A Heuristic Solution Method to a Stochastic Vehicle Routing Problem

Lars Magnus Hvattum* Arne Løkketangen* Gilbert Laporte†

*Molde University College
6411 Molde, Norway
{Lars.M.Hvattum,Arne.Lokketangen}@himolde.no

†Canada Research Chair in Distribution Management
HEC Montréal, Canada
gilbert@crt.umontreal.ca

1 Introduction

In order to solve real-world vehicle routing problems, the standard Vehicle Routing Problem (VRP) model usually needs to be extended. Here we consider the case when customers can call in orders during the daily operations, i.e., both customer locations and demand may be unknown in advance. Our heuristic approach attempts to minimize the expected number of vehicles and their travel distance in the final execution, under the condition that the unknown parameters have a known (approximated) distribution.

The problem discussed in this paper is based on a real-world case, where a freight company is serving pick-up customers, i.e., customers who need goods to be picked up at their location and brought back to a central depot. The goods may then be sent to another location from the depot. The company is the largest freight company in Norway, and has provided us with real-world data. Although the company also handles delivery customers, we focus only on the pick-up part in this work.

Unfortunately, only a portion of the pick-up customers is known at the beginning of the day. During operation additional customers can call in their orders, and the freight company will modify its routes in order to serve the new customers. In current practice this often implies sending out additional vehicles.

A customer has a location and a time-window where service is accepted, but the time-windows are rather wide, often spanning the entire working day. In addition, each order has a weight, and the vehicle that services the customer must have sufficient capacity for supporting the pick-up load. For the deterministic customers we assume that the time-window and the weight of the order are known, whereas the number of stochastic customers, their call-in time, as well as the location, time-window, and capacity-demand of each individual stochastic customer, are only known probabilistically. The location, time-window and capacity-demand will, however, be revealed at call-in time.

2 Problem description

In the Vehicle Routing Problem (VRP) a set of customers require some kind of service, which is offered by a fleet of vehicles. The goal is to find routes for the vehicles, each starting from a given depot to which they must return, such that every customer is visited exactly once. Usually there is also an objective that needs to be optimised, e.g. minimizing the travel cost or the number of vehicles needed.

The standard VRP is an NP-hard problem, and much work has been done in order to find efficient algorithms and heuristics to solve it. In real-world situations, however, the VRP usually fails to capture all essential aspects of the problem at hand. Hence, a large number of extensions have been studied, e.g. adding capacities, time windows, different kinds of customers, split deliveries, and more, which all try to make richer and more useful formulations.

Another way of extending the VRP is to focus on problems with uncertainty, where some of the parameters to the model are initially unknown or only known probabilistically. These approaches can roughly be divided in two classes: Dynamic Vehicle Routing Problems (DVRPs) and Stochastic Vehicle Routing Problems (SVRPs). Both classes have recently received an increasing amount of attention, and reviews can be found in Gendreau, Laporte and Séguin (1996a), Gendreau and Potvin (1998), and Ichoua et al. (2001).

The difference between static and dynamic VRPs is described in Psaraftis (1995), where the terms are defined as follows; a vehicle routing problem is *static* if the inputs to the problem do not change, neither during execution of the algorithm that solves it nor during the execution of the solution. On the other hand, a problem is considered *dynamic* when inputs to the problem are made known or updated to the decision maker concurrently with the determination of the solution (i.e., the set of routes). In dynamic solution approaches no plan is generated a priori, and new events are handled as they are revealed over time.

A Stochastic Vehicle Routing Problem (SVRP) arises when some of the elements of the problem are stochastic. This could be relevant for many of the components that may be included in a standard VRP, for instance travel times, demands, or the presence of each customer. SVRPs are often formulated as two-stage stochastic programming problems. Then probabilistic information is used to construct an a priori plan, and recourse actions are defined to handle the situations that occur when the random variables are realized.

We would like to treat our problem as a mix between a DVRP and a SVRP. In our multi-stage formulation, we divide the time horizon into intervals. The goal is then to make a plan for how to serve customers within the next time interval, using known, deterministic information, as well as knowledge about distributions for the stochastic elements. This plan should be feasible over the whole (remaining) time horizon. Assuming that the routes followed in subsequent time-intervals are optimal, the plan should minimize the expected cost of serving all customers, stochastic as well as deterministic, before the end of the day. The cost of a plan must be calculated after all stochastic variables are realized, and is dependent on the number of vehicles used as well as the total travel distance.

The deterministic, capacitated VRP with time-windows and a homogeneous fleet of vehicles can be formulated as a mixed integer program (MIP), see e.g. Cordeau et al. (2002). This deterministic formulation can then be extended to two- and multi-stage stochastic formulations. We direct the reader to Kall and Wallace (1994) for more on stochastic formulations.

3 The heuristic

Exact solution methods for stochastic VRPs currently fail to consistently solve problems with more than just a few customers (see Gendreau, Laporte and Séguin, 1995 and 1996b). In our case, even evaluation of the recourse cost function can be extremely difficult, depending on the distribution of the random variables. Hence, the need for practical heuristic solution methods is evident.

Our approach is based on solving sample-scenarios, where the idea is to use common features from sample-solutions to build a hopefully good plan. The term sample-scenario is here used to describe one of many potential future situations, i.e. actual realizations of customer call-ins. A sample-solution is then a plan that would be possible to execute given that the scenario would actually happen. The approach requires that we are able to generate possible future scenarios based on current deterministic information as well as approximations of the distributions of the stochastic variables involved. The idea of sampling in stochastic problems can be traced to Jagannathan (1985).

