Dynamic Vehicle Routing with Stochastic Requests

Russell Bent and Pascal Van Hentenryck
Brown University
{rbent,pvh}@cs.brown.edu

Abstract

This paper considers vehicle routing problems (VRP) where customer locations and service times are random variables that are realized dynamically during plan execution. It proposes a multiple scenario approach (MSA) that continuously generates plans consistent with past decisions and anticipating future requests. The approach, which combines AI and OR techniques in novel ways, is compared with the best available heuristics that model long-distance courier mail services [Larsen *et al.*, 2002]. Experimental results shows that MSA may significantly decrease travel times and is robust wrt reasonably noisy distributions.

1 Introduction

The vehicle routing problem (VRP) is a difficult combinatorial optimization problem with many important applications in distribution and transportation systems. It has received considerable attention from the operations research (OR) and AI communities for many years and sophisticated local search methods are quite effective at finding good quality solutions in reasonable amounts of time. In more recent years, technology has advanced so that it is now possible and practical to address dynamic and/or stochastic versions of the problem. These new versions are motivated by the inherent uncertainties that arise in many everyday VRPs and advances in onboard computers and communications systems that allow modification of routing plans even after vehicles are deployed.

Most of the existing work has focused solely on stochastic or dynamic versions of the problem exclusively. In stochastic optimization, the expected cost of a solution is optimized with respect to a recourse function which restores feasibility during plan execution. In dynamic optimization, various data items, such as customer requests, are unknown and are only revealed after some decisions are taken. A more detailed overview of the methods used in both approaches can be found in [Bent and Van Hentenryck, 2002]. How to combine the two approaches is a research topic that is often mentioned (e.g., [Gendreau *et al.*, 1999], among others).

This paper considers the multiple scenario approach (MSA) recently proposed in [Bent and Van Hentenryck, 2002]. MSA combines a variety of AI and OR techniques to address dynamic vehicle routing in the presence of stochastic information.

2 Problem Formulation

In the dynamic VRP, a number of customer requests are available initially, while others become available during the plan execution. In the applications considered in this paper, a request consists of the location of a customer and a service time, both of which are random variables. We assume that the distribution of the requests, or some approximation thereof, is available, which is typically the case in practical applications. For each incoming request, a dynamic algorithm must decide whether to accept or reject it. Once a request is accepted, it must be serviced. Problems are generally characterized by their degree of dynamism (DOD), i.e., the ratio of *unknown customers*/*total customers* [Larsen *et al.*, 2002]. A more in depth formalization of the VRP model can be found in [Bent and Van Hentenryck, 2002].

3 The Approaches

The Multiple Plan Approach The Multiple Plan Approach (MPA) is a fundamental generalization of many modern approaches. MPA generalizes the approach of [Gendreau *et al.*, 1999] by making it independent of the search procedure. In short, MPA continuously generates plans that are compatible with the current state of information and removes those that are not. In addition, since decisions must be made with respect to a specific plan to guarantee service, a *distinguished* plan is maintained via a ranking function.

More precisely, MPA handles four types of events (1) customer requests, (2) vehicle departures, (3) plan generations, (4) timeouts. A more complete description of the events and MPA can be found in [Bent and Van Hentenryck, 2002].

The Multiple Scenario Approach The Multiple Scenario Approach (MSA) generalizes MPA by considering both existing and potential future requests during plan generation. Future requests are obtained by sampling their distributions. Once a routing plan σ is discovered, MSA stores the routing plan σ^- obtained by removing future requests from σ. As a result, plan σ^- leaves room to accommodate future requests, should they actually materialize. This ability to anticipate the future is the strength of MSA.

Ranking Functions Both MPA and MSA are parametrized by a ranking function f_t, which selects the *distinguished* plan at each time t. We will evaluate two ranking functions for nominating σ^* in this paper. The obvious first choice for f_t would be to select the plan with the smallest travel cost (MPA^d). In [Bent and Van Hentenryck, 2002], it was shown that it is possible to do substantially better in practice on highly-constrained problems by using a consensus function which selects the plan in that most resembles all the stored plans.

Optimization Both MPA and MSA use large neighborhood search (LNS) for optimization, which has been shown to be very effective for vehicle routing [Shaw, 1998; Bent and Van Hentenryck, 2001]. LNS combines the advantages of branch and bound, constraint propagation, and local search.

4 Experimental Results

The Models The starting point of this research was the experimental model in [Larsen *et al.*, 2002], where customers are uniformly distributed in a 10km×10km region and must be served by a single vehicle with uniform speed of 40 km/h. Customers inter-arrival times are governed by an exponential distribution. Problem sets are derived with expected DOD of $\{0, 5, \ldots, 100\}$. More details about the model are found in [Larsen *et al.*, 2002]. We generated 15 instances for each DOD configuration, which gives 315 problems for each model described.

