

# Diversion Issues in Real-Time Vehicle Dispatching

SOUMIA ICHOUA, MICHEL GENDREAU, AND JEAN-YVES POTVIN

*Centre de recherche sur les transports and Département d'informatique et de recherche opérationnelle, Université de Montréal, C.P. 6128, succ. Centre-ville, Montréal, Québec, H3C 3J7, Canada*

*Recent technological advances in communication systems now allow the exploitation of real-time information for dynamic vehicle routing and scheduling. It is possible, in particular, to consider diverting a vehicle away from its current destination in response to a new customer request. In this paper, a strategy for assigning customer requests, which includes diversion, is proposed, and various issues related to it are presented. An empirical evaluation of the proposed approach is performed within a previously reported tabu search heuristic. Simulations compare the tabu search heuristic, with and without the new strategy, on a dynamic problem motivated from a courier service application. The results demonstrate the potential savings that can be obtained through the application of the proposed approach.*

In the past few years, there has been a rapid growth in communication and information technologies (e.g., global positioning satellites, cellular phones, geographic information systems, geosynchronous satellite-based systems, etc.). These recent advances afford opportunities for using real-time information to enhance the performance of decision systems in the area of vehicle routing. According to the fraction of requests that are known in advance, vehicle routing problems may be classified as static or dynamic. In the static case, all data are known before the routes are constructed and do not change afterward (e.g., location of transportation requests, demand, etc.). However, in the dynamic case, all or a fraction of all requests are revealed as the routes are executed. Hence, dispatchers are forced to react to events that occur in real time, such as new service requests, unexpected delays, accidents, etc. Numerous examples of dynamic vehicle routing and dispatching problems may be found in practice, like ambulance or police services, courier services, dial-a-ride problems (e.g., transportation-on-demand for the handicapped) and many others.

Due to the new information technologies mentioned above, real-time information is now available at lower costs. This explains the new interest that dynamic vehicle routing problems have gained recently. However, there is still a lack of methodologies that can efficiently solve dynamic vehicle routing

problems through a judicious integration of real-time information.

In dynamic vehicle routing problems, one potential use of real-time information is to divert a vehicle away from its current destination to serve a request that just occurred in the vicinity of its current position. In this work, we propose an approach for the dynamic assignment of new requests, which includes diversion, and we examine different ways of implementing it. An empirical evaluation is performed within the tabu search heuristic reported in GENDREAU et al. (1996b).

The problem considered is motivated from a courier service application found in the local operations of long-distance shipping services, where the mail is collected at different customers' locations and brought back to a central office for further processing and shipping. This problem belongs to the class of pick up (or delivery) only problems where a set of requests must be transported to (or from) a central depot. The goal is to design a set of minimum cost routes, originating and ending at the central depot, to satisfy the transportation requests. In a dynamic context, each new request is inserted in real time in the current set of planned routes, where a planned route is the sequence of requests that have been assigned to a vehicle but have not been serviced yet.

The paper is organized as follows. Section 1 presents a brief literature review dedicated to dynamic

vehicle routing problems. Then, Section 2 describes the problem considered in this study. Section 3 briefly introduces the original tabu search heuristic reported in Gendreau et al. (1996b) that we subsequently modified to explore the benefits of a new dynamic assignment strategy. Section 4 presents and discusses several issues related to it. Section 5 explains how the original tabu search heuristic was modified to include it. Section 6 reports computational results based on different ways of implementing the proposed strategy within the tabu search heuristic. Finally, section 7 summarizes our findings.

## 1. LITERATURE REVIEW

THE EARLIEST PAPERS in the literature on dynamic vehicle routing and dispatching were presented in the seventies and were either application-oriented (WILSON and COLVIN, 1977) or analytical (DAGANZO, 1978; STEIN, 1978).

By the end of the eighties, dynamic vehicle routing gained an increasing attention. Two major factors explain this tendency: new developments in information technologies and the need for decision systems that could exploit this information to better represent the real world. Interesting survey articles on dynamic vehicle routing can be found in PSARAFTIS (1988, 1995), POWELL (1995) and LUND, MADSEN, and RYGAARD (1996).

Because real-time vehicle routing problems are NP-hard and quick response times are required, exact algorithms are not yet capable of handling problems of realistic sizes (PSARAFTIS, 1980, 1983; et al. 1985 DIAL, 1995). This justifies the use of heuristics in real-time environments. According to how they deal with the dynamic aspects of the problem, the problem solving approaches reported in the literature can be classified in three major categories. They are reported in the following subsections.

### 1.1 Adaptation of Static Algorithms

This approach is based on the notion of a rolling horizon. As time unfolds, static problems are solved repeatedly over events found within an horizon of length  $L$  that extends from the current time  $t$  to  $t + L$ . Different strategies emerge when the length  $L$  of the horizon is modified. If  $L$  is very small, a myopic near-term strategy is observed. In some cases this strategy can provide near-optimal solutions (BELL et al., 1983; Dial, 1995; Psaraftis, 1985; Powell, 1988). In contrast, if  $L$  is very large, the problem considered is richer but, because it contains long-term events, the solution obtained is typically weaker unless a fast and powerful solution procedure is

used (TRUDEAU et al., 1989; GENDREAU et al., 1996a, 1996b). Note also that longer-term events may be postponed, because they are likely to lead to useless calculations and updates.

Adaptation of static procedures to dynamic vehicle routing problems can be divided in two classes:

- a. A sophisticated static problem-solving procedure, which typically involves a re-optimization of the routes, is applied each time an input update occurs. Several researchers have used this approach like Bell et al. (1983), HILL et al. (1988), BROWN et al. (1987), Psaraftis (1980, 1983), Psaraftis et al. (1985), POWELL et al. (1988) and Dial (1995). The drawback of this approach is the amount of computation time resulting from repeatedly executing the static algorithm. This disadvantage is more dramatic when new events occur frequently and when the execution of the static algorithm requires more time.
- b. Fast local operations (e.g., insertion) are used for reacting to any input revision (Trudeau et al., 1989; Wilson and Colvin, 1977; ROUSSEAU and ROY, 1988; Solonki, 1991; MADSEN, RAVN, and RYGAARD, 1995). This approach is easy to implement and is appropriate for a dynamic environment where time pressure is important (e.g., Lund, Madsen, and Rygaard, 1996). However, it is myopic because solutions are produced through consecutive insertions (whereas a complete reordering of the routes may lead to better solutions). To overcome this weakness, some authors combine local operations with re-optimization procedures. This is often achieved by executing a set of successive insertions followed by a local search (e.g., exchange procedures like 2-opt, see LIN, 1965). For different applications reported in the literature, see RIVARD (1981), ROY et al. (1985), Gendreau et al. (1996a, 1996b).

All studies mentioned above ignore the potential benefits of considering the stochastic aspects of the problems and trying to forecast the future. Stochastic methods are aimed at overcoming this weakness.

### 1.2 Stochastic Methods

Real-time dispatching problems have a stochastic nature (e.g., accidents, congestion, unexpected changes in meteorological conditions, etc.). Stochastic methods can be viewed as a natural way to judiciously address these issues. The goal is to react properly to an event to insure a good quality of service to the customers disturbed by these events, while minimizing their undesirable impact on the whole system. Two major classes of stochastic ap-

proaches are reported in the literature: stochastic programming and Markov decision processes.

