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

Interuniversity Research Centre on Enterprise Networks, Logistics and Transportation

Multicriteria Optimization of a Long-Haul Routing and Scheduling Problem

Marie-Ève Rancourt
Julie Paquette

July 2012

CIRRELT-2012-32

Multicriteria Optimization of a Long-Haul Routing and Scheduling Problem

Marie-Ève Rancourt[1,2], Julie Paquette[1,3,*]

[1] Interuniversity Research Centre on Enterprise Networks, Logistics and Transportation (CIRRELT)
[2] Department of Management Sciences, HEC Montréal, 3000 Côte-Sainte-Catherine, Montréal, Canada H3T 2A7
[3] Department of Logistics and Operations Management, HEC Montréal, 3000 Côte-Sainte-Catherine, Montréal, Canada H3T 2A7

Abstract. Long-haul carriers are facing a shortage of drivers in North American countries. To reduce turnover rates and improve driver retention, trucking companies are making more efforts to improve their drivers' quality of life. The aim of this paper is to introduce and solve a multi-objective vehicle routing and truck driver scheduling problem under the legislative requirements on work and rest hours in the United States (US MOVRTDSP). We present a tabu search algorithm that solves the US MOVRTDSP and provides heuristic non-dominated solution sets from which tradeoffs between operating costs and driver inconvenience are evaluated. Results show that decreasing driver inconvenience increases carrier operating costs in a smaller proportion. The tradeoffs between the number of vehicles used and operating costs are also estimated. This information is used to evaluate the efficacy of the rule of thumb generally applied in the industry of focusing on minimizing the number of vehicles sent on the road. We show that for most cases, operating costs increase when the number of vehicles exceeds the minimum fleet size. However, for a real-life instance which is an open routing problem, the rule of thumb is not necessarily cost efficient. Overall, interpretations of the computational results on artificial and real-life instances provide meaningful information to long-haul carriers.

Keywords. Multicriteria analysis, vehicle routing, truck driver scheduling, driver working conditions, HOS regulations, tabu search heuristic.

Acknowledgements. This work was partly funded by HEC Montréal. This support is gratefully acknowledged. We also thank Groupe Robert for their kind cooperation.

---

* Corresponding author: Julie.Paquette@cirrelt.ca

Dépôt légal – Bibliothèque et Archives nationales du Québec
Bibliothèque et Archives Canada, 2012

1 Introduction

For many years, North American countries have been facing a shortage of drivers for long-haul transportation services which imply long distances traveled over several days by drivers (Johnston et al. 2009; Doss 2012). In that context, trucking companies increasingly seek to improve driver working conditions as a way to reduce the turnover rates and enhance the retention of this labor. For example, carriers now make greater efforts to adjust schedules and routes so that drivers won't be away from their families for too many days in a row. Trucking companies are thus pursuing two different objectives simultaneously: minimizing operating costs and minimizing driver inconvenience. Moreover, to ensure road safety, different governments around the world have established working hour rules to be taken into account when planning vehicle routing and scheduling. These rules typically stipulate that truck drivers should not work for more than a certain number of hours during successive days without taking a rest. This stands in contrast to classical routing problems where vehicles leave and come back to the depot on the same day.

In the trucking industry, long-haul routes are valued by the distance and number of clients visited. The variable costs for a fixed number of customers, referred to as "operating costs" in this project, are thus proportional to the distance traveled by the trucks. Because drivers are not remunerated for resting time or waiting time, these elements account for "driver inconvenience" in this project. In fact, the larger the number of resting hours on the road, the more time a driver is away from home. We also consider waiting time to be an inconvenience given that it is idle time away from home in addition to being unpaid. According to a study by Johnston et al. (2009), time away from home is the most disagreeable aspect of being a long-distance truck driver, which is consistent with findings derived from discussions with experts from the trucking industry.

Another commonly pursued goal of trucking companies is to minimize the number of vehicles used to provide the service. Generally, dispatchers adopt the rule of thumb of using as few vehicles as possible in order to minimize operating costs while seeking to maintain the capacity to provide service in the case of last-minute demands. In this study, rather than including a fixed cost associated with the number of vehicles used in the objective function, we will sequentially solve the problem using the minimum number of vehicles and then adding one, two and three vehicles to the minimum fleet.

The aim of this paper is to present and solve a multicriteria vehicle routing problem that has emerged in the less-than-truckload industry by taking into account the legislative requirements on the hours of service in the United States (US). This problem is referred to as the US multi-objective vehicle routing and truck driver scheduling problem (US MOVRTDSP). The contribution of this paper is threefold. First, we solve the problem by means of a scheduling algorithm and a multicriteria scheme, both embedded within a tabu search heuristic. Our study, although generally applicable for long-haul transportation in the United States, was first motivated by the case of the US Operations Division of Groupe Robert Inc., one of the largest and most widely known third-party logistics providers in Canada. In addition to the US regulations and other constraints sometimes considered in vehicle routing problems (VRPs), such as multiple time windows and a heterogeneous vehicle fleet, our study takes into account a bi-objective function representing carrier operating costs and driver inconvenience. Second, this heuristic helps to enhance the understanding of the tradeoffs between the carrier costs and the driver inconvenience through the analysis of non-dominated solution sets. Drivers are an important asset for carriers. Improving their working conditions can therefore reduce their turnover rate, thus enhancing the level of service and reducing training costs. Third, our results also improve the understanding of the tradeoffs between the number of vehicles used and the operating costs, which allows assessing the efficacy of the above-mentioned rule of thumb employed by dispatchers. Overall, interpretations of the computational results on artificial instances and on a real-life instance can provide meaningful information to carriers operating transportation services in the United States and elsewhere.

The remainder of this paper is organized as follows. Section 2 presents a literature review on long-haul vehicle routing and scheduling problems and on the different algorithms used to solve multicriteria combinatorial problems. Section 3 presents the multicriteria objective function and the constraints used to model the problem. The tabu search heuristic developed to solve this US MOVRTDSP is then described in Section 4. Computational results are presented and analyzed in Section 5, followed by the managerial implications in Section 6. Finally, conclusions are summarized in Section 7.

2 Literature review

In long-haul transportation worldwide, truck driver fatigue is recognized to be an important factor of accidents (see Williamson et al. 2001; European Transport Safety Council 2001; Federal Motor Carrier Safety Administration 2008, as sample studies for Australia, Europe and the United States). Governments are making efforts to improve driver safety and working conditions by adopting hours of service legislation in road freight transport. These regulations typically impose restrictions on the length of driving and working times during which drivers are permitted to drive without taking a rest or a break over given periods. Trucking companies must take into account the rules that apply to their regions of operation when constructing routes and planning truck driver schedules. Even though routing problems have attracted a lot of attention in the literature, it is only until recently that working hour regulations in road freight transport are being taken into consideration in the vehicle routing and scheduling problems. Archetti and Savelsbergh (2009), Goel and Kok (2012) and Goel and Vidal (2012) proposed models and exact methods for finding a feasible schedule in compliance with US regulations for a fixed sequence of customers with time windows. Goel and Kok (2012), Goel and Rousseau (2012) and Goel et al. (2012) have studied this problem in the context of the European, Canadian and Australian regulations, respectively. The problem that combines this type of scheduling component with a routing component is called the vehicle routing and truck driver scheduling problem (VRTDSP). The scheduling component consists of the planning of the driver rest periods and the customer service times, and the routing component consists of determining the sequence of customers visited by each vehicle. Solving the VRTDSP thus involves determining a set of routes that minimize operating costs for a fleet of vehicles subject to limited capacities such that each customer is visited within given time windows and that each truck driver schedule complies with the applicable working hour regulations.

Heuristic approaches for solving the VRTDSP have been presented by Goel (2009), Kok et al. (2010) and Prescott-Gagnon et al. (2010) with regard to EU regulations and by Rancourt et al. (2012) with regard to US regulations. Other variants of the VRTDSP have been addressed by Xu et al. (2003), Zäpfel and Bögl (2008), Ceselli et al. (2009), Bartodziej et al. (2009) and Kok et al. (2011). Goel and Vidal (2012) were the first to propose an approach that is not restricted to a particular set of rules. They conducted extensive experiments to assess the impact of the different regulations in the United States, Canada, the European

Union and Australia. With regard to EU regulations, their method has proven to outperform all other proposed methods for tackling the VRTDSP. They also show that the risk levels of the different countries' regulations correlate negatively with the associated operating costs, except for the Australian regulations, which rank highest in terms of risk level and operating costs. The European Union legislation provides for the highest level of safety, while the Canadian regulations are the most competitive in terms of costs.

Table 1 summarizes the papers published on the VRTDSP. We can observe that all heuristics developed to solve the VRTDSP are single-criterion, with the primary objective being to minimize the number of vehicles used, and the secondary objective typically being to minimize operating costs, which are generally based on the total distance traveled or the total duration of routes. Rancourt et al. (2012) have considered alternative secondary objectives in conjunction with the primary objective of minimizing the number of vehicles. To do so, they have used a weighted objective function and applied a lexicographic process. However, to our knowledge, no method has been proposed to solve the VRTDSP as a multicriteria optimization problem that yields a non-dominated solution set.

Table 1: Summary of the references related to the VRTDSP

Article	Regulations				Objectives		
	EU	US	CAN	AUS	Distance	Duration	Cost
Xu et al. (2003)		√					√
Ceselli et al. (2009)		√[a]					√
Goel (2009)	√				√		
Kok et al. (2010)	√					√	
Prescott-Gagnon et al. (2010)	√				√		
Kok et al. (2011)	√				√		
Rancourt et al. (2012)		√			√	√	√
Goel and Vidal (2012)	√	√	√	√	√		

[a] The legislation considered in this study is not specified, but the applied rules are very similar to the US regulations

Decision-makers generally pursue more than one objective when making a decision. Subsequently, when modeling practical problems encountered in real-life applications, considering several criteria in the optimization process has become common practice. The approaches to solving these problems can be divided into three categories: (1) a priori decision-making – when the decision-maker's preferences are known before the optimization phase starts; (2) a posteriori decision-making – when the decision-maker provides his/her preferences after the optimization phase by interpreting the set of resulting solutions; (3) dynamic decision-making – when the decision-maker interacts iteratively during the optimization phase, namely by reacting to the solutions presented to him/her, to provide more information about his/her preferences.

The first category of multicriteria problem-solving is very popular. Indeed, many multicriteria studies are using weighted functions of preferences as an objective function to drive their algorithm. This enables the resolution of a single-criterion problem after scalarization, which in turn enables the use of single-criterion algorithms that already exist. A field of research on multicriteria optimization has emerged from the development of specific methods to solve multicriteria problems. The most common approaches are the ϵ-constraint method, goal programming and ranking methods. The ϵ-constraint method consists of transforming a problem's objective functions into constraints as a way to reduce their number. Here, information on the acceptable level of the different objective functions must be known in advance. Goal programming, for its part, aims to minimize the distance to an ideal point. The ideal point values must be provided a priori by the decision-maker, as well as the weights used to put more or less emphasis on the different objectives. Finally, ranking methods, also called lexicographic methods, sequentially optimize the objective functions in the order determined by the decision-maker. Some information on the decision-maker's preferences is thus required in advance. These methods usually apply a priori known information and yield a unique optimal solution.

Other research to solve multicriteria problems has focused on extending single-criterion methods to the multicriteria context. These methods can be exact such as dynamic programming or branch-and-bound, or metaheuristics such as tabu search (Gandibleux et al. 1997) and simulated annealing (Alves and Clímaco 2000), which provide an approximation of the set of solutions. Moreover, evolutionary algorithms (Bäck 1996), which are global search optimization algorithms based on the mechanics of natural selection and reproduction, have proven to be very effective in solving complex multicriteria optimization problems (Cvetkovic and Parmee 2002; Jaszkiewicz 2003; Refanidis and Vlahavas 2003). These methods usually provide the decision-maker with a set of non-dominated solutions. Further details on the methods for solving multicriteria combinatorial optimization problems can be found in Hansen (2000), Coello Coello et al. (2002), and Ehrgott and Gandibleux (2002).

The specific case of multicriteria routing problems has been discussed and reviewed in Boffey (1996) and Jozefowiez et al. (2008). The resolution methods mentioned earlier are similar to those used for this specific group of problems. Most research is concentrated on the bi-criteria traveling salesman problem. Some research has also been done on different vehicle routing problems (Park and Koelling 1986; Jozefowiez et al. 2002; Murata and Itai 2005; Ombuki et al. 2006; Tan et al. 2006; Murata and Itai 2007; Jozefowiez et al. 2009; Parragh et al. 2009; Repoussis et al. 2009; Garcia-Najera and Bullinaria 2011; Paquette et al. 2012). Table 2 summarizes the problem, objectives, groups of constraints and type of algorithm presented for each article.

From this table we can see that the most common objectives in multicriteria vehicle routing problems are to (1) minimize the number of vehicles, (2) minimize the total cost (or distance, or time) and (3) equilibrate the length of routes among the vehicles. The first two objectives are also those that are most commonly used to solve the VRTDSP in a lexicographic manner. Moreover, different algorithms are used to optimize multicriteria vehicle routing problems such as evolutionary algorithms, variable neighborhood search and tabu search heuristics.

This paper is the first study that proposes a methodology that exploits the strengths of multicriteria optimization to generate a set of non-dominated solutions for a variant of the VRTDSP. Our results will allow to measure the tradeoffs between the operating costs and the driver inconvenience and to assess the impact of using more vehicles than the minimum fleet size.

