European Journal of Operational Research 236 (2014) 736–747

Contents lists available at ScienceDirect

European Journal of Operational Research

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

Innovative Applications of O.R.

Fleet routing position-based model for inventory pickup under production shutdown

N.C.P. Edirisinghe [a],*, R.J.W. James [b]

[a] College of Business Administration, University of Tennessee, Knoxville, TN 37996, USA
[b] College of Business and Law, University of Canterbury, Christchurch, New Zealand

ARTICLE INFO

Article history:
Received 14 August 2012
Accepted 25 December 2013
Available online 7 January 2014

Keywords:
(B) Inventory routing
Vendor-managed inventory
Pickup and delivery optimization
Off-shore barge scheduling

ABSTRACT

This paper addresses the problem of collecting inventory of production at various plants having limited storage capacity, violation of which forces plant shutdowns. The production at plants is continuous (with known rates) and a fleet of vehicles need to be scheduled to transport the commodity from plants to a central storage or depot, possibly making multiple pickups at a given plant to avoid shutdown. One operational objective is to achieve the highest possible rate of product retrieval at the depot, relative to the total travel time of the fleet. This problem is a variant (and generalization) of the inventory routing problem. The motivating application for this paper is barge scheduling for oil pickup from off-shore oil-producing platforms with limited holding capacity, where shutdowns are prohibitively expensive. We develop a new model that is fundamentally different from standard node-arc or path formulations in the literature. The proposed model is based on assigning a unique position to each vehicle visit at a node in a chronological sequence of vehicle-nodal visits. This approach leads to substantial flexibility in modeling multiple visits to a node using multiple vehicles, while controlling the number of binary decision variables. Consequently, our position-based model solves larger model instances significantly more efficiently than the node-arc counterpart. Computational experience of the proposed model with the off-shore barge scheduling application is reported.

1. Introduction

This paper considers the routing problem for a fleet of vehicles dispatched to collect a perishable (or time-sensitive) commodity which is available (or produced) at a known set of locations (or nodes), and the commodity must be delivered to a collection facility (or central depot). Such problems occur frequently in practice, for example, collecting blood at various locations to be delivered to a central blood bank, collecting milk at various farms to be delivered to a central cold storage/warehouse, or in scheduling multiple barges to pick up oil from various offshore platforms and delivered to an onshore depot. Nodal production of the commodity may be continuous and at differing rates across nodes and, depending on the commodity produced, may have a predetermined lifetime beyond which the product may become obsolete. In a different situation, a nodal production facility may have limited on-site storage for the commodity and the inability to pick up (or remove) the product may possibly force a prohibitively expensive production shutdown at the node. Therefore, several vehicles may have to visit the same location possibly multiple times in order to avoid production shutdowns within a given time horizon. Considering the repeatability aspect of such an operation into the future, pre-specifying a time horizon, or cycle time, with optimal initial conditions is difficult. The focus of this paper is solving the problem with deterministic nodal production rates.

The above problem is a generalization of the standard vehicle routing problem (VRP), e.g. Bodin, Golden, Assad, and Ball (1983) and Magnanti (1981), and the VRP specified with time windows of service, termed VRPTW, e.g. Solomon (1987), Solomon and Desrosiers (1988), Kolen, Kan, and Trienekens (1987), Desrochers, Desrosiers, and Solomon (1992), Desaulniers, Lessard, and Hadjar (2008), and Jepsen, Petersen, Spoorendonk, and Pisinger (2008). In VRPTW, time windows are static because the earliest and latest time within which a node must be serviced is known ‘a priori’. There are several variants of static time windows; for instance, imposing strict time limits on route duration (when perishable goods are transported) as in Azi, Gendreau, and Potvin (2010), allowing traveling times and costs to be time-dependent functions as in Hashimoto, Yagiura, and Ibaraki (2008), and the case of non-static and sliding time windows in Ferland and Fortin (1989).

In our problem, multiple service calls at a node in a horizon may become mandatory due both to the schedule repeatability requirement, as well as to guarantee that neither the product is time-expired nor the nodal storage limit is exceeded prior to the

---

* Corresponding author. Tel.: +1 865 974 1684.

E-mail addresses: chanaka@utk.edu (N.C.P. Edirisinghe), ross.james@canterbury.ac.nz (R.J.W. James).

0377-2217/$ - see front matter © 2014 Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.ejor.2013.12.039

subsequent cycle of operation. Each successive visit-time is, then, required to be within a time window that is dependent on the nodal inventory level, which depends on the prior visits (and inventory removed thus far from the node) and initial stocks. Therefore, each nodal visit window is tour-dependent, hence dynamic.

The problem addressed in this paper is, therefore, more closely related to the class of problems known as Inventory Routing Problem (IRP) where the nodes are consumers rather than producers. See Moin and Salhi (2007), Andersson, Hoff, Christiansen, Hasle, and Lokketangen (2010), and Coelho, Cordeau, and Laporte (2013) for recent reviews of the IRP literature. The basic IRP is concerned with the distribution of a single product from a central facility (or vendor) to a set of customers (or retailers) over a given planning horizon. Customer consumption rates and maximum inventory capacities are known. Product is distributed from the vendor using a fleet of homogeneous (and known) capacity to the customers who do not specify an order size or when to deliver. Instead, the vendor determines when it is best to deliver and in what quantities, possibly multiple times a day, with an intent to minimize distribution and inventory costs, while incurring no stockouts at any customer. Thus, the general setup in IRP is that of a vendor managed inventory (VMI) system, which is becoming popular in logistics where a vendor manages the inventory of its customers, resulting in distribution cost savings to the supplier by better coordinating deliveries to customers, see e.g. Campbell, Savelsbergh, Clarke, and Kleywegt (1998) and Campbell and Savelsbergh (2002, 2004).

