A Pickup-and-Delivery Routing Problem with Stochastic Demands

F.V. Louveaux J.J. Salazar

October 30, 2006

1 Introduction

The *One-Commodity Pickup-and-Delivery Travelling Salesman Problem* (1-PDTSP) is defined on a directed graph as follows. A node of the graph represents a depot and the other nodes represent customers. Every customer has a positive or negative demand of a product. The product is the same for all customers. Customers with positive demand are called *pickups* and customers with negative demand are called *deliveries*. A *tour* is a route for a vehicle which departures from the depot and visits each customer exactly once before returning back to the depot. The vehicle has a given capacity and leaves the depot with some initial load. Given the vehicle capacity and the initial load, a tour is said to be *feasible* if it completely satisfies the demand of each customer without violating the capacity limits of the vehicle. Given a cost associated with each arc of the graph, an *optimal* solution is a feasible tour with minimum total cost. The 1-PDTSP is the combinatorial optimization problem of finding such an optimal solution. For low capacities, the problem may be infeasible. If the problem is feasible then the 1-PDTSP also includes finding the appropriated initial load of the vehicle that makes an optimal tour feasible. The initial load may be different from zero and the vehicle capacity. Although the capacity of the vehicle is typically an input of the 1-PDTSP, one could also address the different problem of finding the minimum capacity such that the 1-PDTSP is feasible. While the 1-PDTSP was first introduced in [1], the latter problem has not been approached in the literature (as far as we know). For the 1-PDTSP, instead, optimal and near-optimal algorithms have been proposed.

This paper considers a stochastic variant of the 1-PDTSP where we relax the hypothesis that the demand of each customer is deterministic and known in advance. Instead, we assume that the demand of each customer is a random variable with a discrete probability distribution. The particular situation where each random value is replaced by one of each realizations leads to what is called *scenario*. Therefore, unlike in the 1-PDTSP where only one scenario defining the demands of the customers is possible, in this paper we consider that many scenarios are possible, each one with a known probability.

With this stochastic input, the meaning of "feasible tour" and "optimal solution" has to be adapted. Indeed, a tour can turn out to be feasible or unfeasible depending on the scenario. A tour is *feasible for a given scenario and a given capacity* when it is possible to find an initial load such that all customer requests will be satisfied following this tour. This corresponds to the deterministic situation. Assume first that the capacity has to be decided before the scenario is given and the initial load can be decided when the scenario is given. A tour and a vehicle capacity are said to be *adaptable* when for each scenario one can find an initial load such that the tour is feasible with the given capacity. Assume now that the vehicle capacity and the initial load have to be decided before the scenario is given. A tour is said to be *survivable* if there exists a vehicle capacity and an initial load that make the tour feasible for all scenarios. Thus, the deterministic concept of feasibility is replaced by adaptability or survivability, depending on the moment the information on the scenario is made available. In both cases, vehicle capacity has to be decided in advance, which is nothing else than deciding which vehicle to purchase and is typically done under incomplete information. When the information on the scenario may be obtained in advance (e.g., on a daily basis) the concept of adaptability applies. If not, the survivable case is needed. Clearly adaptability is achieved with a smaller capacity than survivability. In this paper we study the question of how to find the smallest vehicle capacity needed to obtain adaptability or survivability. It turns out that the answer to these questions only depends on the two extreme scenarios, i.e. the one with smallest demands for every customer and the other with largest demand for every customer.

In most cases scenarios are not known in advance and capacity needed for survivability of a tour may be too large in practice. This raises questions related to a limited given vehicle capacity. One of these questions is probability of survivability of a given tour. Another question appears when we are also given penalties associated with not satisfying customer demands. This leads to the problem of finding an initial load which maximizes the probability of survivability or minimizes the expected penalty.

Depending on the capacity of the vehicle, a survivable tour may not exist. Still, in this situation it is of interest to find a (non-survivable) tour feasible for "most" of the scenarios. Because we know the probability distribution associated with each customer demand, instead of counting the number of scenarios it is more useful to measure the probability of a set of scenarios. Thus, given a vehicle capacity and a tour, another interesting question is to compute, for each possible initial load of the vehicle, the probability of the set of scenarios for which this tour starting with this initial load is feasible. This measure associated to each capacity, each tour and each initial load, will be shortly referred as *survivable probability*. Once this survivable probabilities are known, it is trivial to select the initial load that maximizes the survivability of a given tour when using a vehicle with a given capacity. These probabilities are also of interest to compute the minimum vehicle capacity such that a given tour keeps a survivable probability greater than a pre-specified threshold. Finally, given a vehicle capacity, finding a tour with minimum total costs or with a survivable probability greater than a threshold is also a challenging question.