3 Problem description

The problem considered in this paper is a bi-objective VRTDSP in compliance with the current hours of service (HOS) regulations for commercial vehicle drivers (Part 395 of the US Department of Transportation) as well as other VRP constraints. In this section, we first present the HOS regulations, followed by a description of the ensuing VRTDSP and its constraints and objectives.

Table 2: Summary of the main studies on multi-objective vehicle routing problems

Article	Problem	Objectives	Algorithm
Park and Koelling (1986)	CVRP	(1) Minimize the travel distance of vehicles, (2) minimize the total deterioration of goods, and (3) maximize the total fulfillment of emergent services and compliance with conditional dependencies of stations.	Two-phase heuristic: (1) clustering heuristic and (2) goal programming (decision-maker provides a priority structure)
Jozefowiez et al. (2002)	CVRP	(1) Minimize total distance, and (2) minimize the difference between maximum route length and minimum route length.	Evolutionary algorithm with parallel and hybrid models for intensification and diversification tasks.
Murata and Itai (2005)	VRP	(1) Minimize the number of vehicles, and (2) minimize the maximum routing time among vehicles.	Evolutionary algorithm.
Ombuki et al. (2006)	VRPTW	(1) Minimize the number of vehicles, and (2) minimize total cost (distance).	Genetic algorithm.
Tan et al. (2006)	TTVRP	(1) Minimize routing distance, and (2) minimize the number of trucks.	Hybrid multi-objective evolutionary algorithm.
Murata and Itai (2007)	VRP	(1) Minimize the maximum routing time among vehicles, (2) minimize the number of vehicles, and (3) maximize the similarity of solutions.	Memetic evolutionary algorithm.
Jozefowiez et al. (2009)	VRPRB	(1) Minimize total length of routes, and (2) minimize the difference between maximum route length and minimum route length.	Evolutionary algorithm.
Parragh et al. (2009)	DARP	(1) Minimize cost, and (2) minimize mean user ride time.	Variable neighborhood search algorithma and path relinking technique.
Repoussis et al. (2009)	VRPTW	(1) Minimize the number of routes (vehicles), and (2) minimize the total travel time.	Evolutionary algorithm combined with tabu search heuristic and lexicographic evaluation in a first phase; with a guided local search in a second phase.
Garcia-Najera and Bullinaria (2011)	VRPTW	(1) Minimize the number of routes (vehicles), (2) minimize total travel distance, and (3) minimize total travel time.	Evolutionary algorithm.
Paquette et al. (2012)	DARP	(1) Minimize cost, (2) minimize waiting time within the destination time window for an outbound trip, (3) minimize waiting time within the origin time window for an inbound trip, and (4) minimize ratio of actual ride time over direct ride time.	Tabu search heuristic and reference point method.

3.1 Hours of service regulations

In the United States, the Federal Motor Carrier Safety Agency is responsible of publishing the HOS regulations. These regulations mainly restrict the amount of cumulated driving and working time between rest periods. The regulation distinguishes between on-duty time and off-duty time. On-duty time consists of the time when a driver is performing work or is required to be available for work. On-duty time includes driving activities as well as any other type of work such as loading and unloading. In this project, we adopt a conservative definition of idle time in that we consider waiting time to be on-duty time, mainly because we assume that drivers may have to remain in the vehicle during these waiting periods. Off-duty time refers to time during which drivers have no obligation to perform work, in other words, time during which they are free to pursue any activity and leave the vehicles where they are parked. For example, resting time is considered to be off-duty time.

The HOS regulations essentially impose limits on the duration of times during which drivers of commercial motor vehicles can drive or be on-duty without taking off-duty periods of prescribed durations. Table 3 describes the three fundamental rules: the 60/70-hour on-duty limit, the 11-hour driving limit and the 14-hour limit. The first rule prescribes a limit on on-duty time during a defined number of consecutive days. In our study, we use the 70-hour on-duty limit to be consistent with Groupe Robert's practice. The second rule limits to 11 hours the driving time after an off-duty period of at least 10 hours, which means that a driver must take an off-duty period of at least 10 hours to regain the right to drive after cumulating 11 hours of driving. The third rule forbids driving after an on-duty period of 14 hours that begins right after an off-duty period of 10 hours. However, a driver may perform work other than driving beyond this 14-hour limit. We point out that an off-duty period of at least 10 consecutive hours but less than 34 hours will be considered a rest period in the remainder of this paper. A driver can also use the sleeper berth provision to get the equivalent of 10 consecutive off-duty hours with two separate split rest periods of at least eight and two hours. However, neither this special provision, nor a new provision to come into effect starting July 2013, will be taken into consideration in this paper. The interested reader can obtain more information on HOS regulations from the U.S. Department of Transportation (2012) website.

Table 3: Hours of service regulations

60/70-hour on-duty limit	A driver may not drive after 60/70 hours on-duty in 7/8 consecutive days. He/she may again start a 7/8 consecutive day period after 34 or more consecutive off-duty hours.
11-hour driving limit	A driver may only drive a maximum of 11 hours after 10 consecutive off-duty hours.
14-hour limit	A driver may not drive beyond the 14^{th} consecutive hour after coming on-duty following a period of 10 consecutive hours off-duty. Off-duty time does not delay the 14-hour period.

In summary, a driver may drive at most h^{drive} hours and may be on-duty at most h^{on_duty} hours before taking a rest of at least h^{rest} hours to regain the right to drive. In other words, a rest period must be taken at the latest when the driving limit or the on-duty limit is reached. Duty time consists primarily of driving, waiting and service time (loading and unloading activities). Moreover, a driver may be on-duty for at most h^{total_duty} during a horizon of H consecutive days. We can thus set the following parameters:

H = eight days, the number of consecutive days during a planning horizon;

$h^{total\_duty} = 70$ hours, the maximum cumulated on-duty time during the planning horizon H;

$h^{drive} = 11$ hours, the maximum cumulated driving hours between two prescribed off-duty periods;

$h^{on\_duty} = 14$ hours, the maximum cumulated on-duty hours after which it is prohibited to drive without resting;

$h^{rest} = 10$ hours, the minimum duration of a rest period in order to be granted the right to drive again after the 11-hour or the 14-hour limits are reached.

3.2 Vehicle routing and scheduling constraints

We now describe the constraints considered in the VRTDSP we are studying in this paper. Given a fleet of m vehicles based at a depot, the problem consists of determining a set of feasible routes to serve a set V of customers in order to minimize a given objective (in the standard version of the problem) or two objectives (in this multicriteria version of the problem described hereafter). The depot is denoted by 0 and we define V_0 as being $V \cup \{0\}$. Let c_{ij} be the travel distance between i and j, where $i, j \in V_0$, and let d_{ij} be the driving time to reach j from i. The problem is solved over a planning horizon of length H. A unique time window $[\tau_{1,0}^{\min}, \tau_{1,0}^{\max}]$ of length H is associated with the depot, where $\tau_{1,0}^{\min}$ and $\tau_{1,0}^{\max}$ represent the earliest possible departure time from the depot and the latest admissible arrival time at the depot, respectively.

An ordered set of time windows $T_i = \{[\tau_{ti}^{\min}, \tau_{ti}^{\max}], t = 1, \ldots, \bar{t_i}\}$, with $\tau_{t-1,i}^{\max} \leq \tau_{ti}^{\max}$ ($t = 2, \ldots, \bar{t_i}$), is also associated with each customer i. These sets of time windows restrict the time intervals within which delivery is allowed. Each customer $i \in V_0$ must be visited once by a vehicle during l_i time units, without service interruption. Consequently, a single time window must be chosen for each customer delivery. If a vehicle arrives before the opening of this time window, it has to wait. It must also arrive before the closing of the selected time window, otherwise the driver will have to wait until the opening of the next available time window.

Groupe Robert seeks to meet customer requirements by using a heterogeneous fleet of vehicles suited to various functions (dry and heated vehicles in this case). Consequently, every customer $i \in V_0$ has a non-negative demand q_i and can be served only by a subset of the vehicle fleet. Every vehicle type has a given load capacity that cannot be exceeded by the total demand it carries.

The routing and scheduling constraints just described consist of a vehicle routing problem with multiple time windows and a heterogeneous fleet. By including the HOS regulations in the scheduling component of this problem, we obtain the resulting decision problem that can be denoted the US VRTDSP.

3.3 Vehicle routing and scheduling objectives

Carriers try to arrange schedules and routes to accommodate their drivers' preferences, in particular with regard to reducing the number of consecutive days they are away from their families. However, while pursuing the objective of minimizing driver inconvenience, the trucking companies also seek to minimize operating costs. Given the aim to optimize both of these objectives simultaneously, the problem in question is referred to as the multi-objective vehicle routing and truck driver scheduling problem (US MOVRTDSP). One objective aims to minimize the operating costs, which are proportional to the total distance traveled by vehicles. The other objective aims to minimize the driver inconvenience, considered to be the sum of the resting time and waiting time of all drivers. Waiting time is measured by the difference, if it is positive, between the beginning of the time window and the arrival time of the vehicle at the customer.

4 Multicriteria solution approach

In this section, we present the multicriteria tabu search algorithm we have developed to solve the US MOVRTDSP. The method we are proposing is a generalization of the algorithm proposed by Rancourt et al. (2012) for the single-criterion US VRTDSP and aims to determine a set of heuristic non-dominated solutions. To do so, we implemented a search mechanism that consists of a single thread in which the weighting attributed to the two objectives, namely operating costs and driver inconvenience, are dynamically modified and in which dominated solutions are discarded along the search. The principle used to adapt the objective function and changes in the weighting have been proven to be efficient by Paquette et al. (2012). In the following, we first explain the outlines of the tabu search. We then describe the scheduling algorithm embedded within the tabu search that enables handling the HOS regulations. Finally, we present how the tabu search has been adapted to manage the bi-objective function and yield a non-dominated solution set.

4.1 Tabu search principles

The tabu search heuristic is first initialized with a solution x_0. This solution is constructed by selecting the customer located the farthest away from the depot and by adding the customer that is closest to the current route until the thresholds with regard to capacity or duration have been met. This process is sequentially repeated until all the customers are assigned to a route. The obtained solution then forms the initial solution x_0, which is not necessarily feasible. The algorithm then moves at each iteration from a solution x to another solution x' in the neighborhood $N(x)$ of x for a fixed number of iterations. The neighborhood $N(x)$ is composed of all the solutions that can be obtained by applying a given type of transformation to x. A solution x is identified with an attribute $B(x) = \{(i, k) : i \in V, k = 1, ..., m\}$ that indicates on which route k each customer i is visited. A neighbor solution is thus obtained by replacing a pair $(i, k) \in B(x)$ with another non-tabu pair $(i, k') \notin B(x)$. The move that best minimizes the increase of the objective function is performed at each iteration and the attribute (i, k) is declared tabu for a fixed number of iterations. However, an aspiration criterion allows the search process to accept a solution x containing a tabu attribute (i, k) if this solution is the best known solution for this attribute.

The search is broadened by accepting intermediate infeasible solutions. This is achieved through the use of a penalized objective function with self-adjusting penalties applied to the constraint violations. Relaxing the constraints allows the algorithm to explore infeasible solutions in order to reach better solutions that may not be possible to obtain otherwise. For a given solution x, let $q(x)$ be the violation of capacity constraints, $w(x)$ the violation of on-duty constraints emerging from the 70-hour limit, and $v(x)$ the violation of time window constraints. The objectives of the carrier are represented by a function $f(x)$ and the total duration of the solution x is represented by $d(x)$. The global objective function $F(x)$ corresponds to:

$$F(x) = f(x) + \alpha q(x) + \beta w(x) + \gamma v(x) + \lambda d(x),$$

where the parameters α, β, and γ are dynamically updated throughout the search. After each iteration, the values of these parameters are modified by a factor of $1 + \delta$, where δ is chosen randomly in the interval $[0, 1]$ at each iteration. If the current solution is feasible with respect to a constraint, the value of the parameter associated to this constraint is divided by $1 + \delta$; if not feasible with respect to a constraint, it is multiplied by $1 + \delta$. A perturbation mechanism has also been implemented by modifying parameter λ during the tabu search. If no solution improving $F(x)$ or one of the carrier objectives in $f(x)$ has been found for 1,000 iterations, the parameter λ will be set equal to 10 for a fixed number of iterations. Otherwise, this parameter is equal to zero. This mechanism helps to diversify the search by increasing the weight on the value of the total duration in the global cost function. The classic diversification and intensification schemes proposed by Cordeau et al. (2001) are also used to improve the search. The goal of the diversification is to penalize solutions containing frequently encountered attributes, whereas intensification is a process aimed at deepening the search around good solutions.

To evaluate the costs and the time constraint violations of a solution, a schedule complying with the HOS regulations has to be determined. To do so, a rest scheduling algorithm that tackles the HOS regulations and the multiple time windows is embedded within the tabu search heuristic.

4.2 Scheduling heuristic

The routes constructed during the course of the search are evaluated using an algorithm proposed by Rancourt et al. (2012), referred to as Algorithm 1 in their paper. This procedure allows determining a schedule complying with the HOS regulations for a given sequence of customers so that the attributes of a route, such as customer service time, driver waiting and resting time can be computed. For a route $r = (i_0 = 0, i_1, ..., i_p, i_{p+1}, ..., i_{n_r} = 0)$ of n_r customers, the forward labeling procedure iteratively generate a set S_p of schedules for each partial route $r_p = (i_0, i_1, ..., i_p)$, with $p \in \{0, ..., n_r\}$.

