CIRRELT

Centre interuniversitaire de recherche sur les réseaux d'entreprise, la logistique et le transport

Interuniversity Research Centre on Enterprise Networks, Logistics and Transportation

Impact of Online Tracking on a Vehicle Routing Problem with Dynamic Travel Times

Jean Respen
Nicolas Zufferey
Jean-Yves Potvin

January 2014

CIRRELT-2014-05

Bureaux de Montréal :
Université de Montréal
Pavillon André-Aisenstadt
C.P. 6128, succursale Centre-ville
Montréal (Québec)
Canada H3C 3J7
Téléphone : 514 343-7575
Télécopie : 514 343-7121

Bureaux de Québec :
Université Laval
Pavillon Palasis-Prince
2325, de la Terrasse, bureau 2642
Québec (Québec)
Canada G1V 0A6
Téléphone : 418 656-2073
Télécopie : 418 656-2624

www.cirrelt.ca

Impact of Online Tracking on a Vehicle Routing Problem with Dynamic Travel Times

Jean Respen[1,2], Nicolas Zufferey[1,2], Jean-Yves Potvin[1,3,*]

[1] Interuniversity Research Centre on Enterprise Networks, Logistics and Transportation (CIRRELT)

[2] HEC Genève, University of Geneva, UNIMAIL, 40, Boulevard du Pont-d'Arve, 1211 Geneva 4, Switzerland

[3] Department of Computer Science and Operations Research, Université de Montréal, P.O. Box 6128, Station Centre-Ville, Montréal, Canada H3C 3J7

Abstract. This paper analyzes the impact of vehicle tracking devices, such as global positioning systems, on a vehicle routing problem with time windows in the presence of dynamic customer requests and dynamic travel times. It is empirically demonstrated that substantial improvements are achieved over a previously reported model which does not assume such tracking devices. We also analyze how the system handles dynamic perturbations to the travel times that lead to earliness or lateness in the planned schedule.

Keywords: Dynamic vehicle routing, travel times, diversion, time windows.

Acknowledgements. Financial support for this work was provided by the Natural Sciences and Engineering Research Council of Canada (NSERC). This support is gratefully acknowledged.

---

* Corresponding author: Jean-Yves.Potvin@cirrelt.ca

Dépôt légal – Bibliothèque et Archives nationales du Québec
Bibliothèque et Archives Canada, 2014

1. Introduction

Dynamic vehicle routing is attracting a growing attention in the research community. In these problems, some data are not know in advance, but are rather revealed in real-time while the routes are executed. Dynamically occurring customer requests have often been considered, but also dynamic customer demands and dynamic travel times. Since our work deals with dynamic customer requests and dynamic travel times, we provide a non exhaustive review of these variants in the following. Note that general considerations as well as exhaustive surveys on different types of dynamic vehicle routing problems can be found in [3, 13, 15].

In [4], the authors propose a parallel tabu search heuristic for a vehicle routing problem with soft time windows in the presence of dynamic customer requests. In this work, a central dispatch office manages the planned routes. Furthermore, the vehicles are not aware of their planned routes and are informed of their next destination only when they have reached their current customer location. The optimization procedure runs in background and is interrupted when a vehicle reaches a customer or when a new customer request is received. At this point, the best known solution is returned and updated, based on the new information received, and a new optimization task is launched on the updated solution. An adaptive memory is also combined with the parallel tabu search to maintain a pool of interesting solution alternatives. It is shown that this algorithm improves over simple greedy heuristics when the optimization tasks can run long enough before they are interrupted. This work was later extended in [5] to address a courier service application where each new customer request is made of a pick-up and a delivery location, with a precedence constraint between the two locations.

The impact of diversion has also been studied in the literature. It consists in diverting a vehicle to a newly occurring customer request, close to the vehicle's current location, while en route to another destination. In [7], diversion is integrated within the tabu search heuristic reported in [4], and is shown to provide substantial improvements. Diversion is also considered in [6] where two different approaches are compared. The first approach, called sample-scenario planning, provides high-quality solutions, but at the expense of large computation times. At each step, a sample of likely-to-occur future customer requests is generated to obtain a number of scenarios. Robust planned routes are then computed based on these scenarios. The second method, called anticipatory-insertion heuristic, incorporates information about expected future customer requests when each new request is inserted into the current planned routes.

As illustrated by the last method, the myopy of methodologies developed for static problems can be alleviated by exploiting any probabilistic knowledge about the occurrence of future customer requests, either implicitly or explicitly. Different approaches are based on waiting and relocation strategies. In [1], for example, the vehicles can either wait at their current customer location or at any other site, to answer customer requests that are likely to occur in their vicinity. A similar idea is also found in [9]. Here, dummy customers in the planned routes stand for future, likely-to-occur, customer requests which are replaced by true requests when they occur. Another approach reported in the literature uses a short-term and a long-term objective, where the latter tends to introduce waiting times in the planned routes to facilitate the inclusion of future requests [12].

Dynamic travel times, where times can change due to road congestion, have also raised the attention of the research community. In [2], for example, a traffic management system forecasts the travel times, based on road conditions, and transmits this information to the dispatch office. The latter then takes appropriate actions in the context of a pickup and delivery problem, assuming that the communication between the dispatch office and the drivers is possible at all time. The authors also describe a general framework to account for dynamic travel times and report results based on traffic information from the city of Berlin, Germany.

