



Innovative Applications of O.R.

Multi-objective microzone-based vehicle routing for courier companies: From tactical to operational planning [☆]



Jochen Janssens ^{a,*}, Joos Van den Bergh ^a, Kenneth Sörensen ^a, Dirk Cattrysse ^b

^a Universiteit Antwerpen, Faculteit Toegepaste Economische Wetenschappen, Prinsstraat 13, 2000 Antwerp, Belgium

^b Centre for Industrial Management, KU Leuven Celestijnenlaan 300A, 3001 Leuven (Heverlee), Belgium

ARTICLE INFO

Article history:

Received 29 January 2014

Accepted 21 September 2014

Available online 17 October 2014

Keywords:

Metaheuristics

Variable neighborhood tabu search

Workload balancing

Multi-objective optimization

Vehicle routing

Courier companies

ABSTRACT

Distribution companies that serve a very large number of customers, courier companies for example, often partition the geographical region served by a depot into zones. Each zone is assigned to a single vehicle and each vehicle serves a single zone. An alternative approach is to partition the distribution region into smaller *microzones* that are assigned to a preferred vehicle in a so-called tactical plan. The moment the workload in each microzone is known, the microzones can be reassigned to vehicles in such a way that the total distance traveled is minimized, the workload of the different vehicles is balanced, and as many microzones as possible are assigned to their preferred vehicle.

In this paper we model the resulting microzone-based vehicle routing problem as a multi-objective optimization problem and develop a simple yet effective algorithm to solve it. We analyze this algorithm and discuss the results it obtains.

© 2014 Elsevier B.V. All rights reserved.

1. Introduction

Following the example of FedEx, many courier companies now organize their distribution system as a hub-and-spoke model (see FedEx, 2006; Lyster, 2012; Rodrigue, Comtois, & Slack, 2006; Waldheim & Berger, 2008). In this model, vehicles collect parcels and bring them to the nearest depot where they are consolidated and transported, through a chain of hubs, to the depot closest to their final destination. From this final depot, the parcels are brought to their destination by vehicles performing so-called milk-runs.

This paper focuses on the last step: the delivery process from the final depot of a courier company to the end customers. On a daily basis, courier companies face a vehicle routing problem (VRP) that is not only among the largest in the logistics industry, but also highly dynamic. It is not uncommon for a courier company to execute several thousand stops from a single depot, using several hundred vehicles. Additionally, the list of pick-up locations changes throughout the day as customers call to order the pick-up of a parcel, and the list of drop-off locations is only available a very short time before the vehicles leave the depot. In order to deliver these parcels in an efficient manner, route planning software could in theory provide efficient

solutions, but experience shows that such software is rarely used in practice.

The fact that the distribution planning process in many courier companies is not automated can be ascribed to a number of reasons. Firstly, correct data on exact drop locations are often lacking, i.e., customer addresses can be incomplete due to sloppy writing, carelessness, or lack of information on behalf of the sender. Because deciphering these addresses is best done by people with an understanding of the geographical distribution region, this task is often left to the drivers or to other people in the final depot. Couriers therefore lack the possibility to generate the routes algorithmically without considerable investment in a centralized address-scanning-and-correcting service. Secondly, conventional route planning software requires all data to be available and correct *before* a final routing plan can be built. This implies that processes that depend on the final route plan, such as sorting and loading parcels into the correct vehicles, have to wait until all order entries are known. As this would extend the delivery time, one of the main competitive factors for a courier, it is considered an unattractive option.

As the surplus in route efficiency achieved by route planning software would not compensate for the increased delivery time and the loss of flexibility, couriers rely on more stable organizational structures for the distribution of their parcels. A popular technique to deal with this flexibility requirement is to partition the area serviced from the depot into *regions*, usually one per vehicle. All customers that fall within a region are then assigned to the vehicle this region belongs to, which implies that the sorting of parcels per vehicle can start as

[☆] This research has been supported by the Inter-university Attraction Poles Program initiated by the Belgian Science Policy Office (COMEX project).

* Corresponding author. Tel: +32 (0)32654017.

E-mail address: jochen.janssens@uantwerpen.be (J. Janssens).




Fig. 1. The parcel distribution process from the arrival of the parcels, until the departure of the vehicles.

soon as the parcel arrives at the depot. The final routing, i.e., determining the order in which to visit the customers in his¹ region is then left to the driver. This way of working has the additional advantage that a driver has the possibility to get acquainted with his region.

Fig. 1 shows the parcel distribution process from the arrival of the parcels, until the departure of the vehicles. The left image shows the time consumed if all parcels are processed based on their destination address. The right image shows the time consumed if parcels are processed based on zones. The sorting of the parcels can start much earlier, which results in a reduction of time from the arrival of the parcels until the departure of the vehicle.

In case of workload imbalances or vehicle capacity violations, drivers are sometimes allowed to swap individual packages among themselves. For this reason, the physical layout of the distribution depot sometimes resembles that of the geographical delivery region, in the sense that vehicles that serve adjacent routes are stationed close to each other.

To diminish the main disadvantage of zone-based organization – the limited possibility to automate the planning process – several authors have proposed an alternative method, in which the geographical region is divided into smaller zones (“microzones”). In practice, each microzone corresponds to a bin or container into which the parcels can be dropped at the moment they arrive at the final depot.

To ensure stability in the planning, the preferred assignment of microzones to vehicles is fixed in a *tactical plan*. When an estimate of the workload in each microzone becomes available, the microzones are reassigned to the different available vehicles, with a preference for the vehicle they are assigned to, in the tactical plan. The exchange between vehicles no longer happens on the level of individual parcels, but on the microzone level. In this paper, we present an approach to assign microzones to vehicles and determine the order in which the microzones should be traversed. We model this problem as a tri-objective vehicle routing problem and develop an efficient multi-objective metaheuristic to solve it.

This paper is organized as follows: Section 2 gives a short literature overview, followed by the problem delineation in Section 3. Section 4 elaborates on the approach we used, the neighborhoods that are explored and the algorithm we developed. Finally Section 5 presents the results of our experiments. Our findings and conclusions as well as our ideas on future research are posed in Section 6.

2. Literature review

Although the microzone concept has been described by different names (“zones”, “cells”, “core areas”, “flex zones”), the advantages of using microzones as building blocks for routes in large-scale vehicle routing problems have been recognized by several authors (Mourgaya, 2004; Mourgaya & Vanderbeck, 2006, 2007; Zhong, Hall, & Dessouky, 2007). A complete application of this concept requires solving several challenging problems (Expósito Izquierdo, Rossi, & Sevaux, 2013): (1) creating robust microzones, (2) assigning microzones to routes in a tactical plan, (3) estimating the workload in each microzone and each route (Daganzo, 1984), (4) reassigning micro-

zones to vehicles in balanced routes on an operational level, and (5) determining the final route planning within the microzones.