This paper provides some answers to these questions.

2 Notation

The directed graph is represented by $G = (N, A)$, where $N = \{0, 1, \ldots, n\}$ is the node set and A is the arc set. Node 0 represents the depot, and each other node represents a customer. The random variable d_i defining the demand of customer i takes a value in the set $D_i = \{d_{i,1}, \ldots, d_{i,k_i}\}$. We assume $d_i = d_{i,k}$ with probability $p_{i,k}$, for each $k \in K_i = \{1, \ldots, k_i\}$. The smallest and largest values are respectively denoted by l_i and u_i, i.e. $l_i = \min\{d_i : d_i \in D_i\}$ and $u_i = \max\{d_i : d_i \in D_i\}$. Note that $l_i \leq u_i$, but one or both may be positive or negative. Let $a^+ = \max\{a, 0\}$ and $a^- = \min\{a, 0\}$.

A scenario can be represented by a vector d in $D = D_1 \times \ldots \times D_n$. The scenario associated with $l = (l_1, \ldots, l_n)$ is named *lower scenario*, and the scenario associated with $u = (u_1, \ldots, u_n)$ is named *upper scenario*. These two "extreme" scenarios play a fundamental rule when answering some of the questions presented in Section 1.

The capacity of the vehicle is denoted by Q. It may be an input for some questions and an output for other questions in this paper.

A tour is a permutation σ of the n customers, starting and ending at the depot. The customer in position i in the permutation is denote by $\sigma(i)$, $i = 1, \ldots, n$, with $\sigma(0) = \sigma(n+1) = 0$. The initial load of the vehicle when leaving the depot through a tour is represented by s. For a given scenario $d \in D$, the load of the vehicle leaving customer in position i in the tour is denoted by f_i. Clearly $f_0 = s$ and f_i for $i = 1, \ldots, n$ is a function of $d_1, \ldots, d_i$, s, Q and σ, i.e. $f_i(d, s, Q, \sigma)$. By abuse of notation, we will use shorter notation which only includes those parameters needed for understanding. As an example, we may be using the notation $f_i(d)$ when only the scenario d is of relevance, or $f_i(d, s)$ when the scenario d and the initial load s are of relevance, or even f_i if all parameters are clear from the context. The default rule used in this paper is

$$f_i = \max\{\min\{f_{i-1} + d_{\sigma(i)}, Q\}, 0\} \qquad \text{for all } i = 1, \ldots, n, \tag{1}$$

meaning that the vehicle should always serve as much as possible a customer request (split is allowed). The function f_i calculated on the extreme scenario l is denoted by g_i, and on the extreme scenario u is denoted by h_i.

In accordance with Section 1, given a capacity Q and a scenario d, a tour σ is *feasible* if there exists s such that if $0 \leq f_i \leq Q$ for all $i = 0, 1, \ldots, n$. Under feasibility, condition (1) simplifies into

$$f_i = f_{i-1} + d_i \qquad \text{for all } i = 1, \ldots, n. \tag{2}$$

Let us denote $\underline{f} = \min\{f_i : i = 0, 1, \ldots, n\}$ and $\overline{f} = \max\{f_i : i = 0, 1, \ldots, n\}$. Note that f_i depends on the value s, but $\overline{f} - \underline{f}$ is independent of s, if rule (2) is applied. For a given tour σ and a given scenario d, the minimal required capacity to make σ feasible is denoted by $Q(\sigma, d)$, i.e.

$$\begin{aligned} Q(\sigma, d) := \qquad & \text{minimize } Q \\ f_i - f_{i-1} = d_{\sigma(i)} \qquad & \text{for all } i = 1, \ldots, n \\ 0 \leq f_i \leq Q \qquad & \text{for all } i = 0, 1, \ldots, n. \end{aligned}$$

Clearly this optimization problem is solvable in linear time on n.

