Notes on Dynamic Vehicle Routing
- The State of the Art -
Technical Report IDSIA-05-01

Leonora Bianchi -IDSIA -

December 20, 2000

Contents

1 Introduction

The aim of this document is (i) to introduce dynamic optimization problems and in particular the class of Dynamic Vehicle Routing Problems (DVRP), (ii) to review the main methodology advances and (iii) to point out the main properties of these problems.

Psaraftis' review article on DVRP [Psa95] focusses on three aspects: a) real-world examples of DVRP, b) technological advances which enhance the interest on DVRP c) methodological advances in the solution techniques developed for DVRP.

Contents of point a) and b) are fully valid today, therefore we will be quite schematic in order to tell the main concepts, without being redundant. Point c) is integrated here with the main contributions of the last years.

The attribute 'dynamic' in DVRP refers to distribution problems such that the information that is needed to come up with a set of good vehicle routes and schedules is dynamically revealed to the decision maker. The first question we want to answer about this class of problems is:

Why are Dynamic Vehicle Routing Problems interesting?

Real-world experience suggests an answer to this question:

- efficient vehicle routing becomes more important as markets tend to become increasingly open;
- the economic benefits that will be realized if the efficiency of these logistical systems increases are very significant;
- real-time distribution scenarios, in which the problem information is dynamically revealed to the decision maker, are becoming more and more common;
- processing of real-time data is more and more feasible and affordable, thanks to revolutionary advances in information and communication technologies.

The first two points of this list address to Vehicle Routing in general, while the second two justify why the attribute of dynamism is of interest to us.

From the point of view of Operation Research methodology, dynamic problems represent an interesting issue, since they are different in many fundamental aspects from static ones. Good methodology requires therefore non-trivial adaptation of static solution strategies, and new ideas.

Trying to classify the several types of DVRP can be useful, since no general framework exists. Such an effort could also be useful when comparing different optimization strategies and algorithms.

Besides the search of good solution techniques, it can be very useful to understand the nature an properties of Dynamic Routing Problems. This is often done for instance by borrowing techniques and results from queuing theory.

Let us now take a look to the

Context of Dynamic Vehicle Routing:

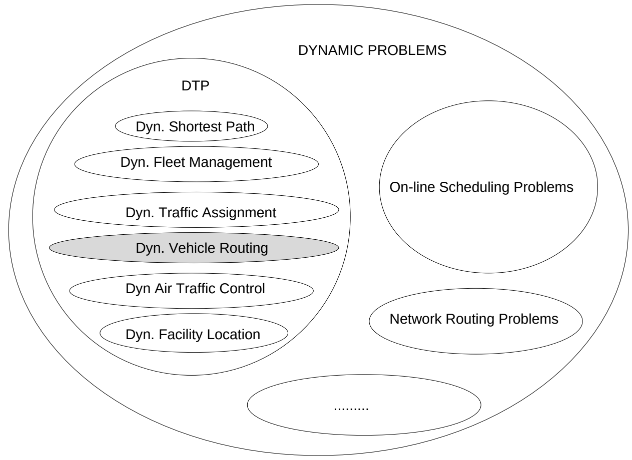

Figure 1: The context of Dynamic Vehicle Routing Problem

DVRP 'from inside': Real-world distribution problems that generally involve DVRP are well described in [Psa95] with representative references, therefore we skip this list.

DVRP 'neighbors': DVRP are a subset of the broader family of "Dynamic Transportation Problems" (DTP). The label "stochastic" is often tagged to this class of models. These include Dynamic Transportation, Dynamic Shortest Paths, Dynamic Traffic Assignment, Dynamic Fleet Management, Dynamic Air Traffic Control, and Dynamic Facility Location. Bibliographic references about these several DTP can be found in [Psa95].

DVRP 'surroundings': in domains different from DVRP, dynamic problems are also called 'on-line' or 'real-time', depending also on the methodology adopted. Some examples of these domains are On-line Scheduling Problems and Network Routing Problems. For recent review articles about these classes of on-line problems see [AL99] and [Sga98].

Figure 1 represents the whole context behind DVRP.

