European Journal of Operational Research 198 (2009) 748–759

Contents lists available at ScienceDirect

European Journal of Operational Research

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

Production, Manufacturing and Logistics

A stochastic dynamic traveling salesman problem with hard time windows

Tsung-Sheng Chang [a,*], Yat-wah Wan [a], Wei Tsang OOI [b]

[a] Institute of Global Operations Strategy and Logistics Management, National Dong Hwa University, Hualien, Taiwan
[b] Department of Computer Science, National University of Singapore, Singapore

ARTICLE INFO

Article history:
Received 25 October 2006
Accepted 12 October 2008
Available online 1 November 2008

Keywords:
Traveling salesman problem
Just-in-time
Time windows
Stochastic, dynamic network

ABSTRACT

Just-in-time (JIT) trucking service, i.e., arriving at customers within specified time windows, has become the norm for freight carriers in all stages of supply chains. In this paper, a JIT pickup/delivery problem is formulated as a stochastic dynamic traveling salesman problem with time windows (SDTSPTW). At a customer location, the vehicle either picks up goods for or delivers goods from the depot, but does not provide moving service to transfer goods from one location to another. Such routing problems are NP-hard in deterministic settings, and in our context, complicated further by the stochastic, dynamic nature of the problem. This paper develops an efficient heuristic for the SDTSPTW with hard time windows. The heuristic is shown to be useful both in controlled numerical experiments and in applying to a real-life trucking problem.

1. Introduction

1.1. Motivational application: JIT transport service

The JIT operating concept is introduced by Toyota for precise co-ordination in production, distribution, inventory, and scheduling. Its goal is to deliver, collect, or deliver and collect goods in the right place at the right time to minimize the total inventory cost. Experience shows that the effective implementation of the JIT concept can dramatically reduce inventories in different stages of a supply chain. With the adoption of cross docking, merge-in-transit, and e-fulfillment, the demand for scheduled trucking service has spread from the manufacturing to the service sector, especially in developed countries with keen concern in process improvement. For example, in Osaka and Kobe, Japan, as early as 1997, 52% (by weight) of cargo deliveries and 45% of cargo pickups had designated time windows or specified arrival times (Taniguchi et al., 2001). The trend of scheduled trucking service is likely to prevail and all trucking operators should prepare for it.

Ideally, goods should be handled at exact time when they are needed. In applying this to freight transportation, the JIT pickup/delivery in fact specifies time windows of jobs, possibly as narrow as 15 minutes or half an hour (Groenevelt, 1993). Such time windows can further be classified as *soft* with penalties on earliness and tardiness, or as *hard* with earliness and tardiness generally forbidden. In cases of hard time windows, unlike in manufacturing settings where buffer space may not be available in production facilities for goods arriving early, some pickup and delivery services permit early vehicles to wait till the beginning of time windows (Cordeau et al., 2002).

Despite its importance, to implement JIT pickup/delivery is very challenging. With traffic volumes varying throughout a day, the travel time of a route segment generally depends on the time of the day. Moreover, service times at customers are often of random duration. The routing decisions are affected by the stochastic, dynamic nature of the problem. Improper application of routing or scheduling policy on any such stochastic dynamic transportation network often leads to dissatisfied customers for late service or idle vehicles for early arrivals.

In this paper, we consider precisely the JIT pickup/delivery service: A vehicle starting from a depot needs to serve a given number of customers at different locations (nodes) and returns to the depot at the end. The customers can be served in any sequence, though there is a time window for the earliest and the latest service start times of each customer. Each processing sequence of customers defines a route of the vehicle. On each route, the vehicle moves from one customer to another, spending time both in traveling among customers and in serving customers at their locations. At each location, the vehicle may pick up goods for or deliver goods from the depot; however, it does not provide moving service of goods among customers, which sets precedent constraints among them. The travel and service times of the

* Corresponding author. Fax: +886 3 8633160.
E-mail addresses: ts@mail.ndhu.edu.tw (T.-S. Chang), ywan@mail.ndhu.edu.tw (Y.-w. Wan), ooiwt@comp.nus.edu.sg (W.T. OOI).

0377-2217/$ - see front matter © 2008 Elsevier B.V. All rights reserved.
doi:10.1016/j.ejor.2008.10.012

vehicle are random with known distributions that possibly change with the time period in the problem horizon. The nature of the problem makes it a *stochastic dynamic traveling salesman problem with time windows* (SDTSPTW).

