European Journal of Operational Research 247 (2015) 60–71

Contents lists available at ScienceDirect

European Journal of Operational Research

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

Discrete Optimization

A comparison of column-generation approaches to the Synchronized Pickup and Delivery Problem

Timo Gschwind*

Chair of Logistics Management, Johannes Gutenberg University Mainz, Jakob-Welder-Weg 9, Mainz D-55128, Germany

ARTICLE INFO

Article history:
Received 29 April 2014
Accepted 8 June 2015
Available online 14 June 2015

Keywords:
Vehicle routing
Pickup and delivery
Temporal synchronization
Labeling algorithm
Branch-and-cut-and-price

ABSTRACT

In the Synchronized Pickup and Delivery Problem (SPDP), user-specified transportation requests from origin to destination points have to be serviced by a fleet of homogeneous vehicles. The task is to find a set of minimum-cost routes satisfying pairing and precedence, capacities, and time windows. Additionally, temporal synchronization constraints couple the service times at the pickup and delivery locations of the customer requests in the following way: a request has to be delivered within prespecified minimum and maximum time lags (called ride times) after it has been picked up. The presence of these ride-time constraints severely complicates the subproblem of the natural column-generation formulation of the SPDP so that it is not clear if their integration into the subproblem pays off in an integer column-generation approach. Therefore, we develop four branch-and-cut-and-price algorithms for the SPDP based on column-generation formulations that use different subproblems. Two of these subproblems are considered for the first time in this paper have not been studied before. We derive new dominance rules and labeling algorithms for their effective solution. Extensive computational results indicate that integrating either both types of ride-time constraints or only the maximum ride-time constraints into the subproblem results in the strongest overall approach.

1. Introduction

In the family of one-to-one *Pickup-and-Delivery Problems* (PDPs), customer requests consist of transporting goods or people between paired origin and destination points: for each request a specific good or person has to be picked up at one location and to be transported to the corresponding delivery location. Typically, the task is to design a set of minimum-cost routes satisfying all customer requests subject to pairing and precedence, and other problem-specific constraints. For details on different PDP-variants we refer to the recent surveys (Berbeglia, Cordeau, Gribkovskaia, & Laporte, 2007; Cordeau, Laporte, & Ropke, 2008; Parragh, Doerner, & Hartl, 2008).

A well-studied one-to-one PDP is the *Pickup-and-Delivery Problem with Time Windows* (PDPTW) (e.g., Dumas, Desrosiers, & Soumis, 1991; Ropke & Cordeau, 2009; Baldacci, Bartolini, & Mingozzi, 2011) in which vehicle routes must respect pairing and precedence, capacities, and time windows. In this article, we introduce the *Synchronized Pickup and Delivery Problem* (SPDP). It extends the PDPTW by imposing additional constraints that couple the service times at the pickup and delivery locations of the customer requests in the following way: a delivery node has to be serviced within prespecified *minimum* and *maximum* time lags (called *ride times*) after the service at the corresponding pickup node has been completed. Because both pickup and delivery are performed by the same vehicle, these additional constraints are temporal *intra-route synchronization* constraints. As a generalization of the PDPTW the SPDP is clearly $\mathcal{NP}$-hard.

As pointed out, e.g., by Dohn, Rasmussen, and Larsen (2011) or Drexl (2012), synchronization aspects are highly relevant in routing practice and there is a growing interest on *Vehicle Routing Problems* (VRPs) with synchronization constraints in the research community. We see the SPDP as the prototypical VRP with temporal intra-route synchronization in the sense that synchronization takes place only within disjunctive pairs of nodes and that there are no other non-standard constraints present. In this respect, the development of an effective algorithm for solving the SPDP constitutes a central building block for the solution of richer VRPs with synchronization constraints.

A special case of the SPDP is the so-called *Dial-a-Ride Problem* (DARP) in which only a maximum ride time is specified for each pickup-and-delivery pair. The DARP mainly arises in door-to-door transportation services for school children, handicapped persons, or the elderly and disabled (see, e.g., Russell and Morrel, 1986; Madsen, Ravn, & Rygaard, 1995; Toth & Vigo, 1997; Borndörfer, Klostermeier, Grötschel, & Küttner, 1997). In this context, maximum

---

* Tel.: +49 6131 39 22085; fax: +49 6131 39 22097.
E-mail address: gschwind@uni-mainz.de

http://dx.doi.org/10.1016/j.ejor.2015.06.017

ride times are used to guarantee a certain service level by limiting the time a passenger is on board of the vehicle. A similar service-related use of maximum ride-time constraints is described by Plum, Pisinger, Salazar-González, and Sigurd (2014) in the context of liner shipping service design. When there is a limit on the total working hours of drivers (Ceselli, Righini, & Salani, 2009) or when transporting perishable goods (Azi, Gendreau, & Potvin, 2010), the time a vehicle is away from the depot has to be restricted. This can be modeled by imposing a maximum ride-time constraint on a dummy request originating and destinating at the depot. Similarly, one might want to have a limit on both the minimum and maximum duration of the routes in order to achieve an even work-distribution of the drivers.

Other applications of temporal intra-route synchronization in which minimum and maximum ride times are relevant include the planning of security guards where locations have to be inspected repeatedly within given time intervals (Bredström & Rönnqvist, 2008). There, no actual pickup at one location followed by a delivery at another location takes place. Instead, just a pairing and precedence relation between the services at the nodes forming a customer request is given. Similar planning problems arise in home health care, e.g., when patients have to be monitored by a nurse several times a day (Eveborn, Flisberg, & Rönnqvist, 2006; Rasmussen, Justensen, & Dohn, 2012). Note that in many health care applications, including those considered in Eveborn et al. (2006) and Rasmussen et al. (2012), it is not mandatory that the patients are always treated by the same nurse, i.e., these problems are of a more general nature than the one considered in this paper. The temporal aspects of this more general synchronization constraints are considered in (Dohn et al., 2011). However, personnel consistency often plays an important role in health care problems (Kovacs, Golden, Hartl, & Parragh, 2014; Rasmussen et al., 2012) so that it may be reasonable to have specific patients monitored by a single staff member only and, hence, to require pairing and precedence of the corresponding services.