We call our method DSHH – the Dynamic Stochastic Hedging Heuristic, and it is in part inspired by progressive hedging as found in Haugen, Løkketangen and Woodruff (2000). It is outlined in Box 1. Assume we have divided the time horizon into n intervals, $I_1,...,I_n$. At the point in time corresponding to the beginning of interval I_u, a plan is constructed in the following manner:

Let C be the set of known customers at the start of I_u, and let S be a subset of C, initially empty.

1. Create p sample scenarios with the same known customers as in C, including stochastic customers drawn from the given distributions.
2. Repeat the following:
   2.1. Solve the sample scenarios, while forcing all customers in S to be served during I_u.
   2.2. Find the customer $c \in C - S$ that is most frequently visited in the time interval I_u in the p sample scenarios.
   2.3. If some stop criteria is met then go to 3, otherwise let $S = S \cup \{c\}$.
3. Repeat until *S* is empty:
   3.1. Solve the sample scenarios
   3.2. Count the number of times each customer $c \in S$ is served first by each vehicle in each scenario solution, disregarding customers that have already been placed.
   3.3. Choose the customer/vehicle-pair that has the highest count. Lock the customer to be served before all remaining customers by the given vehicle, and remove it from S.

Terminate with a plan for the given time interval, I_u.

Box 1. The DSHH heuristic

After the end of I_u has passed the procedure is repeated, possibly including new customers that have called in their orders. Customers served during the previous time intervals can now be ignored in the solution process, as long as the position and spare capacity of each vehicle is remembered.

In phase 1 we exploit knowledge about the probability distributions of the stochastic variables of the model, i.e. distributions of call-in times, demands, and geographical locations, to create possible future events. When including the known customers and their attributes, this constitutes our sample scenarios. They have the property that events that are likely to happen are present in a scenario with a high probability.

The second phase produces solutions of the sample scenarios, as if they were deterministic VRPs. Customers that are frequently served during I_u are then identified, and a decision is made to visit these customers during the next time interval in our final plan. Sample scenarios where these customers were not already served in I_u are then solved again, before the selection procedure continues. The selection is considered complete when there are no unselected customers that are frequently handled during I_u in solutions to the sample scenarios. We require that a customer must be handled early in at least half of the sample solutions in order to be selected, which in this case constitutes our stop criterion.

After phase two is completed we have found which customers to serve in the next time interval, but still have to decide in what sequence they should be visited. A tempting approach is to simply solve a deterministic VRP with the selected customers; however, we would rather exploit the information that is contained in the sample scenario solutions. Therefore we insert customers into the final plan depending on their positions in the sample plans; the customer most frequently visited first by some vehicle is locked to always be visited first by that vehicle. This corresponds to phase 3 in the procedure.

The heuristic has been implemented in C++, using a commercially available VRP-solver called SPIDER, produced by SINTEF, to generate heuristic solutions to the sample scenarios. In our heuristic we do not solve the sample scenarios very accurately, but use a simple insertion-neighbourhood to insert customers consecutively into the plan at the locally best location.

Note that this heuristic could be easily adapted to different types of stochastic VRPs, and that the underlying model could be extended in many directions. Extensive computational results will be reported, comparing our method with both the deterministic and dynamic approaches, showing the value of our method.

References

Cordeau, J.-F., G. Desaulniers, J. Desrosiers, M.M. Solomon and F. Soumis, “The VRP with Time Windows,” in *The Vehicle Routing Problem*, P. Toth and D. Vigo (eds), 157-193, SIAM, Philadelphia, 2002.

Gendreau, M., G. Laporte and R. Séguin, “An Exact Algorithm for the Vehicle Routing Problem with Stochastic Demands and Customers”, *Transportation Science* 29, 143-155 (1995).

Gendreau, M., G. Laporte and R. Séguin, “Stochastic Vehicle Routing”, *European Journal of Operational Research* 88, 3-12 (1996a).

Gendreau, M., G. Laporte and R. Séguin, “A Tabu Search Heuristic for the Vehicle Routing Problem with Stochastic Demands and Customers”, *Operations Research* 44, 469-477 (1996b).

Gendreau, M. and J.-Y. Potvin, “Dynamic Vehicle Routing and Dispatching,” in *Fleet Management and Logistics*, T.G. Craininc and G. Laporte (eds), 115-126, Kluwer Academic Publishers, 1998.

Haugen K.K., A. Løkketangen, and D.L. Woodruff, “Progressive Hedging as a Meta-Heuristic Applied to Stochastic Lot-Sizing”, *European Journal of Operational Research* 132, 103-109 (2000).

Ichoua, S., M. Gendreau and J.-Y. Potvin, “Exploiting Knowledge about Future Demands for Real-Time Vehicle Dispatching”, technical report CRT-2001-45, Centre for Research on Transportation, Montreal, Canada, 2001.

Jagannathan, R., “Use of Sample Information in Stochastic Recourse and Chance-Constrained Programming Models”, *Management Science* 31, 96-108 (1985).

Kall, P. and S. W. Wallace, *Stochastic Programming*, Wiley, New York (1994).

Psaraftis, H.N., ”Dynamic vehicle routing: Status and prospects”, *Annals of Operations Research* 61, 143-164 (1995).