



Contents lists available at ScienceDirect

European Journal of Operational Research

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

Production, Manufacturing and Logistics

Faster rollout search for the vehicle routing problem with stochastic demands and restocking

Luca Bertazzi^{a,*}, Nicola Secomandi^b^a Department of Economics and Management, University of Brescia, Contrada Santa Chiara, 50, Brescia 25122, Italy^b Tepper School of Business, Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, PA 15213, USA

ARTICLE INFO

Article history:

Received 22 April 2017

Accepted 26 March 2018

Available online xxx

Keywords:

Routing

Rollout algorithms

Restocking

Stochastic vehicle routing problem

ABSTRACT

Rollout algorithms lead to effective heuristics for the single vehicle routing problem with stochastic demands (VRPSD), a prototypical model of logistics under uncertainty. However, they can be computationally intensive. To reduce their run time, we introduce a novel approach to approximate the expected cost of a route when executing any rollout algorithm for VRPSD with restocking. With a sufficiently large number of customers its theoretical speed-up factor is of big-o order $1/3$. On a set of instances from the literature, our proposed technique applied to a known rollout algorithm and three variants thereof achieves speed-up factors that range from 0.26 to 0.34 when there are more than fifty customers, degrading only marginally the quality of the resulting routes. Our method also applies to the a priori case, in which case it is exact.

© 2018 Elsevier B.V. All rights reserved.

1. Introduction

Given a set of geographically dispersed customers, a quantity to deliver to each customer, and a fleet of capacitated vehicles located at a depot, the vehicle routing problem consists of determining a set of minimal cost routes, each starting and ending at the depot, such that the demand of all the customers is satisfied without exceeding the capacity of the vehicles. Since its introduction by Dantzig and Ramser (1959), this problem and variants thereof have been well studied (see Fisher, 1995; Laporte, 1992; Toth & Vigo, 2014; and Laporte, 2009 for reviews).

In the vehicle routing problem with stochastic demands (VRPSD), given probability distributions describe the customer demands and the realization of the demand of a customer becomes known upon the first visit to this customer. If the realized demand of a customer exceeds the remaining capacity of a vehicle when this customer is visited then a route failure occurs and a recourse action must be taken. VRPSD is relevant in both strategic distribution planning, when only estimates of customer demands are typically available, and tactical and operational decision making, when there remains residual uncertainty about the demands of the customers.

The extant literature includes three VRPSD solution strategies: a priori, restocking, and reoptimization (see Bertsimas & Simchi-Levi, 1996; Dror, 2002; Dror, Laporte, & Trudeau, 1989; Gendreau, Laporte, & Séguin, 1996a; Stewart & Golden, 1983; and Gendreau, Jabali, & Rei, 2016 for reviews). Under the a priori strategy, vehicles follow a given set of routes and back-and-forth replenishment trips to the depot are performed when a failure occurs (Bertsimas, 1992; Bertsimas, Chervi, & Peterson, 1995; Bertsimas, Jaillet, & Odoni, 1990; Gendreau, Laporte, & Séguin, 1995; 1996b; Goodson, Ohlmann, & Thomas, 2012; Gupta, Viswanath, & Ravi, 2012; Hjorring & Holt, 1999; Jabali, Rei, Gendreau, & Laporte, 2012; Laporte, Louveaux, & Van Hamme, 2002; Rei, Gendreau, & Soriano, 2010; Secomandi, 2003). The restocking strategy modifies the a priori approach by allowing preventing replenishment trips to the depot to avoid potentially costly route failures (Yee & Golden, 1980; Bertsimas et al., 1995; Secomandi, 2003; Yang, Mathur, & Ballou, 2000). With reoptimization, the decisions of which customer to visit next or whether to replenish depend on the demand observed and served so far (Goodson, Ohlmann, & Thomas, 2013; Goodson, Thomas, & Ohlmann, 2016; Novoa & Storer, 2009; Secomandi, 2000; 2001; Secomandi & Margot, 2009). In other words, as discussed by Secomandi and Margot (2009), both routing and replenishment decisions are static in the a priori case, routing decisions are static and replenishment decisions are dynamic in the restocking case, and both types of decisions are dynamic in the reoptimization case. Although dominated by reoptimization, the a priori and restocking strategies are appealing in practice because static routing creates regular service that is appreciated by both

* Corresponding author.