Several models are used to evaluate the various approaches. The first two models, M1 and M2, use a single vehicle. Model M1 is the basic model with 40 customers. Model M2 is similar to M1, except that the region is now 40km×40km. The objective function minimizes the travel distance. The next model is a multiple-vehicle model. Model M3 uses 4 vehicles, 160 customers, and a 20km×20km space. Each vehicle can serve at most 50 customers and the vehicle must return to the depot by the time horizon. For the multiple-vehicle model, the objective function consists in minimizing the number of missed customers and minimizing the travel distance. It is possible that some customers be left unserviced, since M3 has capacity constraints, as well as a hard deadline.

Single Vehicle Results Reference [Larsen *et al.*, 2002] tested various heuristics on Model M1. Their best heuristic is nearest neighbor (NN), where a pool of unserviced customers is maintained and the vehicle is sent to the nearest customer in the pool once it served its current request. Interestingly, in Model M1, the vehicle is able to service customers faster than they arrive. As a consequence, all "reasonable" heuristics converge towards a first come, first serve (FCFS) strategy as the DOD converges to 100%. This same behavior is also exhibited by our approaches. In Model M2, the vehicle is not able to service customers as quickly and the heuristics do not converge to FCFS. Figure 1 shows that MPA^d produces a 4.6% improvement on average and the benefits can be as high as 11.1%. Stochastic information (MSA^d) improves these results slightly. Moreover, MSA^d is never worse than NN. These results essentially indicate that the LNS-based MPA approach provides clear benefits over traditional heuristics. Stochastic information helps, but only marginally.

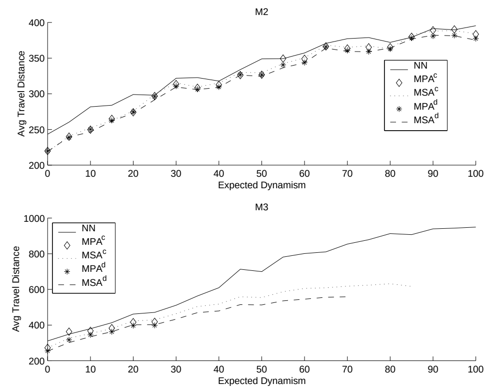

Figure 1: Travel Distance Results

Multiple Vehicle Results First, the NN heuristic was generalized to provide guarantees on servicing customers. Whenever a request arrives, the NN algorithm is simulated to determine if it can accommodate the new request. Surprisingly, the MPA approaches perform very poorly, especially for higher DODs. The "optimized" plans of MPA always use as few vehicles as possible since this generally improves travel distance and feasibility is not a issue until late in the search. This is in sharp contrast with problems with time-windows where feasibility is a major issue and MPA does not experience such problems [Bent and Van Hentenryck, 2002]. The MSA approaches perform roughly the same as NN as far as missed customers are concerned, but the true benefit of MSA is with respect to travel costs as shown in Figure 1. (The figure only shows the results for similar number of unserviced requests). Here, MSA decreases travel distances significantly, especially for high DODs. *These models clearly show the benefits of MSA and the value of stochastic information over traditional heuristics and MPA.* These results are even more dramatic when non-uniform distributions are used. Significant gains are observed for DODs as low as 25%. In addition to these results, the algorithm is robust under reasonable amounts of noise, for example when MSA is given incorrect customer arrival rates.

References

[Bent and Van Hentenryck, 2001] R. Bent and P. Van Hentenryck. A Two-Stage Hybrid Local Search for the Vehicle Routing Problem with Time Windows. To appear in Transportation Science

[Bent and Van Hentenryck, 2002] R. Bent and P. Van Hentenryck. Scenario-Based Planning for Partially Dynamic Vehicle Routing with Stochastic Customers. To appear in *Operations Research* (pending minor revisions).

[Gendreau *et al.*, 1999] M. Gendreau, F. Guertin, J. Y. Potvin, and E. Taillard. Parallel Tabu Search for Real-Time Vehicle Routing and Dispatching. *Transportation Science*, 33 (4):381–390, 1999.

[Larsen *et al.*, 2002] A. Larsen, O. Madsen, and M. Solomon. Partially Dynamic Vehicle Routing-Models and Algorithms. *Journal of Operational Research Society*, 53:637–646, 2002.

[Shaw, 1998] P. Shaw. Using Constraint Programming and Local Search Methods to Solve Vehicle Routing Problems. In *Principles and Practice of Constraint Programming*, 417–431, 1998.