The algorithm starts with a partial route containing nothing other than the depot. The only schedule in S_0 is the one departing from the depot at the beginning of its time window. In the subsequent iterations $p = (1, ..., n_r)$, each schedule of S_p is extended into a new schedule for the partial route r_p by successively appending driving, waiting, working and rest periods to the end of the schedule that is currently expanded. Once the subsequent customer of the current route is reached, the algorithm generates two possible schedules, namely the direct execution of the service without a rest and the insertion of an anticipated rest before the execution of the service. To minimize the waiting time and to ensure that the driver is as flexible as possible for subsequent activities, unnecessary waiting time is prohibited and the duration of the previously scheduled rest periods is extended (or the departure from the depot is delayed) by computing the forward time slack since the last rest.

By applying this process sequentially, beginning with the depot and continuing with each succeeding node in a route, multiple schedules are generated to reach each customer. To reduce the number of possibilities and to speed up the process, some schedules are pruned during the enumeration by using a heuristic principle of dominance. The retained schedule is the one that allows to complete the trip within the shortest amount of time. For more details on how the schedules are extended and on how the alternative schedules are managed within the labeling process, see Rancourt et al. (2012).

4.3 Adaptation of the tabu search to the bi-objective problem

In this section, we explain how the tabu search heuristic was adapted to handle the bi-objective problem and find a non-dominated solution set. Each solution is represented by a vector, the elements of which correspond to the values of each objective. The reference point method of Clímaco et al. (2006), which was first designed for generic multicriteria problems, is now adapted to our case. According to this method, potential solutions identified during the search process are evaluated by their distance from an ideal point. This point, which is generally infeasible, is defined by a vector featuring the minimum value of each objective optimized separately. The distance between a potential solution and the ideal point is computed with the Manhattan metric for its ease of implementation in the tabu search heuristic. The objective function we are using is defined as:

$$f(x) = \omega_1 |c(x) - c^*| + \omega_2 |i(x) - i^*|,$$

where $c(x)$ are the operating costs, $i(x)$ the driver inconvenience, ω_z the weight put on criterion z ($z = 1$ and 2), and (c^*, i^*) the ideal point. Because we are seeking a set of solutions and not only one optimal solution, the weights are changed dynamically and every feasible solution x' found during the search is compared to the solutions kept in a pool X of non-dominated solutions, even if $F(x') > F(x^*)$, with x^* the value of x yielding $\min\{F(x) \mid x \in X\}$. Indeed, having $F(x') > F(x^*)$ does not necessarily mean that solution x' is dominated when considering the two objectives $c(x)$ and $i(x)$.

The weights ω_1 and ω_2 play two distinct roles. First, since the first objective is measured in units of distance and the second objective in units of time, the weights are used to put the two objectives on a comparable scale. As recommended by Clímaco et al. (2006), we use the scaling factors $\sigma_z = 1/(UB_z - LB_z)$, where UB_z and LB_z are the upper and lower bounds on the value of criterion z. The bounds have to be computed empirically given the absence of a priori information in this case. Second, the weights also help exploring different parts of the solution space, which could be seen as a form of diversification procedure. To this end, three strategies that differ in the way the weights are changed dynamically have been tested.

Strategy 1 is the one developed by Paquette et al. (2012), which updates the weights whenever κ iterations have elapsed since the last feasible solution was encountered, or whenever the algorithm has not identified any feasible solution for more than ι consecutive iterations. The procedure to update the weights is as follows:

$\omega_1 = \sigma_1 \omega_1 |c(x) - c^*|$;

$\omega_2 = \sigma_2 \omega_2 |i(x) - i^*|$;

$\omega_z = \omega_z / (\omega_1 + \omega_2)$, with $z = 1$ and 2.

This procedure puts more emphasis on the term of the objective function for which the worst value was obtained compared to the ideal point at the last iteration. At the same time, the the weights are normalized by setting their sum equal to one. Whenever the weights are updated, the tabu list is emptied.

Strategy 2 has four distinct phases. In the first phase, λ is set equal to 1, and ω_1 as well as ω_2 are fixed equal to zero for the first 10,000 iterations. Thus, the objective function is oriented toward finding solutions with the best duration, which speeds up the process of finding feasible solutions at the beginning of the search. In the second phase, which lasts 20,000 iterations, the procedure described in Strategy 1 is applied. In the third and fourth phases, which last 10,000 iterations each, ω_1 and then ω_2 are respectively set equal to 1. These phases are intensification procedures applied to help finding the two tails of the Pareto front.

Strategy 3 implemented to change the weights is similar to the first one. The only difference is that λ is set equal to 1, and ω_1 as well as ω_2 are fixed to zero for the first 1,000 iterations. Thus, the objective at the beginning of the search is to find solutions with the best duration. Afterwards, the same procedure described for Strategy 1 is applied. Calibrations of the algorithm will first enable us to establish the best values of the parameters ι and κ, and to then determine which strategy of the three is the best.

In the literature, studies on the VRTDSP have solved the problem by minimizing the number of vehicles as the primary objective by adding a large fixed cost on the number of vehicles used. In order to evaluate the impact of increasing the number of vehicles on the operating costs and the driver inconvenience, the tabu search heuristic is sequentially applied for $k = 0, 1, 2,$ and 3, where k corresponds to an increment of vehicles made to the minimum number of vehicles in the fleet. In other words, k=0 means that only the minimum fleet is used and k=3 means that three vehicles are added to the minimum fleet. By applying this process, a set of non-dominated solutions can be determined for the four different fleet sizes. Note that the minimum fleet size m_{min} of each instance was extracted from preliminary experiments and it might not correspond to the optimal value because it was obtained using a heuristic approach.

5 Computational results

The algorithm was first calibrated using 12 of the 56 instances developed by Rancourt et al. (2012) for the US VRTDSP. Using a representative sub-group of the artificial instances enables us to determine the best values for parameters ι and κ as well as the best weight-changing strategy. The calibrated multicriteria algorithm was then applied to solve all 56 artificial instances to provide results before solving the real-life instance of Groupe Robert. In this section, we first describe the different groups of instances. We then explain the indicators used to compare the different parameterizations of the algorithm, followed by a presentation of the results of the different calibrations. Finally, the tradeoffs between the number of vehicles, operating costs and driver inconvenience are evaluated for the artificial instances and the real-life instance.

5.1 Artificial instances

The artificial instances were borrowed from Rancourt et al. (2012), who had generated them based on the benchmark instances of Solomon (1987) by adding multiple time windows for visiting the customers and changing the travel time matrices. They contain 100 customers each and are divided in six classes that differ with regard to the geographical distribution of the customers and their time window tightness. In the C1 and C2 classes, the customers are clustered, unlike in the classes R1 and R2, where the customers are uniformly distributed. Some customers are clustered while others are uniformly distributed in classes RC1 and RC2. The C1, R1 and RC1 classes have narrow time windows and smaller load capacity per vehicle than the C2, R2 and RC2 classes. These latter classes are more difficult to solve since the number of customers per route increases, which considerably increases the number of possibilities evaluated in the scheduling process.

All customers have eight time windows during which they can be visited. The available time of all vehicles is set to 192 hours (or 8 days). The traveling speed of a vehicle is set to five distance units per hour as suggested by Goel (2009). Furthermore, in order to adjust the travel matrix of the customers visited during an eight-day horizon, each value is multiplied by 12. The geographic coordinates of the customer and the depot, the distance matrices and the vehicle capacities remain the same as in the Solomon instances.

5.2 Real-life Groupe Robert instance

The instance of Groupe Robert corresponds to a typical week of distribution of goods to 162 customers in the United States as depicted in Figure 1. Up to eight time windows are associated with each customer depending on the service quality level guaranteed to the customer and the delivery schedule. In order to handle three types of goods with mutual incompatibilities, two vehicles types are considered and a set of compatible vehicles is associated to each customer. Because the routes do not include the trips back to the depot, this real-life instance corresponds to an open vehicle routing problem. Indeed, planning does not consider backhaul loads, as these are only known a posteriori. On-line requests for backhaul opportunities are thus be handled after the deliveries have been completed. As a result, the distances and traveling times from any customer back to the depot are set to zero, as suggested by Pisinger and Ropke (2007).

Figure 1: Customer locations in the Groupe Robert instance (Source: Rancourt et al. 2012)

5.3 Calibration of the algorithm

To assess the quality of the solutions, we have used five performance measures, two of which are indicators frequently used in the literature on evolutionary algorithms (Zitzler et al. 2003, 2010). The first measure is the *hypervolume indicator*, which computes the volume in the objective space formed between the nadir point and all points in the set of non-dominated solutions. The hypervolume indicator thus measures the diversity of the solutions found, which is of great value to decision-makers. Indeed, solutions must be diverse in order to obtain meaningful estimates of the tradeoffs between operating costs and driver inconvenience in different parts of the objective space. The algorithm developed by Fonseca et al. (2006) was used to compute the hypervolume indicators.

The second measure is the *multiplicative unary epsilon indicator*. It represents the smallest value of epsilon by which each point in a reference set has to be multiplied in order to obtain a set that is dominated by the set of non-dominated solutions. Ideally, the reference set should be the Pareto front. However, since our instances cannot be solved exactly, we will use the combination of all the solutions found for the same instance across all the different runs performed on this instance to form a heuristic reference set. For this indicator, a small value is preferable because it means that the evaluated set of non-dominated solutions is close to the reference set.

The other three measures considered to assess the quality of the solutions are the *number of non-dominated solutions found*, the *best encountered value of the operating costs* and the *best encountered value of the driver inconvenience.*

The calibration of the algorithm was realized with 12 of the 56 artificial instances. Two instances were randomly chosen in each class (C101, C108, R103, R110, RC104, RC105, C207, C208, R205, R209, RC202 and RC206). The parameterization has two goals. First, it aims to identify the best values of the number of iterations ι and κ after which the weights are dynamically changed. Second, the calibration will enable determining the best strategy to adopt for changing the weights.

Each of the twelve artificial instances were solved for 50,000 iterations taking the minimum number of vehicles necessary for finding a feasible solution (k=0). Then, each instance was solved by using a fleet that exceeds the minimum fleet size by three vehicles (k=3). The algorithm was first tested using the following values of ι and κ: (1,000; 100), (1,000; 200), (1,000; 500), (2,000; 100), (2,000; 200), (2,000; 500). Tables 4 and 5 show the five solution quality measures obtained with the different parameterizations of ι and κ with an increment of k trucks to the minimum fleet size ($k = 0$ and 3).

Table 4: Average solution quality measures with different ι and κ values and minimum fleet size (k=0)

ι	κ	Epsilon indicator	Hypervolume indicator	Number of solutions	Best cost	Best inconvenience	Average ranking
1,000	100	1.0984	139,979,621.8	29.17	856.05	12,904.94	3.2
1,000	200	1.0899*	139,981,765.8	27.75	858.97	12,530.51	3.2
1,000	**500**	**1.0973**	**140,191,091.0**	**30.00***	**857.16**	**12,769.48**	**2.4***
2,000	100	1.1184	139,756,096.3	25.83	855.17*	13,135.00	5.0
2,000	**200**	**1.1008**	**140,342,369.9***	**28.33**	**856.26**	**12,499.13***	**2.4***
2,000	500	1.1152	139,891,209.3	27.17	858.00	12,785.81	4.8

Table 5: Average solution quality measures with different ι and κ values and minimum fleet size incremented by three vehicles (k=3)

ι	κ	Epsilon indicator	Hypervolume indicator	Number of solutions	Best cost	Best inconvenience	Average ranking
1,000	100	1.1796	137,189,434.4	18.33	860.35	14,934.65	4.0
1,000	200	1.1559*	136,891,235.7	20.25	864.70	14,689.70	3.0
1,000	500	1.2155	137,454,893.0	22.58*	858.05*	14,890.22	2.6
2,000	**100**	**1.1848**	**137,478,455.5***	**20.25**	**860.25**	**14,828.98**	**2.4***
2,000	200	1.1582	137,112,368.6	20.08	865.15	14,632.73*	3.4
2,000	500	1.2158	136,645,404.0	19.58	863.03	15,151.33	5.4

The best values of ι and κ were then selected according with the measures employed to evaluate the quality of the resulting solution sets. Each measure yields a different best parameterization depending on whether we consider the results with the minimum fleet or with a fleet containing three additional vehicles. To analyze these tables, the different parameterizations were ranked for each of the five measures, followed by a calculation of the average ranking for all measures. If we use this ranking method for k=0, the best values of ι and κ are (1,000; 500) and (2,000; 200). For k=3, the best values are (2,000; 100).

However, if we aggregate the results as shown in Table 6, we observe that three of the measures identify (1,000; 500) as the best values. Moreover, if we rank the assessed parameterizations according to the different measures, we obtain once again (1,000; 500) as the best calibration, followed closely by (2,000; 200). The parameterization (2,000; 100), which was best for the fleet with three additional vehicles, is high in this ranking because it performs very poorly in tests using the minimum fleet size. For these reasons, we chose to further compare the parameterizations (1,000; 500) and (2,000; 200) in the evaluation of the three weight-changing strategies.

