European Journal of Operational Research 237 (2014) 802–813

Contents lists available at ScienceDirect

European Journal of Operational Research

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

Discrete Optimization

An iterated local search algorithm for the single-vehicle cyclic inventory routing problem

Pieter Vansteenwegen [a,b,*], Manuel Mateo [c]

[a] Centre for Industrial Management, KU Leuven, Celestijnenlaan 300a Box 2422, B-3001 Leuven, Belgium
[b] Department of Industrial Management, UGent, Technologiepark 903, B-9052 Zwijnaarde, Belgium
[c] Departament Organització d'Empreses, Universitat Politècnica Catalunya, Avda Diagonal, 647, 7°, Barcelona, Spain

ARTICLE INFO

Article history:
Received 4 July 2013
Accepted 9 February 2014
Available online 17 February 2014

Keywords:
Routing
Inventory
Single-vehicle cyclic inventory routing problem
Iterated local search
Metaheuristic

ABSTRACT

The Single-Vehicle Cyclic Inventory Routing Problem (SV-CIRP) belongs to the class of Inventory Routing Problems (IRP) in which the supplier optimises both the distribution costs and the inventory costs at the customers. The goal of the SV-CIRP is to minimise both kinds of costs and to maximise the collected rewards, by selecting a subset of customers from a given set and determining the quantity to be delivered to each customer and the vehicle routes, while avoiding stockouts. A cyclic distribution plan should be developed for a single vehicle.

We present an iterated local search (ILS) metaheuristic that exploits typical characteristics of the problem and opportunities to reduce the computation time. Experimental results on 50 benchmark instances show that our algorithm improves the results of the best available algorithm on average with 16.02%. Furthermore, 32 new best known solutions are obtained. A sensitivity analysis demonstrates that the performance of the algorithm is not influenced by small changes in the parameter settings of the ILS.

© 2014 Elsevier B.V. All rights reserved.

1. Introduction

A class of routing problems that currently gets a lot of attention are the Inventory Routing Problems (IRP), e.g. Dror and Ball (1987), Andersson, Hoff, Christiansen, Hasle, and Lokketangen (2010), Schmid, Doerner, and Laporte (2013) and Zhong and Aghezzaf (2012). A typical characteristic of these routing problems is that inventory and handling costs at the customers are also taken into account. This corresponds to logistic providers not only distributing goods from a central depot to their customers, but also managing the inventories at these customers. In the literature, the IRP is discussed as a "Vendor Managed Inventory" (VMI) problem, where the supplier coordinates both the routing of a fleet of vehicles distributing the goods and the timely replenishment of inventories at the customers. This strategy allows reducing costs compared to the traditional strategy of the separated management of inventory, done by the customers, and vehicle routing, done by the supplier (Campbell, Clarke, & Savelsbergh, 1998; Christopher, 1998; Waller, Johnson, & Davis, 1999). The Cyclic Inventory Routing Problem (CIRP) is a well-known variant of the general IRP (Aghezzaf, Raa, & Van Landeghem, 2006; Raa & Aghezzaf, 2009). The CIRP is an appropriate optimisation model for a VMI policy when customer demand rates are stable and the planning horizon is infinite. For this class of problems, the objective function is to minimise the long term transportation and inventory costs.

In this paper, we deal with a specific variant of a cyclic inventory routing problem: the Single-Vehicle Cyclic Inventory Routing Problem (SV-CIRP) (Aghezzaf, Zhong, Raa, & Mateo, 2012; Zhong & Aghezzaf, 2011, 2012). In this case, the demand rate is considered constant and a cyclic distribution plan should be developed for a single vehicle starting and ending at a single depot (Andersson et al., 2010). The goal of the SV-CIRP is to minimise the total cost, i.e. the addition of transportation and inventory costs, by determining the quantity to be delivered to the selected customers and the vehicle routes, while avoiding stockouts. Most papers in the literature (e.g. Campbell et al., 1998; Federgruen & Simchi-Levi, 1995; Moin & Salhi, 2007; Ribeiro & Lourenço, 2005) consider an unlimited number of vehicles for the fleet or determine the minimum number of needed vehicles.