The authors in [14] consider a vehicle routing problem with time windows and dynamic travel times. The latter have three different components: static long-term forecasts (often referred to as time-dependent travel times in the literature), short-term forecasts, where the travel time on a link is modified with a random uniform value to account for any new information available when a vehicle is ready to depart from its current location, and dynamic perturbations caused by unforeseen events that might occur while traveling on a link (e.g., an accident causing sudden congestion). A modification to a planned route is only possible when the vehicle is at a customer location. Hence, a planned route cannot be reconsidered while the vehicle is traveling on a link between two customer locations. An extension to this model is proposed in [10]. In this work, the position of each vehicle can be obtained when a vehicle reaches its lateness tolerance limit or when a new customer request occurs. Based on this information, the planned route of each vehicle is reconsidered, including the possibility of diversion (i.e., redirecting a vehicle en route to its current destination). The results show that the setting of an appropriate lateness tolerance limit can provide substantial improvements. Here, we propose a further extension by assuming that the position of each vehicle is known at all time. This assumption allows the system to detect perturbations to the travel times and take appropriate actions much earlier.

This paper is organized as follows. A description of the problem is provided in Section 2. Then, Section 3 describes the two models in [10, 14] and explain the extension proposed here. Section 4 introduces travel time perturbations that lead to earliness in the planned schedule. The results obtained with the model in [10] and the new extension are then compared in Section 5. Finally, Section 6 concludes the paper and proposes future research avenues.

2. Problem description

The description of the problem is based on [10] where a fleet of vehicles performs routes, starting from and ending at a central depot, to collect goods at customer locations. Each customer must be visited exactly once by a vehicle within a (soft) time window. Some customer requests are said to be static, because they are known in advance and can be used to create initial planned routes. Other requests occur dynamically through the day and must be incorporated in real-time into the current solution. The ratio between the number of static requests and the total number of requests (static plus dynamic) is known as the degree of dynamism and is denoted d_{od} in the following [11]. Additional details on this topic can be found in [13].

Formally, let us consider a complete undirected graph $G = (V, E)$ with a set of vertices $V = \{0, 1, 2, \ldots, n\}$, where vertex 0 is the depot, and a set of edges E. Each edge $(i, j) \in E$ is characterized by a travel time t_{ij}. Also, each vertex $i \in V \setminus \{0\}$ has a time window $[e_i, l_i]$. A vehicle can arrive before the lower bound e_i but must wait to start the service. Conversely, a vehicle can arrive after the upper bound l_i, but a (linear) penalty is incurred in the objective. We assume that K vehicles of virtually infinite capacity are available. Each vehicle performs a single route which must end before an upper bound l_0, otherwise another penalty is incurred in the objective.

The objective function f takes into account (1) the travel time, (2) the sum of lateness at customer locations and (3) the lateness at the depot. Denoting t_{i^k} the arrival time of vehicle k at customer $i \in V \setminus \{0\}$ (assuming that customer i is served by vehicle k) and by t_0^k the return time of vehicle k at the depot 0, the objective can be written as:

$$\begin{aligned} f(S) &= \sum_{k\in K} f(S^k) \\ &= \sum_{k\in K}\left(\alpha\sum_{p=1}^{m_k} t_{i^k_{p-1},i^k_p} + \beta\sum_{p=1}^{m_k-1} max\{0, t_{i^k_{p-1}} - l_{i^k_{p-1}}\} + \gamma\max\{0, t_0^k - l_0\}\right) \quad (1) \end{aligned}$$

where $S = \bigcup_{k\in K} S^k$ represents a solution (a set of routes) and $S^k = \{i_0^k, i_1^k, \ldots, i_{m_k}^k\}$ is the route of vehicle $k \in K$, with $i_0^k = i_{m_k}^k = 0$. The weights α, β and γ are used to put more or less emphasis on travel time or lateness.

With regard to the static, time-dependent, component of the travel time, we do as in [8] and split the operations day in three time periods for the morning, lunch time and afternoon. With each period is associated a coefficient that multiplies the average travel time (namely, 1.25 for the morning, 0.5 for the lunch time, and 1.25 for the afternoon). To guarantee that a vehicle leaving earlier from some customer location also arrives earlier at destination, which is known as the FIFO property, the travel times are adjusted when a boundary between two time periods is crossed.

The travel times also suffer dynamic perturbations due, for example, to unexpected congestion. A dynamic perturbation is thus included based on a normal probability law with mean 0 and different standard deviations σ. Perturbations with negative values, leading to earliness, are reset to 0 in the first implementation, so that only lateness in the planned schedule can occur (as it is done in [10, 14]). In a second implementation, perturbations with negative values are also considered.

3. Models

Three related models are presented in this section. The third model is an extension of the two previous ones.

3.1. Model 1

In Model 1 [14], a central dispatch office manages the planned route of each vehicle. It is assumed that communication between the drivers and the dispatch office takes place only at customer locations. When a driver has finished serving a customer, he communicates with the dispatch office to know his next destination. Hence, the drivers are not aware of their planned route, but only of the next customer to be served. The static requests are first used to construct initial routes through an insertion heuristic where, at each iteration, a customer is selected and inserted at the best possible place in the current routes (i.e., with minimum increase in the objective value). At the end, a local search-based improvement procedure is triggered using CROSS exchanges [18], where sequences of customers are moved from one route to another. Finally, another local search-based improvement procedure is applied to each individual route, based on the relocation of each customer. Whenever a new dynamic request is received, the same insertion and reoptimization procedures are applied to update the planned routes.