Table 6: Average on the two fleet sizes of the measures representing the quality of solutions obtained with the different parameterizations of ι and κ

ι	κ	Epsilon indicator	Hypervolume indicator	Number of solutions	Best cost	Best inconvenience	Average ranking
1,000	100	1.1390	138,584,528.10	23.75	858.20	13,919.80	3.6
1,000	200	1.1229*	138,436,500.76	24.00	861.84	13,610.10	3.4
1,000	**500**	**1.1564**	**138,822,991.99***	**26.29***	**857.61***	**13,829.85**	**2.2***
2,000	100	1.1516	138,617,275.88	23.04	857.71	13,981.99	4.2
2,000	**200**	**1.1295**	**138,727,369.21**	**24.21**	**860.70**	**13,565.93***	**2.4**
2,000	500	1.1655	138,268,306.67	23.38	860.51	13,968.57	5.2

The calibration also aimed to identify which of the three strategies can best execute the weight changes. To evaluate the three weight-changing strategies, we used the same five quality measures to compare the results obtained on the twelve instances for 50,000 iterations of our tabu search. Table 7 presents the results obtained with the three tested strategies using the parameterizations (1,000; 500) and (2,000; 200). The results are listed for two different fleet sizes ($k = 0$ and $k = 3$), and for the average of these. The bold numbers represent the best values obtained for each measure across the three strategies. From this table, we observe that, for the average on the two fleet sizes, Strategy 3 using parameters (2,000; 200) is the best strategy according to all solution quality measures except for *best encountered value of the operating costs*. This parameterization was thus selected for the tests described in the next subsection.

5.4 Results for artificial instances

In this sub-section, we analyze the results of the tests run on the 56 modified Solomon instances using the best calibration, which is Strategy 3 with ι equal to 2,000 and κ equal to 200. Our analysis has two goals. First, we want to find out if the dispatchers' rule of thumb of using as few vehicles as possible is effective for minimizing the operating costs. To do so, we will identify the best solutions obtained in terms of costs using different fleet sizes, which allows us to better understand the tradeoffs between the number of vehicles used and the operating costs. Second, we want to better understand the tradeoffs between the carrier's

Table 7: Average solution quality measures with the different weight-changing strategies implemented during the search with parameterizations (1,000; 500) and (2,000; 200)

Strategy	Number of vehicles	ι	κ	Epsilon indicator	Hypervolume indicator	Number of solutions	Best cost	Best inconvenience
1	0	1,000	500	1.17	139,841,144.8	28.17	**851.56**	13,602.28
	3	1,000	500	1.21	**137,445,744.6**	17.25	**857.56**	15,105.05
	Average	1,000	500	1.19	138,643,444.7	22.71	**854.56**	14,353.67
	0	2,000	200	1.17	138,811,734.3	26.75	858.76	13,502.37
	3	2,000	200	1.41	135,357,701.2	**20.08**	865.30	15,954.80
	Average	2,000	200	1.29	137,084,717.8	23.42	862.03	14,728.59
2	0	1,000	500	1.19	138,136,620.5	21.75	861.20	13,992.54
	3	1,000	500	1.23	136,478,627.5	18.00	864.92	15,081.12
	Average	1,000	500	1.21	137,307,624.0	19.88	863.06	14,536.83
	0	2,000	200	**1.11**	138,972,083.6	27.00	872.05	12,161.50
	3	2,000	200	1.28	136,122,565.9	19.00	867.26	15,063.57
	Average	2,000	200	1.19	137,547,324.8	23.00	869.66	13,612.54
3	0	1,000	500	1.11	140,170,100.5	**29.25**	860.77	**12,068.57**
	3	1,000	500	1.23	135,524,870.4	19.17	876.11	15,334.09
	Average	1,000	500	1.17	137,847,485.4	**24.21**	868.44	13,701.33
	0	2,000	200	1.13	**140,342,369.9**	28.33	856.26	12,499.13
	3	2,000	200	**1.16**	137,112,368.6	20.08	865.15	**14,632.73**
	Average	2,000	200	**1.14**	**138,727,369.2**	**24.21**	860.70	**13,565.93**

operating costs and the driver inconvenience for a given fleet size. This information is useful for carriers who are looking for ways to improve the working conditions for their drivers while keeping increases of operating costs within reasonable limits. Indeed, increasing the costs of transportation means increasing the price charged to customers, who may well be price sensitive. To estimate the impacts that a reduction of driver inconvenience has on operating costs, we will analyze the results using an approximation of the operating cost variations generated by decreasing driver inconvenience by 1%, 5%, 10%, 25% and 50%.

To accomplish our first goal, we ran tests sequentially on all 56 instances using four different sizes of available vehicle fleets. Solution sets X_k were thus found using k vehicles in addition to the minimum fleet size, where $k = 0, 1, 2$ and 3. As an example, Figure 2 shows X_k, with $k = 0, 1, 2$ and 3, for instance RC106. Because the four runs of the heuristic were performed sequentially using different fleet sizes, some solutions found with a certain k could be dominated by another solution found with a different fleet size. To eliminate dominated solutions for each instance, the four solution sets X_k were first combined to form the solution set X, subsequent to which the dominance principles were applied. The number of vehicles represents a third objective for eliminating the dominated solutions in X. Thus, the vector representing a solution x becomes $(c(x), i(x), m_{min} + k)$. The solutions were compared, and the dominated solutions in X were discarded to form the global set of non-dominated solutions X'. One set of non-dominated solutions X'_k was also determined for each fleet size $m_{min} + k$, with $k = 0, 1, 2$ and 3, by extracting the solutions using $m_{min} + k$ vehicles in X'. Figure 3 shows the non-dominated solution sets X'_k, which consists of the global set of non-dominated solutions X' for instance RC106. This procedure was applied to all instances. The detailed results are presented in the Appendix and the aggregated results in Table 8.

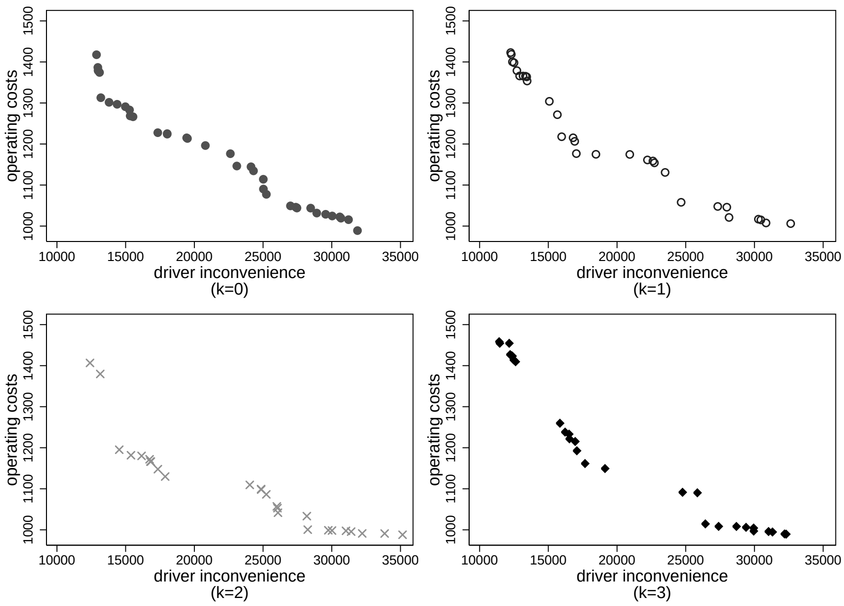

Figure 2: Solution sets X_k for instance RC106, with $k = 0, 1, 2$ and 3

The first rows of Table 8 contain the results aggregated by classes of instances (C1, R1, RC1, C2, R2 and RC2) for the tests run sequentially using the different fleet sizes. The last eight rows represent the results aggregated a level further and show the averages for Class 1 (C1, R1 and RC1) and Class 2 (C2, R2 and RC2). The first column indicates the class of instances; the second column shows the number of k vehicles added to the minimum fleet size; the third column indicates the numbers of solutions in X'_k after the dominance rules were verified on the global solution set X; and the fourth column represents the best encountered value of operating costs $c(x^*_k)$ in solution set X'_k, with x^*_k the value of x yielding $\min\{c(x) \mid x \in X'_k\}$. If a non-dominated solution exists with the corresponding augmented fleet sizes, we calculate the variations of operating costs between the best encountered value obtained with the minimum fleet size and the best encountered objective value obtained with the augmented fleet sizes. These variations are expressed as the percentages of the

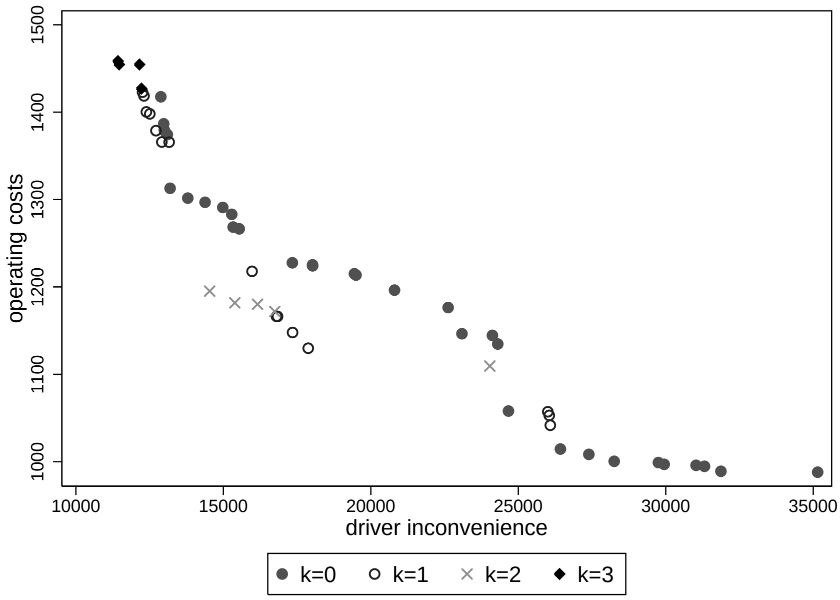

Figure 3: Non-dominated solution sets X'_k for instance RC106, with $k = 0, 1, 2$ and 3

cost increases or decreases that resulted from augmenting the fleet size to serve customers ($100(c(x^*_k) - c(x^*_0))/c(x^*_0)$, with $k = 1, 2$ and 3) and are shown in the fifth column of Table 8. The last column provides the CPU times in minutes.

From the results in Table 8, we observe that for instances in Class 1 and RC2, the number of vehicles used should not exceed the minimum fleet size, even if vehicles are available, as this would lead to a considerable increase in operating costs. For class R2, it can be beneficial to use one additional vehicle depending on the instance. However, for this class of instances, operating costs increase when using two or three additional vehicles. For class C2, adding one or more vehicles can reduce operating costs, but it can also have the opposite effect. Which effect takes place probably depends on the location of the customer clusters relative to the depot. In general, however, for artificial instances, using more vehicles than the minimum fleet size increases the operating costs. From the analysis of the computational results on artificial instances, we can conclude that the rule of thumb applied by most dispatchers of using as few vehicles as possible is appropriate to minimize operating costs.

To accomplish the second goal, we performed the following analyses in order to enhance our understanding of the tradeoffs between operating costs and driver inconvenience. To do so, all the solutions contained in the set of non-dominated solutions were first proportionally

Table 8: Average results of the impact of fleet sizes on operating costs

Instances	k	$\|X'_k\|$	Best cost ($c(x^*_k)$)	Cost variation (%)	CPU (min)
C1	0	20.3	819.56	–	96.98
	1	2.1	944.21	15.21	92.09
	2	1.2	960.42	17.19	89.43
	3	0.6	1,043.66	27.34	83.79
C2	0	23.0	960.35	–	209.98
	1	23.6	930.62	–3.61	177.34
	2	7.7	1,035.60	4.20	178.64
	3	1.1	1,147.78	–6.41	166.22
R1	0	35.2	863.32	–	144.85
	1	8.7	1,020.14	17.75	135.53
	2	2.2	1,093.76	27.93	121.17
	3	1.8	1,133.82	30.95	122.53
R2	0	25.3	738.05	–	288.81
	1	17.0	720.86	–1.54	253.22
	2	8.3	880.98	17.95	225.98
	3	2.7	925.24	25.84	208.12
RC1	0	23.5	995.54	–	115.19
	1	8.9	1,089.77	9.43	104.44
	2	7.8	1,161.24	16.43	95.84
	3	1.9	1,309.50	31.55	93.53
RC2	0	26.6	841.60	–	219.10
	1	17.1	882.51	1.65	195.10
	2	3.0	1,061.72	22.13	182.06
	3	1.6	1,133.80	40.71	175.65
Class 1	0	26.3	892.80	–	115.19
	1	6.6	1,018.04	14.13	110.69
	2	3.7	1,071.81	20.52	102.14
	3	1.4	1,162.33	29.95	99.95
Class 2	0	25.0	846.67		239.30
	1	19.3	844.66	–1.17	208.55
	2	6.3	992.77	14.76	195.56
	3	1.8	1,068.94	20.05	183.33

transformed according to a scale ranging from zero to one for each instance. We then estimated the Pareto front using a piecewise linear function where two consecutive points are linked with a straight segment. This allowed to extrapolate the percentage increases in operating costs obtained by improving driver inconvenience for each instance. Figure 4 shows, on the left side, an example of the set of non-dominated solutions obtained for instance RC106, while on the right side, it shows its corresponding scaled set in which each solution $(c(x), i(x)) \in X'_1$, has been transformed according to the following rule:

$$(c'(x), i'(x)) = ((c(x) - 1,042)/(1,423 - 1,042), (i(x) - 12,261)/(26,085 - 12,261)),$$

where 1,042 and 12,261 are the minimum cost and inconvenience values observed, and where 1,423 and 26,085 are the maximum cost and inconvenience values observed. The according piecewise linear function is shown on Figure 4. The corresponding estimated increase in operating costs obtained by improving the driver inconvenience by 25% is also highlighted with the straight grey lines. These estimates were calculated for all instances. The detailed results are presented in the Appendix and the aggregated results of each class in Table 9.