The SV-CIRP considers a set of potential customers, each with a demand rate, inventory and handling costs and a fixed reward. Other distribution aspects to be considered are the travel times between customers, the travel cost, the vehicle cost, the average speed and the capacity of the vehicle. The main variable is the cycle

* Corresponding author at: KU Leuven, Centre for Industrial Management, Celestijnenlaan 300a Box 2422, B-3001 Leuven, Belgium. Tel.: +32 16 321669; fax: +32 16 322986.

E-mail addresses: Pieter.vansteenwegen@cib.kuleuven.be (P. Vansteenwegen), manel.mateo@upc.edu (M. Mateo).

http://dx.doi.org/10.1016/j.ejor.2014.02.020
0377-2217/© 2014 Elsevier B.V. All rights reserved.

time, i.e. the time between two deliveries to each customer. An important property of the SV-CIRP is that the single vehicle is allowed to make multiple trips from the depot within one cycle.

It is worthwhile to study the SV-CIRP since it arises as a sub-problem of the CIRP (Aghezzaf et al., 2006). In this case, the proposed solution procedures are branch-and-price or column generation. The SV-CIRP should also be considered as the inventory routing variant of the Orienteering Problem (OP), the "Inventory Orienteering Problem" (IOP). The goal of the regular OP (Vansteenwegen, Souffriau, & Van Oudheusden, 2011) is to maximise the total score, collected by visiting a selection of customers, without violating a time constraint. Every customer can be visited at most once and the travel times between the customers are given. If inventory and handling costs are added to the standard OP, the result will be the SV-CIRP, except that for the SV-CIRP no upper bound for the cycle time is fixed beforehand. Actually, we think the IOP would better model reality than the SV-CIRP, since the IOP can limit the cycle time to a given upper bound, e.g. a number of hours or a working day. It should be noted that Federgruen and Simchi-Levi (1995) mention this constraint as an extension of what they called inventory-routing models. However, this upper bound would only make the problem easier to solve, due to a smaller number of possible solutions. Moreover, benchmark test instances are available for the SV-CIRP, but not for the IOP. Therefore, we limit our research to the SV-CIRP in this paper, knowing that the proposed solution technique can also solve the IOP.

When solving the SV-CIRP, a number of questions need to be answered in an integrated way: Which customers should be selected? Which cycle time would minimise the combination of inventory, handling and routing costs? How should the selected customers be divided in multiple trips for the single vehicle?

Due to the complexity of this problem, a metaheuristic approach was designed to tackle it. The Iterated Local Search (ILS) metaheuristic is used as a framework to solve this problem. The general structure of ILS is described by Lourenço, Martin, and Stützle (2010): a sequence of local search solutions is built up, instead of repeating random local search trials starting from an empty solution each time. During each iteration of the algorithm, a different part of the solution is removed and then local search moves are used in order to try to reach a better solution. This framework was successful before in dealing with problems where customers need to be selected (e.g., Ribeiro & Lourenço, 2005).

The next section presents an overview of the related literature. In Section 3 a detailed definition of the SV-CIRP is presented and Section 4 discusses some insights in the complexity of the problem. These insights can be considered as a significant contribution of this paper and will be exploited by our solution approach to tackle the SV-CIRP efficiently. This approach is described in detail in Section 5. Experimental results and parameter settings are discussed in Section 6. Thanks to its specific implementation, our approach outperforms the solution approaches available in the literature and obtains many new best known solutions. Conclusions and future work are presented in Section 7.