Since travel times are dynamic, a lateness tolerance limit TL is defined, which is the maximum acceptable delay to a vehicle's planned arrival time at its current destination before some reassignment action is considered. For example, if we assume that s^k is the current destination of vehicle k and its planned arrival time is t_{s^k}, then $t_{s^k} + TL$ defines the tolerance time limit TTL^k of vehicle k. That is, if vehicle k has not reached customer s^k at time TTL^k, s^k is removed from its planned route and inserted in the planned route of some other vehicle l (note that vehicle k is not aware of this change and will continue toward s^k, as communication between the dispatch office and vehicles only take place at customer locations). If it happens that vehicle k still reaches s^k before vehicle l and while l is en route to s^k, then vehicle k serves s^k, but vehicle l will only know when reaching s^k. A major drawback of this model thus relates to the limited communication scheme between the dispatch office and the vehicles.

3.2. Model 2

In [10], Model 1 was extended by adding diversion to allow any vehicle to be redirected to another customer, while en route to its current destination (if it provides some benefit with regard to the objective). When (a) a new customer request is received or (b) some vehicle k has reached its tolerance time limit, it is assumed that the dispatch office can obtain the current location of each vehicle to evaluate the benefits of a diversion. In case (a), a pure diversion of vehicle k to serve the newly occurring customer request is considered whereas, in case (b), the current destination of vehicle k is reassigned to another vehicle $l \neq k$. Vehicle k is then redirected to the customer that immediately follows (what was) its current destination in the planned route. Figure 1 illustrates these two cases. In Figure 1 (a), a new customer request s occurs while vehicle k is located at position x between vertices i_{p-1}^k and i_p^k. In this case, vehicle k will serve s before i_p^k if it is beneficial to do so. In Figure 1 (b), vehicle k has reached its tolerance time limit. Thus, its current destination $s = i_p^k$ is removed from its planned route and is reassigned to another vehicle l, while vehicle k is redirected to i_{p+1}^k. The results reported in [10] show that Model 2 significantly outperforms Model 1. Also, empirical results demonstrated that the best TL value is 0. That is, an appropriate action must be considered as soon as a perturbation to the planned schedule is detected.

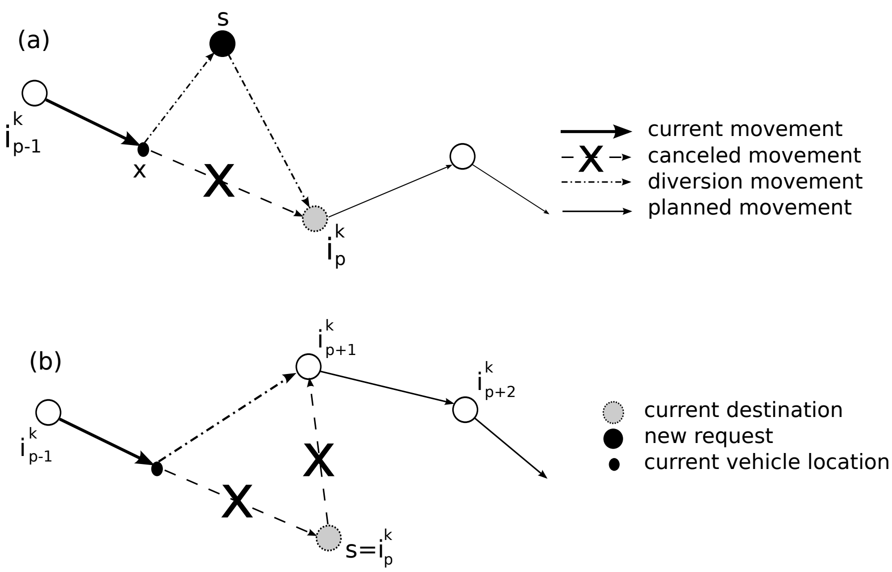

Figure 1: Diversion and reassignment actions

3.3. Model 3

The new model proposed here extends Model 2 by assuming that the position of each vehicle is known at all time, not only when the two types of situations described above for Model 2 occur. To this end, we first assume that the dynamic perturbation component of the travel time is distributed uniformly along a link. Then, as soon as it is impossible for vehicle k to arrive at its current destination at time TTL^k, a reassignment action is considered. For example, let us assume that vehicle k departs from i to j at time t, with a travel time $t_{ij} = 5$ and a dynamic perturbation $\Delta t_{ij} = 5$. That is, the vehicle is planned to arrive at j at time $t{+}5$, but will in fact arrive only at time $t{+}10$. If $TL = 0$, then $TTL^k = t{+}5$ and Model 2 will consider a reassignment at time $t + 5$ when it is observed that vehicle k has not yet reached j. On the other hand, by tracking the current position of each vehicle, Model 3 can detect the problem much earlier. For example, at time $t+1$, vehicle k has still to cover $\frac{9}{10}$ of the distance, so that the planned arrival time at location j could be updated to $t + 1 + \frac{9}{10} \cdot 5 = t + 5.5$ (assuming no more perturbation on the remainder of the link) which already exceeds TTL^k. This assumption holds if we assume the availability of tracking devices, which are now widely available at competitive prices, as well as a mobile network coverage of the service area (note that preliminary results regarding this model have been presented in [16]).