The first rows of Table 9 contain the results aggregated by classes of instances (C1, R1, RC1, C2, R2 and RC2) for the tests run sequentially using the different fleet sizes. The last eight rows represent the average results for Class 1 and Class 2. The first column indicates the instance class and the second column shows the number of k vehicles added to the minimum fleet size. The other columns indicate the operating cost variations in percentages for decreases in driver inconvenience of 1%, 5%, 10%, 25% and 50%. Similarly, Figure 4 shows the average variations in operating costs for decreases in driver inconvenience of 1%, 5%, 10%, 25% and 50% for the instances of Class 1 and Class 2.

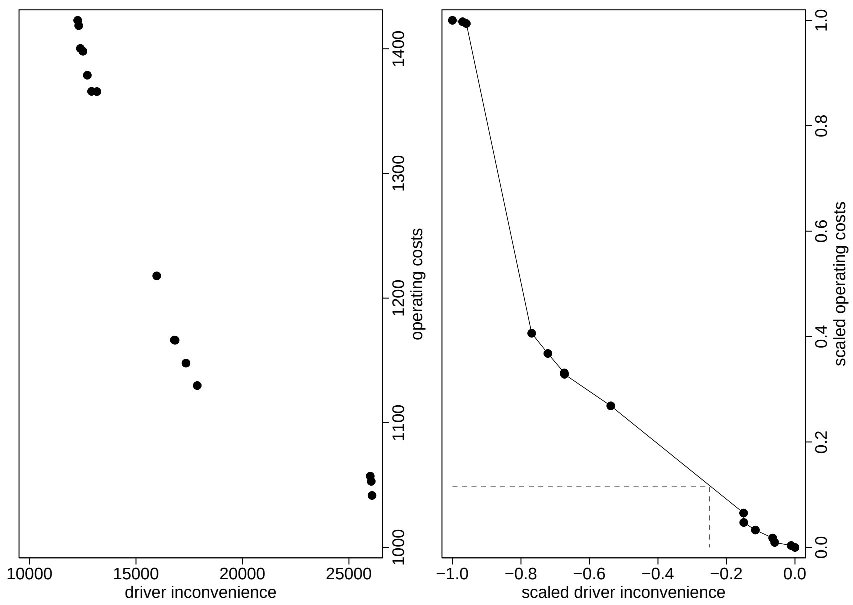

Figure 4: Evaluations of the tradeoffs between operating costs and driver inconvenience for instance RC106

We observe from Figure 5 and Table 9 that for the calculated average value of the instances in Class 1, a 0.38% cost increase corresponds to a 1% reduction of driver inconvenience when using the minimum fleet size. Improving driver inconvenience by 5% and 10% will result in 0.87% and 1.26% cost increases, respectively. Thus, we can decrease driver inconvenience up to 10% by increasing costs by less than 1.5%. However, a decrease of 50% in driver inconvenience would cause a 16.8% increase in operating costs, which is not acceptable. Similarly, we observe that for the calculated average value of the instances in Class 2 using the minimum fleet size, a 0.66% cost increase corresponds to a 1% reduction of driver inconvenience. Improving driver inconvenience by 5% and 10% will result in 2.38% and 4.28% cost increases, respectively. Thus, we can decrease driver inconvenience up to 10% by increasing costs by less than 5%. Finally, a decrease of 50% in driver inconvenience would cause a 22.34% increase of operating costs, which is also not acceptable.

The results presented by Figure 5 clearly show a negative correlation between operating costs and driver inconvenience. The more the driver inconvenience decreases, the more costly it is. The relationship between these two variables seems to be linear in general. Moreover, we can observe that it is more costly to improve driver inconvenience with a larger fleet size then it is with the minimum fleet size, this without taking into account a fixed cost related to vehicles. This figure also depicts the minimum and maximum variations observed on all instances and vehicle fleet sizes, values which represent the best and worst case scenarios that were observed.

Table 9: Tradeoffs between operating costs and driver inconvenience (average over all instances in the same class)

Instances	k	Variation in operating costs (%) for decreases in driver inconvenience of				
		1%	5%	10%	25%	50%
C1	0	0.48	1.16	1.53	4.42	12.03
	1	5.87	8.66	12.15	22.61	39.78
	2	0.93	4.63	9.26	23.16	46.32
	3	1.00	5.00	10.00	25.00	50.00
C2	0	0.47	1.93	2.80	6.30	20.65
	1	0.64	2.71	5.94	16.00	29.81
	2	0.15	0.69	1.84	9.51	22.56
	3	8.33	41.64	52.79	54.03	54.85
R1	0	0.40	0.95	1.57	4.28	17.15
	1	0.68	3.32	7.03	17.11	34.39
	2	1.34	5.79	9.90	26.62	49.48
	3	30.80	44.25	45.44	48.04	64.42
R2	0	0.57	1.61	3.00	7.19	19.12
	1	0.51	1.23	2.36	8.62	19.67
	2	0.74	2.69	4.22	13.30	35.41
	3	1.24	5.13	8.65	18.92	40.97
RC1	0	0.26	0.51	0.67	2.93	21.10
	1	2.59	8.64	13.75	23.30	41.00
	2	2.18	4.61	7.45	22.31	46.61
	3	3.24	16.21	28.75	39.05	56.22
RC2	0	0.92	3.60	7.04	11.47	27.25
	1	0.34	0.85	1.91	5.60	19.21
	2	1.44	7.21	14.42	26.97	46.19
	3	2.45	9.53	11.47	17.31	32.95
Class 1	0	0.38	0.87	1.26	3.88	16.76
	1	3.05	6.87	10.97	21.01	38.39
	2	1.48	5.01	8.87	24.03	47.47
	3	11.68	21.82	28.06	37.36	56.88
Class 2	0	0.66	2.38	4.28	8.32	22.34
	1	0.50	1.60	3.40	10.08	22.90
	2	0.78	3.53	6.82	16.59	34.72
	3	4.01	18.77	24.30	30.08	42.92

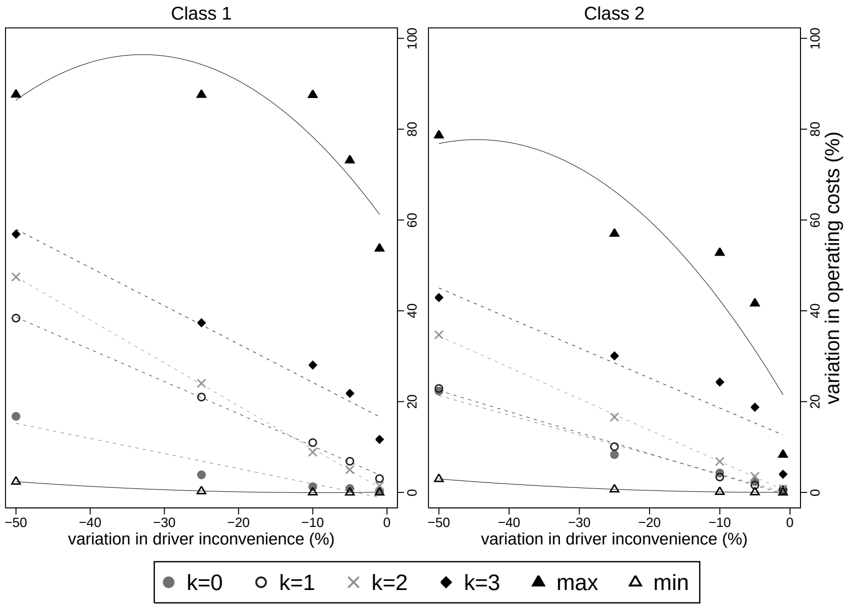

Figure 5: Average over all instances of Class 1 (left side) and Class 2 (right side) of variations in operating costs for different decreases in driver inconvenience

5.5 Results for the real-life instance

Tests for five different fleet sizes have also been sequentially run on the instance provided by Groupe Robert. Here, the same parameterization and weight-change strategy as used in the artificial instances were implemented in the tabu search heuristic. The algorithm finds non-dominated solutions in terms of operating costs and driver inconvenience for $k = 0, 1, 2, 3$ and 4 as shown in Figure 6. The solution established by the dispatchers of Groupe Robert is provided in the last graph on this figure. When dominated solutions are eliminated considering the number of vehicles as a third objective, we obtain non-dominated solution sets X'_k only for k=0, 1 and 4. Figure 7 shows these three sets of non-dominated solutions and the solution applied by Groupe Robert. When compared to the solution that Groupe Robert used for this instance, our solutions reduce operating costs and driver inconvenience simultaneously. The solution from Groupe Robert is dominated by solutions that have been identified by the multicriteria tabu search heuristic. The results presented in Table 10 show the best value encountered in terms of operating costs for different fleet sizes. We observe from this table and Figure 7 that costs decrease when fleet size increases. Thus, the rule of thumb does not seem to apply to the real-life instance.

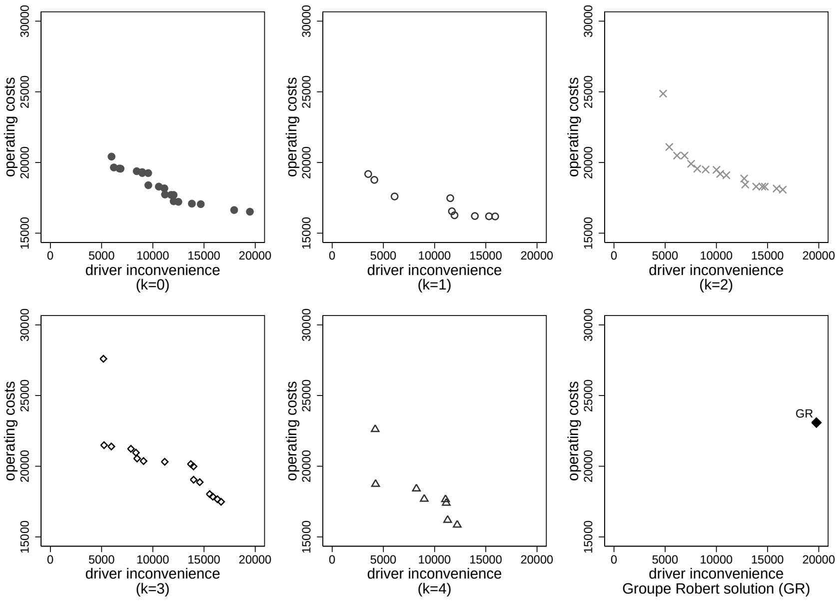

Figure 6: Solution sets X_k for Groupe Robert instance, with $k = 0, 1, 2, 3$ and 4, and Groupe Robert solution

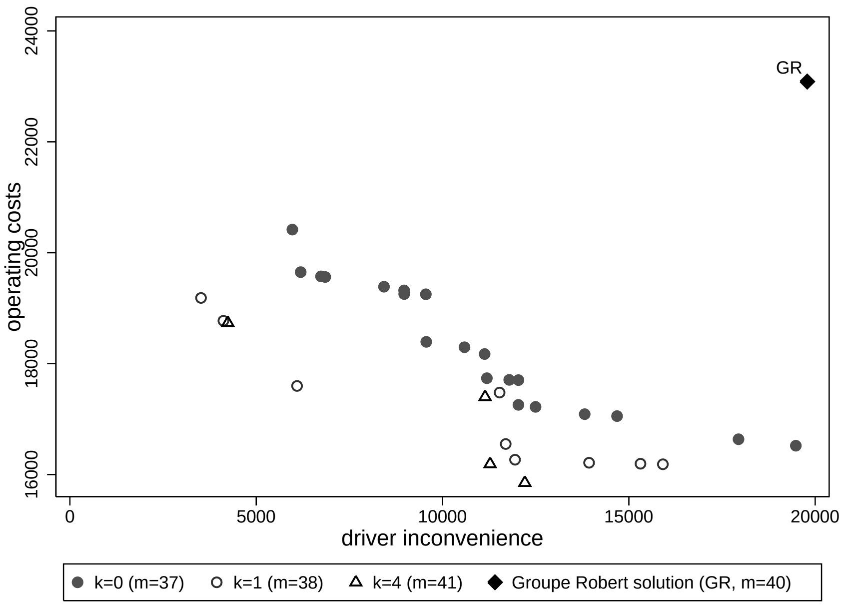

Figure 7: Non-dominated solution sets X'_k for Groupe Robert instance, with $k = 0, 1$ and 4, and Groupe Robert solution

Table 10: Results for the real-life instance provided by Groupe Robert

m	Number of solutions	Best cost	Variation in cost (%)	CPU (min)
37	20	16,520.50		74.38
38	9	16,185.30	–2.03	45.83
39	0	–	–	90.34
40	0	–	–	89.39
41	4	15,853.10	–4.04	44.81

Figure 8 and Table 11 present the variations in operating costs for decreases in driver inconvenience of 1%, 5%, 10%, 25% and 50% for fleets of 37, 38 and 41 vehicles. We observe that to decrease driver inconvenience by 1%, operating costs will increase by 0.26% when using the minimum fleet size. Decreasing driver inconvenience by 5% and 10% will lead to 1.31% and 2.62% cost increases, respectively. Thus, we can decrease driver inconvenience by up to 10% by increasing costs by less than 3%. Finally, a 50% decrease in driver inconvenience would result in a 17.38% increase of operating costs, which is not acceptable. We also observe that when 41 vehicles are used, the cost of decreasing driver inconvenience increases significantly. Moreover, a decrease of driver inconvenience between 1% and 30% appears to be less costly with a fleet of 38 vehicles than with the minimum fleet size, while decreasing driver inconvenience by more than 30% becomes more costly with a fleet of 38 vehicles than with the minimum fleet size.

