Transportation Research Part E 95 (2016) 95–123

Contents lists available at ScienceDirect

Transportation Research Part E

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

An adaptive large-neighborhood search heuristic for a multi-period vehicle routing problem

Iman Dayarian [a,]*, Teodor Gabriel Crainic [b,d], Michel Gendreau [c,d], Walter Rei [b,d]

[a] School of Industrial & Systems Engineering, Georgia Institute of Technology, Atlanta, GA, USA
[b] School of Management, Université du Québec à Montréal, Montreal, QC, Canada
[c] MAGI, École Polytechnique, Montreal, QC, Canada
[d] CIRRELT, C.P. 8888, succ. Centre-ville, Montreal, QC H3C 3P8, Canada

ARTICLE INFO

Article history:
Received 23 February 2016
Received in revised form 10 September 2016
Accepted 12 September 2016
Available online 30 September 2016

Keywords:
Multi-period vehicle routing problem with seasonal fluctuations
Tactical planning
Seasonal variation
Adaptive large neighborhood search

ABSTRACT

This problem involves optimizing product collection and redistribution from production locations to a set of processing plants over a planning horizon. This horizon consists of several days, and the collection-redistribution is performed on a repeating daily basis. A single routing plan must be prepared for the whole horizon, taking into account the seasonal variations in the supply. We model the problem using a sequence of periods, each corresponding to a season. We propose an adaptive large-neighborhood search with several specifically designed operators and features. The results show the excellent performance of the algorithm in terms of solution quality and computational efficiency.

1. Introduction

The vehicle routing problem (VRP) is a difficult combinatorial optimization problem that appears in many practical applications relating to the design and management of distribution systems. Studies of the classical VRP and its many variants and extensions, starting with the seminal work of Dantzig and Ramser (1959), represent a significant portion of the operations research literature (Toth and Vigo, 2002, 2014). The classical VRP, referred to as the capacitated vehicle routing problem (CVRP), concerns the determination of routes for a fleet of homogeneous vehicles, stationed at a central depot, that must service a set of customers with known demands (supplies). The goal is to design a collection of least-cost routes such that: (1) each route, performed by a single vehicle, begins at a depot, (2) each customer is visited once by exactly one vehicle, and (3) the quantity of goods delivered (collected) on each route does not exceed the vehicle capacity (Golden et al., 2008).

In many settings, e.g., the CVRP, the routing plan is executed repeatedly over a long planning horizon. The parameters of the problem, such as the quantities to be delivered (collected) at each customer location, are assumed fixed over the horizon and known a priori. However, in many real-life applications, this assumption may result in poor-quality routing plans. This occurs, for instance, in settings that display significant seasonal fluctuations in the level of supply/demand throughout the considered planning horizon. The class of problems addressed in this paper, inspired by milk collection and redistribution in the dairy industry of Quebec (see Dayarian et al., 2015b), has several problem-specific attributes and characteristics. The routing corresponds to the collection of a given product from producers' facilities followed by the distribution of the product

---

* Corresponding author.
E-mail address: iman.dayarian@isye.gatech.edu (I. Dayarian).

http://dx.doi.org/10.1016/j.tre.2016.09.004
1366-5545/

to a set of processing plants. The routes must be designed in such a way that the plant demands are completely satisfied, while every producer is visited by exactly one vehicle and each vehicle delivers to just one plant per day. We assume that the total daily quantity produced satisfies the total plant demand. Because of contractual and service-consistency requirements, a single routing plan must be designed for a given horizon.

For service consistency, each producer should always be included in the same route and serviced by the same vehicle. The drivers also use this routing plan to schedule their daily operations.

