Robust Duration-Constrained Tours for Vehicle Routing Problems with Stochastic Demands

Juan C. Morales
The Logistics Institute
Georgia Institute of Technology

Alan L. Erera
The Logistics Institute
Georgia Institute of Technology

Martin W.P. Savelsbergh
The Logistics Institute
School of Industrial and System Engineering
Georgia Institute of Technology, Atlanta, GA 30306, U.S.A.
fax.: +1 – 404 – 894 – 2301, e.mail: mwps@isye.gatech.edu

1 Introduction

This research considers a single-period, single-depot vehicle routing problem in which a fleet of homogeneous capacitated vehicles is dispatched to serve the demands of customers. When customer load demands are not known with certainty when planning, the routing problem is referred to as the *Vehicle Routing Problem with Stochastic Demands* (VRPSD). Traditional solution methodologies for the VRPSD can be grouped roughly into two classes: chance-constrained approaches and two-stage expectation minimization approaches; see (3) for more discussion. Most research focuses on the latter approach, using either exact stochastic integer programming methods (*e.g.,* (4) and (6)) or heuristics (*e.g.,* (5) and (1)).

To date, however, the VRPSD research literature has not explicitly addressed time considerations, and thus its utility in practice is limited. In this work, we address this deficiency by focusing on vehicle duration considerations, including hard tour duration constraints. Duration constraints are usually critical in practice due to driver hours-of-service regulations and deadlines inherent in time-definite transportation services, and it is therefore important to explore solution approaches that treat such constraints explicitly. With this motivation, we propose and study a two-stage solution approach for VRPSD problems with duration considerations motivated by ideas drawn from the field of robust optimization (see, *e.g.*, (2)). The approach relies on efficient methods for determining the maximum duration of each *a priori* vehicle tour given a recourse policy for any joint realization of customer demands. With such methods, it is then possible to solve problems that include hard duration constraints or those with minimax objective functions.

2 Problem definition and notation

The VRPSD is defined on a directed graph $G = (V_0, A)$ with $V_0 = \{v_0\} \cup V$, where v_0 represents the depot and $V = \{v_1, \cdots, v_n\}$ represents the set of customers, and with $A = \{(v_i, v_j) \mid i \neq j,\, v_i, v_j \in V_0\}$. A travel time $t_{ij}(= t_{ji})$ is defined for each arc $(i, j) \in A$, and the vector $\{t_{ij}\}$ is assumed to satisfy the triangle inequality. Vehicle capacity is denoted by Q, and customer demands are integer-valued random variables with known distributions. Customer demands are denoted by vector $\widetilde{d} \in \mathbb{Z}_+^{|V|}$. It is assumed that $\underline{d} \leq \widetilde{d} \leq \overline{d}$

for known $\underline{d}, \overline{d} \in \mathbb{Z}_+^{|V|}$, and that element $\overline{d}(i) \leq Q$ for each $i \in V$. The actual customer demand becomes known only when the vehicle arrives at a customer's location. A tour specifies the *a priori* sequence in which a set of customers is visited by a single vehicle. Let $\mathcal{T}_k$ denote the tour of vehicle k. Assume that vehicles carry inventory to be delivered to customers.

A *recourse policy* $\mathcal{P}$ determines uniquely for any given demand instance a set of recourse actions to be undertaken by vehicles to ensure that all customers are fully served. We consider only recourse policies that do not jointly replan vehicle operations; *i.e.*, each vehicle k will serve all customers in $\mathcal{T}_k$ without assistance from other vehicles. Let $\mathrm{L}(\mathcal{T}_k)$ be the total travel time required by vehicle k to complete its *a priori* tour if no recourse were required, and let $\phi(\mathcal{T}_k, \mathcal{P}, d)$ be the total additional travel time required given policy $\mathcal{P}$ and demand realization d. Then, the maximum duration of tour k is:

$$\mathcal{L}(\mathcal{T}_k, \mathcal{P}) \;=\; \mathrm{L}(\mathcal{T}_k) + \max_{\underline{d} \leq d \leq \overline{d}} \phi(\mathcal{T}_k, \mathcal{P}, d) \quad . \tag{1}$$