In this overview, we focus on the fourth challenge, the operational assignment of microzones to vehicles and the sequencing of the microzones for each vehicle, taking into account the tactical plan. Previous research has focused on the use of a tactical assignment of zones to routes with the purpose of balancing these routes and allowing drivers to familiarize themselves with certain parts of their tours. Less variability allows drivers to gain experience in their region and improve the relationship with regular customers (see Kunkel & Schwind, 2012).

In this research, we tackle this problem, by using a multi-objective approach. More specifically, we address the problem of transforming unbalanced tactical route plans into balanced and feasible routes, where the imbalance in the tactical plan is due to the stochastic nature of the problem.

Both Zhong et al. (2007) and Mourgaya (2004) try to improve package delivery under uncertainty by building a robust dispatching system, which determines flexible routes, with more familiar drivers in order to provide better service and lower costs. Ouyang (2007) has done research concerning the design of zones for large-scale distribution systems.

The method proposed in Zhong et al. (2007) tries to find a balance between driver familiarity with a zone, this can be accomplished by assigning the same driver to the same service area each day, and the flexibility to optimize dispatch plans on a daily basis, which is accomplished by adjusting the number of vehicle routes, and maximizing driver utilization. They propose using core areas, which are serviced by the same driver every day, in combination with a flex zone and unassigned zones, where customers are serviced by a different driver on daily basis depending on the utilization of the vehicles in the neighborhood of that cell. The authors further develop a constructive heuristic in which cells are added to the partial routes, minimizing time duration of the routes, and taking into account driver familiarity with a cell.

Similar to portfolio management, where combining uncorrelated assets leads to a lower portfolio risk (see Bodie, Kane, & Marcus, 2005), to decrease the imbalance of the work load for each vehicle in the tactical plan, a combination of zones will be taken (as the work load for the zones is not correlated).

Mourgaya and Vanderbeck (2006) deal with a multiperiodic problem. They propose a model for tactical planning that tries to optimise both the work load balance and the regionalization of the routes. They cluster clients based on the distance between them. The clustering is based on well known classification models.

Several papers have tackled the issue of workload balancing between the vehicles of a vehicle routing problem (Liu, Chang, & Huang, 2006; Nikolakopoulou, Kortesis, Synefaki, & Kalfakakou, 2004; Santiago, Melián-Batista, Álvarez, & AngelBello, 2013).

Like many related problems found in the literature, the problem developed in this paper is a multi-objective optimization problem. Unlike single objective optimization, multi-objective optimization leads to a set of solutions known as the Pareto optimal set. This set consists of non-dominated solutions, i.e., solutions for which no other solution exists, that has at least the same performance on all objectives and a better performance on at least one. For a more formal definition of dominance we refer the reader to Jozefowiez, Semet, and Talbi (2008a).

For multi-objective combinatorial optimization problems (MOCO) like the problem solved in this paper, finding the efficient frontier by means of exact methods is computationally very expensive and sometimes even impossible (Caballero, Gonzalez, Guerrero, Molina, & Parolera, 2007; Zhiping & Yuxing, 2010). For that reason, metaheuristic algorithms become an attractive choice (Ehrgott & Gandibleux, 2000). For courier companies, finding a small set of good solutions

¹ To avoid awkward formulations like his/her, we will refer to drivers in the masculine form and decision-makers/planners in the feminine form.

in short computing time is more important than finding the optimal Pareto set.

A common approach to approximate the Pareto set in multi-objective optimization problems is to optimize the weighted sums of objectives for varying weights (Konak, Coit, & Smith, 2006; Marler & Arora, 2004). The main drawback with this varying weights approach is that not all Pareto optimal points can be found when the Pareto front is not convex, i.e., it only finds the solutions on the convex hull of the optimal Pareto set (see Geiger, 2008; Jozefowicz, Semet, & Talbi, 2008b). Nevertheless, with regard to our goal, it remains an attractive technique as it is straightforward, guarantees that the solutions converge towards the Pareto-front, and is computationally efficient (Konak et al., 2006). Another important advantage is that it allows employing single-objective heuristics or metaheuristics. Using different weight combinations will guide the search in different directions, leading to a more diversified result set, which is either the true shape of the Pareto-optimal set, or a convex approximation of it (see Caramia & Dell'Olmo, 2008; Jozefowicz et al., 2008b). The use of a different multi-objective optimization method, or a true multi-objective metaheuristic could be helpful to give a better estimate of the real Pareto set (Geiger, 2008).

The methods defined in Zhong et al. (2007) and Mourgaya and Vanderbeck (2006) use a constructive heuristic to create the tactical and operational plan. The focus, however, is more on the creation of the tactical plan. In this paper we present a variable neighborhood tabu search heuristic which is focused on rebalancing an existing tactical plan. Variable neighborhood search (VNS) is a well-known technique in which different neighborhoods are used systematically (see e.g., Fleszar, Osman, & Hindi, 2009; Hansen & Mladenović, 2001; Imran, Salhi, & Wassan, 2009). Tabu search is a technique used to diversify the search by excluding previously found solutions or previously executed moves (see e.g., Brandão, 2004).

To compare different Pareto optimal sets, we will make use of the hypervolume (see Okabe, Jin, & Sendhoff, 2003). The hypervolume for two or three objectives can be seen respectively as the area or volume between a reference point and the Pareto-front. Zitzler defined it as the size of the dominated space (see Zitzler & Thiele, 1998, 1999). This means that the larger the hypervolume is, the closer the found solution set is to the theoretical optimal solution set.

3. Problem definition

The vehicle routing problem developed in this paper is defined on a graph $G = \{S_o, E\}$. The nodes of this graph represent the *microzones*, i.e., geographically distinct regions that are part of the distribution area. The set of microzones is labeled S . An extra node represents the depot o , from which all vehicles depart and where they all arrive. Together, the microzones and the depot form the set $S_o = S \cup \{o\}$. A set of vehicles V is also given, where $|V| = R$ is the number of available vehicles. A tactical plan is given, in which each microzone is assigned to exactly one vehicle. This is represented by the parameter \bar{y}_{ik} which is 1 if microzone $i \in S$ is assigned to vehicle $k \in V$ in the tactical plan and 0 otherwise. The *working time* c_i in a microzone i is a number that represents an estimate of the time to deliver all parcels in that microzone and includes the (estimated) intra-zone driving times. The edges in the edge set E of graph G represent the travel between zones. The time a vehicle takes to drive between microzones i and j is represented as c_{ij} . We use the travel time between the center points of micro zones as an estimate for this driving distance. It is a rough estimate but it will stay constant as long as the same microzones are considered. The real driving distance can vary a lot based on the distribution of the points inside a microzone, and based on the position where the microzone is entered. We make the assumption that individual addresses of customers in a microzone are not known, and are volatile over time, and therefore the distance between the nearest

points of the microzones is not known and is hence not a suitable measure for the driving distance between microzones.

