Multi-objective microzone-based vehicle routing for courier companies: from tactical to operational planning

Jochen Janssens¹, Joos Van den Bergh¹, Kenneth Sörensen¹, and Dirk Cattrysse²

¹Universiteit Antwerpen, Faculteit Toegepaste Economische Wetenschappen, Prinsstraat 13, 2000 Antwerp, Belgium,
jochen.janssens@ua.ac.be

²Centre for Industrial Management, KU Leuven Celestijnenlaan 300A, 3001 Leuven (Heverlee), Belgium

Jan 2014

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. When 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 that it obtains.

Key words: Variable neighborhood tabu search, workload balancing, meta-heuristics, multi-objective optimization, vehicle routing, courier companies.

1 Introduction

Following the example of FedEx, many courier companies now organize their distribution system as a hub-and-spoke model (see FedEx (2006), Rodrigue et al. (2006), Waldheim and Berger (2008) and Lyster (2012)). 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 is 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 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

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

driver. This way of working has the additional advantage that a driver has the possibility to get acquainted with his region.

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

Figure 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. 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 and Vanderbeck, 2006, 2007; Zhong et al., 2007). A complete application of this concept requires solving several challenging problems (Expósito Izquierdo et al., 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 microzones 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 and 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 et al. (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) approach a multiperiodic problem by using a tactical plan optimizing both the work load balance as 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 et al., 2006; Nikolakopoulou et al., 2004; Santiago et al., 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 et al. (2008b).

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 et al., 2007; Zhiping and Yuxing, 2010). For that reason, metaheuristic algorithms become an attractive choice (Ehrgott and Gandibleux, 2000). For courier companies, finding a small set of good solutions 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 (convex) weighted sums of objectives for varying weights (Konak et al., 2006; Marler and 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) and Jozefowiez et al. (2008a)). 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 meta-heuristics. 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 and Dell’Olmo (2008) and Jozefowiez et al. (2008a)). The use of a different multi-objective optimization method, or a true multi-objective meta-heuristic 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 will present a heuristic where the focus is more on rebalancing an existing tactical plan.

To compare different Pareto optimal sets, we will make use of the hypervolume (see Okabe et al. (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 and Thiele (1998, 1999)). This means that the larger the hypervolume, 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.

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.

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

where l_k is the workload of vehicle k and is defined as $l_k = \sum_{i \in S_o} \sum_{j \in S_o} y_{ik} x_{ij} c_{ij} + \sum_{i \in S_o} y_{ik} c_i \quad \forall k \in V$.

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

$$\forall k \in V : l_k \leq C \quad (4)$$

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 Equation 4 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:

$$f(x) = \left[w_1 f_1^n(x)^P + w_2 f_2^n(x)^P + w_3 f_3^n(x)^P \right]^{\frac{1}{P}} \quad (5)$$

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 non-convex Pareto points (Marler and Arora, 2004; Messac et al., 2000a,b). 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 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 k 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.

Figure 2: Delaunay triangulation to find natural neighbors

Krasnogor et al. (1995), 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 Figure 2.

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

- 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.
- The swap move exchanges two microzones that are natural neighbors and in different tours.
- The 2-opt move cuts two arcs, each arc between two natural neighbors, and reconnects the solution.

Figure 3: Examples of the different moves used in the variable-neighborhood tabu search algorithm

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 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.