Similarly, the expected duration of tour k is given by:

$$\mathcal{L}_E(\mathcal{T}_k, \mathcal{P}) \;=\; \mathrm{L}(\mathcal{T}_k) + \mathbf{E}_{\widetilde{d}}[\phi(\mathcal{T}_k, \mathcal{P}, d)] \quad . \tag{2}$$

The two problems considered are now defined. Given a fleet of m vehicles, the Robust Vehicle Routing Problem with Stochastic Demand (RVRPSD) develops a set of m vehicle tours with minimum maximum duration:

$$\textbf{RVRPSD} \qquad \min_{\mathcal{T}_k} \; \max_{k} \; \mathcal{L}(\mathcal{T}_k, \mathcal{P}) \tag{3}$$

Note that the RVRPSD does not require knowledge of the distribution of $\widetilde{d}$.

Alternatively, the Vehicle Routing Problem with Stochastic Demands and Duration Constraints (VRPSDDC) finds the set of tours with minimum total expected duration subject to a hard constraint on individual vehicle duration:

$$\begin{aligned} \textbf{VRPSDDC} \qquad \min_{\mathcal{T}_k} \qquad & \sum_k \mathcal{L}_E(\mathcal{T}_k, \mathcal{P}) \\ \text{s.t.} \qquad & \mathcal{L}(\mathcal{T}_k, \mathcal{P}) \;\leq\; \Delta \quad \forall\, k \end{aligned} \tag{4}$$

Here it is assumed that each vehicle faces the same maximum allowable duration Δ.

3 Evaluating $\mathcal{L}(\mathcal{T}, \mathcal{P})$

Given $\mathcal{P}$ and $\mathcal{T}$ where $n = |\mathcal{T}|$, the problem of determining $\mathcal{L}(\mathcal{T}, \mathcal{P})$ is central to the solution of any RVRPSD or VRPSDDC. The *adversarial problem* is to determine the demand realization d^* that maximizes $\phi(\mathcal{T}, \mathcal{P}, d)$, enabling computation of $\mathcal{L}(\mathcal{T}, \mathcal{P})$. Note that if ϕ were non-decreasing in d, it is easy to see that $d^* = \overline{d}$. For many recourse policies, ϕ does not have this property and alternative approaches are needed.

Consider policies in which vehicles serve customers in their *a priori* tours in order, and may make a return trip(s) to the depot (a *detour*) to reload when appropriate. The recourse actions during tour $\mathcal{T}$ can then be described by a triplet (i, r, I) defined for each customer $i \in \mathcal{T}$. For customer i, let $r = j$ if the j-th detour to the depot is initiated immediately prior to serving the customer, or $r = 0$ if no detour occurs. Let I represent the inventory carried by the vehicle after serving customer i. While I is clearly bounded above by Q, it is also easy to show that r is bounded above by a maximum value $R \leq n$ which results from the scenario where $d = \overline{d}$.

For most practical recourse policies $\mathcal{P}$, it is possible to define *recourse conditions* $C^1(i, I_i)$ that provide simple necessary and sufficient conditions for the existence of some demand realization in which the vehicle first detours to the depot before serving customer i and then departs with inventory I_i, and $C^{r,r+1}(i, I_i, k, I_k)$ that provide conditions in which the vehicle detours for the $r+1$-th time before customer k and departs with inventory I_k if the r-th detour occurred before i leaving inventory I_i. The problem of determining the demand vector d which maximizes $\mathcal{L}(\mathcal{T}, \mathcal{P})$ is then equivalent to solving a longest (s, t) path problem on an acyclic network $\mathcal{G}(\mathcal{P}, \mathcal{T}) = (\mathcal{N}, \mathcal{A})$. Node set $\mathcal{N}$ includes s and t, and a node for each (i, r, I) for each $i \in \mathcal{T}$, $r \in \{1, ..., R\}$, and $I \in \{0, 1, ..., Q\}$. Arc set $\mathcal{A}$ includes:

1. arcs connecting s to nodes $(i, 1, I_i)$ when warranted by conditions C^1, with costs equal to the net cost of the recourse action before serving node i;
2. arcs connecting (i, r, I_i) to $(k, r+1, I_k)$ when warranted by conditions $C^{r,r+1}$, with costs equal to the net cost of the recourse action before serving node k; and
3. arcs connecting all nodes $(i, 1, I_i)$ to t with no cost.