As the IRP problem is a generalization of the VRP, solution of these problems is very difficult. As such, numerous heuristic techniques, ranging from simple heuristics to more complex metaheuristics, are primarily used to solve the IRP, see Coelho et al. (2013). More recently, however, some exact techniques have been proposed for solving smaller problems. One exact technique based on a branch-and-cut approach for the multi-vehicle IRP was proposed by Coelho and Laporte (2013). They were able to solve problems with up to 50 customers over 6 time periods to optimality using their technique, while problems with up to 200 customers and 6 time periods were attempted, but could not be solved in the 24 hour (cpu) time limit imposed. The use of mathematical programming techniques to solve all or part of these types of problems is currently a fertile area of research.

One extension of the problem that is of particular relevance is the addition of a periodic or repeatability requirement to the schedule. See, for instance, Francis, Smilowitz, and Tzur (2008) for a review of the Periodic Vehicle Routing Problem (PVRP) and its extensions. They identify that the key difference between the VRP and the PVRP is that the PVRP requires the selection of a "visit schedule" over a given planning horizon, thus not all nodes are visited every day but all nodes will be visited the required number of times within the planning horizon. Solution methods proposed in the literature include heuristics, metaheuristics and mathematical programming approaches based on Lagrangian relaxation. Moreover, the Periodic Inventory Routing Problem (PIRP) was considered by Gaur and Fisher (2004) in the context of a supermarket chain. They formulate the multiple visits to nodes as a set of shortest-path problems to form clusters, clusters are chosen using a set-partitioning formulation which is solved using a heuristic. Aksen, Kaya, Salman, and Akca (2012) address the Selective and Periodic Inventory Routing problem (SPIRP) in a reverse logistics context, whereby part of the problem is to determine which customers to service realizing that some customers may not be visited at all if it is not profitable to do so. They propose two MILP formulations based on a single commodity flow formulation, one based on selecting a node to be visited in each period, while the second is based on predefined schedules which are allocated to each node. With a two-hour CPU time limit, solutions were obtained to 25 node problems within 3.28% of optimality.

While related to the IRP, our problem has a few notable differences from the standard IRP and the periodic version of the problem. In IRP or PIRP, the objective is to minimize the total travel and inventory costs. This is not a useful objective for the problem in this paper as the central facility is concerned with the quantity of oil recovered, and labor staffing the fleet is a major variable cost. Therefore, the objective is to recover the product at the depot at the highest possible rate relative to the time the fleet spends away from the depot, i.e., achieve the maximum fleet utilization while avoiding shutdowns in all future periods. Thus, we are dealing with a fractional objective involving the volume of oil collected (which is directly related to the cycle period) and the total time the fleet are at sea. Furthermore, the minimal number of vehicles needed to guarantee no nodal shutdowns is unknown ahead of time and the fleet may be heterogeneous. Additionally, with the required repeatability of fleet schedules for long-term operational success, nodes may need to be visited multiple times within a given schedule (possibly by different vehicles) for inventory removal. This implies that both the total number of times a node is visited by the fleet and the required length of the time horizon to guarantee repeatability of fleet routes are unknown 'a priori', and they must be determined optimally within the model. In this sense, the modeling situation in this paper is more general than that of IRP or PIRP.

Our routing problem is also close to the so-called Pickup and Delivery Problem (PDP). The class of PDPs has a growing interest due to its applications in planning situations in transportation logistics and public transit. In the PDP, see Savelsbergh and Sol (1995), objects, people, or commodities must be collected and transported among locations of origins and destinations without any transshipment. The schedule for each vehicle is a sequence of pickup and delivery locations along with associated arrival and departure times at each location. In practice, most pickup and delivery problems have to deal with restrictions requiring pickup and delivery within certain time intervals, the class of problems known as PDP with time windows, or PDPTW, see Dumas, Desrosiers, and Soumis (1991). Applications of such models include airline scheduling, or the so-called *dial-a-ride* problem, in which requests are received to pick up personnel by a cab company or letters/parcels by a courier company.

In an earlier attempt, Edirisinghe, Bowers, and Agarwal (2010) used the conventional node-arc (NA) approach to model the problem of scheduling barges to pickup oil from offshore oil platforms and delivering to a central depot. In an NA formulation, binary decision variables are defined for each vehicle corresponding to the traversal of a given arc linking two nodes, with identifiers for the visit numbers at each of the two nodes since multiple nodal visits may be needed to avoid shutdowns. However, since the number of such binary variables in the NA model grows with the visit frequency, the optimal value of which is unknown 'a priori', the latter reference limited (and pre-specified) the maximum number of such visits per node within a cycle by all vehicles. Close variants of the problem for ammonia pickup reported in Christiansen (1999) and the ship scheduling problem to deliver multi-commodities between supply and demand ports reported in Al-Khayyal and Hwang (2007) are all node-arc formulations as well. These node-arc formulations grow in problem size quite dramatically as the fleet size, number of nodes, and the maximum number of visits to a node are increased. These models cannot determine an optimal visit frequency for nodes or an optimal horizon length of repeatable operation. Although NA models are generally easier to formulate, the price to be paid is the astronomical growth of the model-size as well as the sub-optimality of schedules so-determined; in Edirisinghe et al. (2010), the model is computationally intractable