European Journal of Operational Research 256 (2017) 462–472

Contents lists available at ScienceDirect

European Journal of Operational Research

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

Stochastics and Statistics

Worst-case demand distributions in vehicle routing

John Gunnar Carlsson[a,*], Mehdi Behroozi[b]

[a] Department of Industrial and Systems Engineering, University of Southern California, United States
[b] Department of Mechanical and Industrial Engineering at Northeastern University, United States

ARTICLE INFO

Article history:
Received 11 May 2015
Accepted 28 March 2016
Available online 22 April 2016

ABSTRACT

A recent focal point in research on the vehicle routing problem (VRP) is the issue of robustness in which customer demand is uncertain. In this paper, we conduct a theoretical analysis of the demand distributions whose induced workloads are as undesirable as possible. We study two common variations of VRP in a continuous approximation setting: the first is the VRP with time windows, and the second is the capacitated VRP, in which regular returns to the vehicle's point of origin are required.

1. Introduction

Since its original formulation in 1959, two of the primary features that have distinguished the *vehicle routing problem* (VRP) from the *travelling salesman problem* (TSP) have been the introduction of *capacities* on vehicles that originate from a central depot (Dantzig & Ramser, 1959) and the imposition of *time windows* that constrain the times when customers can be visited (Orloff, 1976). Not surprisingly, the imposition of such constraints presents a major obstacle in obtaining solutions to a problem instance, both in terms of the added computational burden of finding solutions and in the overall quality of the solution itself. As identified in Bertsimas and Simchi-Levi (1996), one useful feature of the capacitated VRP is that one can actually describe the additional cost somewhat concretely:

> Any solution for the capacitated VRP has two cost components; the first component is proportional to the total "radial" cost between the depot and the customers. The second component is proportional to the "circular" cost; the cost of traveling between customers. This cost is related to the cost of the optimal traveling salesman tour. It is well known (Beardwood, Halton, & Hammersley, 1959) that, for large N, the cost of the optimal traveling salesman tour grows like $\sqrt{N}$, while the total radial cost between the depot and the customers grows like N.... Therefore, it is intuitive that when the number of customers is large enough the first cost component will dominate the optimal solution value.

The additional cost due to time windows is more difficult to quantify, although insights can be made under certain assumptions (Daganzo, 1987a):

> Imagine an extreme case, where only a tiny fraction of the customers have very stringent...time window constraints. Because, as we shall see, the distance travelled increases with [the square root of the number of time windows], the total system cost may be large because of the requirements of very few customers.

This paper addresses VRP from the perspective of a *continuous approximation model*: we assume that a fleet of vehicles must provide service to a contiguous planar geographic region, and our goal is to quantify precisely the role that vehicle capacities and time windows play in the worst-case workloads of the vehicles. We assume that customer demands are independently sampled from a (possibly unknown) demand distribution, and study the asymptotic behavior of the worst-case distributions as the number of customers becomes large. In this sense, our paper is philosophically similar to (for example) Burns, Hall, Blumenfeld, and Daganzo (1985), which analytically determines trade-offs between transportation and inventory costs, (Huang, Smilowitz, & Balcik, 2013), which shows how to route emergency relief vehicles to beneficiaries in a time-sensitive manner, and Jabali, Gendreau, and Laporte (2012), which describes a simple geometric model for determining the optimal mixture of a fleet of vehicles that perform distribution. The basic premise of the continuous approximation paradigm is that one replaces combinatorial quantities that are difficult to compute with simpler mathematical formulas, which (under certain conditions) provide accurate estimations of the desired quantity (Campbell, 1992; Geunes, Shen, & Emir, 2007; Novaes, de Cursi, & Graciolli, 2000; Ouyang, 2007). Such approximations exist for many combinatorial problems, such as the travelling

* Corresponding author. Tel.: +1 213 740-3858; fax: +1 213 740 1120.
E-mail address: jcarlsso@usc.edu (J.G. Carlsson).

http://dx.doi.org/10.1016/j.ejor.2016.03.047
0377-2217/

salesman problem (Beardwood et al., 1959; Few, 1955), facility location (Haimovich & Magnanti, 1988; Hochbaum, 1984; Papadimitriou, 1981), and any *subadditive Euclidean functional* such as a minimum spanning tree, Steiner tree, or matching (Redmond & Yukich, 1994; Steele, 1987, 1981). In our computational districting experiment, an approximation of this kind is used as the first level of an optimization problem in which we design service zones that are associated with different vehicles. for example. Our study of the VRP with time windows adopts similar assumptions to those of Daganzo (1987a,1987b), namely, that the service period is divided into a collection of pre-specified intervals of equal duration. One might contrast this model with other approaches like (Figliozzi, 2009), which assumes that time windows are independently drawn from an arbitrary probability measure. Our study of the capacitated VRP makes extensive use of upper and lower bounds derived in Bertsimas and van Ryzin (1993); Daganzo (1984); Haimovich and Rinnooy Kan (1985), as well as seminal results on the TSP that can be found in Beardwood et al. (1959); Redmond and Yukich (1994); Steele (1981).