The aim of the problem is to find an *operational route plan*, i.e., an assignment of microzones to vehicles, as well as an order in which each vehicle should visit the microzones assigned to it. An operational route plan is determined by decision variables y_{ik} which take a value 1 if microzone i is assigned to vehicle k and 0 otherwise, and x_{ij} , which is 1 if a vehicle travels directly between microzones i and j and 0 otherwise.

This problem has three objectives:

- **Objective 1—Minimize total transportation cost:** measured by total travel time.

$$\min f_1 = \sum_{i \in S_o} \sum_{j \in S_o} c_{ij} x_{ij} \quad (1)$$

- **Objective 2—Minimize deviation from tactical plan:** measured by counting the regions that are assigned to a different vehicle.

$$\min f_2 = \frac{1}{2} \sum_{i \in S} \sum_{k \in V} |y_{ik} - \bar{y}_{ik}| \quad (2)$$

- **Objective 3—Minimize workload imbalance:** measured by the square root of the sum of the squared deviations from the average route time.

$$\min f_3 = \sqrt{\frac{1}{R} \sum_{k \in V} \left(l_k - \frac{1}{R} \sum_{k \in V} l_k \right)^2} \quad (3)$$

where l_k is the workload of vehicle k and is defined as the sum of the (estimated) inter-zone driving times plus the working time for the assigned microzones.

$$l_k = \sum_{i \in S_o} \sum_{j \in S_o} y_{ik} x_{ij} c_{ij} + \sum_{i \in S} y_{ik} c_i, \quad \forall k \in V \quad (4)$$

For each vehicle k , its total working time in the operational route plan is limited to a pre-specified value.

$$\forall k \in V : l_k \leq C \quad (5)$$

For reasons of simplicity, we assume in this paper that both the preferred working time and the maximum allowed working time are the same for all vehicles (drivers), something which is definitely not the case in real life. However, the model can be easily adapted to such situations. The constraint in Eq. (5) can be changed to have a different maximal working time per vehicle by replacing the global constant C with a maximal allowed working time per vehicle C_k . Objective 3, the minimization of the workload imbalance, can be modeled in different ways when working times are variable. One way is to minimize the workload imbalance separately for different values of expected workload (e.g., 8h and 4h). Another way would be to normalize the data, i.e., divide l_k by C_k and minimize the sum or squared sum of deviations of l_k/C_k from 1.

Another assumption of our model is that the vehicles are not capacity-constrained. This is a realistic assumption, given that many courier companies only distribute relatively small parcels, letters, etc.

Except for the two additional objectives, the problem developed in this paper is a distance-constrained vehicle routing problem (without capacity constraint). In the next section, we develop an efficient algorithm for this problem.

4. Methodology

In this section, we develop a simple yet effective metaheuristic to solve the multi-objective vehicle routing problem described in Section 3. The algorithm presented in this paper is able to produce a

set of non-dominated solutions, and it is left to the decision maker to pick the final solution she decides to implement.

Our algorithm constructs a new routing plan for the multi-objective VRP by iteratively applying a multi-neighborhood tabu search heuristic to a weighted exponential sum of the objective functions discussed in Section 3. This heuristic uses a set of problem-specific operators to efficiently search for high-quality solutions.

To efficiently evaluate intermediate solutions generated using the tabu search heuristic, a weighted exponential sum of normalized values of the tree objective functions is used:

$$f(\mathbf{x}) = \left[w_1 f_1^n(\mathbf{x})^P + w_2 f_2^n(\mathbf{x})^P + w_3 f_3^n(\mathbf{x})^P \right]^{\frac{1}{P}} \quad (6)$$

where $w_1 + w_2 + w_3 = 1$.

By varying the weights of the weighted exponential sum, the tabu search heuristic is forced to explore different areas of the search space, resulting in diverse solutions. Moreover, the value of P allows to adjust the curvature of the objective function, which results in a better fit with the Pareto front, and which allows to capture efficient points on the non-convex part of the Pareto front (Messac, Renaud, Sundararaj, & Tappeta, 2000b; Messac, Puemi-Sukam, & Melachrinoudis, 2000a; Marler & Arora, 2004). This type of objective function is easy to implement, and allows to prioritize certain types of solutions by changing the values of the weights.

If the value of P is not large enough, however, not all Pareto-optimal points can be found when the Pareto front is not convex. Nevertheless, with regard to our goal, it remains an attractive technique as it is straightforward and computationally efficient.

4.1. Neighborhoods

Like a majority of efficient algorithms to solve vehicle routing problems, our variable neighborhood tabu search uses several different neighborhoods to escape from local optima and overcome the myopic behavior of a single neighborhood.

To speed up the search process, the moves are restricted to operate on so-called “natural neighbors”. Natural neighbors are used to limit the search space by excluding moves that could never yield any positive results. It should be clear that microzones that are not near to each other (e.g. on the opposite side of the area that is serviced) can be changed, but that said change will give an undesirable result, as it is not efficient at all. By limiting the search space to only natural neighbors, we might exclude the optimal solution, but will reduce the time for exploring the search space and finding a very good solution. Natural neighbors are defined as nodes that are connected in a Delaunay triangulation of the original fully connected graph on which the problem is defined. The Delaunay triangulation was first defined by Delaunay (1934) as a triangulation of the points in the system, in such a way that the circumcircle of any triangle does not contain any other points of the system.

We consider only euclidean instances in this paper. If other instances are to be considered, a projection can be made into euclidean space and the Delaunay triangulation can be performed on that projection. If no projection can or should be made, another method to find nearest neighbors is to be used. One possibility is to use the distance matrix to select n nearest neighbors to evaluate.

An alternative procedure to restrict the search space of the neighborhood operators would be to limit the moves to pairs of microzones that are not more than a certain maximum distance apart. The main drawback of this approach is that this maximum distance needs to be determined, an arbitrary decision which may be difficult when the nodes are not evenly spaced. In such situations, microzones corresponding to remote areas may feature in much fewer moves than microzones in busy urban regions. The Delaunay triangulation does not have this drawback.




Fig. 2. Delaunay triangulation to find natural neighbors.

