Short term strategies for a dynamic multi-period routing problem

E. Angelelli[a], *N. Bianchessi*[a], *R. Mansini*[b], *M.G. Speranza*[a]
[a]*Dipartimento Metodi Quantitativi, University of Brescia, Italy*
[b]*Dipartimento di Elettronica per l'Automazione, University of Brescia, Italy*

1 Introduction

In this paper we consider the case of a courier company facing orders which arrive continuously over time. The courier service is organized in geographical areas. Each area has a central depot which provides the vehicles for serving the customers within the same area. Each request issued by a customer is made of two distinct and consecutive services, a *pick-up* to be accomplished at a specified location and a subsequent *delivery* at a destination, which is assumed to be in a different area. As a consequence each request is split in two (almost) independent requests. Pick-up requests imply the carriage of parcels to a central depot, while delivery requests imply the carriage of parcels from a central depot to different destinations. As the pick-up and delivery services are typically run at day-time, collected parcels are transported to the central depot of the destination area during the night. In this study we are interested in the problem faced by a single depot that has to face independent requests of pick-up and delivery by means of a fixed fleet of vehicles. When a request is issued, a deadline of k days is fixed. This means that if $k = 1$ the service has to be performed on the same day, if $k = 2$ the service has to be performed either today or tomorrow, if $k > 2$ the service has to be performed on any of the next k days. The time horizon of the company is potentially unlimited in the number of days, but decisions have to be taken day by day as new requests arrive. With respect to a fixed day in the time horizon we call *off-line* all the requests that are known in the morning before vehicles leave the depot, while we call *on-line* all the requests that arrive at the call center of the company during the day when the vehicles are already moving in the area. An on-line request which arrives today with deadline $k = 3$, if it is not served today, tomorrow it will be considered as off-line and *postponable*, while on the day after tomorrow it will be considered as off-line and *unpostponable*. There are two fundamental differences between pick-up and delivery requests. None of the delivery requests is directly issued by a customer, but it comes from some other depot during the night; in this sense all delivery requests are off-line. On the other hand a pick-up request is directly issued by a customer by a phone call during day-time. Thus, a pick-up can be either on-line or off-line (e.g. a postponable request which was not served yesterday). Another difference concerns the opportunity to exchange tasks between vehicles. Typically, after a parcel has been loaded on a vehicle, it must be delivered to the destination by the same vehicle. Thus, while delivery tasks are assigned to the vehicles in the morning and no change is permitted during the day, the assignement of a pick-up service to a vehicle may be changed until the very last moment before the service itself. According to the urgency of the request and the service tariffs the customer chooses the deadline for the delivery and implicitly defines the deadline for the pick-up. It is possible that a very urgent pick-up request is issued during the day with deadline $k = 1$. However, it is common practice that after a fixed time L, say noon or 1:00PM, no unpostponable pick-up requests are accepted. This rule has the consequence that at a convenient time in the day the company knows the set of unpostponable requests. In the case the company is not able to guarantee the service to all requests, it still has the opportunity to forward, at a higher cost, some tasks to a backup service. According to the operational context many different constraints may apply. The most common ones include constraints on traveling-time, vehicle capacity and time windows. The objective of the company is to guarantee the accomplishment of all the requests received by the depot while minimizing the average cost per day.

2 Literature

Dynamic routing problems have been attracting the interest of many researchers in the last years. Some of the applications cover dynamic fleet management ([10] and [11]), couriers ([5], [3] [6] and [7]) and dial-a-ride problems [4]. See [8], [9] and [2] for surveys.

The issues typically addressed in the literature concern problems with one day horizon. New requests are accepted only if they can be feasibly inserted in the service plan, otherwise they are rejected. Heuristic strategies are thus developed in order to optimize one or more of the following objectives: the number of customers served, the average waiting time and the traveled distance. A common approach to the problem is to solve an off-line optimization problem on the basis of the known information when a new request arrives.

This is especially important when the customer must know within a short time whether his request is accepted or rejected. The dynamic multi-period routing problem (DMPRP) was introduced in [1] where every day a set of customers to be served today and a set of customers that can be served either today or tomorrow are available. The company has to decide the set of customers to serve today and the set of customers to serve tomorrow. In that paper the competitive analysis of algorithms in a simple case is presented.

3 The dynamic multi-period routing problem

In the context described above we observe two factors of dynamism. Each day the company has to decide whether to serve postponable requests or not. The decision has to be taken without knowing what the set of new requests will be tomorrow. It is also reasonable to assume that, thanks to modern communication technology, the company can react to on-line requests and possibly make new plans for the service on the day.