E-mail addresses: luca.bertazzi@unibs.it (L. Bertazzi), ns7@andrew.cmu.edu (N. Secomandi).

customers and drivers (Bertsimas & Simchi-Levi, 1996). Further, the restocking approach outperforms the a priori approach (Yee and Golden 1980; Bertsimas et al., 1995; Secomandi, 2003) in terms of expected delivery cost. We thus focus our attention on the restocking strategy, but our methodological development is also relevant for the a priori case.

Heuristics are widespread in the VRPSD literature because exact methods can be used only on moderately-sized instances (Dror et al., 1989; Gendreau et al., 1995; Hjorring & Holt, 1999; Laporte et al., 2002; Bianchi et al., 2006; Jabali et al., 2012). Rollout search is an approximate dynamic programming approach that uses the expected cost of a known (base) heuristic to approximate the optimal value function when making a decision (Bertazzi, 2012; Bertsekas, 2005; Bertsekas & Castanon, 1999; Bertsekas & Tsitsiklis, 1996; Goodson, Thomas, & Ohlmann, 2017). Secomandi (2001, 2003), Novoa and Storer (2009) and Goodson et al. (2013, 2016) develop rollout methods for VRPSD.

The nested execution of a base heuristic, a defining element of rollout algorithms, can make rollout search computationally intensive, especially for instances that feature many customers. Thus, in this paper we develop a method to reduce the computationally requirement of this heuristic search approach applied to VRPSD under the restocking strategy when there is a single vehicle. In contrast to the standard backward dynamic programming evaluation of the expected cost of a restocking route (Yee and Golden 1980; Bertsimas et al., 1995; Secomandi, 2003; Yang et al., 2000), we first derive a novel forward dynamic programming approach to evaluate this cost. We then combine the forward and backward methods in a hybrid fashion when executing any rollout algorithm for the VRPSD version that we consider. This hybrid approach eliminates redundant computations at the expense of additional bookkeeping but approximates the expected cost of a restocking route. Even if it is outside the scope of our research, our proposed technique applies without approximation to the a priori strategy; that is, it yields an exact evaluation of the expected cost of an a priori route. The computational requirement of this method is theoretically appealing when applying VRPSD rollout algorithms to instances with a sufficiently large number of customers, in which case its big-o order speed-up factor equals $1/3$.

We assess the performance of our proposed approach by applying it on the instances of Secomandi and Margot (2009) using as benchmarks the basic rollout algorithm of Secomandi (2003) for VRPSD with a single vehicle and three variants thereof. One of these variants bears some similarities with the rollout algorithms developed by Novoa and Storer (2009) to obtain a rollout policy for the reoptimization version of VRPSD. Hence, our considered rollout algorithms loosely represent extant single vehicle VRPSD rollout algorithms. Consistent with our theoretical analysis, across all the examined rollout algorithms, the observed speed-up factors vary between 0.26 and 0.34 for instances with at least 50 customers. These computational savings are associated with only a marginal degradation of the quality of the resulting routes. Although these results are specific to the rollout algorithms that are the subject of our numerical study, our expected cost computation approach is relevant to other VRPSD rollout algorithms one may design or use.

Other researchers have investigated the possibility of speeding up rollout search. In the context of the traveling salesman problem, Guerriero, Mancini, and Musmanno (2002) propose the pruned and relaxed rollout algorithms, which selectively execute the base heuristic at each iteration. Ciavotta, Meloni, and Pranzo (2016) develop related techniques for parallel machine scheduling problems. Guerriero and Mancini (2005) use parallel computing to reduce the run time of the traditional and pruned rollout algorithms applied to the traveling salesman and sequential ordering problems. In contrast, our approach always executes the base heuristic at each iteration, relies on sequential computing, and deals with a differ-

ent application. The rollout policies of Novoa and Storer (2009) for VRPSD use a base heuristic of the a priori type (see also Birattari, Balaprakash, Stütze, & Dorigo, 2008 for the use of Monte Carlo simulation in local search for stochastic combinatorial optimization). These authors obtain improved computational efficiency by evaluating the expected cost of an a priori route using Monte Carlo simulation, without observing any deterioration in the obtained solutions even with small sample sizes. Monte Carlo evaluation cannot be applied to estimate the expected cost of a restocking route. However, in the a priori case one could combine Monte Carlo simulation and our proposed technique to obtain a possibly even faster method.

