Transportation Research Part E 109 (2018) 1–19

Contents lists available at ScienceDirect

Transportation Research Part E

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

Adaptive orienteering problem with stochastic travel times

Irina Dolinskaya*, Zhenyu (Edwin) Shi, Karen Smilowitz

Department of Industrial Engineering and Management Sciences, Northwestern University, 2145 Sheridan Road, Evanston, IL 60208, United States

ARTICLE INFO

Keywords:
Orienteering problem
Adaptive path
Dynamic programming
Variable neighborhood search
Search and rescue

ABSTRACT

In this paper, we evaluate the extent to which one can increase the likelihood of collecting greater reward in an orienteering problem with stochastic travel times by adapting paths between reward nodes as travel times are revealed. We evaluate whether this adaptivity impacts the choices of reward nodes to visit in a setting where the agent must commit to reward nodes before commencing operations. We explore the computational challenges of adding adaptive consideration in the selection of reward nodes to visit and examine the extent to which one can capture some of the benefits of adaptivity with a simpler model.

1. Introduction

This paper is motivated by search and rescue (SAR) operations in a post-disaster setting, such as one faced by SAR teams following an earthquake, severe storm or flood. The goal of each team is to search damaged or collapsed structures in the affected area to rescue as many survivors as possible within a specified deadline (usually 48–72 h in such settings). The SAR teams can commit to a set of structures to search at the beginning of the time horizon to ensure coverage of the area and avoid redundancy and inefficiency. As a first step in the study of such operations, in this paper, we look at the decisions of a single SAR team. To accurately capture the movement of a SAR team on a road network, we consider a network where some of the nodes are reward nodes, corresponding to collapsed buildings with potential survivors, and the remaining nodes correspond to road intersections. Furthermore, since immediately following a disaster, such as an earthquake, little information about the transportation infrastructure may be known, edge travel times are often stochastic. Thus, one must travel through the network to learn the network and can benefit from dynamically adjusting the paths in response to the learned information.

We model search and rescue operations as a variation of an orienteering problem (OP). OP is closely related to the Selective Traveling Salesman Problem (STSP) and the Traveling Salesman Problem (TSP) with Profits. The objective of the OP is to service a subset of customers within a specified deadline to maximize the total reward collected from the serviced customers. While most existing work studies deterministic versions of the problem, to account for transportation network and infrastructure uncertainty, problems with stochastic travel times have also been considered. In this paper, we study a novel problem setting where multiple paths with stochastic travel times exist between the reward nodes, represented by a more detailed road network (see Fig. 1). The travel time on each edge may be stochastic and, once traversed, is observed and remains constant for the remainder of the problem horizon.

We consider a setting in which one must commit to reward nodes and the visit sequence to these nodes before operations begin; however, multiple paths between reward nodes exist. By fixing the selection and sequence of reward nodes we are able to (1) represent commitment essential in many applications and (2) isolate the impact of path choice in the analysis of adaptive decision making. If the agent can adapt paths between reward nodes based on realized travel time, to what extent could the reward increase? Furthermore, to what extent could the initial set of reward nodes selected change? We show that a *fully adaptive model* greatly

* Corresponding author.
E-mail address: irina.s.dolinskaya@gmail.com (I. Dolinskaya).

https://doi.org/10.1016/j.tre.2017.10.013
Received 13 January 2016; Received in revised form 19 September 2017; Accepted 16 October 2017
1366-5545/ Published by Elsevier Ltd.

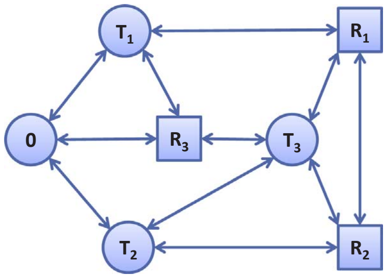

Fig. 1. Example of detailed network from the depot node (0) to reward nodes (R_1,R_2,R_3) through possible transition nodes (T_1,T_2,T_3).

increases the computational burden of selecting reward nodes to service. Therefore, we also consider a *less adaptive model* that chooses reward nodes to service and the sequence of service based on expected travel time (ignoring adaptivity) in the planing phase and allows adaptivity only during operations. Both models are compared against a *base model with no adaptivity* in which nodes to service, the sequence and the specific paths between nodes are chosen base on expectation and no adaptation is allowed during operation.

The key contribution of this work to the transportation routing literature is the **integration of the adaptive path finding component** as information about the traversed part of the network is updated. To achieve this, we present **adaptive path planning methods** to take advantage of **dynamically updating data**; combine the orienteering problem and optimal path finding into a single model; and study a new class of problems, which we call the Adaptive Orienteering Problem with Stochastic Travel Times (AOPST). The AOPST has the following unique characteristics.

1. Multiple feasible paths may exist between a pair of reward nodes, which should be considered in an adaptive approach.
2. Travel time distributions for paths between reward nodes are often not independent of each other, since paths may share common edges.
3. Since paths traversed to reach earlier reward nodes affect knowledge of the travel time distributions of the paths to the future nodes, the objective function is not additive in the subset of reward nodes selected to be serviced.
4. The selection of reward nodes to be serviced and the adaptive pathfinding policy to be followed between those nodes are made simultaneously.

The rest of the paper is organized as follows. Section 2 positions our problem in the existing literature and identifies outstanding gaps that are addressed in this paper. In Section 3, we present in detail our problem setting, introduce multiple models corresponding to levels of adaptivity and provide an overview of the solution approach. The following two sections describe the developed solution methods for the Master Problem (Section 4) and the Pathfinding Subproblem (Section 5). Implementation of the developed solution methods and numerical results are presented in Section 6. The concluding remarks and future work are discussed in Section 7.

2. Position in literature and related work

The Adaptive Orienteering Problem with Stochastic Travel Times is related to the orienteering problem, stochastic shortest path finding and adaptive routing problems, as well as stochastic knapsack problem. In this section, we review the existing work in these areas and establish the gaps addressed in this paper.

2.1. Orienteering problem

Most versions of the OP are restricted to deterministic setting (see reviews in Feillet et al. (2005), Sevkli and Erdoan Sevilgen (2006), Vansteenwegen et al. (2011)). Work considering OP with stochastic components is more limited. Teng et al. (2004) consider a variation of an orienteering problem with stochastic travel and service times, where the total travel and service time can be exceeded within a preset tolerance at a penalty. Evers et al. (2014) study OP with stochastic weights (on arc travel times) with a hard travel deadline constraint. In both settings, the problems are formulated as a two stage stochastic programming model with recourse. Tang and Miller-Hooks (2005) study a similar STSP with stochastic service and travel times. The authors use a chance constraint model to bound the probability of finishing the tour before the deadline and find an a priori tour to service a subset of customers that maximizes expected profit. Ilhan et al. (2008) study the orienteering problem with stochastic profits and deterministic travel times with the objective to maximize the probability of collecting reward above some threshold value, subject to finishing the travel within a deadline.

The orienteering problem studied by Campbell et al. (2011) is the most closely related to AOPST. The authors assume stochastic travel and service times on a complete graph, with deterministic rewards and penalties associated with arriving to a customer before and after the deadline, respectively. The problem of finding an a priori route (subset of customers and servicing order) that maximizes the expected profit is modeled with a dynamic programming framework. The authors use state dominance relationships to prune some suboptimal states. Then, a variable neighborhood search (VNS) heuristic is employed to find solutions more efficiently. In the