The dynamic nature of the problem creates new modeling challenges and the need of new solution approaches. We call this class of problems Dynamic Multi-Period Routing Problems (DMPRP) which can be synthetically defined as follows.

A set of requests need to be served by a given fleet of un-capacitated vehicles over a time horizon of T days. Each request has a deadline $k \leq 2$. Every day the vehicles leave the depot in the morning and have to return to the depot at the end of the day. At the beginning of each day a set of requests is already known, while other requests may arrive over time. Each request has a deadline which makes it either unpostponable or postponable. Unpostponable requests can arrive only before a fixed time limit L for each day. At any time during the day the company knows the exact position of the fleet at the current time and is able to forecast the position of the fleet in the near future. The company is also able to route and re-route the vehicles at any point in time. At time L no more unpostponable requests are accepted and the company is able to decide the set of unpostponable requests which it is able to service with its fleet and the set of requests that will be forwarded to the back-up service. The objective is to minimize the average operational costs per day. These operational costs includes a very high cost paid for each request forwarded to the back-up service.

In such a dynamic context a number of questions naturally arises about the service of each request:

– Shall the company service the request?

– Shall the requests be serviced today or be postponed?

– At what time the request will be serviced?

Given the dynamicity of the context, at least in principle, the third question may be answered only when the vehicle actually stops at the pick-up or delivery location. The second question is obviously answered when the third is answered, but the decision to serve or to postpone could be taken earlier; for example, at the time the requests are issued the company decides on which day to perform the service and inform the customer. Similarly, the first question could be answered at the time the requests are issued either accepting or rejecting service.

In the current literature on dynamic routing problems the second question is not addressed as it is assumed that the customers require service for the current day. The first question is addressed as soon as the requests arrive, if the company is able to guarantee a profitable and feasible insertion of service in the currently planned routes the request is accepted, otherwise the request is rejected. In our context the company never rejects any requests and postpones as much as possible the decision about the service. That is, the decision to serve or do not the postponable requests can be taken late in the evening, while the decision to serve or do not the unpostponable requests can be taken up to the time limit L after which no more unpostponable requests are taken into account for that day. Unpostponable requests that cannot be served directly by the company can be forwarded to an expensive back-up service.

4 The approach to the problem

The objective of the DMPRP is the minimization of the operational costs and we formalize this with two hierarchical objectives. The first one is the maximization of the requests directly served by the company and the second one is the minimization of the length of the routes traveled. It is well known that each decision taken in absence of information on the future can be regretted when new requests arrive. The common approach to apply short term strategies by solving an off-line version of the problem with known data and follow the resulting solution as far as possible (at least until new requests arrive) is not guaranteed to lead to good solutions in the long term. Thus, the first question which need to be addressed is what problem should be solved, if any, in order to make plans for the future. Should we optimize for the long term objective? If not, what is the right objective to optimize? Should we forget about optimization and seek for some other goals like feasibility?

Since in our case no request will be rejected, there is another question to address. How often it is beneficial to re-optimize? Indeed while on one hand it is clear that updating the routes on the basis of the new information may be beneficial, on the other hand, since we allow diversion of the vehicles, a higher number of re-optimizations might imply a higher number of diversions and thus higher cost. In this paper we propose a heuristic framework where each day an off-line optimization problem P is solved first at the beginning of the day and then every Δt time till the end of the day. At each optimization all the known requests are taken into account. The definition of the re-optimization problem is crucial to the success of the approach. The first characteristic we consider is how long we look ahead in making a service plan. We may consider the remaining part of the current day only or we may also consider the day after. We cannot go too far because no information is available on the future. On the other hand, it seems that all the remaining part of the current day should be considered in order to fully exploit the availability of the vehicles. The second characteristic we consider is the criterion to measure the quality of a solution, that is the objective function. On one hand the number of served requests is important, on the other hand the length of the routes cannot be forgotten. We call *1-day look-ahead(f)* problem and *2-day look-ahead(f)* problem the optimization problem that 'works' on the routes until the end of the current day or until the end of the day after, respectively, by using f as objective function to evaluate solutions. The objective functions f considered in the *1-day look-ahead(f)* problem try first to maximize the number of un-postponable requests served, then different objectives are considered: (1) the length of the routes traveled, (2) the number of postponable requests served and the length of the routes, (3) the average distance traveled by the vehicles to serve a request. For the *2-day look-ahead(f)* problem we consider objective functions that share the common objective to maximize first the number of un-postponable requests served and secondly to maximize the total number of postponable requests to be served before the end of the day after. These functions differ in the minimization of a weighted sum of either the traveled distance or the average distance traveled per request.