Dayarian et al. (2015b) modeled this problem as a Multi-Period VRP with Seasonal Fluctuations (MPVRPSF) and proposed an exact solution method based on a branch-and-price approach. Their solution approach provides optimal solutions for instances with up to sixty producers. However, real-life instances may have several hundred producers. Therefore, we need solution approaches that can find good but not necessarily optimal solutions to larger instances. Furthermore, since the solutions obtained from the model can be the basis for negotiations with carriers in some settings, such as the one described in Dayarian et al. (2015b), it may be imperative to be able to solve different versions of an instance repeatedly and within a short time span. Therefore, it is critical to have a fast and efficient solution approach. The main goal of this paper is to derive such a solution approach for the MPVRPSF based on the adaptive large-neighborhood search (ALNS) framework (Pisinger and Ropke, 2007; Ropke and Pisinger, 2006).

Our main contributions are:

- We design an ALNS-based metaheuristic for the MPVRPSF, a rich vehicle routing problem. The proposed solution procedure includes a set of novel algorithmic features, including several new operators based on the inherent structure of the problem. These are detailed in Section 4.
- To evaluate the quality of the solution, we compute a series of lower and upper bounds on the value of the multi-period solution. We compare the solutions obtained through the ALNS with these bounds.
- We extensively analyze the performance of the method and its components in terms of computational time and solution quality, through a series of numerical tests on a large set of randomly generated instances.

The remainder of this paper is organized as follows. In Section 2, we describe the problem and the notation that we use. Section 3 discusses the state-of-the-art in this field. In Section 4, we present our ALNS-based approach to the problem. In Section 5, we propose a series of bounds that allow us to evaluate the performance of the algorithm. The experimental results are reported in Section 6, and Section 7 provides concluding remarks.

2. Problem statement

In this section, we describe more precisely the Multi-Period VRP with Seasonal Fluctuations (MPVRPSF). As mentioned earlier, the purpose of the problem is to design a routing plan that will serve to organize transportation between a set of producers and a set of processing plants for a given horizon H of several days (typically, several months). A plan consists of a set of routes, each performed by a single vehicle on every collection day of horizon H. An unlimited fleet of identical vehicles is assumed to be available at multiple depots. On every collection day, each vehicle departs from a depot, collects a single product type from a subset of producers, delivers the collected product to a single plant, and then returns to its depot. This can be seen as an extension of the VRP with additional deliveries to multiple plants, and it is therefore NP-hard (Lenstra and Rinnooy Kan, 1981).

There are many application settings, in which a routing plan must be designed to be operated repetitively on several collection "days" over a long horizon: the collection of dairy products, poultry and eggs, beverage distribution, waste collection, etc. We are interested in environments in which the supply (demand) exhibits seasonal variations significant enough to have a major impact on the routing. Moreover, we focus on situations in which producers' (customers') supply (demand) are sufficiently strongly correlated that we can make the assumption that they are perfectly correlated.

To treat this correlation, we assume that a year can be divided into several periods, each representing a seasonal production level. We take inter-period production variations into account; the potential intra-period fluctuations are neglected. Intra-period fluctuations can often be handled by leaving a spare capacity of 5–10% on each vehicle when designing the routes. In most applications of the MPVRPSF, this correlation is expected to arise because almost all the producers/customers in a given geographical region are exposed to similar seasonal cycles. The plants must adjust their seasonal demands according to the supply so that the total supply always meets the total demand.

The proposed multi-period model has strong similarities with an *a priori* optimization framework in the context of the vehicle routing problem with stochastic demand (VRPSD). In a two-stage formulation of a stochastic problem, the solution from the first stage is updated at the second stage as the exact values of the stochastic parameters are revealed. One seeks a solution that minimizes the total expected cost of the original plan and the potential adjustments in the second stage. Similarly to algorithms for the VRPSD, in the context of our multi-period problem we design a single plan for the planning horizon in the first stage, taking into account possible supply changes between periods. In the second stage, the plan is adjusted based on the specificities of each period. In seasons with higher supply levels, a vehicle may have insufficient residual capacity to collect the supply at a given producer location. We refer to this as a *failure*. Following a failure, the vehicle usually trav-