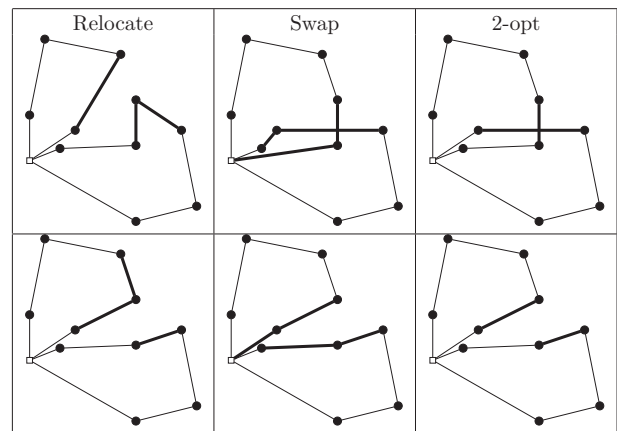


Fig. 3. Examples of the different moves used in the variable-neighborhood tabu search algorithm.

Krasnogor, Moscato, and Norman (1995) and Lau and Shue (2001) are examples of papers in which the Delaunay triangulation is used in the same manner as in this paper, i.e. to restrict the search space of the neighborhood operators. For an example of a Delaunay triangulation see Fig. 2.

The different moves that are used to explore the solution space are illustrated in Fig. 3:

- The **2-opt** move cuts two arcs, each arc between two natural neighbors, and reconnects the solution.
- The **swap** move exchanges two microzones that are natural neighbors and in different tours.
- The **relocate** move transfers a microzone from one tour to the position immediately after one of its natural neighbors in another tour. Natural neighbors inside the same tour are not regarded.

4.2. A variable neighborhood tabu search algorithm

A general outline of our algorithm is presented in Algorithm 1. As mentioned, the aim of this paper is not to develop a state-of-the-art multi-objective metaheuristic, but to provide a simple and

transparent heuristic that can be used to illustrate the multi-objective vehicle routing problem.

Algorithm 1: Diversification and storage sequence.

Let $w_i \in [0, 1]$, be the weight for objective i , with step size $\Delta t = 0.1$;
 Initialization - calculate an initial solution given the assignment of microzones to vehicles of the tactical plan;
for all combinations of $w_i \mid \sum_{i=1}^n w_i = 1$ **do**
 Do variable neighborhood tabu search;
 Store solution in the solution archive A ;
 Filter A ;
 Return A ;

The tactical plan is used as a starting point for the multi-objective variable neighborhood tabu search in the sense that the initial solution is determined by assigning all microzones to their preferred vehicle and then calculating a traveling salesperson solution per vehicle.

For each combination of weights w_i which sum up to 1, and which vary in steps of 0.1, an iteration of the variable neighborhood tabu search is executed, starting from the initial solution. The varying weight combinations will guide the search in different directions, leading to a more diversified approximation of the Pareto set.

The result retrieved from the variable neighborhood tabu search routine is stored in a solution archive A , which is kept in memory until the end of the program. This solution archive is filtered before it is returned to the decision maker. The filtering step will remove all solutions from the archive that are dominated by another solution, only accepting non-dominated solutions.

Moves are executed using a “best-improvement” strategy. The tabu list is a list of microzones on which moves have been executed in the last t iterations of the variable neighborhood tabu search procedure, and each neighborhood has its own tabu list. In a limited pilot experiment, it was established that a value of $t = 6$ robustly resulted in high-quality solutions.

The variable neighborhood tabu search algorithm (see Algorithm 2) is executed for every weight combination that sums up to one. A solution vector V is created, and the algorithm keeps track of the neighborhood it is exploring. Depending on the neighborhood, a certain move is executed, and the respective tabu list is

Algorithm 2: Variable neighborhood tabu search.

Let V be the solution vector ;
 Let N be the neighborhood tracker, initiate $N = 0$;
while $N < 3$ **do**
 if $N = 0$ **then**
 Do 2-opt and update tabu list ;
 else
 if $N = 1$ **then**
 Do swap and update tabu list ;
 else
 if $N = 2$ **then**
 Do relocate and update tabu list ;
 if no improvement found then
 Increase N
 else
 Store the current solution in V

Find best solution in V using weighted exponential sum of the objective values using w_i ;

updated. If no improvement can be found, the algorithm moves on to the next neighborhood. Otherwise, all solutions are evaluated using the weighted exponential sum (see Eq. (7)) of the objective functions, and the best solution is used as the starting solution for the next iteration.

$$f(\mathbf{x}) = \left[w_1 f_1^n(\mathbf{x})^2 + w_2 f_2^n(\mathbf{x})^2 + w_3 f_3^n(\mathbf{x})^2 \right]^{\frac{1}{2}} \quad (7)$$

As it is not necessary, nor desirable to spend time capturing all Pareto optimal solutions, P is set to 2. This value for P gives a curvature of the objective function, and can enable the algorithm to find more points on the Pareto front.

5. Computational results

In this section, a variable neighborhood tabu search algorithm to solve the multi-objective vehicle routing problem, and balance a tactical plan with microzones, is proposed. The experiment is designed to investigate if there is a correlation between any of the three objective functions, and to attempt to find an explanation for this correlation. Further, the results of the experiment are used to investigate if the amount of deviation in workload from the expected workload in a microzone has a large impact on the Pareto front found by the algorithm. A second experiment will show that the variable neighborhood tabu search algorithm is robust with respect to its initial solution.

5.1. Test instances

As stated before, the model assumes that a tactical plan is given, with known work times for the microzones and known microzone-vehicle allocations. To test the algorithm described in this paper extensively, a set of realistic test instances is generated. The test instances comprise: zones that are assigned to vehicles in a tactical plan, an initial routing for the tactical plan, a limit in amount of work a vehicle (driver) can perform, the amount of work in the microzones and the coordinates of the microzones in the distribution zone.

To construct the instances, microzones are generated in randomly distributed locations on a square canvas (of size 1000 by 1000). To replicate city centers and industrial areas, or other areas with a denser customer base than rural areas, the instances are generated in such a way that 90 percent of the microzones are concentrated in clusters. The locations of the remaining 10 percent are uniformly generated.

To mimic the typical layout of a distribution area, microzone clusters are generated in such a way that the center of the cluster is denser than the boundaries. This models the fact that houses in city centers are built closer to each other, and apartment blocks are more common. To this end, the distance between a microzone and the center of the cluster it belongs to is generated as an exponentially distributed variate with mean λ^{-1} , a parameter which can vary across different clusters.

To obtain a tactical plan from this generated instance, the open-source VRP solver, VRPH (Groer, 2008), is used to build routes using the microzones as customers and the distances between those microzones as travel times. These routes determine the tactical assignment of microzones to drivers.

In the generation of the tactical plan, the microzones are of equal workload. This is a realistic assumption as the generation of microzones for use in a tactical plan in a previous step in the “tactical plan-operational plan” approach, would mean that they are robust and as close as possible to equal workloads for known (historical) pick-up and delivery locations. Of course, for the daily operational plan the workload in each microzone can vary. For this reason, the operational demand in each microzone is generated as a single draw from a triangular distribution. The median of this triangular distribution is set equal to the tactical demand of the microzone, its minimum and

Table 1
Parameter values.