4. Earliness

As stated earlier, only positive perturbations to the travel times that lead to lateness in a vehicle schedule were considered in [10], by resetting any negative value to 0. Negative perturbation values, leading to earliness in a vehicle schedule (i.e., the vehicle will arrive earlier than expected at its current destination) have been tested here for both Models 2 and 3. If some vehicle l is late, then the earliness in the schedule of another vehicle k will automatically be exploited by the optimization procedure. That is, the current destination of vehicle l will likely be transferred to vehicle k. The benefits of Model 3 over Model 2 in this situation are the same as those mentioned for positive perturbation values: it will be possible to detect the earliness and lateness in the vehicle schedules before the vehicles reach their current destination and, consequently, react more promptly. Figure 2 illustrates this capability. In the figure, vehicle k is currently traveling between customers i^k_{p-1} and i^k_p and is ahead of its schedule. Similarly, another vehicle l is traveling between customers i^l_{p-1} and i^l_p and is late. Then, vehicle k can be redirected to i^l_p and vehicle l to i^l_{p+1} while both vehicles are en route.

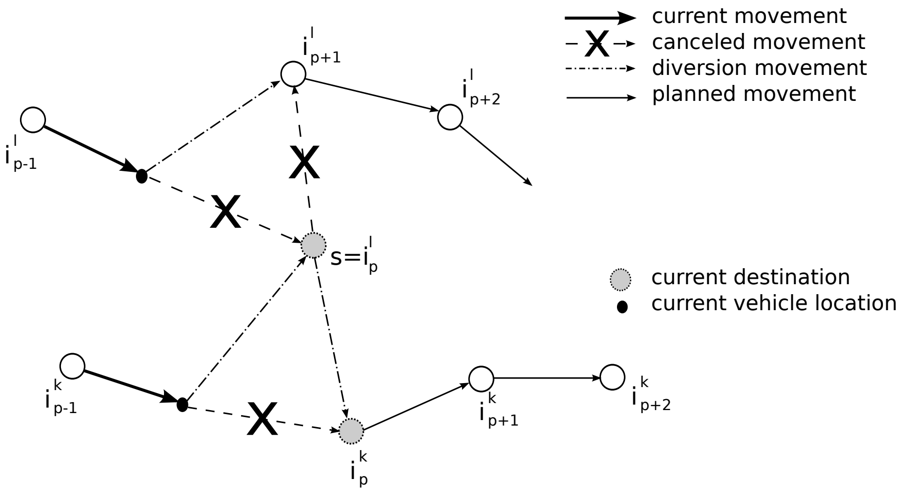

Figure 2: Integrating earliness into the model

5. Computational results

Tests were performed on a 3.4 GHz Intel Quad-core i7 with 8 GB of DDR3 RAM memory. The Euclidean 100-customer Solomon's benchmark instances [17] were used to compare Models 2 and 3. Any dynamic customer request i was set to occur at time $e_i \cdot r$, where e_i is the lower bound of the time window at customer i and r is a random number between 0 and 1. Parameters α, β and γ were set to 1 in the objective. For these experiments, only the three classes of instances $R2$, $C2$ and $RC2$ with 11, 8 and 8 instances, respectively, were considered due to their large time horizon which allows for many customers per route. Note that customers are randomly generated in $R2$, clustered in $C2$ or both clustered and randomly generated in $RC2$. Note also that the computing times are not commented given that the optimization takes place within a fraction of a second.

Tables 1 to 3 show the results on classes $R2$, $C2$, and $RC2$ respectively, for various tolerances TL and σ values in Euclidean units (where σ is the variance of the dynamic perturbations to the travel times). Each entry in these tables correspond to the average objective value over ten runs, using ten different seeds, and over each instance of a given class. There is also a pair of numbers between parentheses: the first number is the average number of times a reassignment action was considered (per instance) and the second number is the percentage of reassignments that were undertaken because they proved to be beneficial. The last row with $TL = 1000 \cdot \sigma$ is an extreme case where no action is taken. That is, the planned routes are followed whatever the perturbation. The degree of dynamism d_{od} was set to 0.5 and only lateness with regard to the current schedule was allowed (i.e., negative perturbations to the travel times were reset to 0).

Table 1: Results of Model 3 on class R2

TL	$\sigma = 1$	$\sigma = 4$	$\sigma = 16$	$\sigma = 32$
0	1548.57 (52.76,8.86%)	2099.08 (54.65,15.77%)	4779.91 (67.24,45.86%)	7322.89 (77.35,65.35%)
$0.5 \cdot \sigma$	1595.78 (31.95,8.14%)	2233.27 (32.64,13.59%)	6283.63 (35.56,39.72%)	13491.44 (36.91,56.11%)
$1 \cdot \sigma$	1623.17 (15.96,7.86%)	2288.34 (16.24,13.89%)	6641.03 (16.58,34.10%)	15201.17 (15.05,44.69%)
$2 \cdot \sigma$	1634.21 (2.05,9.73%)	2326.33 (2.04,12.50%)	6945.38 (1.75,18.75%)	15211.18 (1.33,27.40%)
$3 \cdot \sigma$	1640.09 (0.07,25.00%)	2335.04 (0.07,0.00%)	6925.79 (0.04,0.00%)	15260.43 (0.04,0.00%)
$1000 \cdot \sigma$	1641.17 (0.00,-)	2335.04 (0.00,-)	6925.79 (0.00,-)	15260.43 (0.00,-)