2 Properties of dynamic optimization problems

2.1 What do we actually mean with 'dynamic'

There are two conditions which make an optimization problem *dynamic*:

1. Information on the problem is time-dependent. The input instance is made known or updated as time goes on.

2. Solutions must be found while time proceeds, concurrently with incoming informations. This means that no a-priori solutions can be found. The only thing one can do a-priori is determining a 'strategy', or 'policy', which specifies what actions should be taken as a function of the state of the system.

What a dynamic problem is *not*:

- A problem which simply has time varying information, where the way the information changes is completely known in advance. Then in principle one can solve such a problem once and for all ahead of time, with no need to re-optimize. See for example the Time Dependent TSP described in [Psa95] to illustrate this point.

- A problem whose information is simply non-deterministic, which can be solved in principle a-priori with probabilistic assumptions. See for instance the probabilistic TSP, where demands at each node occurs with a known probability p ([RG87, Psa95]).

2.2 Dynamic versus static optimization problems

Dynamic problems are typically derived from static ones, by revealing on-line one or more parameters of those which define the static instance. In the following we point out some important differences between static and dynamic problems.

Time As it should already be clear, the main feature of dynamic problems which is not present in static ones is the dimension of time as integral part of the instance description.

Some problems (problem formulations) are open-ended, and do not have a time limit. These problems are usually solved with the hypothesis that equilibrium is eventually reached by the system, and all the quantities related to the objective function fluctuate around fixed values. In this case queuing considerations may become important. This formulation is used for instance in the Dynamic Traveling Repairman Problem (DTRP) described in the remainder, see also table 1.

Many problem formulations have a finite time horizon, such as a day or a week, and they do not need hypotheses about reaching equilibrium (DVRPTW, DVRPTW and DRIVE in table 1).

Future information In a dynamic problem information about the input instance may be in part given a priori, and in part dynamically revealed or updated. At a given time, there is no uncertainty with regard to the information received up to that instant. On the contrary the future information may be either completely unknown or partially known by means of some probabilistic assumptions.

Objective function Definition of an objective function is a non-trivial aspect of dynamic problems. One may consider the dynamic problem as a series of static problems (*subproblems*), with the goal of tracking the time-progression of the objective function extremes (maximum or minimum) as close as possible.

Sometimes optimizing only over known past input may be myopic if some information about the future is known (for instance probabilistic informations). Such informations should be explicitly considered by the objective function.

Strategy Previously we told that a strategy specifies what actions should be taken as a function of the state of the system. Less in general, a strategy is also the way a dynamic problem is split into several static sub-problems. One can for instance consider a new sub-problem each time that a new customer appears, or waiting a certain time period for a whole set of new customers, and also considering several spatial grouping of customers and so on. Strategies for some significant DVRP will be discussed in section 3.

The way one solves each sub-problem is a separate issue. For the same strategy, several optimization algorithms can be applied to sub-problems.

In the literature there are two types of works: those whose aim is more the study of strategies ('strategy oriented papers' as in section 3.1), and those who fix a strategy but investigate the performance of some optimization algorithm to solve the single sub-problems ('heuristic oriented papers' as in section 3.2).

2.3 Main DVRP in the literature

For every static VRP it is possible to imagine *many* different DVRP. Consider each parameter defining the instance of the static problem, such as for instance customer demand, customer position, time windows, number of vehicles, vehicle capacity and so on. Now each one of these parameters may be dynamically revealed. This means that in principle we can define for each static problem a number of dynamic problems equal to the number of possible parameter combinations.

In practice however, only certain DVRP are of interest, be the reason the simplicity of a model, its resemblance with a real-world problem or other reasons.

Table 1 shows a selection of significant DVRPs in the literature.

3 Main results - Methodologies

3.1 Strategy oriented papers - DTRP

The DTRP was introduced by Bertsimas and van Ryzin in [BVR91], and studied extensively in [BVR93a, BVR93b, Pap96]. The simplest version of this problem consists of a 'traveling repairman' or vehicle which has to service customers which dynamically appear over time. More complicated versions have m vehicles and finite capacity q (q=maximum number of customers before going back to depot). The objective function is the expected system time,