Parameter	Min	Max	Step size
Number of microzones (N)	100	200	50
Number of clusters (n)	0	3	1
Maximum workload deviation (w)	0	30	5
Maximum distance (M)	1000	1000	0
Lambda (L)	20	20	0
Case number (case)	0	2	1

maximum are set to $(100 - w)$ percent and $(100 + w)$ percent of the tactical demand respectively, where w is allowed to vary between 0 and 30.

The instances used in this paper, which can be found on <http://antor.ua.ac.be/ZVRP>, use the following naming syntax:

ZVRP_<unique name>_N<number of microzones>_n<number of clusters>_w<workload deviation>_M<maximum distance>_L<lambda>_case_<case number>.csv

ZVRP refers to “zone-based vehicle routing problem”. The unique name is used to differentiate between instances where all other variables have identical values. The maximum distance refers to the size of the problem instance. The values we have used for the different parameters in our generated instances can be found in Table 1. An example of an instance name generated with a particular combination of values in the table and the unique name “full_test” would be:

ZVRP_full_test_N100_n0_w0_M1000_L20_case_0.csv

5.2. Results

The variable neighborhood tabu search algorithm was run on 252 test instances, with number of zones varying from 100 to 200 in steps of 50, the number of clusters from 0 to 3 in steps of 1 and the maximum deviation in the work per zone from 0 to 30 percent in steps of 5 percent. Three different instances of all combinations of the parameters are created. Since this is a novel problem, no algorithms are available to compare the results. We therefore discuss the results of our algorithm, which provide some insight into the structure of the problem. The average computational time needed for one instance is 4 minutes.

An example of a Pareto-front obtained using our method can be found in Fig. 4. For an increase in maximum workload deviation there is no visible influence on the shape of the Pareto-front. In Table 2, we can see that the overall average percentage difference between the best hypervolume for an instance and the other hypervolumes for that instance is only 3 percent, with a maximum of only 5 percent. This shows that there are no big differences between the different Pareto sets we find for each instance. We can conclude that the maximum deviation in workload (w_x for $x \in [0, 30]$, with step size $\Delta_t = 5$) has nearly no impact on the final Pareto front obtained. Two possible explanations for this behavior come to mind. The first one being that the workload imbalance in microzones gets resolved by a good combination of microzones (which would prove the usefulness of combining microzones for balancing workload between vehicles). The other possible explanation, that the amount of work inside a microzone is very

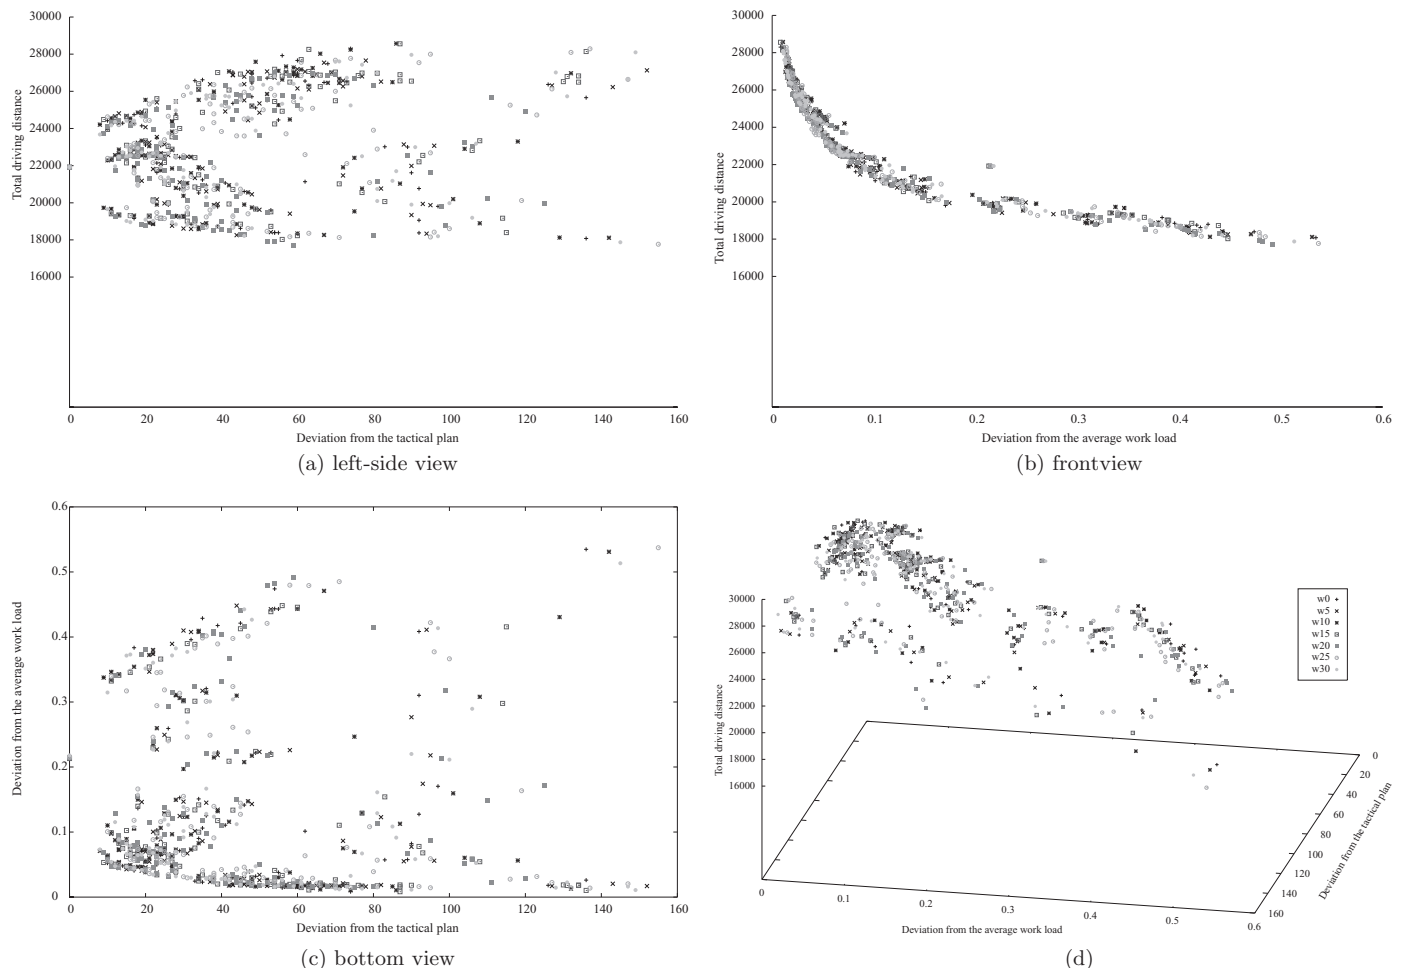


Fig. 4. Pareto-front for an instance with 200 microzones and 1 cluster, for different maximum deviations of the amount of work in the microzones. It clearly shows the relation between driving distance and maximum deviation from the average workload.