Table 2: Results of Model 3 on class C2

TL	$\sigma = 1$	$\sigma = 4$	$\sigma = 16$	$\sigma = 32$
0	2145.41 (51.65,6.53%)	2746.84 (51.68,7.21%)	6047.75 (53.70,12.90%)	10865.41 (58.20,26.29%)
$0.5 \cdot \sigma$	2362.23 (31.60,7.83%)	2972.81 (31.53,8.88%)	6432.07 (32.15,13.61%)	11914.13 (32.68,25.78%)
$1 \cdot \sigma$	2419.65 (15.80,8.39%)	3054.61 (15.75,9.52%)	6722.66 (15.48,15.51%)	12808.44 (12.50,28.60%)
$2 \cdot \sigma$	2696.53 (2.03,11.11%)	3311.12 (2.03,9.88%)	7119.54 (1.75,18.57%)	13263.69 (0.90,33.33%)
$3 \cdot \sigma$	2725.66 (0.10,50.00%)	3349.47 (0.10,25.00%)	7197.18 (0.10,25.00%)	13326.74 (0.00,-)
$1000 \cdot \sigma$	2728.99 (0.00,-)	3359.17 (0.00,-)	7200.55 (0.00,-)	13326.74 (0.00,-)

Table 3: Results of Model 3 on class RC2

TL	$\sigma = 1$	$\sigma = 4$	$\sigma = 16$	$\sigma = 32$
0	1515.47 (52.03,6.78%)	1878.14 (53.98,13.06%)	3975.85 (63.75,36.39%)	6205.70 (73.58,56.81%)
$0.5 \cdot \sigma$	1547.14 (31.58,6.25%)	1952.08 (32.45,11.86%)	4815.76 (34.40,30.31%)	10187.55 (34.33,47.20%)
$1 \cdot \sigma$	1572.37 (15.85,4.73%)	2006.65 (16.03,10.30%)	5313.23 (15.38,27.80%)	11370.48 (13.75,36.91%)
$2 \cdot \sigma$	1590.77 (1.98,3.80%)	2037.86 (1.98,7.59%)	5490.76 (1.88,13.33%)	11641.22 (1.25,26.00%)
$3 \cdot \sigma$	1591.66 (0.10,0.00%)	2041.68 (0.10,0.00%)	5495.05 (0.10,25.00%)	11685.73 (0.08,66.67%)
$1000 \cdot \sigma$	1591.66 (0.00,-)	2041.68 (0.00,-)	5501.21 (0.00,-)	11734.49 (0.00,-)

These results indicate that small TL values lead to better results, with the best value being $TL = 0$ in all cases. In other words, a reactive action should be considered as soon as a perturbation to the current schedule is detected. This observation is in line with the results reported in [10]. Also, the percentage of reassignments that provide an improvement increases with σ. This is not surprising, given that the current plan is likely to be improved when large perturbations are encountered.

Table 4 shows the objective values as well as the percentage of improvement of Model 3 over Model 2 when d_{od} ranges from 0.1 to 0.9 with $TL = 0$. Although we show only these results, additional experiments with other TL values led to the same observation, namely, that Model 3 is clearly superior to Model 2 due to its ability to detect perturbations to the current plan much earlier. The improvement is quite substantial in the case of $R2$ and $RC2$, and can even reach 30% for large σ values. The results are less impressive in the case of $C2$ (Model 3 is even worse than Model 2 for $d_{od} = 0.9$ and $\sigma = 16$). This observation can be explained by the geographical clustering of customers which seriously limits the benefit of redirecting a vehicle, for example to serve a distant customer in another cluster. Table 5 summarizes the improvements obtained over all d_{od} values for each class of instances, as well as over all classes of instances.

Table 4: Improvement of Model 3 over Model 2 with $TL = 0$ for different d_{od} values

d_{od}			$\sigma = 1$	$\sigma = 4$	$\sigma = 16$	$\sigma = 32$
0.1	**R2**	Model 2	1183.59	1653.41	4827.73	8606.15
		Model 3	1168.82	1598.71	4097.30	6545.20
		% Imprv.	1.26%	3.42%	17.83%	31.49%
	C2	Model 2	1195.68	1536.89	4073.00	8578.14
		Model 3	1171.77	1533.58	3920.30	8230.87
		% Imprv.	2.04%	0.22%	3.89%	4.22%
	RC2	Model 2	1267.60	1557.40	4002.04	7499.71
		Model 3	1259.96	1524.96	3512.06	5805.93
		% Imprv.	0.61%	2.13%	13.95%	29.17%
0.3	**R2**	Model 2	1381.69	1933.86	5334.02	9228.19
		Model 3	1346.87	1853.73	4520.57	7097.35
		% Imprv.	2.59%	4.32%	17.99%	30.02%
	C2	Model 2	1645.09	2014.47	5269.25	9705.95
		Model 3	1642.23	1970.29	5103.47	9466.89
		% Imprv.	0.17%	2.24%	3.25%	2.53%
	RC2	Model 2	1460.75	1857.16	4521.94	8105.48
		Model 3	1442.85	1775.30	3815.26	6338.51
		% Imprv.	1.24%	4.61%	18.52%	27.88%
0.5	**R2**	Model 2	1604.03	2245.95	5891.80	9636.36
		Model 3	1548.57	2099.08	4779.92	7322.89
		% Imprv.	3.58%	7.00%	23.26%	31.59%
	C2	Model 2	2193.26	2769.02	6221.36	11652.32
		Model 3	2145.41	2746.84	6047.76	10865.42
		% Imprv.	2.23%	0.81%	2.87%	7.24%
	RC2	Model 2	1551.26	1949.16	4621.90	8426.49
		Model 3	1515.47	1878.14	3975.86	6205.70
		% Imprv.	2.36%	3.78%	16.25%	35.79%
0.7	**R2**	Model 2	2014.79	2823.41	6553.96	10366.10
		Model 3	1932.65	2599.74	5292.18	8037.42
		% Imprv.	4.25%	8.60%	23.84%	28.97%
	C2	Model 2	2444.12	2806.64	6075.49	11937.20
		Model 3	2340.94	2790.25	5809.04	10929.91
		% Imprv.	4.41%	0.59%	4.59%	9.22%
	RC2	Model 2	1805.32	2343.49	5311.32	9153.11
		Model 3	1747.96	2185.97	4708.65	6948.25
		% Imprv.	3.28%	7.21%	12.80%	31.73%
0.9	**R2**	Model 2	2182.42	2994.81	6961.94	11386.24
		Model 3	2044.20	2769.36	5728.86	8902.86
		% Imprv.	6.76%	8.14%	21.52%	27.89%
	C2	Model 2	3498.62	4337.93	8702.07	15191.83
		Model 3	3414.02	4216.22	8941.81	14441.63
		% Imprv.	2.48%	2.89%	-2.68%	5.19%
	RC2	Model 2	2097.39	2782.93	6743.54	10730.58
		Model 3	2008.42	2554.65	5759.39	8285.90
		% Imprv.	4.43%	8.94%	17.09%	29.50%