Table 1: A selection of DVRP in the literature

problem	*a-priori info.*	*dynamic info.*	*method*	*refs*
DSVRP	m vehicles, finite capacity	amount of request	Markov decision process	[DMT89]
DTRP	1 vehicle, infinite capacity	new request location	Poisson process, Euclidean space, queuing theory	[BVR91] [Pap96]
DTRP	m vehicles, finite and infinite capacity	new request location	Poisson process, Euclidean Space, queuing theory	[BVR93a]
DTRP	m vehicles, finite and infinite capacity	request location	General renewal process, Euclidean Space, queuing theory	[BVR93a]
DVRPTW	m vehicles, infinite capacity, time windows	new requests (location, time window)	heuristics, numerical calculations, simulation	[GGPT99]
DVRPTWD&P	m vehicles, infinite capacity	new requests (pick-up & delivery location, time window)	heuristics, numerical calculations, simulation	[GGPS98]
DRIVE	m heterogeneous vehicles, more depots	new requests (pick-up & delivery locations)	real world problem, branch-and-price algorithm	[SS91]
DTSP	1 vehicle, infinite capacity	deletion and insertion of cities	ACO	[GBH00]
DVRP			review article	[Psa95]

i.e. the expected waiting time plus on-site service time. The system time is the average time a customer must wait before its request is completed.

The problem formulation chosen is one that allows queuing theory analysis. Precisely, demands for service arrive according a renewal process with frequency λ, and are independently distributed on an Euclidean region of area A, r is the expected distance from a location in A to the depot. Vehicles travel at constant velocity v and spend at each location an amount of time that is independently and identically distributed with finite first (s) and second moments. The fraction of total vehicle time spent in on-site service is ρ.

The aim of the above cited papers is that of finding the best routing strategy for the service vehicle(s) that minimize the expected system time (it is assumed that every sub-problem is optimaly solved). This is in contrast with other works on DVRP, such as [GGPT99, GGPS98] analyzed in section 3.2, whose aim is not to find the best strategy, but to find a good optimization algorithm for the time changing sub-problems, given the strategy.

3.1.1 DTRP properties

In papers on DTRP authors derive general properties of the problem itself, such as stability condition, behavior of the optimum expected system time T^* by varying the spatial distribution of requests, lower bounds for the optimum, asymptotic behavior of the optimum in light and heavy traffic.

- Light traffic: for $\lambda \to 0$ the optimum expected system time T^* has a very simple form,

$$T^* \to \frac{r}{v} + s \text{ as } \lambda \to 0. \tag{1}$$

- The stability condition guarantees that at equilibrium the system time is finite. Bertsimas and van Ryzin [BVR93b] showed that this condition is

$$\rho + \frac{2\lambda r}{mvq} < 1 \tag{2}$$

  The second term can be interpreted as the fraction of time a vehicle spends in radial travel, since $\frac{2r}{v}$ 'is essentially the average time required to reach a set of q customers'. If $q = \infty$ (i.e. no capacity constraints), the stability condition (2) becomes simply $\rho = \lambda s < 1$ ([BVR91]).

- Heavy traffic: asymptotic behavior of the expected optimal system time T^* has been derived for $\rho \to$ limit stability value,

$$T^* \sim \frac{\gamma^2 \lambda A (1 - \frac{1}{q})^2}{m^2 v^2 (1 - \rho - \frac{2\lambda r}{mqv})^2}, \text{ as } \rho + \frac{2\lambda r}{mvq} \to 1, \tag{3}$$

  where γ is a constant that depends only on the strategy.

For infinite capacity (3) becomes

$$T^* \sim \frac{\gamma^2 \lambda A}{v^2(1-\rho)^2}, \text{ as } \rho \to 1, \tag{4}$$

which was derived in [BVR91] before than the more general (3).

- Multi vehicle versus single vehicle DTRP: by comparing (3) and (4) we see that

$$T^*(\text{m vehicles}) \sim \frac{1}{m^2} T^*(\text{1 vehicle}). \tag{5}$$

- Spatial distribution of requests has an influence on the optimal system time, as was shown in [BVR93a]. The worst situation is when requests are uniformly distributed on the service region:

$$T^*(\text{uniform distrib.}) \geq T^*(\text{non} - \text{uniform distrib.}). \tag{6}$$

3.1.2 The strategies

The main strategies analyzed in works on DTRP are the Stochastic Queue Median (SQM), the Traveling Salesman Problem strategy (TSP), the modified TSP (mod TSP), the Generation strategy (GEN), the Partition strategy (PART), the Nearest Neighbor (NN) and the Space Filling Curve (SFC). See (table 2). Strategies PART, TSP, mod TSP, GEN are analytically tractable, while SFC and NN must be numerically analyzed.