### Markov Decision Processes

Formulations based on this modeling approach were proposed by Powell (1988), BERTSIMAS and VAN RYZIN (1991, 1993). Unfortunately, Markov decision processes are confronted with the following limitations that often prevent them from being applied to complex real-world problems: (i) the state space grows quickly with problem size; (ii) simplifying assumptions are often made to make the model more tractable.

### Stochastic Programming

The only work that we are aware of in this category is the one done by Powell et al. (1988) in his comparative review of dynamic vehicle allocation problems. The author proposed a hybrid model that combines insights from Markov decision processes and classical network formulations.

### 1.3 Other Methods

A new generation of approaches arises that try to replicate the dispatcher's decision making process. This is achieved by automating the decision procedure based on previous decisions taken by a skilled dispatcher. Within this framework, SHEN and POTVIN (1995) used a neural network to elaborate an expert consulting system for a dispatcher working in a courier service company. In the same context, BENYAHIA and POTVIN (1995) proposed an approach based on genetic programming.

### 1.4 Diversion Strategies

Apart from the study of REGAN, MAHMASSANI, and JAILLET (1994), Regan, Mahmassani, and Jaillet (1995) have proposed different diversion strategies in the context of a truck-load carrier. These studies will be discussed in Section 4.

## 2. PROBLEM DEFINITION

THE PROBLEM CONSIDERED is a dynamic pick up (or delivery) problem motivated from courier service applications. The static version of the problem can be stated as follows. Given a fixed size fleet of  $m$  identical vehicles and a set of requests to be served, the goal is to find a set of minimum cost routes that service these requests. More precisely:

- The vehicle routes must originate from and terminate at a fixed depot;
- Each vehicle services one route and the service point of each request is visited exactly once by exactly one vehicle;

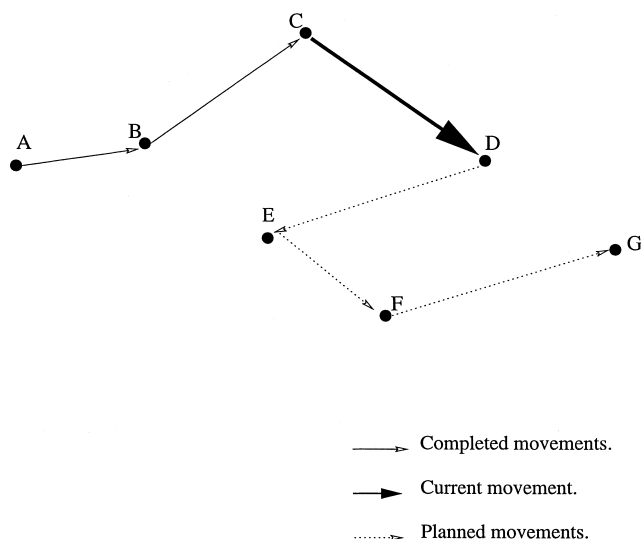


Fig. 1. A vehicle route in a dynamic setting.

- No capacity constraint is considered (because small items are transported);
- Each service point (including the depot) has its own time window  $[e, l]$ , where  $e$  is the earliest service time and  $l$  is the latest service time. Each route must start and end within the time window associated with the depot. However, the service points have “soft” time windows. Thereby, a vehicle can arrive before the lower bound or after the upper bound: if the vehicle arrives too early, it must wait to start its service; if the vehicle is too late, a penalty for lateness is incurred in the objective function.

The objective is to minimize a weighted sum of total distance and total lateness over all customers.

In the dynamic version of the problem, a number of service requests are not known completely ahead of time, but are rather dynamically revealed as time goes by. As illustrated in Figure 1, in a dynamic setting, a vehicle route can be divided into three parts at any instant  $t$ :

- completed movements, which form the part of the route already executed. Thereby, this part cannot be modified anymore;
- current movement to reach the current destination;
- planned movements, which constitute the portion of the routes not yet executed by the vehicle (planned route).

Given a new request at instant  $t$ , the problem is to assign this request to a particular vehicle and include it in its planned route at minimum cost.

In the following section, we briefly describe the original tabu search heuristic developed to tackle this problem.

### 3. A PARALLEL TABU SEARCH ALGORITHM

AT ANY INSTANT  $t$ , a solution to our dynamic problem is a set of planned routes, each beginning with the current destination of the associated vehicle. This current destination is fixed in the procedure to be described and cannot be modified (i.e., the vehicle must reach its current destination).

The algorithm developed in Gendreau et al. (1996b) is a parallel tabu search heuristic with an adaptive memory. Tabu search is an iterative local search technique that starts from some initial solution. At each iteration, a neighborhood is generated around the current solution and the best solution in this neighborhood becomes the new current solution (even if it does not provide an improvement). By allowing a degradation of the objective, it is possible to escape from bad local optima, as opposed to pure descent methods. The interested reader will find more details about this approach in GLOVER and LAGUNA (1997).

The algorithm developed in Gendreau et al. (1996b) can be described as follows:

- Construct  $I$  different initial solutions with a stochastic insertion heuristic (where the choice of the next customer to be inserted is randomized).
- Apply tabu search to each solution and store the resulting routes in the adaptive memory.
- While a stopping criterion is not met do:
  - Use the routes stored in the adaptive memory to construct the initial solution.
  - Decompose the problem into subproblems obtained through a decomposition procedure.
  - Apply tabu search to each subproblem.
  - Add the routes of the resulting solution in the adaptive memory.

In the following, the main points of this algorithm are briefly introduced.

#### Adaptive Memory

An adaptive memory that stores previously found elite solutions is used to generate new starting points for the tabu search. This is achieved by combining routes taken from different solutions in memory (ROCHAT and TAILLARD, 1995). Any new solution produced by the tabu search is included in the memory if it is not filled yet. Otherwise, the new solution replaces the worst solution in memory, if it is better.

#### Neighborhood Structure of Tabu Search

The procedure for generating the neighborhood is called the CROSS exchange. Basically, two segments of variable lengths are taken from two different routes and are swapped.

#### Parallel Implementation

To cope with real-time environments, the algorithm was implemented on a network of workstations. The parallelization of the procedure was achieved at two different levels: (1) different tabu search threads run in parallel, each of them starting from a different initial solution; and (2) within each search thread, many tabu searches run independently on subproblems obtained through a decomposition procedure. Given the available platform, a master-slave scheme was chosen to implement the procedure. The master process manages the adaptive memory and creates initial solutions for the slave processes that run the tabu search.

#### A Dynamic Environment

Given the dynamic context, it is important to maintain the consistency of the adaptive memory with the current environment. Also, the memory itself must be consistent, that is, the current destination of each vehicle (which is fixed) must be the same for all solutions in memory. Thus, whenever a new event occurs, the search threads are first interrupted. Then, after an appropriate update of the adaptive memory, it is possible to restart the search threads with new solutions constructed from this updated memory. A new event may be of two types:

- A vehicle has finished serving its current customer. In this case, its next destination (i.e., the next customer to be serviced) has to be determined. The best solution in memory is used for this purpose. The remaining solutions in the adaptive memory are updated by removing this customer from its current location and by inserting it in first position in the planned route of the vehicle.
- A new service request has just occurred. In this case, the new request is inserted in all solutions in adaptive memory. If there is no feasible insertion position in any solution, the request is rejected. Otherwise, the request is accepted and all solutions with no feasible insertion places are discarded from memory. The best solution is then re-optimized using a local search heuristic based on CROSS exchange to have at least one solution of high quality, which includes the new request.

