European Journal of Operational Research 243 (2015) 52–60

Contents lists available at ScienceDirect

European Journal of Operational Research

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

Discrete Optimization

Sequential and parallel large neighborhood search algorithms for the periodic location routing problem

Vera C. Hemmelmayr*

WU (Vienna University of Economics and Business), Welthandelsplatz 1, 1020 Vienna, Austria

ARTICLE INFO

Article history:
Received 28 September 2013
Accepted 16 November 2014
Available online 26 November 2014

Keywords:
Location routing
Large neighborhood search
Parallel metaheuristics

ABSTRACT

We propose a large neighborhood search (LNS) algorithm to solve the periodic location routing problem (PLRP). The PLRP combines location and routing decisions over a planning horizon in which customers require visits according to a given frequency and the specific visit days can be chosen. We use parallelization strategies that can exploit the availability of multiple processors. The computational results show that the algorithms obtain better results than previous solution methods on a set of standard benchmark instances from the literature.

1. Introduction

The periodic location routing problem (PLRP) is an important problem in supply chain management that combines location decisions with routing decisions and visit day assignment. A set of customers has to be served with a predefined frequency from a set of capacitated depots. Although the frequency is given, the exact service days have to be determined. Furthermore, there is a fixed cost that has to be paid for each vehicle used over the planning horizon. The objective is to minimize travel costs, opening costs of the depots and fixed costs of the used vehicles.

Solving location and routing decisions simultaneously is particularly appealing in problems where these two decisions are on the same decision level. This is frequently the case in problem settings where location decisions are easier to modify, for example, when they involve rented locations or do not require big investments. Furthermore, in many problem settings the depot costs can be broken down to the considered planning horizon so that they are at the same order of magnitude as the routing costs. Location routing problems are also interesting from a strategic point of view where building detailed routes can give a better approximation of future routing costs. An extensive literature review including a description of different applications for location routing and a classification scheme is given in Nagy and Salhi (2007). More recent surveys can be found in Prodhon and Prins (2014), Drexl and Schneider (2014), Drexl and Schneider (2015) and Lopes, Ferreira, Santos, and Barreto (2013).

The PLRP was first introduced in Prodhon (2008) and it is related to the periodic vehicle routing problem (PVRP) and the location routing problem (LRP). The PVRP deals with serving a set of customers over a given planning horizon. Customers have predefined frequencies and a related set of predefined visit combinations.

The PVRP is a well-studied problem and a number of mainly heuristic solution methods have been proposed to solve it. A recent exact method was proposed in Baldacci, Bartolini, Mingozzi, and Valletta (2011). Recent meta-heuristic approaches can be found in Gulczynski, Golden, and Wasil (2011), where integer programming and the record-to-record travel algorithm were combined, and in Cordeau and Maischberger (2012), where a parallel iterated tabu search heuristic was proposed. Finally, in Vidal, Crainic, Gendreau, Lahrichi, and Rei (2012) evolutionary search, local search and elaborate population-diversity management schemes were combined.

When the visit frequency is not given, but must be decided, the problem results in a PVRP with service choice (Francis, Smilowitz, & Tzur, 2006). However, in this paper we assume that the visit frequencies are given and only the particular visit days can be chosen from a set of visit day combinations.

Moreover, the PLRP can be seen as an extension of the LRP to a planning horizon. In the LRP, a set of possible capacitated depots is given to serve customer demand. The goal is to simultaneously select a subset of these depots to open and to solve the corresponding multi-depot vehicle routing problem (MDVRP) such that the routing cost and the opening cost of the depots are minimized. Recent

* Corresponding author. Tel.: +43 1 31336 5980.
E-mail address: vera.hemmelmayr@wu.ac.at

http://dx.doi.org/10.1016/j.ejor.2014.11.024
0377-2217/© 2014 Elsevier B.V. All rights reserved.

heuristic methods for the LRP with capacitated depots and vehicles include Hemmelmayr, Cordeau, and Gabriel Crainic (2012), Contardo, Cordeau, and Gendron (2014), Escobar, Linfati, and Toth (2013) and Ting and Chen (2013). Exact solution methods were developed in Baldacci, Mingozzi, and Wolfler-Calvo (2011), Belenguer, Benavent, Prins, Prodhon, and Wolfler Calvo (2011) and Contardo, Cordeau, and Gendron (2014).

There are only a few papers in the literature that combine location and routing decisions over a planning horizon in which customers have to be served multiple times. Albareda-Sambola, Fernández, and Nickel (2012) tackled a multiperiod location routing problem where location and routing decisions are made at different time scales. In this problem, decisions regarding facility location can only be made in selected time periods of the planning horizon and cannot be changed during time periods between them. Furthermore, unlike in the PLRP, the customer specifies the exact time periods of service.

Previous solution methods for the PLRP include an iterative heuristic (Prodhon, 2008), a memetic algorithm with population management (Prodhon & Prins, 2008) and an ELS with path relinking (Prodhon, 2009). These methods are outperformed in Prodhon (2011) by a later method, a hybrid evolutionary algorithm (HELS). In this algorithm, an evolutionary local search (ELS) works on the selection of visit day combinations. The ELS is hybridized with a randomized extended Clarke and Wright algorithm (RECWA), which was originally designed for the LRP (Prins, Prodhon, & Wolfler Calvo, 2006). It is used for each day separately. Then, the depots that are opened for all days of the planning period are chosen according to different ratios that measure the suitability of opening a depot based on the percentage of daily use and the percentage of total demand satisfied per depot.