Here we only describe the most important ones.

Stochastic Queue Median strategy (SQM) Locate the vehicle at the median and serve the customers in a FCFS order, returning at the median after each service.

Traveling Salesman Problem strategy (TSP)
- While demands arrive, form sets of n consecutive demands, and deposit them in a queue;
- service the queue FCFS with the first available vehicle (or with the only one vehicle, for single-vehicle problems) by following optimal tours which start and end at the depot;
- optimize over n.

modified TSP strategy (mod TSP) Fig.2.

- For some fixed integer $k \geq 1$, divide the service region into wedges of area A/k;
- within each region, form sets with n/k consecutive demands (as in the TSP strategy) and, as sets are formed, deposit them in a queue;
- service the queue FCFS with the first available vehicle (or with the only one vehicle, for single-vehicle problems) by following optimal tours which start and end at the depot;

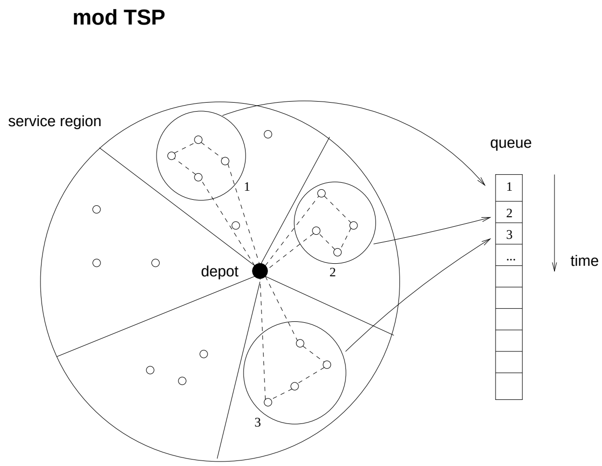

Figure 2: modified TSP strategy for DTRP with $k = 5$ and $n = 20$ (n is to be optimized) [BVR93b, BVR93a]

- optimize over n.

Nearest Neighbor strategy (NN)

- As the first m requests arrive, the m vehicles service them;
- after each service completion, the vehicle goes to the nearest unserviced location.

Generation strategy (GEN) Fig. 3. This has been applied only to single vehicle DTRP.

- Step 1. Initially position the vehicle at the median of the service region;
- Step 2. when the next request arrives, the vehicle moves directly to service it (this first request is the 'first generation');
- Step 3. If, at the completion time of the last generation, there are no demands waiting in queue, go to Step 1. Otherwise, solve the optimal TSP tour on the formed demands (which constitute the new 'generation'), and service them in the order in which they appear in the solution, then return to the median.

  Return to the beginning of Step 3.

Partitioning Policy (PART) Fig. 4. This has been applied only to single vehicle DTRP, on a region A squared.

- Divide the squared region A into m^2 squares (sub-regions).
- Sub-regions are visited by the vehicle in a given order, passing from one sub-region to an adjacent one.
- Inside each sub-region all demands are serviced in a FCFS order.

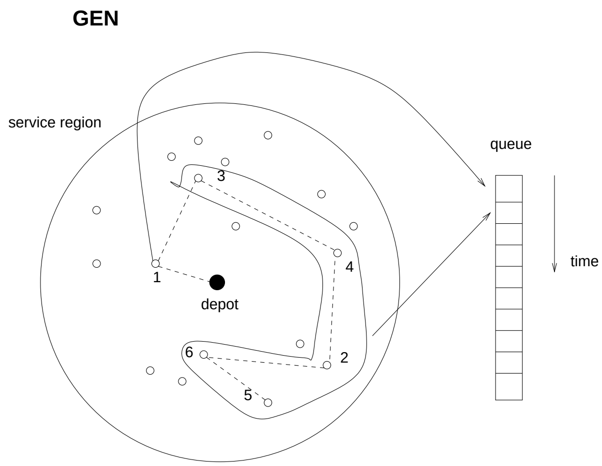

Figure 3: Generation strategy for DTRP [Pap96]