Bianchi et al. (2006) study the effectiveness of approximating the VRPSD objective function under the a priori strategy with the one of deterministic versions of this problem, which is easier to compute than the exact one, for various metaheuristics. Whereas their focus is on the solution quality of the resulting methods, our main attention is on the speed up of the approach that we put forth (although we observe that it has minimal impact on solution value). Moreover, our proposed method does not involve any approximation when applied to VRPSD under the a priori strategy.

We introduce VRPSD under the restocking strategy in Section 2. We present the backward and forward dynamic programming approaches to compute the expected cost of a VRPSD restocking route in Section 3. We discuss our proposed hybrid approach for the computation of the expected cost of such a route for any rollout algorithm in Section 4. We conduct our numerical study in Section 5. We conclude in Section 6. The online supplementary material includes additional information about our numerical results.

2. VRPSD with a single vehicle

We formulate VRPSD with a single vehicle as an optimization model. Let $G(V, E)$ be a given complete graph, where $V := \{0, 1, \dots, n\}$ is the set of $n + 1$ nodes and E is the corresponding set of edges. The depot is located at node 0 and the customers at nodes 1 through n . The cost to travel from node i to node j is denoted by $c(i, j)$. Travel costs are symmetric, $c(i, j) = c(j, i)$ for each i and $j \in V$, and satisfy the triangle inequality, $c(i, j) \leq c(i, k) + c(k, j)$ for each i, j , and $k \in V$. A single and fully loaded vehicle with capacity Q is initially located at the depot. The demand of each customer i is the integer-valued random variable D_i . These random variables are independent. Their probability distributions are known. We denote the probability that the demand of customer i is equal to ℓ_i as $p_i(\ell_i)$, that is, $p_i(\ell_i) := \Pr\{D_i = \ell_i\}$. We let $L_i \geq 1$, and $\bar{L}_i \leq Q$, respectively, be the minimal and maximal values of the support of the random variable D_i . The realization of the demand of a customer becomes known when, and only when, the vehicle arrives at the location of this customer.

A route is feasible if it starts and ends at the depot and visits all the customers exactly once. A route failure occurs when the vehicle does not have enough capacity to fully satisfy the demand of a customer. The VRPSD objective is to find a feasible route with minimal expected cost such that the demands of all the customers are satisfied. Formally, denote a feasible route (tour) by τ and the set of all feasible routes by \mathcal{T} . Let the expected cost of route τ under the restocking strategy be $\mathbb{E}[C_\tau]$. VRPSD is $\min_{\tau \in \mathcal{T}} \mathbb{E}[C_\tau]$.

3. Expected cost computation for a given route

In Section 3.1 we present a known backward recursion to compute the expected cost of a given route for the restocking case. In Section 3.2 we introduce a new forward recursion to compute this cost. In Section 3.3 we illustrate these recursions using a simple example. Given a feasible route τ , for expositional simplicity we

متن کامل مقاله

دریافت فوری ←

ISIArticles

مرجع مقالات تخصصی ایران

- ✓ امکان دانلود نسخه تمام متن مقالات انگلیسی
- ✓ امکان دانلود نسخه ترجمه شده مقالات
- ✓ پذیرش سفارش ترجمه تخصصی
- ✓ امکان جستجو در آرشیو جامعی از صدها موضوع و هزاران مقاله
- ✓ امکان دانلود رایگان ۲ صفحه اول هر مقاله
- ✓ امکان پرداخت اینترنتی با کلیه کارت های عضو شتاب
- ✓ دانلود فوری مقاله پس از پرداخت آنلاین
- ✓ پشتیبانی کامل خرید با بهره مندی از سیستم هوشمند رهگیری سفارشات