1.2. Literature review

With or without time windows, traveling salesman problems are NP-hard in deterministic settings. In fact, even the feasibility problem with time window is NP-complete (Savelsbergh, 1984). See Solomon and Desrosiers (1988) that describe early papers to solve the traveling salesman problems with time windows by dynamic programming, branch and bound, and exploiting the tightness of time windows. In recent years, metaheuristics, local search, and hybrid algorithms of optimization and searching approaches are applied to such problems (e.g., Chiang and Russell (1996), Ibaraki et al. (2005) and Focacci et al. (2002)). Our SDTSPTW is further complicated by its stochastic and dynamic nature. In general, it takes extensive computational effort to identify the optimal route, and is practically impossible to do so even for medium number of customers.

There are very few studies that consider a routing problem with stochastic travel times and time windows. Jula et al. (2006), precisely for the SDTSPTW, is an exception. They also estimate means and variances of arrival times of nodes, and remove routes that are dominated by others. However, the approximations in Jula et al. (2006) and ours are based on different ideas. The estimates in Jula et al. (2006) take the first-order approximation of changes of arrival times in means and variances, with the effect of the second and higher order derivatives dropped. Our approximation provides a quick way to compute means and variances of arrival times under the normal assumption of arrival times. The states in Jula et al. (2006) are vectors of sequences of all nodes visited, and ours are singletons representing nodes just visited.

By definition, our method considers a single vehicle in this traveling salesman problem. If the time-windows are too tight for a single vehicle, our algorithm will report so. One or more vehicles should be added in these cases. There are papers that consider vehicle routing problems with stochastic travel times. For example, the stochastic programming approach in Laporte et al. (1992), the Markov decision process approach in Kim et al. (2005), the insertion heuristic for dynamic travel times in Potvin et al. (2006), the genetic algorithm for time-dependent travel time VRP in Haghani and Jung (2005), and the three-level local search based heuristic for quick route formation in Du et al. (2005). However, this paper deals only with the SDTSPTW and skips its interesting variants in the vehicle routing domain.

1.3. The difficulties of the SDTSPTW

Naturally, the SDTSPTW is an extremely difficult problem. Consider the small four-link network in Fig. 1 such that the travel times $D'_{01},\ D''_{01},\ D'_{12}$, and D''_{12} along the links are random, independent, and time dependent. The service times at nodes are taken to be zero, which do not affect the validity of the inference drawn below. Similarly, the double links between two nodes as shown in Fig. 1 do not affect the generality of its inferences. It is always possible to define dummy nodes to eliminate parallel links between two nodes.

Suppose that the vehicle takes the upper route leading to completion time $D'_{01} + D'_{12}$. Other than some special distributions for D'_{01} and D'_{12}, the distribution of $D'_{01} + D'_{12}$ can only be found numerically from the joint distribution of D'_{01} and D'_{12}. This is true in general, even for static problems with fixed distributions, the evaluation of means and variances of event times at nodes of a network can only be computed numerically.

The dynamic nature of the problem further complicates the computation of event times in three ways. The least effect is that the parameters of joint distributions in computation change with time. More importantly, since the distributions of D'_{12} and D''_{12} change with time, there is no guarantee that an early arrival to node 1 leads to a better route. Thus, one needs to keep track of partial routes of different arrival times to node 1, where the number of these partial routes can be excessive in large networks. The dynamic nature of the problem further eliminates the concept of optimal route in static, deterministic networks: the optimal decision at a node can be a complicated function of time. For example, the best route from node 1 to node 2 may depend on the arrival time at node 1; taking D'_{12} can be the best for arriving at node 1 at a particular moment, though for an infinitesimally short moment later taking D''_{12} can become better, and this flip-flop of optimal links can take place continually as time evolves. The determination of the optimal policy from the node to the end node is equivalent to solving a stochastic dynamic programming problem with uncountable state space and uncountable decision space. For such a problem, the optimal policy can only be found for small-size problems with special travel and service time distributions; see, e.g., Hall (1986), for an elegant example of such a problem.

1.4. Two key ideas for our solution techniques

In this paper, we develop an efficient heuristic, the FAN algorithm, to solve this SDTSPTW when time windows are hard. The key ideas of the FAN algorithm are from the *n-path* relaxation of a deterministic traveling salesman problem (TSP) by Houck et al. (1980) and the convolution–propagation approach (CPA) by Chang et al. (2005) to approximate travel times in stochastic, dynamic networks.

1.4.1. The n-path relaxation and loop elimination

For a traveling salesman problem to minimize the total transportation cost for n customers (i.e., nodes), Houck et al. (1980) find a lower bound of the objective function by solving shortest-path problems of at most n links. Refer to a (partial) route of l links as an l-path, $l \leqslant n$, which is built up by adding links one by one; it is an l-path of node j when the route of l links ends at the node. There is no guarantee that

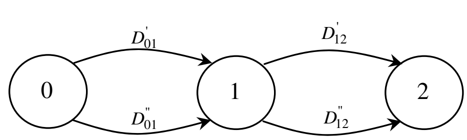

Fig. 1. A small stochastic, dynamic network.