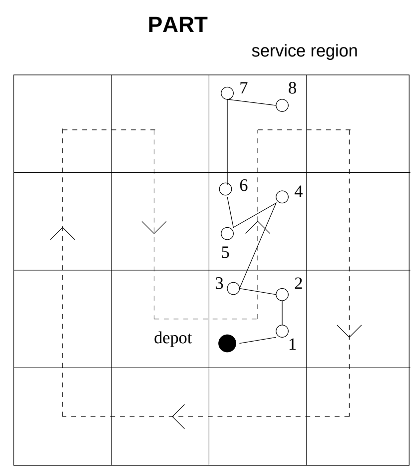

Figure 4: Partitioning strategy for DTRP with $m = 4$ [BVR91].

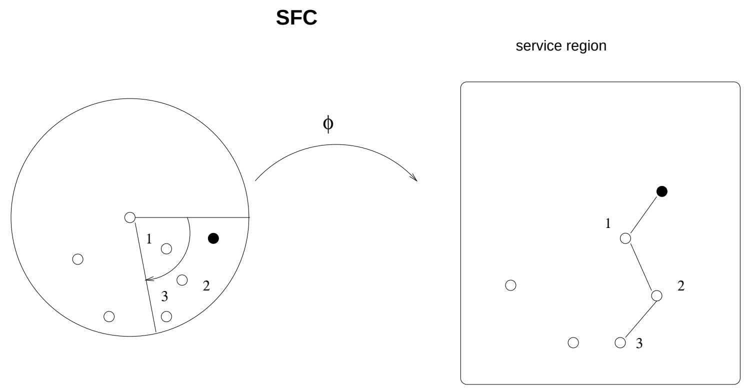

Figure 5: Space Filling Curve strategy for DTRP [BVR91, BVR93a].

- The pattern around A is continuously repeated until all requests have been accomplished.

Space Filling Curve (SFC) Fig. 5. This has been applied only to single vehicle DTRP.

- There is a mapping ϕ between the unit circle and the service region A, which preserves certain 'nearness' properties (see [BVR91]).
- Service demands in A as they are encountered in repeated clockwise sweeps of the circle.

3.1.3 Results on DTRP strategies

- Light traffic ($\lambda \to 0$), $m = 1, q = \infty$: the optimal strategy is the SQM strategy.

In heavy traffic some of the described strategies have been studied for both ($m = 1, q = \infty$) and ($m > 1, q =$ finite) problems, with small adaptations from one type of problem to the other. Here some interesting results about strategy performances in heavy traffic.

- Asymptotic system time behavior for heavy traffic: all strategies have shown the same behavior as in (3) and (4). This was provable for PART, TSP, mod TSP, GEN, and only empirically tested for NN and SFC.
- Best provable strategy in high traffic: the mod TSP strategy, which has a constant $\gamma_{mod\ TSP} \simeq 0.51$ ([BVR93b]), while $\gamma_{PART} = 1.02$, $\gamma_{SFC} \simeq 0.66$ and $\gamma_{NN} \simeq 0.64$ [BVR93a] . For ($m > 1$,q finite) problem, this strategy has the following provable constant factor guarantee: $T/T^* \leq 1.8$.
- Biased and unbiased strategies: in [BVR93a] two interesting variants of the mod TSP strategy have been analyzed. The biased strategy allows different waiting time values

Table 2: Formulations of DTRP and corresponding tested strategies

paper and problem	*strategies*
[BVR91], $m = 1$ $q = \infty$, Poisson process, uniform distrib.	SQM, PART, TSP, SFC, NN
[BVR93b], $m \geq 1$ $q =\leq \infty$, Poisson process, uniform distrib.	SQM, TSP, mod TSP, others
[BVR93a], $m \geq 1$ $q \leq \infty$, General renewal process, general distrib.	biased TSP, unbiased TSP, SFC, NN
[Pap96], $m = 1$ $q = \infty$, Poisson process, uniform distrib.	GEN, mod TSP, PART, NN

for different customers, while the unbiased strategy imposes the same expected waiting time for all. The interesting result is that

