European Journal of Operational Research 241 (2015) 361–369

Contents lists available at ScienceDirect

European Journal of Operational Research

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

Production, Manufacturing and Logistics

A priori policy evaluation and cyclic-order-based simulated annealing for the multi-compartment vehicle routing problem with stochastic demands

Justin C. Goodson*

Department of Operations & Information Technology Management, John Cook School of Business, Saint Louis University, 3674 Lindell Blvd, St. Louis, MO 63108, United States

ARTICLE INFO

Article history:
Received 17 December 2013
Accepted 9 September 2014
Available online 2 October 2014

Keywords:
Logistics
Stochastic multi-compartment
Vehicle routing
Cyclic-order search
Simulated annealing

ABSTRACT

We develop methods to estimate and exactly calculate the expected cost of a priori policies for the *multi-compartment vehicle routing problem with stochastic demands*, an extension of the classical vehicle routing problem where customer demands are uncertain and products must be transported in separate partitions. We incorporate our estimation procedure into a cyclic-order-based simulated annealing algorithm, significantly improving the best-known solution values for a set of benchmark problems. We also extend the updating procedure for a cyclic order's candidate route set to duration-constrained a priori policies.

1. Introduction

Vehicle routing problems with stochastic demands (VRPSDs) underlie operational problems in logistics where routes must be planned without full knowledge of customer demand levels. In such cases, actual demands are first observed upon arrival to customer locations. In some situations, products must be transported in independent vehicle compartments, giving rise to the *multi-compartment vehicle routing problem with stochastic demands* (MCVRPSD), a generalization of the single-compartment case (see the review of Gendreau, Laporte, and Séguin (1996) or the work of Laporte, Louveaux, and Van Hamme (2002) for examples). Applications of the MCVRPSD include the collection and delivery of different types and qualities of milk (Caramia & Guerriero, 2010), distribution of various grades of fuel (Brown, Ellis, Graves, & Ronen, 1987), transport of animal food (El Fallahi, Prins, & Calvo, 2008), pickup and delivery of livestock (Oppen & Løkketangen, 2008), selective waste collection, and transport of groceries requiring different levels of refrigeration.

A common approach to handle uncertainty in customer demands is to restrict attention to a priori policies (see Campbell and Thomas (2008) for a review). An a priori policy requires vehicles to visit customers in the order specified by a set of pre-defined routes, returning to the depot to replenish in the event vehicle capacity is inadequate to fully serve demand, i.e., a *route failure*. A priori routes are routinely used in industry (Erera, Morales, & Savelsbergh, 2010) and create a regularity of service that can be beneficial for both the customers and the drivers – customers may be served at roughly the same time each day they require service and the drivers become familiar with their routes.

We make two main contributions. First, we develop methods to estimate and exactly calculate the expected cost of an a priori MCVRPSD policy. Straightforward procedures exist to calculate the expected cost of an a priori policy when all products are transported in a single compartment (Teodorović & Pavković, 1992). However, in the multi-compartment case, a priori policy evaluation is more difficult. Mendoza, Castanier, Guéret, Medaglia, and Velasco (2008, 2010) and Mendoza, Guéret, Medaglia, and Velasco (2011) derive an expression to calculate the expected cost of an a priori policy, but the expression requires probability calculations to be made over random variables representing partially unserved customer demands. The distributions of these random variables are unknown, thereby making it difficult to implement the expression. Recognizing this challenge, Mendoza et al. (2008) explore methods to approximate the expected cost of an a priori policy for the MCVRPSD. When the ratio of compartment capacities to customer demands is high, Mendoza et al. (2008) find a *take-all* approximation yields good estimates of the expected cost. However, the bias of the estimate increases as the ratio of compartment capacities to customer demands decreases.

The method we propose to calculate the exact expected cost of an a priori MCVRPSD policy builds on Goodson, Ohlmann, and Thomas (2013) and applies to problem instances where customer

* Tel.: +1 314 977 2027; fax: +1 314 977 1483.
E-mail address: goodson@slu.edu

http://dx.doi.org/10.1016/j.ejor.2014.09.031
0377-2217/

demands follow discrete probability distributions with finite support. The method requires exponential time to execute and is not practical for use in optimization procedures. However, the exact method leads to a simulation scheme that provides unbiased and consistent estimates of a policy's expected cost for both discrete and continuous customer demand distributions. Because the simulation scheme can be efficiently incorporated into local search procedures, it provides an attractive alternative when the bias of the take-all approximation is unacceptable.

Our second contribution is the development of a cyclic-order-based simulated annealing procedure for the MCVRPSD. Goodson, Ohlmann, and Thomas (2012) propose cyclic-order neighborhoods as the basis of local search methods for a broad class of routing problems, of which the MCVRPSD is a member. We utilize our simulation scheme for a priori MCVRPSD policies to accelerate the initial iterations of the simulated annealing procedure. We also extend the cyclic-order updating procedure of Goodson et al. (2012) for the candidate route set to the duration-constrained a priori policies considered in this paper.

As of this writing, our method appears to be the most effective heuristic for the MCVRPSD. Without tailoring the cyclic-order search procedure to the MCVRPSD, we improve the best-known solution values for 159 of 180 benchmark problem instances and match the best-known solution values for the remaining 21 instances. These results, in conjunction with the results of Goodson et al. (2012), support cyclic-order-based local search as an effective solution procedure for a variety of routing problems.

The remainder of the paper is organized as follows. In Section 2, we formally state the MCVRPSD. In Section 3, we review related literature. In Section 4, we discuss methods to exactly calculate and estimate the expected cost of an a priori MCVRPSD policy. In Section 5, we present our cyclic-order-based simulated annealing procedure. In Section 6, we discuss the results of computational experiments. We make concluding remarks and suggestions for future research in Section 7. Appendix A provides a guide to the primary notation used in the paper.