Table 11: Tradeoffs between driver inconvenience and operating costs for the real-life instance with different fleet sizes

m	Variation in operating costs (%) for decreases in driver inconvenience of				
	1%	5%	10%	25%	50%
37	0.26	1.31	2.62	9.03	17.38
38	0.07	0.34	0.61	1.95	44.45
39	–	–	–	–	–
40	–	–	–	–	–
41	1.00	5.00	9.99	59.86	73.24

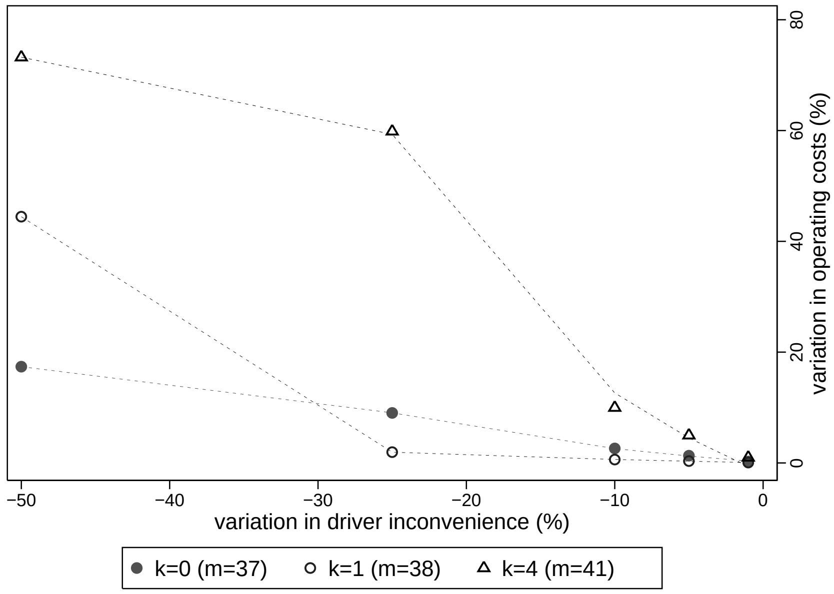

Figure 8: Tradeoffs between operating costs and driver inconvenience for the real-life instance with different fleet sizes

6 Discussion

Carriers in North American countries are facing a driver shortage. To tackle this issue, we propose a decision tool that can provide trucking companies with information for better understanding the tradeoffs between the number of vehicles used and the operating costs. The tradeoffs between operating costs and driver inconvenience on the global non-dominated solution sets have also been analyzed.

6.1 Tradeoffs between fleet sizes and operating costs

Results presented in Section 5 show that for the C1, R1 and RC1 instance classes, which are characterized by tight time windows and lower vehicle capacities, operating costs always increase when more vehicles are used to serve the same pool of customers. This could be explained by the fact that customers are located far from the depot, this being a context of long-haul vehicle routing, and that these instances are highly constrained. Thus, the range of different feasible solutions is rather small, with little possibility for re-optimization even when using a supplement vehicle. In this circumstance, using one more vehicle can add a long-distance stretch to reach a customer that was already served by another vehicle without necessarily significantly reducing the distance traveled between two customers. This phenomenon is illustrated with the example of Figure 9. Carriers that usually have to manage with tight time windows to serve their customers should continue to make use of the rule of thumb that aims to use as few vehicles as possible to minimize their operating costs and to enhance the capability of the carrier to provide service in the case of last-minute demands by keeping vehicles available at the depot if possible.

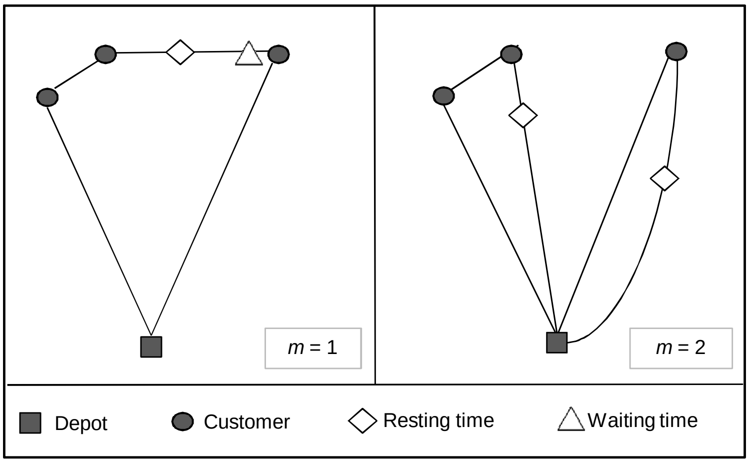

Figure 9: Example where using one more vehicle adds a long-distance stretch

The results for classes C2, R2 and RC2 will be analyzed separately. Results for instances in class RC2 are similar to those from Class 1. Increasing the number of vehicles available increases the operating costs of the carrier. For instances in Class C2, which are characterized by wide time windows and geographically clustered customers, using one or more vehicles decreases operating costs for some instances and increases them for other instances. This could be explained by the fact that some instances have customer clusters located close to the depot whereas for other instances all clusters are farther away. In the first case, adding a vehicle does not significantly increase long distance stretches, but in the second case it does, which would explain the irregular impacts on operating costs. As an exception to the rule, carriers having a cluster of customers close to their depots should consider adding one or more vehicles to the minimum fleet in order to better serve these customers, and determine whether this decreases their operating costs. For instances of class R2, operating costs generally increase when two or more vehicles are added to the minimum fleet, which confirms the rule of thumb. However, sometimes operating costs decrease only when one vehicle is added to the minimum fleet. Even if operating costs sometimes decrease for some instances, the rule of thumb still holds, by and large, for the artificial instances.

The results from the real-life instance stand in contrast to the ones obtained with the artificial instances. Adding vehicles to the minimum fleet of 37 vehicles reduces operating costs. This could be explained by the fact that there is no cost for returning to the depot in this case. Indeed, Groupe Robert's problem is considered as an open routing problem in that returns are not planned or scheduled until after all deliveries are completed. Thus, we do not consider the long distance to be traveled from the last visited customer to the depot when adding a vehicle, a factor that would most likely offset the impact which an additional vehicle has on operating costs. However, sending another vehicle on the road would increase the risk for the carrier because backhaul opportunities are not guaranteed. The 2% decrease in operating costs may not be large enough to compensate for the added risks for the carrier.

6.2 Tradeoffs between operating costs and driver inconvenience

According to the results found in Figures 5 and 8, there exists a negative correlation between operating costs and driver inconvenience for all instances. The more driver inconvenience decreases, the more costly it is. This could be explained by the fact that in order to reduce waiting time or resting time, routes have to be adjusted, which increases the total distance traveled. Thus, to improve driver working conditions, trucking companies have to increase their fuel and maintenance costs. Moreover, we could conclude that it is less costly to improve driver inconvenience with the minimum fleet size than it is with an augmented fleet. This could be explained with the relationship between fleet size and operating costs obtained from previous results on tradeoffs. Because the variations of operating costs increase more than the variation of driver inconvenience when vehicles are added to the fleet, the proportion to improve driver inconvenience becomes higher when more vehicles are used.

Carriers who want to improve the working conditions for drivers will have to (1) increase their prices, which could be difficult in this highly competitive market, (2) reduce their profit margin, which are already low or (3) offer a better service to their customers for an increased price, thereby differentiating themselves with value-added services. Of these three strategies, the last one appears to be the most profitable and sustainable one for these carriers. Overall, drivers are in direct contact with customers more so than any other employees of the carrier. Thus, motivated and dedicated drivers are a good means to improve the service provided to customers. We argue that it is feasible for carriers to offer a better service to their customers for the same price, or for a slightly higher price, by reducing turnover rates, decreasing expenditures for training new drivers and maintaining stability in their workforce.

7 Conclusions

This paper was motivated by the existing shortage of truck drivers in North America, which is partly due to their difficult working conditions. This can be remedied by constructing routes and schedules that not only minimize costs, but also take drivers' welfare into account. We have described a multicriteria tabu search algorithm that can simultaneously account for both these goals.

The multicriteria tabu search heuristic developed for this problem was capable of identifying a set of non-dominated solutions for each artificial instance and for the real-life instance of Groupe Robert within reasonable computational times. Our algorithm provides better solutions in terms of the number of vehicles, operating costs and driver inconvenience, compared to the solution implemented by Groupe Robert. Moreover, our heuristic yields a decision support tool which provides meaningful information to long-haul carriers.

Our results help gain a better understanding of the tradeoffs between the number of vehicles and the operating costs of a solution. Extensive analyses enabled us to verify the efficacy of a rule of thumb which is commonly used in the transportation industry and which is, by and large, also applicable to artificial instances. We show that for most cases, operating costs increase when the number of vehicles exceeds the minimum fleet size. However, in the case of the real-life instance which is an open routing problem, the rule of thumb appears to be less appropriate.

Our results also improve our understanding of the potential tradeoffs between operating costs and driver inconvenience. Drivers are an important asset for carriers. Improving their working conditions can thus reduce their turnover rates, thereby improving the level of service and reducing training costs. Results show that decreasing driver inconvenience increases carriers' operating costs in a smaller proportion. For example, for Groupe Robert, we could decrease driver inconvenience by up to 10% when increasing costs by less than 2.6%.

References

Alves, M. J. and J. Clímaco (2000). An interactive method for 0-1 multiobjective problems using simulated annealing and tabu search. *Journal of Heuristics* 6, 385–403.

Archetti, C. and M. W. P. Savelsbergh (2009). The trip scheduling problem. *Transportation Science* 43, 417–431.

Bäck, T. (1996). *Evolutionary Algorithms in Theory and Practice*. New York: Oxford University Press.

Bartodziej, P., U. Derigs, D. Malcherek, and U. Vogel (2009). Models and algorithms for solving combined vehicle and crew scheduling problems with rest constraints : an application to road feeder service planning in air cargo transportation. *OR Spectrum* 31, 405–429.

Boffey, T. B. (1996). Multiobjective routing problems. *TOP* 3, 167–220.

Ceselli, A., R. Righini, and M. Salani (2009). A column generation algorithm for a rich vehicle-routing problem. *Transportation Science* 43, 56–69.

Clímaco, J. C. N., J. M. F. Craveirinha, and M. M. B. Pascoal (2006). An automated reference point-like approach for multicriteria shortest path problems. *Journal of Systems Science and Systems Engineering* 15, 314–329.

Coello Coello, C., D. Van Veldhuizen, and G. Lamont (2002). *Evolutionary Algorithms for Solving Multi-Objective Problems*. Kluwer Academic/Plenum Publishers.

Cordeau, J.-F., G. Laporte, and A. Mercier (2001). A unified tabu search heuristic for vehicle routing problems with time windows. *Journal of the Operational Research Society* 52, 928–936.

Cvetkovic, D. and I. Parmee (2002). Preferences and their application in evolutionary multiobjective optimization. *IEEE Transactions on Evolutionary Computation 6*(1), 42–57.

Doss, N. (2012). Driver shortage shows gain in U.S. truck cargo: freight markets. *Bloomberg*. URL http://www.bloomberg.com/news/2011-08-25/driver-shortage-shows-gain-in-u-s-truck-cargo-freight-markets.html.

Ehrgott, M. and X. Gandibleux (2002). Multiobjective combinatorial optimization - theory, methodology, and applications. In M. Ehrgott and X. Gandibleux (Eds.), *Multiple Criteria Optimization*, pp. 369–444. Norwell, MA: Kluwer.

European Transport Safety Council (2001). The role of driver fatigue in commercial road transport crashes. ETSC Review. URL www.etsc.eu.

Federal Motor Carrier Safety Administration (2008). Hours of service of drivers. *Federal Register 73*(224), 69569–69586.

Gandibleux, X., N. Mezdaoui, and A. Fréville (1997). A tabu search procedure to solve multiobjective combinatorial optimization problems. In R. Caballero, F. Ruiz, and R. Steuer (Eds.), *Advances in Multiple Objectives and Goal Programming*, Volume 455 of *Lecture Notes in Economics and Mathematical Systems*, pp. 291–300. Berlin: Springer-Verlag.

Garcia-Najera, A. and J. Bullinaria (2011). An improved multi-objective evolutionary algorithm for the vehicle routing problem with time windows. *Computers & Operations Research 38*(1), 287–300.

Goel, A. (2009). Vehicle scheduling and routing with drivers' working hours. *Transportation Science* 43, 17–26.

Goel, A., C. Archetti, and M. W. P. Savelsbergh (2012). Truck driver scheduling in Australia. *Computers & Operations Research* 39, 1122–1132.

Goel, A. and A. L. Kok (2012). Truck driver scheduling in the United States. *To appear in Transportation Science*.

Goel, A. and L.-M. Rousseau (2012). Truck driver scheduling in Canada. *To appear in Journal of Scheduling*.

Goel, A. and T. Vidal (2012). Hours of service regulations in road freight transport : an optimization-based international assessment. *Technical Report CIRRELT* 2012-08.