This shortest-path network leads directly to the following pseudopolynomial complexity result:

Lemma 3.1 *Given a tour $\mathcal{T}$ with n customers operating using recourse policy $\mathcal{P}$, $\mathcal{L}(\mathcal{T}, \mathcal{P})$ can be determined in $O(n^3Q^2)$ time.*

For many recourse policies, determination of $\mathcal{L}(\mathcal{T}, \mathcal{P})$ is possible in polynomial time using a simplified acyclic network. One class of such policies are called *history independent*:

Definition 3.1 (History-independent recourse policy) *A recourse policy $\mathcal{P}$ is history independent, if for every $i, k \in \mathcal{T}$ such that $i < k$, recourse conditions $C^{r,r+1}$ can be evaluated without considering $d(1), \cdots, d(i-1)$.*

Lemma 3.2 *Given a tour $\mathcal{T}$ with n customers operating using history-independent recourse policy $\mathcal{P}$, $\mathcal{L}(\mathcal{T}, \mathcal{P})$ can be determined in $O(n^3)$ time.*

4 Recourse Policies

We consider a set of recourse policies where customer demand is not splittable across vehicles, or across multiple return trips to the depot for a single vehicle:

Definition 4.1 (Myopic Policy $\mathcal{P}_0$) *A detour recourse occurs at customer i if and only if $d(i)$ is strictly greater than the on-board inventory of the arriving vehicle, I_{i-1}.*

Definition 4.2 (One Lookahead Policy $\mathcal{P}_1$) *A preemptive detour is taken between customer $i-1$ and i if and only if I_{i-1} is strictly less than $\overline{d}(i)$ after serving i.*

Definition 4.3 (All Lookahead Policy $\mathcal{P}_n$) *Prior to departing from customer $i-1$, determine whether to take a preemptive detour before proceeding to customer i by solving an optimal tour partitioning problem on the remaining customers assuming demand $d(j) = \overline{d}(j)$ and using initial vehicle inventory I_{i-1}.*

It turns out that $\mathcal{P}_1$ is history independent, but $\mathcal{P}_0$ is not. Determination of $\mathcal{L}(\mathcal{T}, \mathcal{P}_0)$, however, is possible in general in $O(n^4)$ using a specialized form of network $\mathcal{G}$. Furthermore, we show that determination of $\mathcal{L}(\mathcal{T}, \mathcal{P}_n)$ is simple since $\phi(\mathcal{T}, \mathcal{P}_n, d)$ is non-decreasing in d, and thus the worst-case demand instance is

$d = \overline{d}$. This leads to the following result. Let $L_{OP}(\mathcal{T}, d)$ be the duration of the set of vehicle routing tours formed by using the *optimal tour partitioning* algorithm to determine the best set of vehicle detours given *known* customer demands d.

Theorem 4.1 *For any fixed sequence* $\mathcal{T}$, $\mathcal{L}(\mathcal{T}, \mathcal{P}_n) = L_{OP}(\mathcal{T}, \overline{d}) = \max_d L_{OP}(\mathcal{T}, d)$

Thus, in the worst case, knowing all customer demands *a priori* is no better than using policy $\mathcal{P}_n$ if customers are still to be served in the same order as given by $\mathcal{T}$. Since no such policy can be better than $\mathcal{P}_n$, the following lemma holds.