$$T(\text{unbiased strategy}) \geq T(\text{biased strategy}). \tag{7}$$

- Numerical experiments: [BVR91, Pap96] show that in practice the NN policy is constantly better than the other ones, from moderate to high traffic. This is valid for $(m = 1, q = \infty)$ and $(m > 1, q$ finite) problems.
- A flexible strategy: Papastavrou in [Pap96] analyzed the GEN strategy for the single vehicle, uncapacitated DTRP, which is optimal in light traffic and also behaves well in high traffic. In particular it has a constant factor guarantee respect to the TSP strategy:

$$\frac{T_{GEN}}{T_{TSP}} \leq 2, \quad \rho \to 1. \tag{8}$$

3.2 Heuristic oriented papers - DVRPTW and DVRPTWP&D

The two works of Gendreau et al. [GGPT99, GGPS98] analyze the DVRP with Time Windows and (DVRPTW) DVRPTW with Pick-up & Delivery (DVRPTWP&D). These works aim to test different heuristics embedded in one fixed strategy. This is much different from the approach of strategy-oriented papers, whose aim is to test different strategies, independently from the particular optimization algorithm which lies behind.

The two papers have many common elements, and draw similar conclusions on the performances of the tested algorithms. In the following we describe the strategy adopted, the optimization procedure used and the main results obtained.

3.2.1 Strategy

The strategy is very simple: the dynamic environment is handled by solving a series of static problems, with a new problem being defined each time an input update occurs. The static

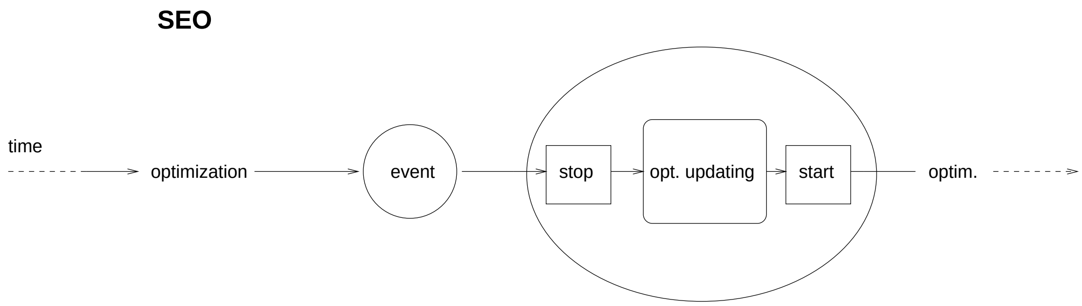

Figure 6: Single Event Optimization strategy for the DVRPTW and DVRTWP&D, [GGPS98, GGPT99]

problem contains the (yet unserviced) customer requests known at the time of the input update (figure 6). We call this strategy 'Single Event Optimization' (SEO).

The optimization procedure which solves the series of static problems runs between couples of "events", and stop each time an event occurs. An event can be either the occurrence of a new request or the end of service at a customer location. In the first case (new request) the static problem being solved is updated by adding the new request to the set of customers to be serviced, and the optimization procedure is restarted from a new starting solution. In the second case (end of service) the next location to visit is chosen by interrogating the stopped optimization algorithm, and then the algorithm is restarted on the previous set of customers minus the chosen one.

Comparison with other strategies: SEO is different from the strategies used for DTRP, where the optimization is done only when a certain new group of events has happened. The 'new group of events' we refer to is for TSP a set of n consecutive demands, for mod TSP a set of n/k consecutive demands, for GEN a new generation of demands.

3.2.2 Optimization procedure

In both papers the main optimization algorithm which the authors want to test is an Adaptive Tabu Search (ATS). Two slightly different versions are respectively used for the DVRPTW and DVRPTWD&P. The main difference lies in the Tabu Search neighborhood structure. In DVRPTW the only constraints are due to the time windows. In DVRPTWD&P the neighborhood must also deal with precedence constraints due to the fact that for every customer request the pick-up location must be visited before the delivery location.

The main features of ATS are:

- Adaptive Memory (diversification);
- Decomposition/Reconstruction procedure (intensification);

- chain exchange-based neighborhood;
- parallel implementation.

For more details on the Tabu Search procedures see [GGPT99, GGPS98].