Given a capacity Q, a tour σ is *adaptable* if every scenario is associated with an initial load for which the tour is feasible. Given a capacity Q, a tour σ is *survivable* if there exists a load s such that σ with the initial load s is feasible for all scenarios. Note that the smallest capacity to make σ survivable may be strictly larger than the maximum $Q(\sigma, d)$ over all the scenarios d since the same initial load must ensure feasibility of the survivable tour in all the scenarios.

3 Results for a given route

In this section we assume that the tour (i.e., the customer sequence) is given. For simplicity in the exposition, we assume that $\sigma(i) = i$ for all $i = 1, \ldots, n$.

In the first result, we search for the smallest vehicle capacity that will make the tour adaptable. Thus, we assume that the capacity is decided under uncertainty while the initial load can be adapted to the scenario, i.e., the information on the scenario becomes known before the tour starts (e.g., on a daily base).

Proposition 3.1 *The smallest vehicle capacity Q' such that the given tour is adaptable can be found in linear time in n by computing the minimal capacity in the two extreme scenarios.*

To be complete we now provide the following lower bound of the smallest capacity valid for all tours.

Proposition 3.2

$$Q' \geq \max\left\{ \left|\sum_{i=1}^{n} l_i\right| , \left|\sum_{i=1}^{n} u_i\right| , \max_{i=1}^{n}\{|u_i|, |l_i|\} \right\}.$$

We now turn to the case when the information on the scenario is not available when starting the tour. It may be desirable to have an initial load and an adequate capacity to guarantee that the given tour is survivable.

Proposition 3.3 *The smallest vehicle capacity Q^* and an initial load s^* such that the given tour is survivable are $Q^* = \overline{f}(u, 0) - \underline{f}(l, 0)$ and $s^* = -\underline{f}(l, 0)$.*

To be complete we now provide the following lower bound of the smallest capacity Q^* valid for all tours:

Proposition 3.4

$$Q^* \geq \max\left\{ Q', \sum_{i=1}^{n}(u_i - l_i) + \min_{j}\{u_j^-, l_j^+\} \right\}.$$

The previous question helps to find a good capacity for the vehicle following a tour. However, in many situations the capacity is fixed and given, and smaller than Q^*. In that case, at some customer the vehicle may be unable to serve the customer demand. This situation is called a *failure*. As an example, if $f_{i-1} + d_i$ exceeds Q then the vehicle is unable to load customer i's demand. Then it will have to take some costly action (e.g., pay some penalty) and leave full (i.e., $f_i = Q$). The load of the vehicle when leaving customer i is still denoted by f_i, although it is no longer always $f_{i-1} + d_i$. The penalty associated with a failure is denoted by $r_i(f_{i-1}, d_i)$. As an example, if a penalty of ϵ is paid for each unloaded unit then $r_i(f_{i-1}, d_i) = \epsilon \cdot (f_{i-1} + d_i - Q)^+$.

The *expected penalty* or *expected recourse* for a given tour and a given initial load is

$$z = E_d\left\{ \sum_{i=1}^{n} r_i(f_{i-1}, d_i) \right\}$$

where the expectation is taken over all the scenarios d.

Problem 3.1 *Given a capacity $Q < Q^*$, find the initial load s which minimizes the expected penalty.*

To solve this problem we propose the following "backward" procedure:

Definition: Let $\pi(i, q)$ be the cumulative expected penalty of the partial tour from i to the depot if the vehicle leave i with load q.

Notation: Let $S_i(q) = \{k \in K_i : 0 \leq q + d_{i,k} \leq Q\}$.

Initialization: set $\pi(n, q) = 0$ for all $0 \leq q \leq Q$

Recursion: for all $0 \leq q \leq Q$ and all $i \in \{n, \ldots, 1\}$ do

$$\pi(i-1, q) \quad = \quad \sum_{k \in S_i(q)} p_{i,k} \cdot \pi(i, q + d_{i,k}) + \sum_{k \notin S_i(q)} p_{i,k} \cdot [\pi(i, f_i) + r_i(f_{i-1}, d_{i,k})]$$

Then, $\pi(0, q)$ is the expected penalty if the vehicle departures from the depot with load q, and $s = \arg\min\{\pi(0, q) : 0 \leq q \leq Q\}$ is the answer to the question.