Table 2

Percentage difference in hypervolume for the same problem instances with different deviations in the amount of work in the zones.

	Average percent difference	Maximum percent difference	Standard deviation
hypervol_N100_n0_case_0	1.39	4.56	1.47
hypervol_N100_n0_case_1	1.40	5.19	1.80
hypervol_N100_n0_case_2	0.655	1.48	0.546
hypervol_N100_n1_case_0	1.62	3.47	1.36
hypervol_N100_n1_case_1	2.81	4.33	1.48
hypervol_N100_n1_case_2	1.57	2.75	0.947
hypervol_N100_n2_case_0	1.27	2.30	0.771
hypervol_N100_n2_case_1	1.13	1.92	0.654
hypervol_N100_n2_case_2	0.680	1.60	0.601
hypervol_N100_n3_case_0	2.14	4.11	1.55
hypervol_N100_n3_case_1	1.03	1.60	0.575
hypervol_N100_n3_case_2	0.697	1.13	0.402
hypervol_N150_n0_case_0	0.529	1.09	0.414
hypervol_N150_n0_case_1	0.837	2.09	0.812
hypervol_N150_n0_case_2	1.44	2.30	0.730
hypervol_N150_n1_case_0	3.20	5.27	1.99
hypervol_N150_n1_case_1	0.890	2.37	0.763
hypervol_N150_n1_case_2	1.24	3.34	1.14
hypervol_N150_n2_case_0	0.775	1.12	0.374
hypervol_N150_n2_case_1	1.09	1.74	0.605
hypervol_N150_n2_case_2	1.30	3.27	1.08
hypervol_N150_n3_case_0	2.80	4.82	2.12
hypervol_N150_n3_case_1	1.21	1.91	0.669
hypervol_N150_n3_case_2	0.872	3.02	1.02
hypervol_N200_n0_case_0	1.87	3.89	1.27
hypervol_N200_n0_case_1	0.877	2.82	0.949
hypervol_N200_n0_case_2	0.968	1.61	0.521
hypervol_N200_n1_case_0	1.14	1.66	0.574
hypervol_N200_n1_case_1	1.18	2.01	0.697
hypervol_N200_n1_case_2	1.16	3.48	1.24
hypervol_N200_n2_case_0	0.731	1.44	0.546
hypervol_N200_n2_case_1	2.51	3.63	1.18
hypervol_N200_n2_case_2	0.531	1.75	0.632
hypervol_N200_n3_case_0	0.988	1.66	0.521
hypervol_N200_n3_case_1	0.321	0.798	0.316
hypervol_N200_n3_case_2	1.54	3.90	1.38




Fig. 5. The Pareto front for different initial solutions for a single instance with 200 microzones, 1 cluster and a maximum work load deviation of 15 percent.

Table 3

Pearson's product-moment correlation for driving distance and deviation of average worktime for instance with 200 microzones, 1 cluster and maximum workload deviation of 20 percent.

Values:	$t = -22.559, df = 108, p\text{-value} < 2.2e-16$
Alternative hypothesis:	True correlation is not equal to 0
95 percent confidence interval:	-0.936; -0.869
Sample estimates:	Correlation -0.908

Table 4

Percentage difference in hypervolume with the best found Pareto set for different start solutions and one tactical plan.

Average difference	2.32
Maximum difference	15.0
Standard deviation	2.41

small in comparison to the time spent driving between microzones and that the impact of different workloads inside the microzones

is therefore very small, can be ruled out. The test instances have been generated to ensure that, on average, more than 40 percent of a vehicle's time is spent inside the microzones.

The initial solution used by the algorithm has little influence on the Pareto front found. In Fig. 5 the Pareto-fronts of the same problem instance, with different start solution and one tactical plan, are shown. The figure presents a few of the Pareto-fronts found for the different start solutions, and demonstrates that they are generally very close to one another. In Table 4, the results are compared in terms of accuracy by using the hypervolume of the found Pareto sets. Eighty four different initial solutions have been tested. The average deviation from the best found Pareto set is 2.3 percent, showing that the method is relatively robust with respect to the initial solution.

Statistical tests on the data of the Pareto-fronts found by the algorithm reveal that there is a clear (inverse) correlation between the total driving distance and the deviation of the average workload (see Table 3 for an example). This result is intuitive as the optimal solution for any instance with a vehicle that has no time or capacity constraint is the optimal TSP tour. If we have a fixed number of vehicles to balance the workload this tour should be split up in parts of roughly the