Section 4 will now consider the case where the next destination of each vehicle is not fixed. To address this issue, a dynamic assignment strategy that includes diversion as a special case is proposed. Its implementation within the tabu search heuristic presented here is reported in Section 5.

**4. A NEW STRATEGY FOR HANDLING CUSTOMER REQUESTS**

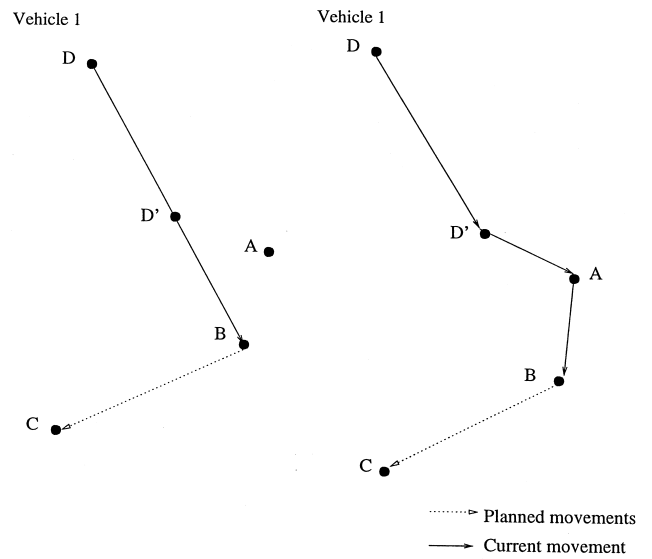
**4.1 Motivation**

Apart from the study of Regan, Mahmassani, and Jaillet (1994, 1995), all approaches that we are aware of fix the current destination of each vehicle. However, diverting a vehicle away from its current destination may be quite beneficial. This could be considered, for example, when a new request occurs in the vicinity of the current position of some vehicle, while the driver is on his way to his current destination. Technically, diversion is now possible due to recent advances in communication technologies. In the following, we present an approach which includes diversion as a special case. We also investigate the trade-off between computation time and solution quality, given that this strategy takes place in a context where time pressure is important (vehicles are moving fast and diversion opportunities may be quickly lost).

In Regan, Mahmassani, and Jaillet (1994, 1995), the authors have proposed different diversion strategies in the context of a truck-load carrier. This problem is a combined pick up and delivery problem with no consolidation (i.e., at any time, a vehicle is either empty or carrying a single load). In our application, in contrast, a vehicle may be used to carry many loads concurrently. Also, in the context of a truck-load carrier, requests are more distant from each other and the activities take place over a longer time horizon (e.g., a few days). Finally, diversion was evaluated in a context where simple local dispatching rules were used, thus resulting in myopic approaches. The results reported in these studies show that substantial benefits can be obtained through the exploitation of diversion.

**4.2 A Broader View**

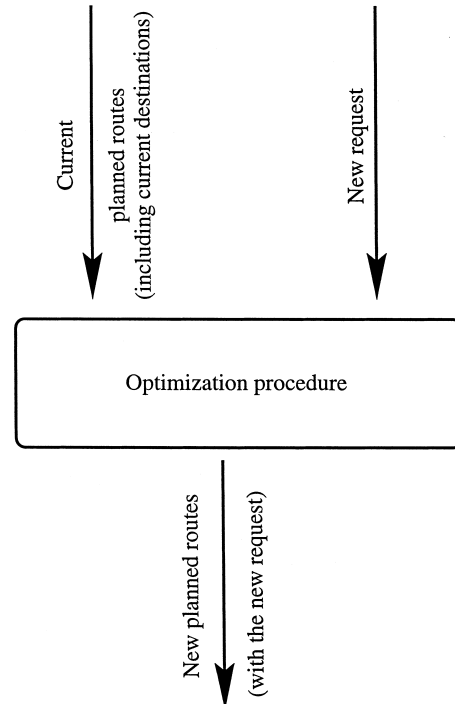
Diversion consists of allowing a vehicle to be diverted away from its current destination to serve a request that just occurred in the vicinity of its current position. Figure 2 shows an example for one vehicle route. We suppose that a new request unfolds at point A at instant  $t$ , while vehicle 1 is at position  $D'$  on his way to service point B. The only modification that can result from en-route diversion is the modification of the current destination of one



**Fig. 2.** Diversion.

vehicle. This occurs if the vehicle services the new request before its current destination B. In this paper, we propose a broader approach, where a change that occurs in the system may lead to redirecting one or more vehicles.

Figure 3 illustrates this new strategy in a schematic way. We first assume that the current destination of each vehicle is part of the planned route



**Fig. 3.** A dynamic assignment strategy that includes diversion.




Fig. 4. Time projection.

(and may thus change). Whenever a new request occurs, this request and the current routes are provided as input to some optimization procedure. The latter then produces a new set of routes with the new request. The optimization procedure may be anything that we can think of, such as:

- A simple insertion procedure. Diversion then happen if the new request is inserted just before the current destination of some vehicle.
- An insertion followed by some re-optimization procedure (e.g., tabu search).
- An exact algorithm (assuming that it is fast enough).

Depending on the optimization procedure, it may well happen that the current destination of each vehicle remains the same after its application (although the planned routes have changed). Conversely, the current destination of one or more vehicles may change, but not necessary for servicing the new request (which appears further away on a planned route). Thus, diversion is simply a particular outcome of the dynamic assignment strategy presented here.

### 4.3 Solution Quality Versus Computation Time Trade-Off

Let us assume that a new request appears at point  $A$  at instant  $t$  and that it is possible to estimate the position  $D'$  of the vehicle (between the point  $D$  just served and the current destination  $B$ ), see Figure 4.

Because new opportunities for including the new request are now offered through the inclusion of the current destination in the planned route, the solution obtained should be better. However, some amount of time  $\delta t$  is required to apply the optimization procedure. Hence, a decision based on the situation at instant  $t$  when a new request appears, does

not correspond to the state of the system at  $(t + \delta t)$  when the decision becomes available (i.e., the vehicle will be at position  $D''$  rather than  $D'$ ). Consequently, the situation should be assessed at instant  $(t + \delta t)$  rather than instant  $t$ . The goal is to provide a decision that will fit in the context found at the end of the time allocated to the optimization procedure.

Clearly, this amount of time is a very important factor. In dynamic settings, finding a value that achieves a good trade-off between execution time and solution quality is very challenging. If  $\delta t$  is too small, not enough time is available for the optimization procedure. Therefore, the situation may be incorrectly assessed, resulting in a bad decision. Otherwise, if  $\delta t$  is too large, opportunities for diversion may be lost.

To assess the benefits that can be gained from the application of our dynamic assignment strategy, the parallel tabu search heuristic developed by Gendreau et al. (1996b) was used as the optimization procedure. The next section shows how the original algorithm was modified to implement it.

## 5. IMPLEMENTATION

### 5.1 Inclusion of the Current Destination

In the original implementation of the tabu search heuristic (Gendreau et al., 1996b) reported in Section 3, the current destination of each vehicle is a fixed node that is not part of the planned route to be optimized. We thus included the current destination of each vehicle in the planned route and used instead the current position of the vehicle as the fixed node. With this modification, the re-optimization procedure based on CROSS exchanges can now move the current destination around (as any other service point on the planned route).