A more recent focal point in research on VRP and its variants is the issue of *robustness* in which one seeks a policy that performs as well as possible against all possible realizations of demand that are compatible with some set of observations or initial conditions. Robust methodologies for the capacitated VRP were first introduced in the paper Sungur, Ordónez, and Dessouky (2008), which adapts the methodology of Ben-Tal and Nemirovski (1998) to solve problems in which customer demands and travel times are uncertain; the goal is to find vehicle routes that meet all feasibility requirements in the worst-case scenario, which occurs precisely when all customer demands and travel times attain their worst-case realizations simultaneously. In most models of the robust VRP, one has a pre-defined ambiguity region and seeks a set of routes that is as good as possible with respect to all of the outcomes; this ambiguity region is usually described as a finite collection of scenarios or a polyhedral set (Agra et al., 2013; Barkaoui & Gendreau, 2013; Gounaris, Wiesemann, & Floudas, 2013; Lee, Lee, & Park, 2011; Solyalı, Cordeau, & Laporte, 2012; Sungur et al., 2008), although the recent paper Allahviranloo, Chow, and Recker (2014) adopts a “robust mean-variance” approach that minimizes a weighted sum of the average cost and the variance of a route when sampled over many scenarios. In our problem, we are concerned with robustness in the *distributional* sense (Calafiore & El Ghaoui, 2006): we seek the distribution of demand for which the expected cost of a tour is as high as possible, while remaining consistent with some observed data samples or some parameters derived thereof. The most closely related result to our paper is (Carlsson & Delage, 2013), which determines the worst-case spatial demand distribution for the TSP when the first and second moments are fixed. Our paper can be seen as a generalization of these principles to the cases where vehicles have capacities and time window constraints.

Our present work uses the notion of robustness to study the negative consequences of fluctuation in demand for delivery services, in either a spatial or temporal sense. Demand fluctuation is of particular concern for emerging delivery services such as Good Eggs, DoorDash, BiteSquad, and Caviar (Caviar, 2014; Door Dash Food Delivery, 2014; Food Delivery & Restaurants Delivery - Order Food Online - BiteSquad.com, 2014; Good Eggs, 2014), which face extremely high volatility in demand due to seasonality and the time-sensitive nature of the requests they satisfy (GrubHub Inc., 2014; Leetaru, 2016):

> Our business is highly dependent on diner behavior patterns that we have observed over time. In our metropolitan markets, we generally experience a relative increase in diner activity from September to April and a relative decrease in diner activity from May to August. In addition, we benefit from increased order volume in our campus markets when school is in session and experience a decrease in order volume when school is not in session, during summer breaks and other vacation periods. Diner activity can also be impacted by colder or more inclement weather, which typically increases order volume, and warmer or sunny weather, which typically decreases order volume. Seasonality will likely cause fluctuations in our financial results on a quarterly basis. In addition, other seasonality trends may develop and the existing seasonality and diner behavior that we experience may change or become more extreme.

In total, this paper makes the following contributions: Section 3 analyzes the vehicle routing problem with time windows, characterizing the worst-case distributions that can arise when demand varies over a specified time horizon. Section 4 deals with the capacitated VRP, and Section 5 extends this analysis to more sophisticated models in which we have information about the mean or covariance of the demand distribution and describes some computational experiments.

2. Preliminaries

We make the following notational conventions in this paper: given a point set X, the *star network*of X is written $\mathsf{SN}(X)$ and consists of the network in which each point in X is connected to some central “depot” point (the location of this central point will be made clear from context). Vehicle capacities are either denoted by the letter c, indicating that a vehicle can visit c destinations before returning to its depot, or by the capacity coefficient t, which satisfies the relationship $c = t\sqrt{|X|}$; this is a standard and useful representation, as can be seen in Section 4.2 of Daganzo (2005) or the paper Daganzo (1984). A TSP tour of a set of points will be denoted by $\mathsf{TSP}(X)$. A capacitated VRP tour of a set of points is written $\mathsf{VRP}(X)$, where we suppress the capacity in the interest of notational brevity. Finally, we say that $f(x) \in o(g(x))$ if $\lim_{x\to\infty} f(x)/g(x) = 0$, we say that $f(x) \sim g(x)$ if $\lim_{x\to\infty} f(x)/g(x) = 1$, we say that $f(x) \in \mathcal{O}(g(x))$ if $f(x) \le \alpha g(x)$ for some $c > 0$ and all x exceeding some threshold x_0, and we say that $f(x) \in \Omega(g(x))$ if $f(x) \ge \alpha g(x)$ for all x exceeding some threshold x_0.

To approximate the length of a TSP tour of a collection of points, we will use the well-known *BHH Theorem* (Beardwood et al., 1959), which says that the length of an optimal TSP tour of a set of points follows a law of large numbers:

Theorem 1. *Suppose that* $X = \{X_1, X_2, \ldots\}$ *is a sequence of random points i.i.d. according to a probability density function* $f(\cdot)$ *defined on a compact planar region* $\mathcal{R}$. *Then with probability one, the length of* $\mathsf{TSP}(X)$ *satisfies*

$$\lim_{N\to\infty} \frac{\text{length}(\mathsf{TSP}(X))}{\sqrt{N}} = \beta \iint_{\mathcal{R}} \sqrt{\bar{f}(x)}\, dA$$

where β *is a constant and* $\bar{f}(\cdot)$ *represents the absolutely continuous part of* $f(\cdot)$.

It is additionally known that $0.6250 \le \beta \le 0.9204$ and estimated that $\beta \approx 0.7124$; see Applegate, Bixby, Chvatal, and Cook (2011), Beardwood et al. (1959). Theorem 1 can also be expressed deterministically, removing any assumptions about the distribution of the points X_i; see for example Goddyn (1990), Karloff (1989):

Theorem 2. *There exists a constant* α *satisfying the following: if* $X = \{X_1, X_2, \ldots\}$ *is any sequence of points contained in a compact planar region* $\mathcal{R}$ *with area* 1, *then*

$$\limsup_{N\to\infty} \frac{\mathsf{TSP}(X_1, \ldots, X_N)}{\sqrt{N}} \le \alpha.$$

Furthermore, it is also true that $(4/3)^{1/4} \le \alpha < 1.392$.