Indipendently of the look-ahead period considered, the re-optimization problems have been addressed by means of a same variable neighborhood search heuristic.

Computational analysis has been carried out on two classes of scenarios that differ for the geographical dispersion of the requests in a service area of $100 \times 100\ km^2$: random and clustered scenarios. The location of the requests are taken from random and clustered Solomon's instances for the VRP. Each scenario considers a planning horizon of 10 days and a daily service period of 10 hours (from 8:00 AM to 6:00 PM) during which requests arrive dynamically according to a Poisson distribution with parameter $\lambda = 100, 200, 300, 400, 500$ per day. In particular, with a probability equals to $\frac{1}{3}$, requests arriving before 1:00 PM are un-postponable. In each random (clustered) scenario, the coordinates of an arriving request are randomly selected among those characterizing customers in c1 and c2 (r1 and r2) Solomon's instances. The depot coordinates are set as in Solomon's instances too. The service is provided by means of a fleet of 3 vehicles traveling at a costant speed of 40 km/h.

Three different values for Δt have been tested: 5h, 2.5h and 1h. Independently of the considered setting, increasing the frequency of the reoptimization allows in general to find better results. When solving a *1-day look-ahead(f)* problem, it is better to consider the objective (3) in case of random requests; while in the case that requests are well structured in clusters, the objective function (1) gives better results. On the other hand, with a 2-day look-ahead strategy, in both of the scenarios' classes, the best results are obtained minimizing the weighted sum of traveled distances, preferring a decrease in the distance traveled in the current day to any decrease in the distance traveled the day after, indipendently of the values of the decreases. The results obtained with a 2-day look-ahead strategy are definitely better than those obtained looking only at the current day. In particular, in the random scenarios, the most difficult ones, the number of not served requests is halved and the length of routes is improved on average by about 2.3% as well.

Future developments of this analysis may involve values of k larger than 2, capacitated vehicles and management of delivery requests.

References

[1] E. Angelelli, M.W.P. Savelsbergh, and M.G. Speranza. Competitive analysis for dynamic multi-period uncapacitated routing problems. *Networks*, to appear, 2006.

[2] G. Ghiani, F. Guerriero, G. Laporte, and R. Musmanno. Real-time vehicle routing: Solution concepts, algorithms and parallel computing strategies. *European Journal of Operational Research*, 151:1–11, 2003.

[3] S. Ichoua, M.Gendreau, and J-Y Potvin. Diversion issues in real-time vehicle dispatching. *Transportation Science*, 34:426–438, 2000.

[4] O.B.G. Madsen, H.F. Ravn, and J.M. Rygaard. A heuristic algorithm for a dial-ride problem with time windows, multiple capacities and multiple objectives. *Annals of Operations Research*, 60:193–208, 1995.

[5] M.Gendreau, F. Guertin, J-Y. Potvin, and E. Taillard. Parallel tabu search for real-time vehicle routing and dispatching. *Transportation Science*, 33:381–390, 1999.

[6] S. Mitrović-Minić, R. Krishnamurti, and G. Laporte. The double-horizon heuristic for the dynamic pickup and delivery problem with time windows. *Transportation Research Part B*, 38:669–685, 2004.

[7] S. Mitrović-Minić and G. Laporte. Waiting strategies for the dynamic pickup and delivery problem with time windows. *Transportation Research Part B*, 38:635–655, 2004.

[8] H.N. Psaraftis. Dynamic vehicle routing: Status and prospects. *Annals of Operations Research*, 61:143–164, 1995.

[9] H.N. Psaraftis. Dynamic vehicle routing problems. In B. L. Golden and A. A. Assad, editors, *Vehicle Routing: Methods and Studies*, pages 223–248. Elsevier Science, Amsterdam, 1998.

[10] M.W.P. Savelsbergh and M. Sol. Drive: Dynamic routing of independent vehicles. *Operations Research*, 46:474–490, 1998.

[11] J. Yang, P. Jaillet, and H. Mahmassani. Real-time multivehicle truckload pickup and delivery problems. *Transportation Science*, 38:135–148, 2004.