Lemma 4.1 *For any fixed sequence* $\mathcal{T}$,

$$\begin{aligned} \mathcal{L}(\mathcal{T}, \mathcal{P}_n) &\leq \mathcal{L}(\mathcal{T}, \mathcal{P}_0) \\ \mathcal{L}(\mathcal{T}, \mathcal{P}_n) &\leq \mathcal{L}(\mathcal{T}, \mathcal{P}_1) \end{aligned}$$

Interestingly, however, the following result only holds for $R \leq 1$; simple counterexamples exist for $R = 2$.

Theorem 4.2 *For any fixed sequence* $\mathcal{T}$ *and* $R \leq 1$,

$$\mathcal{L}(\mathcal{T}, \mathcal{P}_n) \leq \mathcal{L}(\mathcal{T}, \mathcal{P}_1) \leq \mathcal{L}(\mathcal{T}, \mathcal{P}_0)$$

5 Computational Findings

We focus an initial computational study on recourse policies $\mathcal{P}_0$ and $\mathcal{P}_n$. The myopic policy is appealing because of its simplicity, and since it is frequently analyzed in literature on stochastic vehicle routing problems. Tabu search solution heuristics are developed and used for the different problems to follow using the standard $N(p, q, x)$ remove-and-reinsert neighborhood from (5).

5.1 Single-vehicle RVRPSD

We compare robust vehicle tours to those generated by expectation minimization only. Ten geographic instances were generated for each $n \in \{10, 20, 30\}$, with customers located according to a uniform distribution on $[0, 100]^2$ with depot at $(50, 50)$ and Euclidean travel times. Customers have homogeneous demands, with $\underline{d}(i) = 1$. Three different levels for $\overline{d}(i) \in \{5, 11, 17\}$ were used to understand the impact of increasing variance; to compute expectations, $\widetilde{d}(i)$ is assumed to be discrete uniform on $[\underline{d}(i), \overline{d}(i)]$. Finally, Q was varied using a parameter r that estimates the approximate number of required detours to serve a set of customers: $Q = \text{round}\left(\frac{n}{r} \frac{\underline{d}(i)+\overline{d}(i)}{2}\right)$ where $r \in \{1, 2, 3\}$.

For each instance, tabu search was used to determine both a best robust minimax tour $\mathcal{T}_R$ and a best minimum expected duration tour $\mathcal{T}_E$. Next, the expected duration of $\mathcal{T}_R$ and the maximum duration $\mathcal{T}_E$ were computed. Table 1 summarizes the results with using myopic policy $\mathcal{P}_0$, presenting average performance measures over all instances while holding one parameter fixed. Columns WCE give the average percentage difference between the expected and maximum durations of $\mathcal{T}_E$; note that these differences can be quite large. Columns POR give the *price of robustness*, measured as the average percentage difference between the expected duration of $\mathcal{T}_R$ and the expected duration of $\mathcal{T}_E$. For these problems, this price is about 2 to 4 percent. Finally, columns PON give the *price of non-robustness*, measured as the average percentage difference between the maximum duration of $\mathcal{T}_E$ and the maximum duration of $\mathcal{T}_R$. These differences average about 4 to 6 percent, and represent the tour duration penalty one may face in the worst case from not using the robust tour $\mathcal{T}_R$.

n	WCE	POR	PON	$\overline{d}(i)$	WCE	POR	PON	r	WCE	POR	PON
10	48.5	2.2	3.6	5	26.3	3.0	4.8	1	21.5	1.7	2.2
20	32.2	2.2	4.6	11	42.3	3.0	3.7	2	37.1	3.9	5.6
30	30.4	4.3	4.2	17	42.6	2.8	3.8	3	52.5	3.1	4.6

Table 1: Comparative results for single-vehicle RVRPSD

Best robust tours were also generated with lookahead policy $\mathcal{P}_n$. In this case, the minimax duration robust tour $\mathcal{T}_A$ can be determined by solving a deterministic VRP for the customers using demands $\overline{d}(i)$. The expected cost of each tour $\mathcal{T}_A$ is evaluated by simulation. For the test instances, the average penalty in maximum duration for using myopic policy $\mathcal{P}_0$ instead of $\mathcal{P}_n$ was about 14 percent, while the average penalty in expected duration was about 1.5 percent.