### 5.2 Time Projection

When a new request is received at time  $t$ , all solutions in the adaptive memory are updated according to the state of the system at time  $(t + \delta t)$ , where  $\delta t$  is the amount of time allocated to the optimization procedure. The planned routes of the best solution in adaptive memory are used for this purpose. A dummy point  $F_{t+\delta t}^i$  is created for every vehicle route  $i$  of every solution in the adaptive memory. This dummy point  $F_{t+\delta t}^i$  represents the position of vehicle  $i$  at time  $(t + \delta t)$  and is the fixed node associated with route  $i$ .

Time projection of the solutions in adaptive memory is illustrated in Figure 5. In this figure, a fleet of two vehicles is considered: a new request occurs at instant  $t$ , and  $\delta t$  is the amount of time allocated to the optimization procedure. In the best solution  $S^*$

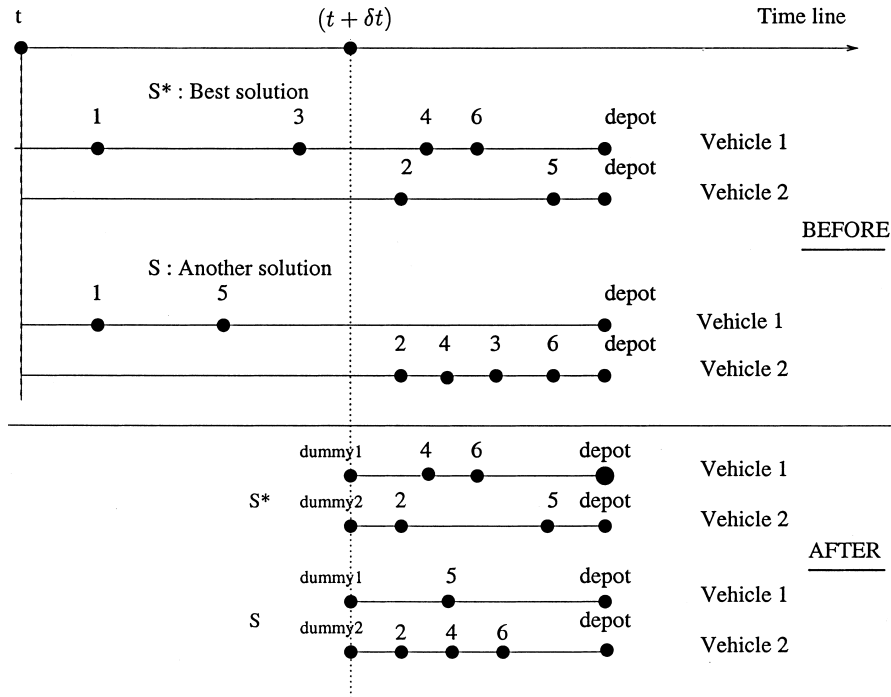


Fig. 5. Time projection in adaptive memory.

at instant  $t + \delta t$ , vehicle 1 will have finished service at customer 1 and customer 3, whereas vehicle 2 will not have yet reached customer 2. Hence, before applying the optimization procedure, customer 1 and customer 3 are removed from  $S^*$  and the dummy points are inserted at the beginning of the two routes. To update solution  $S$ , customer 1 and customer 3 are first removed from it (because they are serviced). Then the dummy fixed points are added at the beginning of the two routes. It is worth noting that the customers that have not been serviced yet are found in the same order on the planned routes before and after the update. Only the service times are likely to change.

### 5.3 Handling New Events

Suppose that a new event occurs at instant  $t$ . The new events that occur during the time interval  $[t, t + \delta t]$  will not interrupt the search threads. Instead, these events are handled as follows:

- a. When a vehicle has finished serving a customer, its next destination is determined by the best solution found in the adaptive memory. That is, in the time interval  $[t, t + \delta t]$ , everything is “frozen” and we simply follow the best solution in memory.
- b. When a new request occurs, we must decide whether to accept or reject it, because the customer may be waiting for a prompt answer about

the ability to handle his request (a request is rejected if the end time of the planned route exceeds the upper bound of the time window at the depot after its insertion). Basically, the new request is inserted in all solutions in a copy of the adaptive memory. Solutions with no feasible insertion places are discarded. If at least one feasible solution remains, the new request is accepted and is kept in a list of pending requests. Otherwise, it is rejected.

At instant  $(t + \delta t)$ , all search threads are interrupted. The adaptive memory is updated with the best solution found by the optimization procedure, and the pending requests are handled one by one, as in the original algorithm (Gendreau et al., 1996b). Finally, the search threads are restarted with new solutions constructed from the adaptive memory.

### 5.4 Setting $\delta t$

The amount of time allocated to the optimization procedure must be related to the rate at which new events occur. When few events occur, a large evaluation time can be allowed. Conversely, in case of a highly dynamic environment, the evaluation must be done quickly. Based on these observations, three rules are considered.

RULE 1.  $\delta t$  is chosen in such a way that the optimization procedure ends before any vehicle begins

its service at its current destination. Let us assume that a service request is received at time  $t$ , and that service at the current destination on route  $i$  begins at time  $t_i$ . Then, choose  $\delta t$  such that  $t + \delta t \leq t_i, \forall i$ , that is,  $\delta t \leq \min_i(t_i) - t$ . By choosing  $\delta t \leq \min_i(t_i) - t$ , we only need to add a dummy fixed node at the beginning of each planned route, which corresponds to the position of the vehicle at time  $t + \delta t$ . Unfortunately,  $\delta t$  is often very small because there is typically one or more vehicles that are close to their current destination. This rule was not retained in our computational experiments.

**RULE 2.** The amount of time  $\delta t$  is fixed a priori and is related to the average time interval between the occurrence of two requests. It is proportional to a moving average of the  $l$  last intervals, that is  $\delta t = \alpha_1 T^-$ , where  $\alpha_1$  is a constant and  $T^-$  is the moving average. By relating  $\delta t$  to the average request inter-arrival times, our choice is based on the intensity of occurrence of new requests.

**RULE 3.** In rule 2, the computation of  $\delta t$  may take into account a service request that is not so urgent (i.e., its time window is much later than the current time). This is true, in particular, in contexts where the calls are received much before the requested service times. Therefore, another rule is proposed. Let  $X$  be the length of some time horizon and  $l_X$  the number of requests on the planned routes found within that horizon. Then the average time per request is defined as  $X/l_X$  and the interval  $\delta t$  is such that  $\delta t = \alpha_2 X/l_X$ , where  $\alpha_2$  is a constant.

Rules 2 and 3 represent two different strategies. In rule 2, intensity is measured using past events, because the moving average is calculated over the  $l$  last request inter-arrival times. In rule 3, intensity is assessed using events to come, because the average time per request is evaluated over requests found on the planned routes. Results obtained with these two strategies for setting  $\delta t$  are reported in the next section.

## 6. COMPUTATIONAL RESULTS

### 6.1 Simulation Framework

Simulations were performed to assess the benefits of the proposed approach. Data were taken from Solomon's 100-customer Euclidean problems (SOLOMON, 1987), as in Gendreau et al. (1996b). In the following, we describe the essential components of the simulator. The interested reader will find more

details about this simulator in Gendreau et al. (1996b).