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  $w_i$ ;

```

The result retrieved from the 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 moves that has been executed in the last t iterations of the 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 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 Equation 6) of the objective functions, and the best solution is used as the starting solution for the next iteration.

$$f(x) = \left[w_1 f_1^n(x)^2 + w_2 f_2^n(x)^2 + w_3 f_3^n(x)^2 \right]^{\frac{1}{2}} \quad (6)$$

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

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% of the microzones are concentrated in clusters. The locations of the remaining 10% 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 maximum are set to $(100 - w)\%$ and $(100 + w)\%$ 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

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

Table 1: Parameter values

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% in steps of 5%. 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 Figure 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%, with a maximum of only 5%. 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

Figure 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.

explanation, that the amount of work inside a microzone is very 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% of a vehicle's time is spent inside the microzones.

	Average % difference	Maximum % difference	standard deviation
hypervol_N100_n0_case_0	1.39120959	4.5593267	1.46926533
hypervol_N100_n0_case_1	1.39908527	5.18476163	1.80136414
hypervol_N100_n0_case_2	0.65542098	1.47838684	0.54603236
hypervol_N100_n1_case_0	1.6154032	3.46909848	1.36114729
hypervol_N100_n1_case_1	2.80846146	4.33350565	1.47778115
hypervol_N100_n1_case_2	1.5729616	2.74987602	0.9470098
hypervol_N100_n2_case_0	1.27253838	2.29765852	0.77140121
hypervol_N100_n2_case_1	1.12459149	1.91531304	0.65385443
hypervol_N100_n2_case_2	0.67997033	1.60294863	0.60087508
hypervol_N100_n3_case_0	2.13858032	4.10520786	1.54878545
hypervol_N100_n3_case_1	1.02914402	1.60015808	0.57465281
hypervol_N100_n3_case_2	0.6970578	1.13002656	0.40211289
hypervol_N150_n0_case_0	0.52908273	1.08699574	0.41417231
hypervol_N150_n0_case_1	0.83707682	2.09173963	0.81233216
hypervol_N150_n0_case_2	1.44316261	2.30062232	0.72971846
hypervol_N150_n1_case_0	3.1983654	5.2709399	1.98687116
hypervol_N150_n1_case_1	0.8900108	2.36564545	0.76347774
hypervol_N150_n1_case_2	1.24379315	3.34203259	1.1387423
hypervol_N150_n2_case_0	0.77501097	1.12065471	0.37423947
hypervol_N150_n2_case_1	1.09003271	1.74020572	0.60512264
hypervol_N150_n2_case_2	1.29529657	3.26659511	1.07563041
hypervol_N150_n3_case_0	2.79944405	4.82379488	2.11847038
hypervol_N150_n3_case_1	1.21411829	1.90636768	0.6690894
hypervol_N150_n3_case_2	0.87218235	3.02312758	1.01994804
hypervol_N200_n0_case_0	1.8662516	3.88906161	1.27226784
hypervol_N200_n0_case_1	0.87717255	2.81684605	0.9489694
hypervol_N200_n0_case_2	0.96784812	1.61127854	0.52096437
hypervol_N200_n1_case_0	1.13679859	1.66438088	0.57401662
hypervol_N200_n1_case_1	1.18150358	2.00969119	0.69749246
hypervol_N200_n1_case_2	1.1568058	3.47853499	1.24272541
hypervol_N200_n2_case_0	0.73140137	1.44253571	0.5456957
hypervol_N200_n2_case_1	2.51040143	3.62735226	1.18090497
hypervol_N200_n2_case_2	0.53068883	1.74728851	0.63170498
hypervol_N200_n3_case_0	0.98817476	1.66111	0.52113997
hypervol_N200_n3_case_1	0.32124655	0.79793064	0.31569055
hypervol_N200_n3_case_2	1.53883286	3.89912511	1.37660352

Table 2: Percentage difference in hypervolume for the same problem instances with different deviations in the amount of work in the zones.

The initial solution used by the algorithm has little influence on the Pareto front found. In Figure 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 3, 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%, showing that the method is relatively robust with respect to the initial solution.

Figure 5: This plot displays the Pareto-front for different initial solutions for a single instance with 200 microzones, 1 cluster and a maximum work load deviation of 15%. It shows that the initial solution only has a limited effect on the Pareto front found. Every different Pareto-set represents the solution set for a different start solution.

A closer look at the Pareto-fronts found by the algorithm reveals that there is a clear (inverse) correlation between the total driving distance, the deviation of the average workload, and the deviation from the tactical plan. 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.

Average difference	2.32464988
Maximum difference	15.03348506
Standard deviation	2.40892581

Table 3: Percentage difference in hypervolume with the best found Pareto set for different start solutions and one 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 meta-heuristic 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., Figure 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.

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.

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

Figure 6: Example solutions found for the instance with 150 zones, zero city centers and a deviation of workload in the zones of maximum 15% of the initial workload

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.

Acknowledgments

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

References

- Z. Bodie, A. Kane, and A. J. Marcus. Investments (International Edition). McGraw-Hill, 2005.
- R. Caballero, M. Gonzalez, F. Guerrero, J. Molina, and C. Parolera. 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(3): 1751–1763, 2007.