5.2 Multiple-vehicle VRPSDDC with $\mathcal{P}_0$

n	Vehicle capacity	$\overline{d}(i)$	$0.75\Delta_k$	$0.85\Delta_k$	$0.95\Delta_k$
20	$Q(r=5)$	5	4.89%	1.89%	1.25%
		11	3.50%	1.66%	1.09%
		17	2.74%	1.69%	1.12%
	$Q(r=4)$	5	5.05%	3.04%	1.05%
		11	1.47%	1.01%	0.81%
		17	1.57%	1.05%	0.83%
	$Q(r=3)$	5	6.63%	3.28%	1.63%
		11	5.85%	2.95%	0.92%
		17	5.54%	2.87%	0.89%
100	$Q(r=5)$	5	11.35%	7.74%	5.19%
		11	15.82%	10.12%	6.08%
		17	14.60%	9.60%	5.60%
	$Q(r=4)$	5	8.87%	6.93%	3.39%
		11	12.55%	8.38%	5.33%
		17	12.37%	9.17%	4.02%
	$Q(r=3)$	5	8.41%	4.62%	2.84%
		11	9.91%	8.18%	4.91%
		17	9.64%	7.87%	6.34%

Table 2: Average percentage increases in expected total travel time for VRPSDDC problems for varying levels of duration constraint tightness.

We study the impact of duration constraints on multiple vehicle routing problems with stochastic demands when using myopic recourse policy $\mathcal{P}_0$. Ten geographic instances were generated for each $n \in \{20, 40, 60, 80, 100\}$, and customer demands are modeled as in the previous section. Parameter Q is varied using the same expression, where $r \in \{3, 4, 5\}$; in this context, r can be interpreted as a lower bound on the required number of vehicles. To understand the potential impact of duration constraints, each instance k is solved first with no duration constraints to determine an unconstrained minimum. Next, the maximum tour duration Δ_k is computed. Then, the instance is resolved with three different duration constraint right-hand sides: $\{0.95\Delta_k, 0.85\Delta_k, 0.75\Delta_k, \}$.

Results in Table 2 summarize the average percentage increase in expected solution cost when enforcing duration constraints for two levels of n. Problems with more customers seem to face larger potential negative impacts from duration constraints, as well as problems with smaller vehicle capacity (larger r). In the largest problems, a 25 percent reduction in allowable maximum tour duration leads to increases in total expected tour cost of 10 to 15 percent.

6 Conclusions

This research has proposed and developed a methodology for incorporating tour duration considerations into vehicle routing problems with stochastic demands. Efficient evaluation methods are presented for typical recourse policies, and are embedded within traditional metaheuristic search approaches. Computational results on VRPSDDC instances demonstrate that tours and expected costs may vary significantly when problems face hard duration tour duration constraints.

References

[1] A. Ak and A.L. Erera. A paired-vehicle recourse strategy for the vehicle routing problem with stochastic demand. *Transportation Science*, submitted, 2006.

[2] D. Bertsimas and M. Sim. The price of robustness. *Operations Research*, 52:35–53, 2004.

[3] A.L. Erera and C.F. Daganzo. A dynamic scheme for stochastic vehicle routing. *Transportation Science*, submitted, 2003.

[4] M. Gendreau, G. Laporte, and R. Séguin. An exact algorithm for the vehicle routing problem with stochastic demands and customers. *Transportation Science*, 29:143–155, 1995.

[5] M. Gendreau, G. Laporte, and R. Séguin. A tabu search heuristic for the vehicle routing problem with stochastic demands and customers. *Operations Research*, 44:469–477, 1996.

[6] G. Laporte, F.V. Louveaux, and L. Van Hamme. An integer l-shaped algorithm for the capacitated vehicle routing problem with stochastic demands. *Operations Research*, 50:415–423, 2002.