The contributions of this paper are the following: First, we introduce the SPDP as the prototypical VRP with temporal intra-route synchronization. This problem has to the best of our knowledge not been considered before. Second, we develop four exact solution approaches to the SPDP based on column-generation formulations whose master programs are formulated on different sets of variables implying different subproblems. Two of these subproblems are considered for the first time in the literature. One of them is the natural subproblem of the SPDP, in which time windows as well as temporal intra-route synchronization with both minimum and maximum ride times have to be dealt with. In the other one, maximum ride-times are relaxed. We derive new dominance rules and labeling algorithms for their solution. The other subproblems are solved with algorithms proposed by Dumas et al. (1991) and Gschwind and Irnich (2015), respectively. Finally, to compare the strength of the different solution approaches, we report extensive computational results over a large number of test instances with different characteristics regarding the number of customer requests and the tightness of capacity, time-window, and minimum and maximum ride-time constraints. The analysis shows that integrating either both types of ride-time constraints or only the maximum ride-time constraints into the subproblem results in the strongest overall approach regarding the number of optimal solutions, computation times, and remaining integrality gap.

Integer column-generation methods have proven to be very successful in solving many VRP-variants including PDPs (e.g., Dumas et al., 1991; Ropke & Cordeau, 2009; Baldacci et al., 2011). The column-generation master programs of such approaches typically are extended set-partitioning models formulated on variables representing feasible routes for the problem at hand. These formulations provide stronger bounds compared to other formulations like, e.g., arc-flow formulations or extended set-partitioning models formulated on a relaxed set of variables, if the respective subproblem does not possess the integrality property (Lübbecke & Desrosiers, 2005). This is the case for many VRPs where the subproblems are typically Elementary Shortest-Path Problems with Resource Constraints (ESPPRC, Desaulniers, Desroisiers, Ioachim, Solomon, Soumis, & Villeneuve, 1998). However, the overall success of an integer column-generation approach for VRP-variants relies not only on strong bounds but also on the effective solution of the subproblem.

This is the main challenge when synchronization comes into play (Drexl, 2012). In the case of inter-route synchronization, additional constraints have to be included in the master programs (Desaulniers et al., 1998). Because of the dual variables associated with these constraints, the resulting subproblems are highly complex (e.g., Christiansen & Nygreen, 1998; Ioachim, Desrosiers, Soumis, & Bélanger, 1999; Dohn et al., 2011) and cannot be solved by standard dynamic-programming labeling algorithms. This is also true for intra-route synchronization where no additional linking constraints are necessary. There, the increased complexity of the subproblems is not caused by additional duals but by the synchronization constraints themselves, which may be hard to incorporate into the subproblem. For the DARP, e.g., Hunsaker, Savelsbergh, and problems (2002) have demonstrated that in the presence of time windows and maximum ride times checking the feasibility of a given route is intricate. Clearly, the effective generation of such routes within a column-generation approach is even more challenging.

In the case of intra-route synchronization, the complexity of the subproblems can be reduced by relaxing one or more types of constraints in the subproblem and handling them in the master programs instead (see, e.g., Ropke & Cordeau, 2005 for the DARP or Cherkesly, Desaulniers, & Laporte, 2014 for the PDPTW with LIFO Loading). The resulting easier-to-solve subproblems come at the cost of weaker lower bounds and, thus, larger branch-and-bound trees. Often, it is a priori not clear what is the best compromise between the strength of the CG formulation and the hardness of the subproblem.

The recent work of Gschwind and Irnich (2015) provides insights regarding this trade-off for the DARP: they proposed a branch-and-cut-and-price algorithm that handles all route constraints of the DARP in the subproblem which is solved by means of an effective labeling algorithm. In a computational study, they compared the strength of their approach to the branch-and-cut-and-price algorithm of Ropke and Cordeau (2005) that uses a subproblem in which the maximum ride-time constraints are relaxed. The results indicated that their approach significantly outperforms the algorithm of Ropke and Cordeau (2005) in terms of computation times and number of solved instances. However, they also tested their approach with a different labeling algorithm that uses a weaker dominance rule and observed that in this case the approach with the relaxed subproblem of Ropke and Cordeau (2005) shows the better overall performance. Decisive for the success of the approach using the stronger formulation, thus, is the availability of an effective pricing procedure for the harder subproblem.

Compared to the DARP, the additional presence of minimum ride times significantly complicates the natural subproblem of the SPDP. As a result, the dominance rule that we are able to derive for its solution is much weaker compared to those that can be used for the subproblems in which one or both types of ride-time constraints are relaxed. Therefore, we propose and compare the efficiency of four column-generation algorithms for the SPDP. Each algorithm uses a different subproblem: one that handles all route constraints of the SPDP, one that relaxes the minimum ride times, one that relaxes the maximum ride times, and one that relaxes both types of ride-time constraints.

The remainder of the paper is organized as follows. Section 2 defines the SPDP and presents column-generation formulations of it. The dominance rules and labeling algorithms we use for solving the different subproblems are detailed in Section 3. In Section 4, we briefly describe our basic branch-and-cut-and-price algorithm and