Hansen, M. P. (2000). Tabu search for multiobjective combinatorial optimization: Tamoco. *Control and Cybernetics* 29, 799–818.

Jaszkiewicz, A. (2003). Do multiple-objective metaheuristics deliver on their promises? a computational experiment on the set-covering problem. *IEEE Transactions on Evolutionary Computation 7*(2), 133–143.

Johnston, J., D. Bristow, D. McClure, and K. Schneider (2009). Long distance truck drivers – their joys and frustrations. *Journal of Transportation Management 20*(1), 1–20.

Jozefowiez, N., F. Semet, and E.-G. Talbi (2002). Parallel and hybrid models for multi-objective optimization: Application to the vehicle routing problem. In J. M. G. et al. (Ed.), *Parallel problem solving from nature VII*, Volume 2439 of *Lecture notes in Computer Science*, pp. 271–280. Springer-Verlag.

Jozefowiez, N., F. Semet, and E.-G. Talbi (2008). Multi-objective vehicle routing problems. 189, 293–309.

Jozefowiez, N., F. Semet, and E.-G. Talbi (2009). An evolutionary algorithm for the vehicle routing problem with route balancing. *European Journal of Operational Research 195*(3), 761–769.

Kok, A. L., E. W. Hans, and J. M. J. Schutten (2011). Optimizing departure times in vehicle routes. *European Journal of Operational Research* 210, 579–587.

Kok, A. L., C. M. Meyer, H. Kopfer, and J. M. J. Schutten (2010). A dynamic programming heuristic for the vehicle routing problem with time windows and european community social legislation. *Transportation Science* 44, 442–454.

Murata, T. and T. Itai (2005). Multi-objective vehicle routing problems using two-fold emo algorithm to enhance solution similarity on non-dominated set. In C. C. Coello, A. Aguirre, and E. Zitzler (Eds.), *Evolutionary multi-criterion optimization*, Volume 3410 of *Lecture notes in Computer Science*, pp. 885–896. Guanajanto, Mexico: Springer.

Murata, T. and T. Itai (2007). Local search in two-fold emo algorithm to enhance solution similarity multi-objective vehicle routing problems. In C. P. T. H. S. Obayashi, K. Deb and T. Murata (Eds.), *EMO 2007*, Volume 4403 of *Lecture notes in Computer Science*, pp. 201–215. Springer.

Ombuki, B., B. Ross, and F. Hanshar (2006). Multi-objective genetic algorithm for vehicle routing problem with time windows. *Applied Intelligence* 24, 17–30.

Paquette, J., J.-F. Cordeau, G. Laporte, and M. Pascoal (2012). Combining multicriteria analysis and tabu search for dial-a-ride problems. *Technical Report CIRRELT* 2012-04.

Park, Y. and C. Koelling (1986). A solution of vehicle routing problems in multiple objective environment. *Engineering Costs and Production Economics* 10, 121–132.

Parragh, S., K. Doerner, R. Hartl, and X. Gandibleux (2009). A heuristic two-phase solution approach for the multi-objective dial-a-ride problem. *Networks 54*(4), 227–242.

Pisinger, D. and S. Ropke (2007). A general heuristic for vehicle routing problems. *Computers & Operations Research* 34, 2403–2435.

Prescott-Gagnon, E., M. Drexl, and L.-M. Rousseau (2010). European driver rules in vehicle routing with time windows. *Transportation Science* 44, 455–473.

Rancourt, M.-E., J.-F. Cordeau, and G. Laporte (2012). Long-haul vehicle routing and scheduling with working hour rules. *To appear in Transportation Science*.

Refanidis, I. and I. Vlahavas (2003). Multiobjective heuristic state-space planning. *Artificial Intelligence 145*(1), 1–32.

Repoussis, P., C. Tarantilis, and G. Ioannou (2009). Arc-guided evolutionary algorithm for the vehicle routing problem with time windows. *IEEE Transactions on evolutionary computation 13*(3), 624–647.

Solomon, M. M. (1987). Algorithms for the vehicle-routing and scheduling problems with time window constraints. *Operations Research* 35, 254–265.

Tan, K., Y. Chew, and L. Lee (2006). A hybrid multi-objective evolutionary algorithm for solving truck and trailer vehicle routing problems. *European Journal of Operational Research* 172, 855–885.

U.S. Department of Transportation (2012). Summary of hours-of-service (hos) regulations. URL www.fmcsa.dot.gov/rules-regulations/topics/hos/index.htm.

Williamson, A., S. Sadural, A. Feyer, and R. Friswell (2001). Driver fatigue: A survey of long distance heavy vehicle drivers in australia. *Information paper* 198. National Road Transport Commission.

Xu, H., Z.-L. Chen, S. Rajagopal, and S. Arunapuram (2003). Solving a practical pickup and delivery problem. *Transportation Science* 37, 347–364.

Zäpfel, G. and M. Bögl (2008). Multi-period vehicle routing and crew scheduling with outsourcing options. *International Journal of Production Economics* 113, 980–996.

Zitzler, E., L. Thiele, and J. Bader (2010). On set-based multiobjective optimization. *IEEE Transactions on Evolutionary Computation* 14, 58–79.

Zitzler, E., L. Thiele, M. Laumanns, C. M. Fonseca, and G. V. da Fonseca (2003). Performance assessment of multiobjective optimizers: An analysis and review. *IEEE Transactions on Evolutionary Computation* 7, 117–132.

Appendix

Table 12: Detailed results for instance class C1

Instance	k	m	$\|X'_k\|$	Best objective values		Variation in operating costs (%) for decreases in driver inconvenience of					CPU
				cost	inconvenience	1%	5%	10%	25%	50%	(min)
C101	0	10	18	819.56	8947.87	0.13	0.67	1.34	4.91	10.37	83.73
C101	1	11	1	875.1	9617.47	–	–	–	–	–	82.62
C101	2	12	0	–	–	–	–	–	–	–	44.94
C101	3	13	0	–	–	–	–	–	–	–	66.63
C102	0	10	21	819.56	8181.93	0.41	1.27	1.45	6.41	8.51	99.45
C102	1	11	0	–	–	–	–	–	–	–	69.51
C102	2	12	0	–	–	–	–	–	–	–	119.36
C102	3	13	0	–	–	–	–	–	–	–	58.03
C103	0	10	35	819.56	9598.7	0.2	0.56	0.7	2.86	6.06	102.14
C103	1	11	13	946.92	8922.29	15.62	15.99	16.45	17.82	19.35	145.36
C103	2	12	0	–	–	–	–	–	–	–	123.79
C103	3	13	0	–	–	–	–	–	–	–	94.53
C104	0	10	20	819.56	7880.76	0.04	0.19	0.39	1.51	7.43	115.71
C104	1	11	2	926.39	7800	1	5	10	25	50	101.61
C104	2	12	1	949.19	7922.75	–	–	–	–	–	81.31
C104	3	13	2	1001.45	6683.7	1	5	10	25	50	111.37
C105	0	10	23	819.56	8819.79	1.11	2.66	2.91	5.76	15.92	90.33
C105	1	11	0	–	–	–	–	–	–	–	66.16
C105	2	12	0	–	–	–	–	–	–	–	81.70
C105	3	13	1	1008.41	8799.45	–	–	–	–	–	70.43
C106	0	10	15	819.56	8306.58	1.18	2.94	3.7	10.21	12.41	87.71
C106	1	11	0	–	–	–	–	–	–	–	83.68
C106	2	12	0	–	–	–	–	–	–	–	65.12
C106	3	13	0	–	–	–	–	–	–	–	81.89
C107	0	10	21	819.56	9433.24	0.2	0.69	1.3	3.1	9.64	98.64
C107	1	11	0	–	–	–	–	–	–	–	95.70
C107	2	12	8	972.34	8518.26	0.85	4.26	8.53	21.32	42.64	95.71
C107	3	13	0	–	–	–	–	–	–	–	87.29
C108	0	10	7	819.56	10124.4	0.09	0.43	0.85	2.46	3.45	95.63
C108	1	11	1	961.87	11102.7	–	–	–	–	–	75.32
C108	2	12	2	959.73	9697.56	1	5	10	25	50	78.94
C108	3	13	2	1121.13	9076.13	1	5	10	25	50	103.16
C109	0	10	23	819.56	8402.76	1.01	1.04	1.13	2.51	34.51	99.51
C109	1	11	2	1010.76	8803.1	1	5	10	25	50	108.90
C109	2	12	0	–	–	–	–	–	–	–	113.98
C109	3	13	0	–	–	–	–	–	–	–	80.77

Table 13: Detailed results for instance class C2

Instance	k	m	$\lvert X'_k \rvert$	Best objective values		Variation in operating costs (%) for decreases in driver inconvenience of					CPU
				cost	inconvenience	1%	5%	10%	25%	50%	(min)
C201	0	8	28	1566.94	25485.8	0.97	2.19	3.96	6.7	20.48	120.08
C201	1	9	42	1278.24	24675.1	0.99	3.17	4.52	5.83	19.57	123.67
C201	2	10	11	1254.38	24594.2	0	0.05	0.79	2.99	26.68	111.42
C201	3	11	7	1205.65	42132.3	8.33	41.64	52.79	54.03	54.85	126.69
C202	0	7	18	1197.21	23222.2	0.09	0.68	1.49	4.38	40.77	231.02
C202	1	8	44	965.25	22678.6	0.06	0.29	0.5	2.97	13.1	179.11
C202	2	9	27	963.05	18959.8	0.14	0.35	1.83	7.83	28.11	204.91
C202	3	10	0	–	–	–	–	–	–	–	180.67
C203	0	6	28	813.08	15989.6	0.29	1.45	2.31	10.63	37.16	215.98
C203	1	7	6	1011.94	18365.3	1.01	5.05	15.21	56.99	69.38	201.77
C203	2	8	10	1017.91	15306.9	0.17	0.85	1.7	4.24	8.48	273.43
C203	3	9	0	–	–	–	–	–	–	–	206.31
C204	0	5	35	676.08	9480.66	0.17	0.89	1.8	4.51	9.05	295.91
C204	1	6	3	760.46	10978	1.08	5.38	10.77	26.92	53.08	234.36
C204	2	7	1	830.23	11191.3	–	–	–	–	–	236.24
C204	3	8	0	–	–	–	–	–	–	–	181.43
C205	0	6	1	870.56	24230.3	–	–	–	–	–	183.46
C205	1	7	44	814.76	16362	0.01	0.05	1.86	3.97	19.09	145.35
C205	2	8	0	–	–	–	–	–	–	–	143.13
C205	3	9	0	–	–	–	–	–	–	–	146.89
C206	0	6	25	988.74	17724.3	0.16	4.72	6.12	12.98	29.98	225.72
C206	1	7	35	879.64	16283.6	0.13	0.65	2.59	5.01	20.95	176.97
C206	2	8	5	1112.43	16211.9	0.3	1.52	3.05	22.97	26.96	145.69
C206	3	9	1	1089.9	20199.5						100.77
C207	0	6	25	705.7	15127.3	0.38	0.5	0.69	1.34	2.97	178.51
C207	1	7	0	–	–	–	–	–	–	–	171.34
C207	2	8	0	–	–	–	–	–	–	–	120.41
C207	3	9	0	–	–	–	–	–	–	–	165.20
C208	0	6	24	864.51	14223.9	1.24	3.08	3.2	3.54	4.11	229.15
C208	1	7	15	804.02	22118.8	1.18	4.42	6.14	10.33	13.53	186.16
C208	2	8	0	–	–	–	–	–	–	–	193.86
C208	3	9	0	–	–	–	–	–	–	–	221.80