Other optimization algorithms were implemented in order to produce comparative test performances of the ATS. Here is the comprehensive list of algorithm being tested:

- Insertion (I): a new request is inserted in the planned routes in the less expensive position. The insertion cost is evaluated with the true objective function in [GGPT99], and with an approximation function in [GGPS98].
- Construction (C): a new solution is rebuilt;
- Insertion followed by a local search (I+);
- Construction followed by local search (C+);
- Adaptive Local Descent (ALD): the ATS stopped at the first local minimum.

3.2.3 Main results

Numerical experiments were run on different scenarios to test the above algorithms. In [GGPS98] scenarios with different degrees of dynamism, number of vehicles and time horizons were simulated. In [GGPT99] additional degrees of freedom in the definition of scenarios were allowed, namely: spacial distribution of requests, short/long depot time windows, short/long on-site time service.

Here the main results. If not explicitly written, the results are valid for both papers.

- Adaptive heuristics (ALD and ATS) perform always better then others, particularly ATS. This means more serviced customers, less total travel time (and distance), and less delay.
- Difference between adaptive heuristics and the others is bigger for less dynamic scenarios. In highly dynamic environments there are more requests per route, which implies that more opportunities for optimization are available. These opportunities cannot be fully exploited by a sophisticated local search algorithm due to the lack of computation time between the occurrence of events.
- ATS improves by increasing the number of parallel processors. This also accounts for the fact that ATS gives its best when enough computation time is available.
- In [GGPT99] a comparison has been done between the dynamic and the static problem. Without any surprise the objective function is bigger in the dynamic case, and the gap between static and dynamic case decreases by decreasing dynamism.

3.2.4 Possible extensions

Possible extensions on the strategy adopted for these DVRPTW are

- allowing diversion of a vehicle which is traveling, by changing on-line its planned destination;
- retaining new service requests for a certain time period, and dispatching them all at once. Note that this is typical of already studied strategies adopted for the DTRP;
- using probabilistic knowledge about the future.

The tendency of Gendreau et al. is to converge in studying more real-world like problems. Therefore they aim in future works to study problems with more and/or other uncertainty sources, such as the travel time between two locations, sudden vehicle breakdown and so on.

4 Possible directions for further research

- The simplest DVRP is in principle the DTRP with one vehicle, or the DTSP (depending on which objective function one chooses). Up to now the graph version of these problem is unsolved. It could be interesting trying to use the insight obtained from the DTRP on a Euclidean region (section 3.1) for its graph formulation.
- What role can Ant System algorithms play in this class of problems? Maybe [DCD98] can suggest something.

  Ant based numerical experiments similarly to [GGPS98, GGPT99] can be done.
- Worst-case analysis: can local search worst-case properties be extended to VRPs?
- 'Competitive analysis': how far from the optimum can a DVRP algorithm deviate by only using information currently available and not the entire information that could be available if the problem were static?

  In the competitive analysis domain, the nearest problem to DVRPs is the on-line *k-server* problem [AL99]. In case there is any link between scheduling and VRP problems, [Sga98] can be useful.
- Global strategy versus specific sub-problem optimization: which is the relative importance of strategy respect to the optimization procedure of sub-problems?

  Trying to find non 'ad-hoc' strategies, but more general ones, can be another task.

References

[AFL+94] Giorgio Ausiello, Esteban Feuerstein, Stefano Leonardi, Leen Stougie, and Maurizio Talamo. Serving requests with on-line routing. In *Proceedings of the 4th Scandinavian Workshop on Algorithm Theory*. LNCS 824 Springer Verlag, 1994.

[AFL+95] Giorgio Ausiello, Esteban Feuerstein, Stefano Leonardi, Leen Stougie, and Maurizio Talamo. Competitive algorithms for the on-line traveling salesman. Workshop on Algorithms and Data Structures, 1995.

[AL99] Susanne Albers and Stefano Leonardi. Online algorithms. *ACM Computing Surveys*, 1999.

[BCP95] Dimitris Bertsimas, Philippe Chervi, and Michael Peterson. Computational approaches to stochastic vehicle routing problems. *Transportation Sciences*, 29(4):342–352, 1995.

[Ber92] Dimitris Bertsimas. A vehicle routing problem with stochastic demand. *Operations Research*, 40(3):574–585, 1992.