- The customer locations are generated within a  $[0, 100]^2$  square.
- The travel times are proportional to the corresponding Euclidean distances.
- Six different sets of problems are defined, namely  $C_1, C_2, R_1, R_2, RC_1$ , and  $RC_2$ . The customers are uniformly distributed in the problems of type  $R$ , clustered in groups in the problems of type  $C$  and mixed in the problems of type  $RC$ . In the problems of type 1, only a few customers can be serviced on each route due to a narrow time window at the depot, as opposed to problems of type 2 where each route may have many customers.
- The set of requests is divided into two subsets. The first subset contains requests that are assumed to be known at the start of the day. They are randomly selected among the entire set of requests with a bias in favor of requests with early time windows. The second subset contains requests that are received in real-time. In this case, the arrival time of each request is adjusted to create two different types of scenarios: scenario 1 with half of the requests known in advance, and scenario 2 with only 25% of the requests known in advance. In the latter case, a higher degree of dynamism is obtained, because there are more requests that occur in real-time within the same time horizon.
- Using minutes as time units for time-related data in Solomon's files, the time horizon is set to 15 minutes. The intent is to have a realistic rate of new unfolding requests (about three requests per minute for scenario 1 and five requests per minute for scenario 2).

### 6.2 Numerical Results

The experiments reported in this section were performed on a network of 9 SUN UltraSparc-III workstations (300 MHz). Each process was programmed in C++ and communication between the processes was handled by the Parallel Virtual Machine Software. In these experiments, the objective function to minimize is the sum of the total distance traveled and the total lateness over all customer locations. The fleet size was set to the number of routes in the best solution reported in the literature for each problem (Gendreau et al., 1996b). The value of the parameters found in the original algorithm were kept as in Gendreau et al. (1996b). Parameters related to our new dynamic assignment strategy are tested in the following.



## 6.2.1 Preliminary Tests

Scenario 1 (with half of the requests known in advance) was first considered for the preliminary experiments. The new algorithm was tested under rule 2 (see Section 5) where the interval  $\delta t$  is computed as a multiple of a moving average  $T^-$  of inter-arrival times. That is,  $\delta t = \alpha_1 T^-$ , where  $\alpha_1$  is a multiplier. The moving average is taken over  $l$  loads, with  $l = \beta_1 K$ , where  $K$  is the number of requests that are known in advance and  $\beta_1$  is a multiplier. Parameters  $\alpha_1$  and  $\beta_1$  are determined empirically.

Preliminary experiments were performed over two problems selected in each of the six categories  $C_1$ ,  $C_2$ ,  $R_1$ ,  $R_2$ ,  $RC_1$ , and  $RC_2$ . In these experiments, the values 0.10, 0.25, 0.50, 0.75, and 0.90 were tested for  $\alpha_1$  whereas the values 0.10 and 0.25 were tested for  $\beta_1$ . The results obtained did not clearly show an improvement over the original algorithm. Thus, more investigations were required, leading to an important refinement in our algorithm.

As mentioned in Section 5, a copy of the adaptive memory is used to decide whether a new request should be accepted or rejected. This can lead to inconsistencies with the true memory. In particular, two important problems were identified:

- A new best solution that is received from a tabu search thread is not inserted in the copy (only in the adaptive memory itself). Consequently, requests that could have been inserted in the solutions found in the true adaptive memory at time  $t + \delta t$  may have been wrongfully rejected.
- A new solution may be received from a tabu search thread while the adaptive memory is full. In this case, the worst solution in adaptive memory is replaced by the new one. At this point, there is an inconsistency with the copy that still includes this worst solution.

To improve the implementation, both the adaptive memory and its copy are updated whenever a new solution is received from a tabu search thread. Table I shows the kind of improvement that can be obtained with the refined implementation, using  $\alpha_1 = 0.9$  and  $\beta_1 = 0.25$ . The four numbers in each entry are the average number of unserved customers, distance traveled, total lateness, and objective value (sum of total distance traveled and total lateness), respectively, for each problem set. The row "Overall" contains averages taken over the entire set of 56 test problems. When the number of unserved customers is maintained, the objective value is improved. In contrast, when the number of unserved customers is

TABLE I  
Comparison of the Refined Implementation (New) with the Old One Using Scenario 1

Problem Set	New $\alpha_1 = 0.90,$ $\beta_1 = 0.25$	Old $\alpha_1 = 0.90,$ $\beta_1 = 0.25$
$C_1$ , 9 problems	0 <sup>a</sup>	0
	835.781 <sup>b</sup>	838.784
	2.025 <sup>c</sup>	1.403
	837.807 <sup>d</sup>	840.187
$R_1$ , 12 problems	0.583	0.75
	1183.00	1182.49
	51.125	52.55
	1234.125	1235.04
$RC_1$ , 8 problems	0.375	0.25
	1317.37	1316.12
	47.058	47.852
	1364.43	1363.97
$C_2$ , 8 problems	0	0.25
	608.321	602.624
	0	0
	608.321	602.624
$R_2$ , 11 problems	0.181	0.181
	1058.76	1073.06
	57.20	66.64
	1115.97	1139.70
$RC_2$ , 8 problems	0	0
	1188.09	1244.00
	45.166	53.276
	1233.26	1297.27
Overall, 56 problems	0.21	0.27
	1040.62	1050.80
	35.69	39.02
	1076.31	1089.82

<sup>a</sup>Number of unserved customers.

<sup>b</sup>Total distance traveled.

<sup>c</sup>Total lateness.

<sup>d</sup>Objective value.

reduced, this is achieved with only a slight increase in the objective value.

## 6.2.2 Parameter Values

To determine the best values for parameters  $\alpha_1$  and  $\beta_1$  under rule 2, tests were performed over a small sample of problems selected in each of the six categories  $C_1$ ,  $C_2$ ,  $R_1$ ,  $R_2$ ,  $RC_1$ , and  $RC_2$ . Several values had to be tested for every problem in the sample. Thereby, a significant amount of computation time is required if the size of the sample is too large. In contrast, this size has to be large enough to get a correct calibration. For these experiments, the size of the sample was set to 4. Because scenarios 1 and 2 represent different degrees of dynamism (requests occur more frequently in scenario 2), the experiments were conducted separately for the two scenarios. Table II reports some results obtained with different combinations of values using scenario 1 and scenario 2. As in Table I, the four numbers in each entry are the average number of unserved customers, distance traveled, total lateness, and objec-