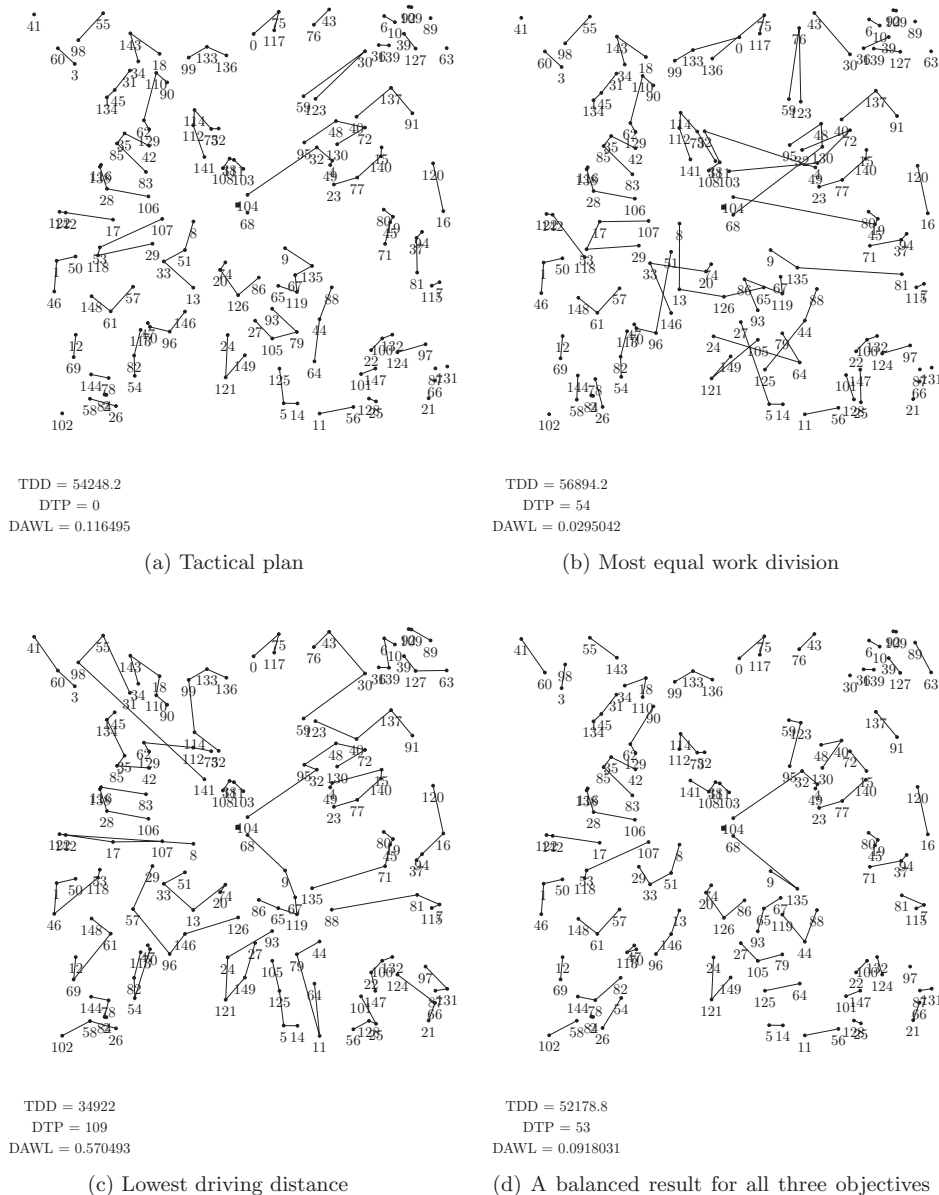


Fig. 6. Example solutions found for the instance with 150 zones, zero city centers and a deviation of workload in the zones of maximum 15 percent of the initial workload.

same amount of work, and be assigned to the vehicles. Next, the travel from and to the depot should be added, which makes the total driving distance longer. The correlation between the deviation of the average workload and the deviation from the tactical plan is not always statistically significant, but logic dictates that there should be some correlation. As we try to minimize the deviation of the average workload, it is obvious that a deviation of the operational solution from the tactical plan is inevitable. In deviating from this tactical plan, the driving distance becomes larger. If the total driving distance is minimized, the routes get more and more unbalanced, and we deviate to a greater extent from the tactical plan.

Because of the technique used to deal with the multi-objective property of the problem, it is not clear if this gives the true shape of the Pareto-optimal set, or just a convex approximation. The use of a different multi-objective optimization method, or a true multi-objective metaheuristic could be helpful to give a better estimate of the real Pareto set.

Taking a look at the routes found by the algorithm proposed in this paper (see e.g., Fig. 6), however, it is clear that the decision maker should keep an eye on the routes she selects, and has to make a keen assessment of the routes proposed. If she pushes the equality of the workload to the lowest point possible, and by doing so relaxes the constraint on driving distance, she has a high probability to find suboptimal routes. The algorithm will try to balance the routes in such a way that it does not regard the driving distance and hence will produce detours, just for the sake of balancing the routes. A healthy balance of the three objective functions should be maintained at all times. It is up to the decision maker to keep this balance, and selecting the optimal route configuration for her needs.

6. Conclusions and future research

In this paper, we have described a multi-objective vehicle routing problem faced by several courier companies, and developed an algorithm to solve it. Afterwards, some results were shown and analyzed. We demonstrated that a variable neighborhood tabu search approach could be used to balance a tactical plan with microzones, but the decision maker must pay attention to the proposed results and assess its usability, and applicability on a case by case basis.

We find a clear (negative) correlation between total driving distance and deviation of average worktime. Between total driving distance and deviation of the tactical plan we sometimes find a significant correlation, but this is not the case for all instances. The same holds for the correlation between deviation of the average worktime and deviation from the tactical plan.

From the results of the experiment, we can conclude that the maximum deviation in workload (w_x for $x \in [0, 30]$, with step size $\Delta_t = 5$) has nearly no impact on the final Pareto front obtained. Furthermore, we have shown that our approach is robust with respect to the initial solution.

A primary point of improvement is the speed of the algorithm. Currently all possible moves with nodes from neighboring tours, tours which contain natural neighbors for the node that is checked, are considered. It might, for example, be possible to distinguish moves that have a higher probability to yield a good result from the ones that are rather unlikely to improve the solution, and decrease the computational time needed to evaluate the entire neighborhood.

A second point in future research is testing different ways to deal with the multi-objective character of the problem. In the current algorithm, a scalarization technique is used, which has the downside of not finding convex-dominated alternatives. In the future more appropriate, but necessarily more elaborate alternatives to this technique should be considered, to make sure the Pareto set is not containing any non-convex parts, and to compare the results of that technique with the results of the scalarization technique. In Caballero et al. (2007)

all visited points are checked for inclusion in the list of Pareto optimal points. Checking each solution visited can let the search sensibly increase the number of efficient solutions found while reducing the number of tabu searches and the computational effort. The insertion of this approach should be considered in future research.

Related work can focus on the development of an algorithm to estimate the amount of work in a microzone. Also, a method to divide the distribution area into robust microzones of equal work load should be developed. These algorithms will provide the possibility to turn data of courier companies into tactical plans with robust microzones, and routes of equal work load. At this point it will be possible to investigate if the creation of the tactical plan can be more effective and accurate.

The results obtained using different maximum working time deviations demonstrate that the algorithm is, to some extent, capable of combining zones in such a way that each vehicle has approximately the same amount of work. An interesting topic of future research could be to investigate whether it is advisable to generate *robust* microzones, i.e., microzones that have approximately the same amount of work every day.

The final step will be the development of an algorithm for the detailed routing within a microzone. Some research has been done on this phase, but an approach that integrates this work into a model like the one proposed in this paper, will certainly be useful.