Table 5: Improvement of Model 3 over Model 2 with $TL = 0$ over all d_{od} values

	$\sigma = 1$	$\sigma = 4$	$\sigma = 16$	$\sigma = 32$
$R2$	3.69%	6.30%	20.89%	29.99%
$C2$	2.27%	1.35%	2.38%	5.68%
$RC2$	2.38%	5.33%	15.72%	30.81%
Overall	2.78%	4.33%	13.00%	22.16%

Tables 6 to 8 are similar to Tables 1 to 3 and report the results of Model 3 for various tolerance TL and σ values with $d_{od} = 0.5$ when negative perturbations to the travel times are allowed (leading to earliness in the schedule). Tables 9 and 10 are also similar to Tables 4 and 5 and report the improvement of Model 3 over Model 2 with $TL = 0$ for d_{od} values between 0.1 and 0.9 when negative perturbations are allowed. Not surprisingly, the trends are the same as those observed previously but are somewhat accentuated, in particular for large σ values.

Table 6: Results of Model 3 on class R2 including negative perturbations

TL	$\sigma = 1$	$\sigma = 4$	$\sigma = 16$	$\sigma = 32$
0	1408.23 (53.13,8.80%)	1576.70 (54.82,14.23%)	2951.06 (66.20,42.85%)	4156.69 (76.67,62.81%)
$0.5 \cdot \sigma$	1451.92 (32.22,7.62%)	1675.40 (32.73,12.89%)	4164.83 (35.77,38.81%)	9962.26 (37.57,55.17%)
$1 \cdot \sigma$	1474.64 (16.20,7.58%)	1719.48 (16.35,14.17%)	4840.45 (16.71,34.33%)	12027.22 (15.36,45.03%)
$2 \cdot \sigma$	1488.51 (2.12,8.58%)	1753.46 (2.10,10.82%)	5176.51 (1.90,25.36%)	12765.64 (1.39,40.52%)
$3 \cdot \sigma$	1490.47 (0.13,14.29%)	1761.16 (0.12,15.38%)	5160.00 (0.09,10.00%)	12485.61 (0.07,12.50%)
$1000 \cdot \sigma$	1491.05 (0.00,-)	1762.37 (0.00,-)	5161.13 (0.00,-)	12489.44 (0.00,-)

Table 7: Results of Model 3 on class C2 including negative perturbations

TL	$\sigma = 1$	$\sigma = 4$	$\sigma = 16$	$\sigma = 32$
0	2208.83 (52.21,6.22%)	2480.35 (52.43,7.20%)	4934.58 (54.14,12.65%)	8706.28 (57.91,24.15%)
$0.5 \cdot \sigma$	2349.66 (31.93,7.01%)	2593.29 (32.06,8.38%)	5153.91 (32.63,13.26%)	9852.66 (32.64,24.01%)
$1 \cdot \sigma$	2433.25 (15.90,8.18%)	2697.11 (15.95,9.01%)	5550.78 (15.54,14.88%)	10736.38 (12.54,26.32%)
$2 \cdot \sigma$	2592.38 (2.09,10.18%)	2877.84 (2.08,12.65%)	5923.32 (1.73,21.01%)	11454.67 (0.89,29.58%)
$3 \cdot \sigma$	2639.34 (0.14,27.27%)	2932.43 (0.14,9.09%)	5999.76 (0.11,22.22%)	11449.34 (0.04,0.00%)
$1000 \cdot \sigma$	2638.97 (0.00,-)	2933.91 (0.00,-)	6018.05 (0.00,-)	11449.34 (0.00,-)