TABLE II  
Searching for the Best Parameter Values under Rule 2

Problem Set	Scenario 1					Scenario 2				
	$\alpha_1 = 0.90$ $\beta_1 = 0.15$	$\alpha_1 = 0.90$ $\beta_1 = 0.25$	$\alpha_1 = 1.20$ $\beta_1 = 0.25$	$\alpha_1 = 0.50$ $\beta_1 = 0.10$	$\alpha_1 = 0.50$ $\beta_1 = 0.15$	$\alpha_1 = 0.50$ $\beta_1 = 0.10$	$\alpha_1 = 0.50$ $\beta_1 = 0.15$	$\alpha_1 = 1.20$ $\beta_1 = 0.15$	$\alpha_1 = 0.90$ $\beta_1 = 0.25$	$\alpha_1 = 0.25$ $\beta_1 = 0.25$
$C_1$ 4 problems	0 <sup>a</sup>	0	0	0	0	0	0	0	0	0
	829.89 <sup>b</sup>	832.35	828.535	852.31	828.94	882.42	871.055	908.478	925.06	881.658
	0 <sup>c</sup>	0	0	0.65	0	0.105	0	19.79	15.027	2.127
$R_1$ 4 problems	829.89 <sup>d</sup>	832.35	828.535	852.96	828.94	882.525	871.055	928.268	940.088	883.785
	0.25	0	0.25	0.25	0	0.25	0	0.25	0	0.50
	1253.02	1284.25	1262.32	1266.03	1281.79	1296.70	1284.24	1303.97	1303.56	1296.28
$RC_1$ 4 problems	63.435	38.307	43.90	47.12	45.385	76.00	63.982	73.315	67.71	53.067
	1316.46	1322.56	1306.22	1313.15	1327.18	1372.70	1348.22	1377.28	1371.28	1349.35
	0.5	0.75	0.5	0.75	0.25	0.25	0.75	0.5	1.00	1.00
$C_2$ 4 problems	1334.86	1324.50	1310.76	1318.16	1330.19	1366.52	1349.24	1405.42	1402.88	1354.10
	74.472	62.985	55.962	67.762	83.022	54.99	80.737	77.917	61.32	64.82
	1409.33	1387.49	1366.72	1385.92	1413.22	1421.50	1429.98	1483.34	1464.20	1418.92
$R_2$ 4 problems	0	0	0	0	0	0	0.25	0	0	0
	642.978	625.628	612.602	616.238	627.51	607.572	619.515	613.62	621.495	612.33
	0	0	0	0	0	0	0	0.002	0	0
$RC_2$ 4 problems	642.978	625.628	612.602	616.238	627.51	607.572	619.515	613.62	621.495	612.33
	0.5	0.25	0	0.75	0	0	0.75	0	0.25	0
	1096.48	1106.52	1106.13	1126.82	1107.19	1113.66	1093.96	1119.40	1140.72	1106.98
$R_2$ 4 problems	24.182	30.155	19.337	53.782	14.645	17.63	93.605	23.73	14.032	22.36
	1120.66	1136.67	1125.47	1180.60	1121.84	1131.29	1187.56	1143.13	1154.76	1129.34
	0	0	0	0	0	0	0	0	0	0
$RC_2$ 4 problems	1226.10	1184.14	1188.22	1198.09	1180.53	1207.95	1235.45	1213.50	1144.22	1218.88
	24.045	38.74	10.745	45.41	17.482	37.50	19.827	35.867	60.822	37.622
	1250.14	1222.88	1198.96	1243.50	1198.01	1245.45	1255.28	1249.37	1205.04	1256.50

<sup>a</sup>Number of unserved customers.  
<sup>b</sup>Total distance traveled.  
<sup>c</sup>Total lateness.  
<sup>d</sup>Objective value.

tive value, respectively, for each problem set. The values  $\alpha_1 = 0.50, \beta_1 = 0.15$  for scenario 1 and  $\alpha_1 = 0.50, \beta_1 = 0.10$  for scenario 2 were finally selected.

Under rule 3, the interval  $\delta t$  is computed as a function of the length of some rolling time horizon, that is,  $\delta t = \alpha_2 X / l_X$ , where  $\alpha_2$  is a constant,  $l_X$  is the number of requests on the planned routes within horizon  $X$ ,  $X = \beta_2 L$ , where  $L$  is the length of the time horizon for the simulation and  $\beta_2$  is a multiplier. The tests were conducted as for rule 2. Table III reports some results obtained with different combinations of values using scenario 1 and scenario 2. The values  $\alpha_2 = 0.50, \beta_2 = 0.25$  for scenario 1 and  $\alpha_2 = 0.25, \beta_2 = 0.25$  for scenario 2 were finally selected. Note that, under rule 2,  $\delta t$  is about 5 seconds for scenario 1, and 2.8 seconds for scenario 2. Under rule 3,  $\delta t$  is about 4 seconds for scenario 1 and 2 seconds for scenario 2.

### 6.2.3 Assessing the New Implementation

Table IV compare the new algorithm (with the best parameter values found under rule 2 and rule 3) with the original algorithm reported in Gendreau et al. (1996b), for scenario 1 and scenario 2. Here, the tests were conducted over all problems in each

category. The results highlight the fact that the new implementation performs well against the original algorithm for the two scenarios. More precisely:

**RULE 2.** For scenario 1, both the number of unserved customers and objective value were improved with regard to the original algorithm (apart from subset  $R_2$ ). The reduction ranged from 33.3 to 100% for the number of unserved customers and from 0.6 to 3.0% for the objective value. In the case of subset  $R_2$ , the increase in the objective value was marginal when compared to the reduction in the number of unserved customers: 0.3% against 60%. Note that the lowest reduction observed in the objective value with the new strategy is still larger than the only increase observed (0.7% against 0.3%). For scenario 2, the results demonstrate that applying our approach under rule 2 reduces both the number of unserved customers and the objective value in all categories. The reduction ranges from 11.2 to 67% for the number of unserved customers, and from 1.0 to 2.7% for the objective value.

**RULE 3.** For scenario 1 and for categories  $C_1$  and  $RC_2$ , where the original algorithm succeeds in

TABLE III  
*Searching for the Best Parameters Values Under Rule 3*

Problem Set	Scenario 1					Scenario 2				
	$\alpha_2 = 0.50$ $\beta_2 = 0.15$	$\alpha_2 = 0.90$ $\beta_2 = 0.15$	$\alpha_2 = 1.20$ $\beta_2 = 0.15$	$\alpha_2 = 0.50$ $\beta_2 = 0.25$	$\beta_2 = 0.25$ $\beta_2 = 0.25$	$\alpha_2 = 0.25$ $\beta_2 = 0.25$	$\alpha_2 = 0.50$ $\beta_2 = 0.25$	$\alpha_2 = 0.90$ $\beta_2 = 0.15$	$\alpha_2 = 1.20$ $\beta_2 = 0.15$	$\alpha_2 = 0.50$ $\beta_2 = 0.15$
$C_1$ 4 problems	0 <sup>a</sup>	0	0	0	0	0	0	0	0	0
	829.26 <sup>b</sup>	834.125	857.005	829.112	839.488	880.675	889.275	880.555	955.958	882.895
	0 <sup>c</sup>	0	0.375	0	0	1.977	0	0.90	1.695	0
	829.26 <sup>d</sup>	834.125	857.38	829.112	839.488	882.652	889.275	881.455	957.652	882.895
$R_1$ 4 problems	0.25	0.25	0.25	0.25	0.25	0	0.25	0.25	0.25	0.25
	1272.71	1274.41	1267.60	1275.08	1265.19	1282.77	1302.15	1293.42	1282.50	1285.20
	51.997	52.532	43.627	39.537	45.217	46.415	51.71	60.865	58.432	56.23
	1324.71	1326.94	1311.23	1314.62	1310.41	1329.18	1353.86	1354.29	1340.93	1341.42
$RC_1$ 4 problems	0.75	0.75	1.00	0.75	0.75	0	1.00	1.00	1.00	0.25
	1339.43	1344.64	1321.06	1325.66	1352.71	1366.24	1358.48	1380.77	1363.36	1334.95
	62.587	74.707	71.185	60.04	64.92	71.422	89.21	60.082	71.322	68.03
	1402.02	1419.34	1392.25	1385.70	1417.63	1437.66	1447.69	1440.86	1434.68	1402.98
$C_2$ 4 problems	0	0	0.5	0	0	0	0.25	0	0	0
	615.518	614.925	616.928	616.878	624.21	601.972	650.912	612.018	630.335	603.122
	0	0	0	0	0	0	0	0	0	0
	615.518	614.925	616.928	616.878	624.21	601.972	650.912	612.018	630.335	603.122
$R_2$ 4 problems	0	0	0.5	0	0	0	0.50	0	0.50	0.50
	1110.10	1095.61	1113.52	1091.81	1105.60	1089.10	1090.41	1083.09	1104.65	1101.28
	60.165	61.09	46.065	56.442	64.455	14.995	36.237	22.917	14.192	40.322
	1170.27	1156.70	1159.58	1148.25	1170.06	1104.09	1126.65	1106.00	1118.84	1141.60
$RC_2$ 4 problems	0	0	0	0	0	0	0	0	0	0
	1186.05	1198.42	1172.54	1177.27	1198.94	1221.02	1207.80	1191.57	1207.86	1218.99
	23.472	29.502	23.862	25.457	32.652	35.692	28.942	40.562	46.522	27.662
	1209.52	1227.92	1196.40	1202.73	1231.60	1256.72	1236.75	1232.13	1254.38	1246.66