Problem 3.2 *What is the smallest capacity for which there exists an initial load such that each customer is served with probability of at least $1 - \alpha$?*

This question is again difficult to solve in general. An upper bound on the vehicle capacity can be found as follows. If each customer is considered separately, serving a customer with probability $1 - \alpha$ is equivalent to request that one can serve a quantity included in a $(1 - \alpha)$ confidence interval of the demand. Let $[l_i^\alpha, u_i^\alpha]$ denote the $1 - \alpha$ confidence interval of customer i's demand.

Proposition 3.5 *Let Q^α be the vehicle capacity and s^α be the initial load obtained by using Proposition 3.3 after replacing each l_i by l_i^α and each u_i by u_i^α. Let $f_0 = s^\alpha$ and*

$$f_i = f_{i-1} + \max\{\min\{d_i, u_i^\alpha\}, l_i^\alpha\}. \tag{3}$$

Then Q^α is an upper bound on the solution of Problem 3.2.

4 Improving routes

In the previous section we addressed the question of finding the minimal vehicle capacity such that a given route is adaptable or survivable. We now address the same question when the tour is not fixed but can be optimized.

4.1 A node formulation

A first tentative is using a Mixed Integer Linear Programming (MILP) model. To this end, let z_{ij} be a 0-1 variable which assumes value 1 if customer i is in position j in an optimal tour, and 0 otherwise. Let g_j be a continuous variable representing the load of the vehicle when leaving node in position j ($j = 0, 1, \ldots, n$) in the lower scenario l, and h_j a continuous variable representing the load of the vehicle when leaving node in position j ($j = 0, 1, \ldots, n$) in the upper scenario u. Then a first MILP model is

$$\text{[Node-Formulation:]} \qquad \min Q$$

subject to

$$\begin{aligned}
g_j = g_{j-1} + \sum_{i=1}^{n} l_i z_{ij} & \qquad \text{for all } j = 1, \ldots, n \\
0 \le g_j \le Q & \qquad \text{for all } j = 0, 1, \ldots, n \\
h_j = h_{j-1} + \sum_{i=1}^{n} u_i z_{ij} & \qquad \text{for all } j = 1, \ldots, n \\
0 \le h_j \le Q & \qquad \text{for all } j = 0, 1, \ldots, n \\
\sum_{i=1}^{n} z_{ij} = 1 & \qquad \text{for all } j = 1, \ldots, n \\
\sum_{j=1}^{n} z_{ij} = 1 & \qquad \text{for all } i = 1, \ldots, n \\
z_{ij} \in \{0, 1\} & \qquad \text{for all } i, j = 1, \ldots, n.
\end{aligned}$$

Observe that this formulation finds the smallest vehicle capacity among all adaptable tours. If one is interested in finding the smallest vehicle capacity among all survivable tours then the extra constraint $g_0 = h_0$ must be added to the above mathematical model.

4.2 An arc formulation

Once the optimal or near-optimal solution of the above problem is available, one could then be interested in finding the minimum-cost tour, where the *cost* of a tour T can be defined as a function $c(T)$ in different ways. An example is $c(T) = \sum_{a \in T} c_a$ where c_a is the known length associated to arc a. Under this cost definition, the optimization problem turns to be a variant of the so-called *two-commodity Pickup-and-Delivery Travelling Salesman Problem* (2-PDTSP). It is a deterministic routing problem where the vehicle has to load two products, and each customer has a (positive or negative) demand associated to each product. In the 2-PDTSP the vehicle has a single container of capacity Q where the two products are transported. In our variant, instead, the vehicle has two containers of capacity Q, one for each product. Briefly speaking, in the 2-PDTSP one has $0 \le g_i + h_i \le Q$ ($i \in N$) while in our variant we have $0 \le g_i \le Q$ and $0 \le h_i \le Q$, which clearly shows that our optimization problem is a relaxed variant of the 2-PDTSP. Still, the two problems are not so different, and indeed known results for the 2-PDTSP can still be adapted to be used in our problem. The m-PDTSP has been considered in [2], and based on their ideas for $m = 2$, let us represent a tour with a 0-1 variable x_a for each $a \in A$ assuming value 1 if and only if $a \in T$. Let g_a and h_a the load of the vehicle when traversing arc a in the extreme scenario l and u, respectively. As standard notation, $\delta^+(i)$ represents the set of arcs leaving node i, $\delta^-(i)$ the set of arcs entering node i, and $\delta^+(S)$ the set of arcs leaving a node subset S. Then, a mathematical model is:

$$\text{[Arc-Fomulation:]} \qquad \min \sum_{a \in A} c_a x_a$$

subject to

$$\sum_{a \in \delta^+(i)} g_a - \sum_{a \in \delta^-(i)} g_a = l_i \qquad \text{for all } i = 1, \ldots, n$$

$$
\begin{aligned}
\sum_{a\in\delta^+(i)} h_a - \sum_{a\in\delta^-(i)} h_a = u_i & \qquad \text{for all } i = 1,\ldots,n \\
0 \le g_a \le Qx_a & \qquad \text{for all } a \in A \\
0 \le h_a \le Qx_a & \qquad \text{for all } a \in A \\
\sum_{a\in\delta^+(i)} x_a = 1 & \qquad \text{for all } i = 0,1,\ldots,n \\
\sum_{a\in\delta^-(i)} x_a = 1 & \qquad \text{for all } i = 0,1,\ldots,n \\
\sum_{a\in\delta^+(S)} x_a \ge 1 & \qquad \text{for all } S \subset \{1,\ldots,n\} \\
x_a\{0,1\} & \qquad \text{for all } a \in A.
\end{aligned}
$$

Note that the tour defined by the x variables of an optimal solution will be feasible for all scenarios, although each scenario may admit a different initial load. For example, the lower scenario l has the initial load $\sum_{a\in\delta^+(0)} g_a$, the upper scenario u has the initial load $\sum_{a\in\delta^+(0)} h_a$, and both values may be different.

4.3 Analysis based on expectations

In Section 3 we proposed a backward procedure to solve Problem 3.1 for a current given route. We now split the backward procedure in two components to obtain properties useful for improving routes.

Definition: Let $\pi_\epsilon(i,q)$ be the cumulative expected excess penalty of the partial tour from i to the depot if the vehicle leave i with load q.

Initialization: set $\pi_\epsilon(n,q) = 0$ for all $0 \le q \le Q$

Recursion: for all $0 \le q \le Q$ and all $i \in \{n,\ldots,1\}$ do

$$
\pi_\epsilon(i-1,q) \quad = \quad \sum_{k\in S_i(q)} p_{i,k}\cdot \pi_\epsilon(i, q+d_{i,k}) + \sum_{k\notin S_i(q)} p_{i,k}\cdot [\pi_\epsilon(i,f_i) + \epsilon(f_{i-1}+d_{i,k})]
$$

Analogously

Definition: Let $\pi_\eta(i,q)$ be the cumulative expected lack penalty of the partial tour from i to the depot if the vehicle leave i with load q.

Initialization: set $\pi_\eta(n,q) = 0$ for all $0 \le q \le Q$

Recursion: for all $0 \le q \le Q$ and all $i \in \{n,\ldots,1\}$ do

$$
\pi_\eta(i-1,q) \quad = \quad \sum_{k\in S_i(q)} p_{i,k}\cdot \pi_\eta(i, q+d_{i,k}) + \sum_{k\notin S_i(q)} p_{i,k}\cdot [\pi_\eta(i,f_i) + \eta(f_{i-1}+d_{i,k})]
$$

Then, $\pi(0,q) = \pi_\epsilon(0,q) + \pi_\eta(0,q)$ is the expected penalty if the vehicle departures from the depot with load q, and $s_0 = \arg\min\{\pi(0,q) : 0 \le q \le Q\}$ is the optimal load if the current route is followed.

References

[1] H. Hernández-Pérez, J.J. Salazar-González, “The one-commodity Pickup-and-Delivery Traveling Salesman Problem,” *Lecture Notes and Computation Science* 2570 (2003) 89–104.

[2] H. Hernández-Pérez, J.J. Salazar-González, “The multi-commodity Travelling Salesman Problem with Pickups and Deliveries”, Working paper, University of La Laguna, 2004.