- M. Caramia and P. Dell’Olmo. Multi-objective Management in Freight Logistics. Springer London, Limited, 2008.
- C. F. Daganzo. The distance traveled to visit n points with a maximum of c stops per vehicle: An analytic model and an application. *Transportation Science*, 18(4): 331–350, 1984.
- B. N. Delaunay. Sur la sphère vide. *Bulletin of Academy of Sciences of the USSR*, 6: 793–800, 1934.
- M. Ehrgott and X. Gandibleux. An annotated bibliography of multi-objective combinatorial optimization. *OR Spectrum*, 22(4):425–460, 2000.
- C. Expósito Izquierdo, A. Rossi, and M. Sevaux. Modeling and solving the clustered capacitated vehicle routing problem. In *Proceedings of the 14th EU/ME Workshop*, pages 110–115, 2013.
- FedEx. FedEx: Powering Global Access. 2006. http://images.fedex.com/us/about/today/access/GreaterAccessChange_global.pdf.
- M. J. Geiger. Genetic algorithms for multiple objective vehicle routing. Number 1, pages 349–354. Program Operational Ciencia, Tecnologia, Inovação do Quadro Comunitário de Apoio III de Fundação para a Ciencia e Tecnologia, 2008. URL <http://arxiv.org/abs/0809.0416>.

- C. Groer. Parallel and Serial Algorithms for Vehicle Routing Problems. PhD thesis, University of Maryland, 2008. URL <http://hdl.handle.net/1903/9011>.
- N. Jozefowiez, F. Semet, and E. Talbi. Multi-objective vehicle routing problems. *European Journal of Operational Research*, 189(9):293–309, 2008a.
- N. Jozefowiez, F. Semet, and E. Talbi. From Single-Objective to Multi-Objective Vehicle Routing Problems: Motivations, Case Studies, and Methods. In B. Golden, S. Raghavan, and E. Wasil, editors, *The Vehicle Routing Problem: Latest Advances and New Challenges*, volume 43 of *Operations Research/Computer Science Interfaces Series*, pages 445–471. Springer US, 2008b.
- A. Konak, D. Coit, and A. Smith. Multi-objective optimization using genetic algorithms: a tutorial. *Reliability Engineering and System Safety*, 91(9):992–1007, 2006.
- N. Krasnogor, P. Moscato, and M. Norman. A new hybrid heuristic for large geometric traveling salesman problems based on the delaunay triangulation. *Anales del xxvii Simposio Brasileiro de Pesquisa Operacional*, 1995.
- M. Kunkel and M. Schwind. Vehicle routing with driver learning for real world cep problems. In *System Science (HICSS)*, 2012 45th Hawaii International Conference on, pages 1315–1322, 2012.
- S. Lau and L. Shue. Solving travelling salesman problems with an intelligent search approach. *Asia-Pacific Journal of Operational Research*, 18:77–87, 2001.
- C.-M. Liu, T.-C. Chang, and L.-F. Huang. Multi-objective heuristics for the vehicle routing problem. *International Journal of Operations Research*, 3(3):173–181, 2006.
- C. Lyster. Learning from fedex: lessons for the city. *Journal of Landscape Architecture*, 7(1):54–67, 2012.
- R. Marler and J. Arora. Survey of multi-objective optimization methods for engineering. *Structural Multidisciplinary Optimization*, 26(6):369–395, 2004.
- A. Messac, C. Puemi-Sukam, and E. Melachrinoudis. Aggregate objective functions and pareto frontiers: Required relationships and practical implications. *Optimization and Engineering*, 1(2):171–188, 2000a.
- A. Messac, J. E. Renaud, G. J. Sundararaj, and R. V. Tappeta. Ability of objective functions to generate points on nonconvex pareto frontiers. *AIAA Journal*, 38(6):1084–1091, 2000b.
- M. Mourgaya. The periodic vehicle routing problem: planning before routing. PhD thesis, Laboratoire de Mathématiques Appliquées de Bordeaux, Université de Bordeaux 1, Bordeaux, France, 2004.
- M. Mourgaya and F. Vanderbeck. Periodic Vehicle Routing Problem: classification and heuristic. *RAIRO Operational Research*, 40:169–194, 2006.

- M. Mourgaya and F. Vanderbeck. Column generation based heuristic for tactical planning in multi period vehicle routing. *European Journal of Operational Research*, 183(3):1028–1041, 2007.
- G. Nikolakopoulou, S. Kortesis, A. Synefaki, and R. Kalfakakou. Solving a vehicle routing problem by balancing the vehicles time utilization. *European Journal of Operational Research*, 152(2):520 – 527, 2004.
- T. Okabe, Y. Jin, and B. Sendhoff. A critical survey of performance indices for multi-objective optimisation. In *The 2003 Congress on Evolutionary Computation, CEC '03*, volume 2, pages 878–885, 2003.
- Y. Ouyang. Design of vehicle routing zones for large-scale distribution systems. *Transportation Research Part B: Methodological*, 41(10):1079–1093, 2007.
- J.-P. Rodrigue, C. Comtois, and S. B. The Geography of Transport Systems. Routledge, Taylor & Francis, New York, 2006.
- A. Santiago, B. Melián-Batista, A. Álvarez, and F. AngelBello. Effect of the initial solutions to balance routes in vehicle routing problem with time windows. In I. Rojas, G. Joya, and J. Cabestany, editors, *Advances in Computational Intelligence*, volume 7903 of *Lecture Notes in Computer Science*, pages 46–52. Springer Berlin Heidelberg, 2013.
- C. Waldheim and A. Berger. Logistics landscape. *Landscape Journal*, 27(2):219–246, 2008.
- Q. Zhiping and Z. Yuxing. Parametric optimization design of aircraft based on hybrid parallel multi-objective tabu search algorithm. *Chinese Journal of Aeronautics*, 23(4):430–437, 2010.
- H. Zhong, R. Hall, and M. Dessouky. Territory planning and vehicle dispatching with driver learning. *Transportation Science*, 41(1):74–89, 2007.
- E. Zitzler and L. Thiele. Multiobjective evolutionary algorithms: A comparative case study and the strength pareto approach. *Parallel Problem Solving from Nature*, 1498:292–301, 1998.
- E. Zitzler and L. Thiele. Multiobjective evolutionary algorithms: A comparative case study and the strength pareto approach. *IEEE Transactions on Evolutionary Computation*, 3(4):257–271, 1999.