<sup>a</sup>Number of unserved customers.

<sup>b</sup>Total distance traveled.

<sup>c</sup>Total lateness.

<sup>d</sup>Objective value.

servicing all customers, the new algorithm improves the objective value. This reduction ranges from 1.3 to 4.3%. In the case of  $R_1$  and  $C_2$ , both the number of unserved customers and objective value were improved. The reduction ranged from 66.8 to 100% for the number of unserved customers, and from 2.6 to 2.7% for the objective value. For  $RC_1$  and  $R_2$ , the number of unserved customers was improved at the cost of a small increase in the objective value (i.e., 67% against 0.04% for  $RC_1$  and 80% against 0.6% for  $R_2$ ). Using rule 3 instead of rule 2 results in a small reduction of the objective value ranging from 0.5 to 2.0% when the number of unserved customers is the same. Rule 3 succeeds in reducing the number of unserved customers obtained by rule 2 for  $RC_1$  and  $R_2$ , at the cost of a small increase in the objective value (i.e. 50% against 1.8%, and 51% against 0.2%).

For scenario 2, the comparison with the original algorithm indicates a reduction of both the number of unserved customers and the objective value over all categories. This reduction ranges from 16.8 to 100% for the number of unserved customers, and from 2.0 to 4.3% for the objective value. The results

also indicate that, apart from  $RC_1$  and  $R_2$ , rule 3 leads to the same number of unserved customers than rule 2, but it reduces the objective value by 0.9 to 2.2%. For  $R_2$  and  $RC_1$ , rule 3 succeeds in reducing both the number of unserved customers (i.e., 12 and 100%, respectively), and the objective value (i.e., 0.6 and 1.1%, respectively).

## 7. CONCLUSION

IN THIS PAPER, a new strategy for the dynamic assignment of new requests, which includes diversion as a special case was proposed. Tests were conducted to assess its benefits within a previously reported tabu search heuristic. Some issues related to the proposed model were addressed and explored. Results show that a reduction in the sum of total distance traveled and total lateness, and in the number of unserved customers is observed when compared with the original heuristic, where the current destination of each vehicle is fixed.

Future work will be aimed at considering additional sources of uncertainty like congestion, accidents, vehicle breakdowns, and others. A taxonomy

TABLE IV  
Comparison with the Original Algorithm for Both Scenarios

Problem Set	Scenario 1			Scenario 2		
	Original Algorithm	New Algorithm		Original Algorithm	New Algorithm	
		Rule 2 $\alpha_1 = 0.50$ $\beta_1 = 0.15$	Rule 3 $\alpha_2 = 0.50$ $\beta_2 = 0.25$		Rule 2 $\alpha_1 = 0.50$ $\beta_1 = 0.10$	Rule 3 $\alpha_2 = 0.25$ $\beta_2 = 0.25$
$C_1$ 9 problems	0 <sup>a</sup>	0	0	0	0	0
	835.098 <sup>b</sup>	830.024	829.628	877.332	843.924	833.061
	5.535 <sup>c</sup>	3.341	0	10.455	11.227	8.027
	840.633 <sup>d</sup>	833.366	829.628	878.378	855.152	841.089
$R_2$ 12 problems	0.5	0.166	0.166	0.4	0.333	0.333
	1219.42	1194.75	1178.54	1226.21	1203.99	1197.30
	43.267	48.625	51.460	55.431	67.167	49.121
	1262.69	1243.38	1230.00	1281.64	1271.16	1246.42
$RC_1$ 8 problems	0.375	0.25	0.125	0.375	0.125	0
	1349.65	1304.58	1333.41	1400.50	1358.78	1346.98
	44.22	64.955	60.965	56.391	73.28	69.002
	1393.87	1369.53	1394.38	1456.89	1432.06	1415.98
$C_2$ 8 problems	0.125	0	0	0	0	0
	612.798	609.232	597.011	610.775	600.796	595.635
	0.326	0	0	0	0	0
	613.124	609.232	597.011	610.775	600.796	595.635
$R_2$ 11 problems	0.454	0.181	0.09	0.818	0.727	0.636
	1028.34	1035.09	1020.55	1043.96	1043.43	1027.82
	31.96	29.053	46.334	216.696	193.885	207.639
	1060.30	1064.15	1066.88	1260.65	1237.32	1235.46
$RC_2$ 8 problems	0	0	0	0	0	0
	1203.75	1175.87	1160.83	1222.22	1189.48	1184.77
	30.811	22.108	20.672	27.441	41.023	19.568
	1234.56	1197.98	1181.50	1249.66	1230.51	1204.34
Overall 56 problems	0.27	0.10	0.07	0.30	0.23	0.19
	1049.83	1034.11	1027.95	1070.75	1048.45	1039.11
	27.20	29.10	31.79	68.10	70.61	65.25
	1077.03	1063.22	1059.74	1138.85	1119.06	1104.36

<sup>a</sup>Number of unserved customers.  
<sup>b</sup>Total distance traveled.  
<sup>c</sup>Total lateness.  
<sup>d</sup>Objective value.

of such events will be developed to help us in the task of determining appropriate responses.

REFERENCES

W. J. BELL, L. M. DALBERTO, M. L. FISHER, A. J. GREENFIELD, R. JAIKUMAR, P. KEDIA, R. G. MACJ, AND P. J. PRUTZMAN, "Improving the Distribution of Industrial Gases with an On-Line Computerized Routing and Scheduling Optimizer," *Interfaces* **13**, 4–23 (1983).  
 I. BENYAHIA AND J. Y. POTVIN, "Decision Support for Vehicle Dispatching Using Genetic Programming," Publication CRT-95-23, Centre de recherche sur les transports, Université de Montréal, Montréal, Canada, 1995.  
 D. BERTSIMAS AND G. VAN RYZIN, "A Stochastic and Dynamic Vehicle Routing Problem in the Euclidian Plane," *Opns. Res.* **39**, 601–615 (1991).  
 D. J. BERTSIMAS AND G. VAN RYZIN, "Stochastic and Dynamic Vehicle Routing in the Euclidian Plane with Multiple Capacitated vehicles," *Opns. Res.* **41**, 60–76 (1993).