2. Problem statement

The MCVRPSD generalizes the classical vehicle routing problem, a NP-hard optimization problem (Toth & Vigo, 2001). The MCVRPSD is characterized by a complete graph $G = (\mathcal{N}, \mathcal{E})$, where $\mathcal{N} = \{0, 1, \dots, N\}$ is a set of $N+1$ nodes and $\mathcal{E} = \{(n, n') : n, n' \in \mathcal{N}\}$ is the set of edges connecting the nodes. Nodes $1, \dots, N$ represent customer locations and node 0 represents a depot from which a set $\mathcal{M} = \{1, 2, \dots, M\}$ of identical vehicles operates. Each vehicle contains a separate compartment for each product in the product set $\mathcal{P} = \{1, 2, \dots, P\}$. The capacity of the compartment assigned to product p in $\mathcal{P}$ is Q_p. Customer demands for products are assumed to be independent random variables. The random demand for product p in $\mathcal{P}$ at customer n in $\mathcal{N}$ is denoted $D_{n,p}$. We denote the distribution function by $F_{D_{n,p}}(\cdot)$ with support $\mathcal{S}(D_{n,p})$, which we require to be a subset of the range $[0, \overline{D}]$, where $\overline{D}$ is a finite number. Prior to arrival at customer locations, customer demands are known only in distribution. Upon arrival, customer demands for each product are observed and served to the maximum extent possible, subject to available vehicle capacity for each product. When capacity in one or more compartments is exhausted (i.e., a route failure occurs), a vehicle returns to the depot and replenishes the capacity of all compartments. Vehicle routes begin and end at the depot and travel times among customers are known. The time to travel from location n to location n' is denoted $t(n, n')$. The objective of the MCVRPSD is to obtain a policy that serves customer demand with minimal expected travel time subject to a route duration limit L (e.g., end of a working day), by which time all vehicles must return to the depot. To be consistent with the literature, we use travel time to represent cost, noting that $t(\cdot, \cdot)$ may be replaced with a general cost function.

We focus on a priori routing policies for the MCVRPSD. A priori policies are characterized by a priori routes, or predetermined sequences of customers. We denote an a priori route for vehicle m in vehicle set $\mathcal{M}$ by the sequence of customers $v^m = (v_0^m = 0, v_1^m, \dots, v_{l^m}^m, v_{l^m+1}^m = 0)$. We denote by $(v^m)_{m\in\mathcal{M}}$ a set containing an a priori route for each vehicle m in $\mathcal{M}$. In a set of a priori routes, each customer appears exactly once on exactly one route. We adopt the classical *detour-to-depot* policy, the same policy employed by Mendoza et al. (2010, 2011). The policy requires vehicles to serve customers in the order they appear in a set of a priori routes. Vehicles must travel directly to the next customer on the route, i.e., preemptive capacity replenishment is not allowed. In the event of a route failure, vehicles must make return trips to the depot until customer demands are fully served. An a priori route is feasible if its expected travel time is less than or equal to the route duration limit L. If we denote by $\mathcal{V}$ the set of all a priori route sets, by $A_{v^m_{l^m+1}}$ the random arrival time of vehicle m to final destination $v^m_{l^m+1}$, and by $\mathbb{E}[A_{v^m_{l^m+1}}]$ the expected arrival time of vehicle m to $v^m_{l^m+1}$, then the problem we seek to solve is

$$\min \left\{ \sum_{m\in\mathcal{M}} \mathbb{E}\left[A_{v^m_{l^m+1}}\right] : \right. \tag{1}$$

$$\mathbb{E}\left[A_{v^m_{l^m+1}}\right] \le L, m \in \mathcal{M}, \tag{2}$$

$$\left. (v^m)_{m\in\mathcal{M}} \in \mathcal{V} \right\}. \tag{3}$$

Throughout the remainder of the paper, because we refer primarily to a single a priori route, to ease the notation we drop the superscript m and refer to a route simply by v and to the ith customer on a route by v_i.

3. Related literature

Research related to multi-compartment vehicle routing appears to begin with van der Bruggen, Gruson, and Salomon (1995), who consider the task of designing delivery routes for gasoline distribution. van der Bruggen et al. (1995) and others model customer demands as deterministic. In this section, we focus on literature treating customer demands as stochastic.

Mendoza et al. (2008) study *take-none* and *take-all* approximations of the recourse cost component of a priori policies. The take-all approximation assumes that when a route failure occurs at a particular customer, all demands for all products at that customer are served in full before replenishing capacity at the depot; the take-none approximation assumes none of the demands are served. Both approximations circumvent the issue of making probability calculations over random variables with unknown distributions. Computational experiments conclude the take-all scheme is superior, but that the quality of both approximations degrades as the ratio of compartment capacities to customer demands decreases.

The methods we propose to estimate and exactly calculate the expected cost of a priori policies serve as alternatives to the take-none and take-all approximations. Our methods differ in three ways. First, the take-none and take-all approximations can be used when demands follow discrete or continuous probability distributions, provided the convolution of random demands can be readily calculated (e.g., Poisson- or normally-distributed demands). In contrast, our exact evaluation procedure can only be used when customer demands follow discrete distributions, but there is no limit imposed by the need to calculate convolutions of random variables. The estimation procedure may be applied regardless of the nature of customer demand distributions. Second, our procedures do not depend on the assumption that demand realizations are less than or equal to compartment capacities, thereby accounting for the possibility of multiple