Pirkwieser and Raidl (2010) developed a variable neighborhood search (VNS) algorithm combined with ILP-based very large neighborhood searches for the LRP and the PLRP. The VNS uses neighborhoods that change visit combinations for customers, exchange segments of customers between routes of the same and of different depots and change the location decisions by opening or closing depots. They introduced three different ILP-based large neighborhood searches. The first neighborhood search, V_1, operates on a route level, in which depots can be opened or closed and routes can be relocated to different depots on the same day. The second approach, V_2 , also allows changing the visit day combinations of customers. The routes used in V_2 come from a set of feasible solutions of the VNS. The third approach is similar to ILP-based refinement techniques (De Franceschi, Fischetti, & Toth, 2006). It removes sequences of customers from given routes and an ILP is used to find the optimal insertion points. This approach is solved for each day. These results were further improved in the thesis (Pirkwieser, 2012) by adapted parameter settings, for instance a longer runtime.

Parallel computing enables the development of fast and robust solution methods for instance by exploiting the availability of multiple processors on computing clusters or multi-core processors. In the following, we will describe the most recent parallelization strategies for vehicle routing problems. For further information on parallel metaheuristics, we refer to the book of Alba (2005) and to the survey of Crainic and Toulouse (2010). Moreover, a survey with a focus on vehicle routing problems that covers heuristic and exact parallel solution methods can be found in Crainic (2008).

Crainic and Nourredine (2005) introduced a classification for parallel metaheuristics along three dimensions. The first dimension indicates whether the search is controlled by a master process (1C) or by several processes together (pC). The second dimension indicates the quantity and the quality of information exchanged. There is rigid (RS) and knowledge synchronization (KS) for synchronous communication and collegial (C) and knowledge collegial (KC) for asynchronous communication, where the respective difference is the amount and quality of information exchanged. The third dimension can be classified as SPSS, SPDS, MPSS or MPDS standing for same or multiple starting point and same or different search strategies used by the processes.

Cordeau and Maischberger (2012) proposed a parallel iterated tabu search heuristic for four different routing problems: the VRP, the PVRP, the MDVRP, and the site dependent VRP with and without time windows. In their parallel algorithm, each process starts from a different initial solution. Some of the parameters of the iterated tabu search are chosen independently in the processes. At given points, knowledge about the solutions is shared. Each process $p \in \{1, 2, \ldots, N\}$ decides whether to accept its working solution with a probability $1 - (\lambda/\eta)^2$, where λ is the current iteration and η is the total number of iterations, or go to the jth best solution, where $j = \lfloor \sqrt{p} \rfloor$. Therefore, most working solutions will be accepted in the beginning, while toward the end the best solutions will be accepted, so that the processes can focus on trying to improve the best solution. The authors can show that the algorithm yields very good results for several variants of vehicle routing problems.

Groër, Golden, and Wasil (2011) developed a parallel algorithm for the VRP that combines integer programming and heuristic search. A master process is used to control the search, while the remaining processes can be either heuristic solvers or set covering solvers. The heuristic solvers run a metaheuristic based on the record-to-record travel algorithm and the set covering solvers solve set covering problems with routes taken from the heuristic solvers.

Jin, Crainic, and Løkketangen (2012) designed an algorithm that achieves very competitive results for the VRP. They use four tabu search threads that each uses different neighborhoods. The best solutions found are exchanged periodically through the use of a solution pool.

Lahrichi et al. (2012) developed integrative cooperative search (ICS) that can tackle multi-attribute combinatorial optimization problems. ICS performs a decomposition of the problem in partial problems that are solved by partial solvers. Integrators select partial solutions and combine them to complete solutions. These complete solutions are sent to the complete solver group. In this work, they apply the method to the MDPVRP and get results that improve upon previous methods.

In the context of parallel computing, it is also important to mention GPU based computing. It takes advantage of the rapid increase in GPU performance. Modern commodity PCs include a multi-core CPU and one or more GPUs. For this parallel, heterogeneous architecture, solution methods are developed. An introduction to modern computer architectures and GPU programming is given in Brodtkorb, Hagen, Schulz, and Hasle (2013), which is part one of a survey in two parts. Part two (Schulz, Hasle, Brodtkorb, & Hagen, 2013) gives a broad overview of the existing literature on parallel computing in discrete optimization aimed at modern PCs. The survey has a strong focus on routing problems. The authors conclude that GPU computing in discrete optimization is still in its infancy. The development of solution methods that exploit the heterogeneity of modern PCs efficiently is still an open research field that is interesting and highly relevant.

In this paper, we propose sequential and parallel variants of a large neighborhood search algorithm (LNS) to solve the PLRP. The computational results show that our algorithm outperforms previous solution methods in terms of solution quality. We can show that for standard benchmark instances from the literature, substantial improvements are possible both in the average solution quality as well as in the quality of best known solutions found. Moreover, we also develop two parallel versions of the algorithm that make use of the availability of clusters of computers. We propose a simple methodology for parallelization that can easily be applied to other problems and algorithms.

The remainder of this paper is organized as follows. Section 2 gives a detailed problem description, in Section 3 the solution method