References

- Bodie, Z., Kane, A., & Marcus, A. J. (2005). *Investments* (International ed.). McGraw-Hill.
- Brandão, J. (2004). A tabu search algorithm for the open vehicle routing problem. *European Journal of Operational Research*, 157, 552–564.
- Caballero, R., Gonzalez, M., Guerrero, F., Molina, J., & Paralera, C. (2007). Solving a multi-objective location routing problem with a metaheuristic based on tabu search. Application to a real case in Andalusia. *European Journal of Operational Research*, 177, 1751–1763.
- Caramia, M., & Dell'Olmo, P. (2008). *Multi-objective Management in Freight Logistics*. Springer London, Limited.
- Daganzo, C. F. (1984). The distance traveled to visit n points with a maximum of c stops per vehicle: An analytic model and an application. *Transportation Science*, 18, 331–350.
- Delanay, B. N. (1934). Sur la sphère vide. *Bulletin of Academy of Sciences of the USSR*, 6, 793–800.
- Ehrgott, M., & Gandibleux, X. (2000). An annotated bibliography of multi-objective combinatorial optimization. *OR Spectrum*, 22, 425–460.
- Expósito Izquierdo, C., Rossi, A., & Sevaux, M. (2013). Modeling and solving the clustered capacitated vehicle routing problem. In *Proceedings of the 14th EU/ME workshop* (pp. 110–115).
- FedEx (2006). *FedEx: Powering Global Access*. http://images.fedex.com/us/about/today/access/GreaterAccessChange_global.pdf Accessed 13.03.13.
- Fleszar, K., Osman, I. H., & Hindi, K. S. (2009). A variable neighbourhood search algorithm for the open vehicle routing problem. *European Journal of Operational Research*, 195, 803–809.
- Geiger, M. J. (2008). *Genetic algorithms for multiple objective vehicle routing*. 1 (pp. 349–354). Program Operational Ciencia, Tecnologia, Inovação do Quadro Comunitário de Apoio III de Fundação para a Ciência e Tecnologia. <http://arxiv.org/abs/0809.0416>.
- Groer, C. (2008). Parallel and serial algorithms for vehicle routing problems (Ph.D. thesis). University of Maryland <http://hdl.handle.net/1903/9011>.
- Hansen, P., & Mladenović, N. (2001). Variable neighborhood search: Principles and applications. *European Journal of Operational Research*, 130, 449–467.
- Imran, A., Salhi, S., & Wassen, N. A. (2009). A variable neighborhood-based heuristic for the heterogeneous fleet vehicle routing problem. *European Journal of Operational Research*, 197, 509–518.
- Jozefowicz, N., Semet, F., & Talbi, E. (2008a). From single-objective to multi-objective vehicle routing problems: Motivations, case studies, and methods. In B. Golden, S. Raghavan, & E. Wasil (Eds.), *The vehicle routing problem: Latest advances and new challenges* (pp. 445–471). Springer USA Operations Research/Computer Science Interfaces Series: Vol. 43.
- Jozefowicz, N., Semet, F., & Talbi, E. (2008b). Multi-objective vehicle routing problems. *European Journal of Operational Research*, 189, 293–309.
- Konak, A., Coit, D., & Smith, A. (2006). Multi-objective optimization using genetic algorithms: A tutorial. *Reliability Engineering and System Safety*, 91, 992–1007.
- Krasnogor, N., Moscato, P., & Norman, M. (1995). A new hybrid heuristic for large geometric traveling salesman problems based on the delaunay triangulation. *Anales del XXVII simposio Brasileiro de Pesquisa Operacional*.
- Kunkel, M., & Schwind, M. (2012). Vehicle routing with driver learning for real world CEP problems. In *45th Hawaii international conference on system science (HICSS)* (pp. 1315–1322).
- Lau, S., & Shue, L. (2001). Solving travelling salesman problems with an intelligent search approach. *Asia-Pacific Journal of Operational Research*, 18, 77–87.

- Liu, C.-M., Chang, T.-C., & Huang, L.-F. (2006). Multi-objective heuristics for the vehicle routing problem. *International Journal of Operations Research*, 3, 173–181.
- Lyster, C. (2012). Learning from fedex: Lessons for the city. *Journal of Landscape Architecture*, 7, 54–67.
- Marler, R., & Arora, J. (2004). Survey of multi-objective optimization methods for engineering. *Structural Multidisciplinary Optimization*, 26, 369–395.
- Messac, A., Puemi-Sukam, C., & Melachrinoudis, E. (2000a). Aggregate objective functions and Pareto frontiers: Required relationships and practical implications. *Optimization and Engineering*, 1, 171–188.
- Messac, A., Renaud, J. E., Sundararaj, G. J., & Tappeta, R. V. (2000b). Ability of objective functions to generate points on nonconvex pareto frontiers. *AIAA Journal*, 38, 1084–1091.
- Mourgaya, M. (2004). The periodic vehicle routing problem: Planning before routing (Ph.D. thesis). Laboratoire de Mathématiques Appliquées de Bordeaux, Université de Bordeaux 1 Bordeaux, France.
- Mourgaya, M., & Vanderbeck, F. (2006). Periodic vehicle routing problem: Classification and heuristic. *RAIRO Operational Research*, 40, 169–194.
- Mourgaya, M., & Vanderbeck, F. (2007). Column generation based heuristic for tactical planning in multi period vehicle routing. *European Journal of Operational Research*, 183, 1028–1041.
- Nikolakopoulou, G., Kortesis, S., Synefaki, A., & Kalfakakou, R. (2004). Solving a vehicle routing problem by balancing the vehicles time utilization. *European Journal of Operational Research*, 152, 520–527.
- Okabe, T., Jin, Y., & Sendhoff, B. (2003). A critical survey of performance indices for multi-objective optimisation. In *The 2003 congress on evolutionary computation (CEC'03)* (Vol. 2, pp. 878–885).
- Ouyang, Y. (2007). Design of vehicle routing zones for large-scale distribution systems. *Transportation Research Part B: Methodological*, 41, 1079–1093.
- Rodrigue, J.-P., Comtois, C., & Slack, B. (2006). *The Geography of Transport Systems*. New York: Routledge, Taylor & Francis.
- Santiago, A., Melián-Batista, B., Álvarez, A., & AngelBello, F. (2013). Effect of the initial solutions to balance routes in vehicle routing problem with time windows. In I. Rojas, G. Joya, & J. Cabestany (Eds.), *Advances in computational intelligence* (pp. 46–52). Springer Berlin/Heidelberg Lecture Notes in Computer Science: Vol. 7903.
- Waldheim, C., & Berger, A. (2008). Logistics landscape. *Landscape Journal*, 27, 219–246.
- Zhiping, Q., & Yuxing, Z. (2010). Parametric optimization design of aircraft based on hybrid parallel multi-objective tabu search algorithm. *Chinese Journal of Aeronautics*, 23, 430–437.
- Zhong, H., Hall, R., & Dessouky, M. (2007). Territory planning and vehicle dispatching with driver learning. *Transportation Science*, 41, 74–89.
- Zitzler, E., & Thiele, L. (1998). Multiobjective evolutionary algorithms: A comparative case study and the strength pareto approach. *Parallel Problem Solving from Nature*, 1498, 292–301.
- Zitzler, E., & Thiele, L. (1999). Multiobjective evolutionary algorithms: A comparative case study and the strength pareto approach. *IEEE Transactions on Evolutionary Computation*, 3, 257–271.