Table 8: Results of Model 3 on class RC2 including negative perturbations

TL	$\sigma = 1$	$\sigma = 4$	$\sigma = 16$	$\sigma = 32$
0	1444.06 (53.05,7.28%)	1570.02 (54.51,12.27%)	2602.30 (65.19,37.03%)	3583.52 (74.21,55.47%)
$0.5 \cdot \sigma$	1472.57 (32.01,6.44%)	1630.52 (32.78,11.94%)	3387.50 (34.94,30.88%)	7333.01 (34.93,48.21%)
$1 \cdot \sigma$	1489.27 (15.96,6.50%)	1669.12 (16.15,12.00%)	3931.03 (15.61,29.78%)	8636.15 (13.38,39.81%)
$2 \cdot \sigma$	1503.83 (2.06,8.48%)	1712.14 (2.10,16.07%)	4122.88 (1.83,24.66%)	9111.62 (1.39,28.83%)
$3 \cdot \sigma$	1508.63 (0.14,0.00%)	1722.23 (0.14,0.00%)	4150.32 (0.14,18.18%)	9185.94 (0.05,25.00%)
$1000 \cdot \sigma$	1508.63 (0.00,-)	1722.23 (0.00,-)	4145.58 (0.00,-)	9184.10 (0.00,-)

Table 9: Improvement of Model 3 over Model 2 with $TL = 0$ for different d_{od} values, including negative perturbations

d_{od}			$\sigma = 1$	$\sigma = 4$	$\sigma = 16$	$\sigma = 32$
	R2	Model 2	1093.38	1238.16	3091.47	5400.52
0.1		Model 3	1082.36	1208.02	2531.94	3723.15
		% Imprv.	1.02%	2.50%	22.10%	45.05%
	C2	Model 2	1180.54	1361.29	3160.62	6975.13
		Model 3	1167.54	1343.24	3095.66	6527.79
		% Imprv.	1.11%	1.34%	2.10%	6.85%
	RC2	Model 2	1205.74	1322.11	2749.25	4765.88
		Model 3	1199.42	1300.70	2298.89	3178.00
		% Imprv.	0.53%	1.65%	19.59%	49.96%
	R2	Model 2	1255.76	1445.00	3417.23	5638.62
0.3		Model 3	1229.63	1377.47	2679.75	3874.91
		% Imprv.	2.12%	4.90%	27.52%	45.52%
	C2	Model 2	1626.13	1767.72	3902.18	7793.65
		Model 3	1607.29	1735.47	3651.05	7159.94
		% Imprv.	1.17%	1.86%	6.88%	8.85%
	RC2	Model 2	1372.78	1541.12	3208.57	5321.58
		Model 3	1356.16	1481.81	2532.72	3579.10
		% Imprv.	1.23%	4.00%	26.68%	48.68%
	R2	Model 2	1454.72	1668.62	3741.13	5996.84
0.5		Model 3	1408.23	1576.70	2951.06	4156.69
		% Imprv.	3.30%	5.83%	26.77%	44.27%
	C2	Model 2	2269.21	2496.09	5110.58	9508.50
		Model 3	2208.83	2480.35	4934.58	8706.28
		% Imprv.	2.73%	0.63%	3.57%	9.21%
	RC2	Model 2	1478.08	1634.76	3338.04	5278.44
		Model 3	1444.06	1570.02	2602.30	3583.52
		% Imprv.	2.36%	4.12%	28.27%	47.30%
	R2	Model 2	1803.15	2053.27	4159.67	6545.16
0.7		Model 3	1728.48	1909.71	3264.42	4409.19
		% Imprv.	4.32%	7.52%	27.42%	48.44%
	C2	Model 2	2310.80	2422.14	4852.43	9672.17
		Model 3	2217.70	2336.18	4573.30	8745.50
		% Imprv.	4.20%	3.68%	6.10%	10.60%
	RC2	Model 2	1669.42	1853.11	3614.90	5818.96
		Model 3	1625.96	1747.68	2805.90	3807.39
		% Imprv.	2.67%	6.03%	28.83%	52.83%
	R2	Model 2	1965.55	2283.00	4631.17	7386.65
0.9		Model 3	1861.09	2099.15	3610.27	4990.11
		% Imprv.	5.61%	8.76%	28.28%	48.03%
	C2	Model 2	3323.86	3621.42	6403.54	12274.79
		Model 3	3261.67	3446.48	6610.18	11123.78
		% Imprv.	1.91%	5.08%	-3.13%	10.35%
	RC2	Model 2	2112.12	2381.04	5101.59	7644.64
		Model 3	2016.67	2195.41	3943.92	5532.53
		% Imprv.	4.73%	8.46%	29.35%	38.18%

Table 10: Improvement of Model 3 over Model 2 with $TL = 0$ over all d_{od} values, including negative perturbations

	$\sigma = 1$	$\sigma = 4$	$\sigma = 16$	$\sigma = 32$
$R2$	3.28%	5.90%	26.42%	46.26%
$C2$	2.22%	2.52%	3.10%	9.17%
$RC2$	2.30%	4.85%	26.55%	47.39%
Overall	2.60%	4.42%	18.69%	34.28%

Finally, Figure 3 illustrates the average objective values of Models 2 and 3 with $TL = 0$ and $d_{od} = 0.5$ for a large number of σ values taken between 1 and 32 using the instances of class $RC2$. In this figure, negative perturbations to the travel times are allowed. This figure clearly shows that the gap between the two models sharply increases with σ.