G. G. BROWN, C. J. ELLIS, G. LENN, W. GRAVES, AND D. RONEN, "Real Time, Wide Area Dispatch of Mobil Tank Trucks," *Interfaces* **17**, 107–120 (1987).  
 C. F. DAGANZO, "An Approximate Analytic Model of Many-to-Many Demand Responsive Transportation Systems," *Transp. Res.* **12**, 325–333 (1978).  
 R. B. DIAL, "Autonomous Dial-a-Ride Transit Introductory Overview," *Transp. Res.* **3C**, 261–275 (1995).  
 M. GENDREAU, P. BADEAU, F. GUERTIN, J. Y. POTVIN, AND E. TAILLARD, "A Solution Procedure for Real-time Routing and Dispatching of Commercial Vehicles," Technical report CRT-96-24, Centre de recherche sur les transports, Université de Montréal, Montréal, Canada, 1996a.  
 M. GENDREAU, F. GUERTIN, J. Y. POTVIN, AND E. TAILLARD, "Tabu Search for Real-Time Vehicle Routing and Dispatching," Technical report CRT-96-47, Centre de recherche sur les transports, Université de Montréal, Montréal, Canada, 1996b.  
 F. GLOVER AND M. LAGUNA, *Tabu Search*, Kluwer, (1997).  
 A. HILL, V. MABERT, AND D. MONTGOMERY, "A Decision Support System for the Courier Vehicle Scheduling Problem," *Omega Int. J. Mgmt. Sci.* **16**, 333–345 (1988).

- S. LIN, "Computer Solutions of the Traveling Salesman Problem," *Bell System Tech. J.* **44**, 2245–2269 (1965).
- K. LUND, O. B. G. MADSEN, AND J. M. RYGAARD, "Vehicle Routing Problems with Varying Degrees of Dynamism," Technical report IMM-REP-1996-1, Institute of Mathematical Modeling, Technical University of Denmark, Lyngby, Denmark, 1996.
- O. B. G. MADSEN, H. F. RAVN, AND J. M. RYGAARD, "A Heuristic Algorithm for a Dial-a-Ride Problem with Time Windows, Multiple Capacities and Multiple Objectives," *Ann. Opns. Res.* **60**, 193–208 (1995).
- A. S. MINKOFF, "A Markov Decision Model and Decomposition Heuristic for Dynamic Vehicle Dispatching," *Opns. Res.* **41**, 77–90 (1993).
- W. POWELL, "A Comparative Review of Alternative Algorithms for the Dynamic Vehicle Allocation Problem," in *Vehicle Routing: Methods and Studies*, B. L. Golden and A. A. Assad (eds.), 249–291, North Holland, Amsterdam, 1988.
- W. POWELL, Y. SHEFFI, K. S. NICKERSON, K. BUTTERBAUGH, AND S. ATHERTON, "Maximizing Profits for North American Van Lines Truckload Division: A New Framework for Pricing and Operations," *Interfaces*, **18**, 21–41 (1988).
- W. B. POWELL, P. JAILLET, AND A. ODONI, "Stochastic and Dynamic Networks and Routing," in *Network Routing*, Handbooks in Operations Research and Management Science, Vol. 8, M. O. Ball, T. L. Magnanti, C. L. Monma, and G. L. Nemhauser (eds.), 141–295, North-Holland, Amsterdam, 1995.
- H. N. PSARAFTIS, "A Dynamic Programming Solution to the Single-Vehicle Many-to-Many Immediate Request Dial-a-Ride Problem," *Transp. Sci.* **14**, 130–154 (1980).
- H. N. PSARAFTIS, "An Exact Algorithm for the Single Vehicle Many-to-Many Dial-a-Ride Problem with Time Windows," *Transp. Res.* **17**, 351–357 (1983).
- H. N. PSARAFTIS, J. B. ORLIN, D. BIENSTOCK, AND P. M. THOMPSON, "Analysis and Solution Algorithms of Sealift Routing and Scheduling Problems: Final Report," Working paper No. 1700-85, Sloan School of Management, M.I.T., Cambridge, MA, 1985.
- H. N. PSARAFTIS, "Dynamic Vehicle Routing Problems," in *Vehicle Routing: Methods and Studies*, B. L. Golden and A. A. Assad (eds.), 223–249, North Holland: Amsterdam, 1988.
- H. N. PSARAFTIS, "Dynamic Vehicle Routing: Status and Prospects," *Ann. Opns. Res.* **61**, 143–164 (1995).
- A. C. REGAN, H. S. MAHMASSANI, AND P. JAILLET, "Real-Time Information for Improved Efficiency of Commercial Vehicle Operations," Proc. First World Congress on Applications of Transport Telematics & Intelligent Vehicle-Highway Systems, (1994a).
- A. C. REGAN, H. S. MAHMASSANI, AND P. JAILLET, "Improving Efficiency of Commercial Vehicle Operations Using Real-Time Information: Potential Uses and Assignment Strategies," *Transp. Res. Rec.* **1493**, 188–198 (1994b).
- A. C. REGAN, H. S. MAHMASSANI, AND P. JAILLET, "Dynamic Vehicle Allocation for Fleet Management: Operational Changes for Improved Efficiency," Proc. 2nd World Congress on Applications of Transport Telematics & Intelligent Vehicle-Highway Systems, (1995a).
- A. C. REGAN, H. S. MAHMASSANI, AND P. JAILLET, "Dynamic Decision Making for Commercial Fleet Operations Using Real-Time Information," *Transp. Res. Rec.* **1537**, 91–97 (1995b).
- R. RIVARD, "Construction des parcours des véhicules et des Horaires des Chauffeurs pour le Transport des Personnes Handicapées," Publication CRT-240, Centre de recherche sur les transports, Université de Montréal, Montréal, Canada, 1981.
- Y. ROCHAT AND E. TAILLARD, "Probabilistic Diversification and Intensification in Local Search for Vehicle Routing," *J. Heuristics*, **1**, 147–167 (1995).
- J. M. ROUSSEAU, AND S. ROY, "RAO Répartition Assistée par Ordinateur: la Description du Prototype," Publication CRT-564, Centre de recherche sur les transports, Université de Montréal, Montréal, Canada, 1988.
- S. ROY, J. M. ROUSSEAU, G. LAPALME, AND J. A. FERLAND, "Routing and Scheduling for the Transportation of Disabled Persons: the Algorithm," Publication CRT-412, Centre de recherche sur les transports, Université de Montréal, Montréal, Canada, 1985.
- Y. SHEN AND J. Y. POTVIN, "A Computer Assistant for Vehicle Dispatching with Learning Capabilities," *Ann. Opns. Res.* **61**, 189–211 (1995).
- R. S. SOLANKI, "An Execution Planning Algorithm for Military Airlift," *Interfaces*, **21**, 121–131 (1991).
- D. M. STEIN, "Scheduling Dial-a-Ride Transportation Systems," *Transp. Res.* **12**, 232–249 (1978).
- P. TRUDEAU, J. M. ROUSSEAU, J. A. FERLAND, AND J. CHOQUETTE, "An Operations Research Approach for the Planning and Operating of an Ambulance Service," *INFOR* **27**, 95–113 (1989).
- N. H. M. WILSON AND N. H. COLVIN, "Computer Control of the Rochester Dial-a-Ride System," Report R77-31, Dept. of Civil Engineering, M.I.T., Cambridge, MA, 1977.

(Received, March 1999; revisions received, January 2000; accepted, March 2000)