[BSL93] Dimitris J. Bertsimas and David Simchi-Levi. A new genaration of vehicle routing research: Robust algorithms, addressing uncertainty. working paper, 1993.

[BVR91] Dimitris J. Bertsimas and Garret Van Ryzin. A stochastic and dynamic vehicle routing problem in the euclidean plane. *Operations Research*, 39(4):601–615, 1991.

[BVR93a] Dimitris J. Bertsimas and Garret Van Ryzin. Stochastic and dynamic vehicle routing with general demand and interarrival time distributions. *Advanced Applied Probability*, 25:947–978, 1993.

[BVR93b] Dimitris J. Bertsimas and Garrett Van Ryzin. Stochastic and dynamic vehicle routing in the euclidean plane with multiple capacitated vehicles. *Operations Research*, 41(1):60–76, 1993.

[DCD98] Gianni Di Caro and Marco Dorigo. AntNet: distributed stigmergetic control for communications networks. *J. Artif. Intell. Res.*, 9:317–365, 1998.

[DMT89] Laporte G. Dror M. and P. Trudeau. Vehicle routing with stochastic demands: properties and solution frameworks. *Transportation Science*, 23:166–176, 1989.

[GBH00] Michael Guntsch, Jürgen Branke, and Schmeck Hartmut. ACO strategies for dynamic tsp. In Marco Dorigo, Martin Middendorf, and Stützle, editors, *Abstract Proceedings of ANTS'2000 From Ant Colonies to Artificial Ants: Second International Workshop on Ant Algorithms*, pages 59–62, 2000.

[GGPS98] Michael Gendreau, François Guertin, Jean-Yves Potvin, and René Séguin. Neighborhood search heuristics for a dynamic vehicle dispatching problem with pick-ups and deliveries. Technical Report CRT-98-10, Centre de Recherche sur les Transports, Université de Montréal, 1998.

[GGPT99] Michel Gendreau, François Guertin, Jean-Yves Potvin, and Eric Taillard. Parallel tabu search for real-time vehicle routing and dispatching. *Transportation Science*, 33(4):381–390, 1999.

[GLS95] Michel Gendreau, Gilbert Laporte, and Reneé Séguin. An exact algorithm for the vehicle routing problem with stochastic demands and customers. *Transportation Sciences*, 29(2):143–155, 1995.

[IGP99] Soumia Ichoua, Michel Gendreau, and Jean-Yves Potvin. Diversion issues in real-time vehicle dispatching. Technical report, 1999. CRT-99-13.

[KPS98] Philip Kilby, Patrick Prosser, and Paul Shaw. Dynamic VRPs: A study of scenarios. Technical report, 1998. Report APES-06-1998.

[LRI00] Xiangwen Lu, Amelia Regan, and Sandra Irani. An asymptotically optimal algorithm for the dynamic traveling repairman problem, 2000.

[Pap96] Jason D. Papastavrou. A stochastic and dynamic routing policy using branching processes with state dependent immigration. *European Journal of Operational Research*, 95:167–177, 1996.

[Psa95] Hilaros N. Psaraftis. Dynamic vehicle routing: status and prospects. *Annals of Operations Research*, 61:143–164, 1995.

[PTM97] Warren B. Powell, Michael T. Towns, and Arun Marar. On the value of optimal myopic solutions for dynamic routing and scheduling problems in the presence of user noncompliance. Technical Report SOR-97-15, Princeton University, August 1997.

[RG87] Francesco A. Rossi and Ilario Gavioli. *Aspects of Heuristic Methods in the Probabilistic Traveling Salesman Problem*, pages 214–227. World Scientific, Singapore, 1987.

[Sga98] Jiří Sgall. *On-Line Algorithms - The state of the art*, pages 196–231. Springer-Verlag, 1998.

[SP99] Michael R. Swihart and Jason D. Papastavrou. A stochastic and dynamic model for the single-vehicle pick-up and delivery problem. *European Journal of Operational Research*, 114:447–464, 1999.

[SS91] Martin Savelsberg and Marc Sol. DRIVE: Dynamic routing of independent vehicles. *Operations Research*, 46(4):474–490, 1991.

[ZQO00] Kenny Zhu Qili and Kar-loon Ong. A reactive method for real time dynamic vehicle routing problem. 12th ICTAI 2000, Vancouver, Canada, 2000.