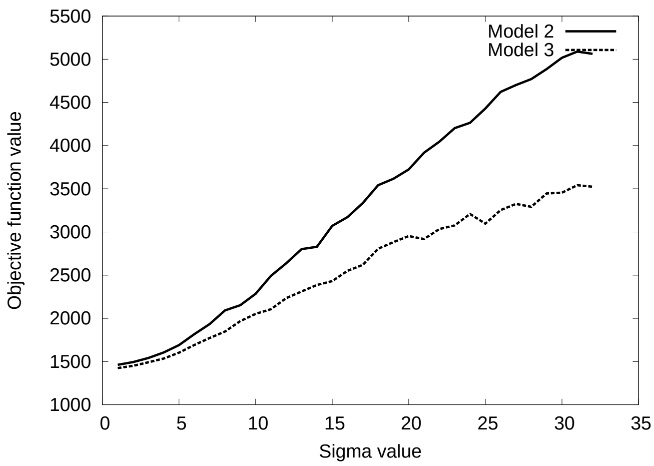

Figure 3: Average objective values of Models 2 and 3 on class $RC2$ with $TL = 0$, $d_{od} = 0.5$ and different σ values

6. Conclusion

This paper has investigated the impact of on-line vehicle tracking devices on solution quality for a dynamic vehicle routing problem with time windows. The dynamic characteristics relate to the occurrence of new customer requests and perturbations to the travel times. We empirically demonstrated that a reactive action should be contemplated as soon as a perturbation is detected in the current planned routes. Our model was also shown to be robust and to behave well under different degrees of dynamism. It also proved to be significantly superior to another model where the location of each vehicle is only known at specific moments during the operations day. In the future, we plan to investigate the impact of other probability distributions to model the travel time perturbations. We also want to consider other ways to distribute a perturbation to the travel time along a link.

References

[1] Bent, R., & Van Hentenryck, P. 2007. Waiting and Relocation Strategies in Online Stochastic Vehicle Routing. *Pages 1816–1821 of: International Joint Conference on Artificial Intelligence*.

[2] Fleischmann, B., Gnutzmann, S., & Sandvoß, E. 2004. Dynamic Vehicle Routing based on Online Traffic Information. *Transportation Science*, **38**, 420–433.

[3] Gendreau, M., & Potvin, J.-Y. 1998. Dynamic Vehicle Routing and Dispatching. *Pages 115 – 126 of:* Crainic, T.G., & Laporte, G. (eds), *Fleet Management and Logistics*. Kluwer.

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

[5] Gendreau, M., Guertin, F., Potvin, J.-Y., & Séguin, R. 2006. Neighborhood Search Heuristics for a Dynamic Vehicle Dispatching Problem with Pickups and Deliveries. *Transportation Research Part C: Emerging Technologies*, **14**, 157–174.

[6] Ghiani, G., Manni, E., & Thomas, B. W. 2012. A Comparison of Anticipatory Algorithms for the Dynamic and Stochastic Traveling Salesman Problem. *Transportation Science*, **46**, 374 – 387.

[7] Ichoua, S., Gendreau, M., & Potvin, J.-Y. 2000. Diversion Issues in Real-Time Vehicle Dispatching. *Transportation Science*, **34**, 426 – 438.

[8] Ichoua, S., Gendreau, M., & Potvin, J.-Y. 2003. Vehicle Dispatching with Time-Dependent Travel Times. *European Journal of Operational Research*, **144**, 379 – 396.

[9] Ichoua, S., Gendreau, M., & Potvin, J.-Y. 2006. Exploiting Knowledge about Future Demands for Real-Time Vehicle Dispatching. *Transportation Science*, **40**, 211–225.

[10] Lorini, S., Potvin, J.-Y., & Zufferey, N. 2011. Online Vehicle Routing and Scheduling with Dynamic Travel Times. *Computers & Operations Research*, **38**, 1086–1090.

[11] Lund, K., Madsen, O.B.G., & Rygaard, J.M. 1996. *Vehicle Routing Problems with Varying Degrees of Dynamism*. IMM-REP. Technical report.

[12] Mitrović-Minić, S., Krishnamurti, R., & Laporte, G. 2004. Double-Horizon Based Heuristics for the Dynamic Pickup and Delivery Problem with Time Windows. *Transportation Research Part B: Methodological*, **38**, 669–685.

[13] Pillac, V., Gendreau, M., Guéret, C., & Medaglia, A. L. 2013. A Review of Dynamic Vehicle Routing Problems. *European Journal of Operational Research*, **225**, 1 – 11.

[14] Potvin, J.-Y., Xu, Y., & Benyahia, I. 2006. Vehicle Routing and Scheduling with Dynamic Travel Times. *Computers & Operations Research*, **33**, 1129 – 1137.

[15] Psaraftis, H.N. 1995. Dynamic Vehicle Routing: Status and Prospects. *Annals of Operations Research*, **61**, 143 – 164.

[16] Respen, J., Zufferey, N., & Potvin, J.-Y. 2014. Online Vehicle Routing and Scheduling with Continuous Vehicle Tracking. *In: Proceedings of the ROADEF 2014, Bordeaux, France. February 26 to 28.*

[17] Solomon, M.M. 1987. Algorithms for the Vehicle Routing and Scheduling Problem with Time Window Constraints. *Operations Research*, **35**, 254–265.

[18] Taillard, É., Badeau, P., Gendreau, M., Guertin, F., & Potvin, J.-Y. 1997. A Tabu Search Heuristic for the Vehicle Routing Problem with Soft Time Windows. *Transportation Science*, **31**, 170 – 186.