2. Literature review

The Inventory Routing Problem (IRP) was first introduced by Bell et al. (1983). However, many variants of the problem exist, since different authors rarely define the problem in exactly the same way (Coelho, Cordeau, & Laporte, 2012). These authors classify a number of papers based on single or multiple customers, stochastic demands, direct deliveries, problems with multiple products or heterogeneous fleets. A number of comprehensive survey papers about IRP are available. Some of them are the following ones. Kleywegt, Nori, and Savelsbergh (2002) and Adelman (2004) present a classification of IRP variants based on four characteristics: finite or infinite planning horizon, deterministic or stochastic demands, limited or unlimited number of vehicles and one or multiple customers visited per trip. In a recent and comprehensive survey, Andersson et al. (2010) extend these four characteristics and focus more on industrial aspects. For instance, they discuss three different topologies: one-to-one, one-to-many and many-to-many and they add the "instant" planning horizon. In the survey of Schmid et al. (2013), inventory routing problems are discussed and situated amongst other rich vehicle routing problems in the section about "inventory management and vendor managed inventory".

The fundamental contributions in the class of the CIRP are that of Anily and Federgruen (1990), Gallego and Simchi-Levi (1990) and Hall (1992). Research about an appropriate solution technique for the CIRP led to the formulation of the SV-CIRP, since it came up as a sub-problem in a solution approach based on column generation (Aghezzaf et al., 2006). In Zhong and Aghezzaf (2011), an in-depth analysis of the SV-CIRP reveals that the objective function is non-smooth and non-convex with many local minima. This makes the problem complicated to solve for both exact and heuristic solution approaches. Recently, Haughton (2014) underlined the complexity considering the possible correlation between customer locations and demand rates.

In his PhD dissertation, Zhong (2012) presents an extensive literature survey about the IRP, the CIRP and the SV-CIRP, together with a detailed discussion about the complexity of the SV-CIRP and an efficient mathematical formulation of the problem. He also discusses the differences and similarities with some closely related problems. None of these related problems integrate the above mentioned aspects that make the SV-CIRP unique and complex to solve.

The first solution technique developed for the SV-CIRP is described by Aghezzaf et al. (2006). They solve the SV-CIRP by combining a savings-based heuristic with an insertion move. Zhong and Aghezzaf (2011) propose a steepest decent hybrid algorithm to solve the SV-CIRP to optimality. This approach reaches the objective, but appears very time consuming, since many mixed integer linear problems need to be solved. The algorithm consists of two major steps: (1) solving the SV-CIRP for a fixed cycle time as a mixed integer linear problem and (2) improving the interval of variation of the cycle time using the Frank–Wolfe method (e.g. Frank & Wolfe, 1956). Due to its time consumption, this exact approach can only be applied to small size instances, with 15 customers or less.

To deal with larger instances of the SV-CIRP, Zhong and Aghezzaf (2012) present an iterated local search (ILS) metaheuristic. This is the best (meta)heuristic currently available to deal with the SV-CIRP. As mentioned above, the general structure of ILS is described in Lourenço et al. (2010). Since we also use the ILS framework for our algorithm, we will now discuss its general structure in more detail. However, our implementation of this framework is totally different from the implementation of Zhong and Aghezzaf (2012) and so will be the performance. We discuss the most important differences when we explain our algorithm in Section 5.

ILS typically starts from an *initial solution*. This solution is further improved by a *local search heuristic*. Every time local search reaches a local optimum, an *acceptance criterion* decides whether or not the new solution is accepted. If it is accepted, the next iteration starts with a *perturbation* of this new local optimum, otherwise the perturbation is applied to the previous local optimum. After the perturbation, the local search heuristic is applied again in the next iteration in order to obtain a new local optimum. The *stopping criterion* determines when the algorithm stops.

In the improvement phase of Zhong and Aghezzaf (2012), four local search moves are considered that insert, remove and/or relocate customers in order to reduce the total cost. In order to

Download English Version:

https://daneshyari.com/en/article/479708

Download Persian Version:

https://daneshyari.com/article/479708

Daneshyari.com