Table 14: Detailed results for instance class R1

Instance	k	m	$\|X'_k\|$	Best objective values		Variation in operating costs (%) for decreases in driver inconvenience of					CPU
				cost	inconvenience	1%	5%	10%	25%	50%	(min)
R101	0	8	27	1012.34	22496.5	0.68	2.79	3.76	8.89	34.83	120.89
R101	1	9	46	934.72	18463.7	0	0	5.36	11.51	28.84	121.98
R101	2	10	0	–	–	–	–	–	–	–	113.00
R101	3	11	0	–	–	–	–	–	–	–	122.78
R102	0	8	38	901.17	20106	0.52	0.99	1.25	1.26	8.67	130.25
R102	1	9	21	1129.61	18637.9	0.78	2.97	4.14	10.06	17.51	127.86
R102	2	10	13	1431.74	17690.1	2.29	12.05	19.93	40.86	59.5	123.35
R102	3	11	1	1063.1	26790.5	–	–	–	–	–	–
R103	0	8	37	853.33	14072.6	0.04	0.71	1	1.04	10.94	132.86
R103	1	9	1	1087.59	15627.6	–	–	–	–	–	–
R103	2	10	0	–	–	–	–	–	–	–	136.16
R103	3	11	4	1240.49	12900.7	53.73	73.16	73.67	75.2	77.75	143.02
R104	0	8	29	837.73	9840.88	0.12	1.25	2.24	3.62	14.15	156.26
R104	1	9	2	1024.83	9876.54	1	5	10	25	50	196.47
R104	2	10	1	1126.99	9731.92	–	–	–	–	–	–
R104	3	11	1	1025.07	9600.61	–	–	–	–	–	–
R105	0	8	33	873.81	13961.4	1.62	2.45	5.62	15.1	43.33	137.71
R105	1	9	9	964.32	13442.7	0.92	4.58	9.16	22.89	45.79	129.20
R105	2	10	5	1065.83	14079.4	0.96	1.41	1.98	17.33	43.84	132.12
R105	3	11	15	1206.62	11823.2	7.88	15.35	17.22	20.88	51.09	122.44
R106	0	8	45	861.63	11725.8	0.2	0.2	0.3	4.39	38.12	156.92
R106	1	9	3	1060.74	19317.7	0.89	4.45	8.91	22.27	44.55	125.47
R106	2	10	1	1013.75	21996.1	–	–	–	–	–	–
R106	3	11	0	–	–	–	–	–	–	–	167.16
R107	0	8	43	841.45	12247.8	0.2	0.95	1.53	3.48	6.4	171.27
R107	1	9	4	992.06	9937.25	0.83	4.14	8.27	20.68	41.35	139.26
R107	2	10	4	1006.96	9083.63	0.78	3.9	7.8	21.66	45.11	108.12
R107	3	11	0	–	–	–	–	–	–	–	132.79
R108	0	8	28	829.16	8634.42	1.13	1.47	1.9	3.94	17.42	141.23
R108	1	9	2	943.25	8160.69	1	5	10	25	50	115.47
R108	2	10	1	983.27	8150.21	–	–	–	–	–	–
R108	3	11	0	–	–	–	–	–	–	–	96.16
R109	0	8	47	840.95	11159.7	0.13	0.22	0.41	2.3	12.32	146.42
R109	1	9	7	1039.05	13226.2	0.65	3.27	6.54	14.38	27.04	141.46
R109	2	10	0	–	–	–	–	–	–	–	148.00
R109	3	11	0	–	–	–	–	–	–	–	120.24
R110	0	8	32	835.75	9816.85	0.01	0.06	0.11	0.28	2.39	151.71
R110	1	9	0	–	–	–	–	–	–	–	126.91
R110	2	10	0	–	–	–	–	–	–	–	122.83
R110	3	11	0	–	–	–	–	–	–	–	92.33
R111	0	8	45	845.08	10378.1	0.07	0.1	0.32	2.46	9.58	156.63
R111	1	9	9	1025.27	9962.95	0.09	0.44	0.88	2.2	4.4	166.01
R111	2	10	0	–	–	–	–	–	–	–	151.66
R111	3	11	0	–	–	–	–	–	–	–	147.47
R112	0	8	18	827.39	7800	0.04	0.2	0.41	4.6	7.64	136.06
R112	1	9	0	–	–	–	–	–	–	–	116.75
R112	2	10	1	1027.8	7242.39	–	–	–	–	–	–
R112	3	11	0	–	–	–	–	–	–	–	91.18

Table 15: Detailed results for instance class R2

Instance	k	m	$\|X'_k\|$	Best objective values		Variation in operating costs (%) for decreases in driver inconvenience of					CPU
				cost	inconvenience	1%	5%	10%	25%	50%	(min)
R201	0	6	17	1024.72	21533.2	0.04	0.22	0.43	6.97	18.32	225.25
R201	1	7	40	877.06	17060.7	2.21	3.58	4.71	25.9	41.29	202.78
R201	2	8	29	1115.87	14682.8	3.41	4.28	4.68	5.87	34.55	165.06
R201	3	9	2	1410.5	14954.5	1	5	10	25	50	169.35
R202	0	6	46	833.66	18234.6	2.31	3.01	3.19	7.46	12.09	243.65
R202	1	7	18	846.77	16259.8	0.07	0.66	1.91	4.65	23.69	220.80
R202	2	8	16	1061.67	14722.5	0.02	0.1	0.2	27.49	50.32	213.39
R202	3	9	17	1224.84	13703.8	0.64	4.05	7.42	15.61	51.96	198.15
R203	0	5	25	790.31	14266.6	2.89	8.47	9.28	11.72	31.9	192.43
R203	1	6	22	762.4	13522.9	1.11	1.16	1.84	5.46	26.11	136.34
R203	2	7	5	1035.16	11484.9	0.07	0.36	0.73	1.82	3.64	112.94
R203	3	8	0	–	–	–	–	–	–	–	89.02
R204	0	5	30	742.08	10010.7	0.01	0.07	0.13	4.8	7.54	301.19
R204	1	6	12	779.33	9537.6	0.77	1.4	1.58	13.92	17.26	278.51
R204	2	7	3	995.84	9155.24	0.53	2.65	5.31	13.27	26.54	267.67
R204	3	8	2	939.26	9138.81	1	5	10	25	50	241.14
R205	0	6	41	821.26	13916.5	0.18	0.39	0.79	7.06	16.17	228.59
R205	1	7	29	809.15	13734.4	0.11	0.58	1.74	3.07	13.45	217.51
R205	2	8	15	1003.38	12383.7	1.6	9.47	10.33	22.09	44.9	193.52
R205	3	9	0	–	–	–	–	–	–	–	195.03
R206	0	5	18	806.57	15975	0.08	0.42	6.02	9.74	17.86	318.74
R206	1	6	38	781.51	11674.2	0.06	0.43	1.38	4.51	11.84	286.92
R206	2	7	11	787.96	12861.9	0.08	0.42	0.59	14.4	40.96	220.75
R206	3	8	7	992.97	11873.3	4.03	11.88	13.11	16.81	34.81	209.63
R207	0	5	33	735.18	8572.78	0.29	1.45	5.5	13.8	28.81	325.88
R207	1	6	5	793.08	11655.3	0.8	4	8	20.01	40.01	267.36
R207	2	7	0	–	–	–	–	–	–	–	246.34
R207	3	8	0	–	–	–	–	–	–	–	223.24
R208	0	4	15	727.43	8420.46	0.76	3.8	7.6	11.62	20.54	393.52
R208	1	5	15	722.53	8400	0.31	0.59	0.64	3.49	7.93	336.68
R208	2	6	2	791.51	8776.27	1	5	10	25	50	314.24
R208	3	7	2	814.27	8017.43	1	5	10	25	50	257.19
R209	0	5	31	812.71	10944.5	0.03	0.14	0.27	0.68	22.28	312.26
R209	1	6	17	773.4	12507.6	0.17	0.86	3.58	7.92	23.55	272.83
R209	2	7	9	795.79	10894.8	0.34	1.6	3.17	7.88	35.35	225.41
R209	3	8	2	1094.86	10765.2	1	5	10	25	50	175.03
R210	0	5	16	811.39	12583.8	0.26	1.28	2.56	11.26	50.2	313.11
R210	1	6	8	784.28	12363	0.05	0.27	0.53	5.92	11.26	273.63
R210	2	7	8	1071.37	11598.1	0.14	0.69	1.37	3.44	53.27	252.61
R210	3	8	0	–	–	–	–	–	–	–	260.30
R211	0	5	32	751.26	9600	0.04	0.1	0.21	1.14	3.71	322.31
R211	1	6	0	–	–	–	–	–	–	–	292.06
R211	2	7	2	1032.26	9000	1	5	10	25	50	273.78
R211	3	8	0	–	–	–	–	–	–	–	271.29

Table 16: Detailed results for instance class RC1

Instance	k	m	$\|X'_k\|$	Best objective values		Variation in operating costs (%) for decreases in driver inconvenience of					CPU
				cost	inconvenience	1%	5%	10%	25%	50%	(min)
RC101	0	9	25	1023.25	15897.2	0.4	1.6	1.7	6.79	48.34	118.80
RC101	1	10	16	1077.53	15504.9	4.67	14.84	18.21	28.34	53.24	110.26
RC101	2	11	21	1200.03	14364.2	0.19	0.94	4.49	40.82	58.03	106.35
RC101	3	12	3	1300.94	21237.7	0.07	0.33	0.67	1.67	3.34	111.14
RC102	0	9	25	1010.04	14079.8	0	0.05	0.14	0.28	11.6	126.36
RC102	1	10	10	1251.8	15032.5	3.45	12.78	14.95	23.34	26.43	130.00
RC102	2	11	5	1184.69	17110.8	11.58	14.3	17.39	33.7	60.9	109.66
RC102	3	12	0	–	–	–	–	–	–	–	133.54
RC103	0	9	22	987.36	11230	1.4	1.64	2.15	8.61	33.24	124.18
RC103	1	10	9	1032.11	11147.9	0.52	2.43	3.7	4.82	33.36	123.06
RC103	2	11	14	1072.5	11011.2	0.87	4.2	7.3	15.5	29.06	128.86
RC103	3	12	3	1355.23	10821.7	1.68	8.39	16.78	41.95	83.91	159.46
RC104	0	9	23	985.67	9533.4	0	0.01	0.03	1.99	10.21	112.60
RC104	1	10	6	1068.65	9054.48	0.25	1.24	20.51	33.04	43.19	113.45
RC104	2	11	0	–	–	–	–	–	–	–	110.50
RC104	3	12	1	1283.25	8478.05	–	–	–	–	–	–
RC105	0	9	23	997.91	14576.8	0.02	0.08	0.16	0.41	12.59	110.37
RC105	1	10	2	1207.09	17093.8	1	5	10	25	50	106.20
RC105	2	11	9	1263	12279.3	0.3	1.51	3.03	18.14	58.4	116.88
RC105	3	12	0	–	–	–	–	–	–	–	107.60
RC106	0	9	32	988.01	12880.8	0.02	0.08	0.16	2.6	36.48	115.86
RC106	1	10	15	1041.68	12260.5	4.19	5.49	7.11	11.98	20.09	90.06
RC106	2	11	5	1109.47	14536.6	0.95	4.75	9.49	23.73	47.46	58.61
RC106	3	12	4	1427.03	11425.2	10.22	51.11	87.56	87.59	87.65	64.79
RC107	0	9	20	994.62	11523.2	0.1	0.52	0.84	1.41	12.16	109.33
RC107	1	10	7	1007.14	13481.1	1.58	7.53	13.89	32.96	64.74	89.14
RC107	2	11	5	1251.58	9399.42	0.66	3.13	3.62	7.2	35.64	74.79
RC107	3	12	2	1181.07	12155.8	1	5	10	25	50	43.26
RC108	0	9	18	977.45	9190.71	0.11	0.13	0.17	1.37	4.15	104.05
RC108	1	10	6	1032.19	8123.68	5.07	19.82	21.59	26.92	36.93	73.38
RC108	2	11	3	1047.44	9570.54	0.68	3.42	6.84	17.11	36.75	61.06
RC108	3	12	0	–	–	–	–	–	–	–	55.59

Table 17: Detailed results for instance class RC2

Instance	k	m	$\lvert X'_k \rvert$	Best objective values		Variation in operating costs (%) for decreases in driver inconvenience of					CPU
				cost	inconvenience	1%	5%	10%	25%	50%	(min)
RC201	0	6	33	1033.55	19581.7	0.47	2.2	8.57	14.42	47.69	230.87
RC201	1	7	45	937.46	18504.7	0.58	2	4.89	9.85	17.64	189.81
RC201	2	8	0	–	–	–	–	–	–	–	179.21
RC201	3	9	0	–	–	–	–	–	–	–	156.26
RC202	0	6	27	885.48	15494.4	0.22	0.89	1.33	2	5.4	235.11
RC202	1	7	23	878.27	14888.5	0.02	0.1	0.27	3.04	19.82	219.12
RC202	2	8	4	1282.87	13838.7	0.18	0.9	1.8	4.51	9.02	202.24
RC202	3	9	0	–	–	–	–	–	–	–	202.31
RC203	0	5	22	834.2	13891.9	0.62	0.93	1.09	1.56	6.95	294.22
RC203	1	6	34	805.62	12205.6	1.07	1.08	1.48	2.55	13.33	244.32
RC203	2	7	2	984.91	12858	1	5	10	25	50	208.59
RC203	3	8	6	1123.48	11382.9	5.18	18.75	19.56	21.99	43.82	219.31
RC204	0	5	24	755.7	9006.52	1.07	1.1	1.35	2.54	26.3	295.57
RC204	1	6	0	–	–	–	–	–	–	–	274.02
RC204	2	7	0	–	–	–	–	–	–	–	252.42
RC204	3	8	0	–	–	–	–	–	–	–	272.35
RC205	0	5	28	707.41	11374.5	0.54	1.87	3.08	12.2	36.3	90.26
RC205	1	6	0	–	–	–	–	–	–	–	103.52
RC205	2	7	0	–	–	–	–	–	–	–	95.45
RC205	3	8	0	–	–	–	–	–	–	–	90.07
RC206	0	6	34	828.85	13211.1	0.32	0.79	1.15	6.48	22.14	230.68
RC206	1	7	8	850.21	16405.7	0.27	1.33	2.67	6.18	11.76	212.03
RC206	2	8	8	1014.11	15258.8	4	19.98	39.96	57	78.66	213.72
RC206	3	9	5	1245.33	11908.8	1.17	4.84	4.86	4.93	5.04	201.40
RC207	0	5	23	934.98	12755.2	3.73	19.01	36.91	47.65	58.07	180.42
RC207	1	6	24	827.01	12400.7	0.01	0.06	1.01	6.19	15.49	140.64
RC207	2	7	10	964.97	11262	0.59	2.95	5.89	21.37	47.08	135.17
RC207	3	8	0								111.77
RC208	0	5	22	752.64	9757.76	0.41	2.04	2.82	4.94	15.16	195.68
RC208	1	6	3	996.49	9008.09	0.11	0.55	1.11	5.8	37.2	177.34
RC208	2	7	0	–	–	–	–	–	–	–	169.66
RC208	3	8	2	1032.6	8411.13	1	5	10	25	50	151.77