a long-horizon solution is developed first and then short-term plans are derived from it. For example, Campbell and Savelsbergh (2004) develop a two phase rolling horizon approach. First, a monthly plan is generated, which is then split into short-term problems for daily scheduling. The plan is implemented only for the first few days of the planned month, after which a new plan will be generated. A similar rolling horizon approach is developed by Jaillet et al. (2000), who build a two weeks schedule but only the first week is implemented. Just as these examples, most of the solving approaches in the literature typically decompose the entire IRP into short term problems and use some method to account for the long term objective.

The majority of models developed for multiple-period IRPs assume deterministic demand as seen in Andersson et al. (2010). However, it is often desirable that a planning system is able to cope with stochastic processes, especially when considering that the different realizations of real-life demand might prevent the plan of being executed as desired (Ronen, 2002). According to the classification scheme of Andersson et al. (2010), our problem can be characterized as a Dynamic and Stochastic Inventory–Routing Problem (DSIRP), with a finite horizon, one-to-many deliveries, multiple customer visits per route, order-up-to level inventory policy, using back-ordering, with a fleet of multiple homogeneous trucks. In the DSIRP, customer demand is known only in a probabilistic sense and revealed over time. Frequently, stochastic IRPs are modelled as Markov Decision Processes (Kleywegt et al., 2002; Adelman, 2004; Hvattum et al., 2009). However, this approach might easily