Large Neighborhood Search with Constraint Programming for a Vehicle Routing Problem with Synchronization Constraints

Hossein Hojabri †§, Michel Gendreau ‡§,
Jean-Yves Potvin †§, Louis-Martin Rousseau ‡§

†Département d'informatique et de recherche opérationnelle
Université de Montréal,
C.P. 6128, succ. Centre-Ville, Montréal, Québec, Canada H3C 3J7.

‡Département de mathématiques et de génie industriel
École Polytechnique de Montréal,
C.P. 6079, succ. Centre-Ville, Montréal, Québec, Canada H3C 3A7.

§Centre interuniversitaire de recherche
sur les réseaux d'entreprise, la logistique et le transport,
Université de Montréal,
C.P. 6128, succ. Centre-Ville, Montréal, Québec, Canada H3C 3J7.

Abstract

This paper considers an extension of the vehicle routing problem with time windows, where the arrival of two vehicles at different customer locations must be synchronized. That is, one vehicle has to deliver some product to a customer, like a home theater system, while the crew on another vehicle must install it. This type of problem is often encountered in practice and is very challenging due to the interdependency among the vehicle routes, but has received little attention in the literature. A constraint programming-based adaptive large neighborhood search is proposed to solve this problem. The search abilities of the large neighborhood search and the constraint propagation abilities of constraint programming are combined to determine the feasibility of any proposed modification to the current solution. Numerical results are reported on instances derived from benchmark instances for the vehicle routing problem with time windows with up to 200 customers.

Keywords: Vehicle routing; synchronization; large neighborhood search; constraint programming.

1 Introduction

The Vehicle Routing Problem with Multiple Synchronization constraints (VRPMS) can be used to model many real-world applications, like home care delivery, aircraft fleet assignment, vehicle routing and scheduling, and forest operations, as reported in Rousseau et al. (2013). The problem considered in this paper comes from a software development company that produces vehicle routing solutions for large retailers involved in the delivery of large items to their customers, like furniture, appliances, home theater systems, etc. In the case of electronics equipment, the crew in the delivery vehicle simply unload the goods, while another crew with the required capabilities takes charge of the installation. In this case, the routes of two different types of vehicles must be synchronized at a certain number of customer locations. A dynamic version of a similar problem for pickup vehicles (i.e., goods are collected at customer locations rather than being delivered) is addressed in Rousseau et al. (2013). The authors consider a problem where new customer requests are inserted one by one into the current routes, as they occur. Constraint programming (CP), a computational paradigm based on constraint propagation, is used to determine if an insertion is feasible or not. Since only one customer is considered at a time, this approach proved to be viable. Surprisingly and to the best of our knowledge, no solution approach, neither exact nor heuristic, has been proposed for the static version of the problem where all customers are known in advance. This a priori knowledge is required in the case of delivery applications, because the vehicles are loaded before they depart from the depot. The goal of this paper is thus to provide a contribution in this regard. In particular, we exploit the constraint propagation capabilities of CP in a context where many customers must be handled at the same time, through a proper integration within the Adaptive Large Neighborhood Search (ALNS) framework.

Apart from the paper of Rousseau et al. (2013), a problem quite similar to ours is addressed in Bredström and Rönnqvist (2008) in the domain of home care crew scheduling. The authors formally state their problem as a mixed integer programming (MIP) model that accounts for temporal precedence and synchronization constraints. The problem is then solved through an optimization-based heuristic. Rasmussen et al. (2012) solve the same problem exactly with a branch-and-price algorithm. Due to the application context, there is no capacity constraint and specific issues about home care crews are taken into account, in addition to the synchronization requirements, like care giver preferences, customer priority and ability of a particular care giver to serve a given customer. Also, there is no need to serve all customers (i.e., visits can be rescheduled or canceled).

Schmid et al. (2010) report a MIP for the delivery of concrete to construction sites from several plants by a fleet of heterogeneous vehicles. Two hybrid approaches are introduced to solve the problem by combining an exact algorithm with a variable neighborhood search (VNS). El Hachemi et al. (2011) consider an application in the

forest industry where trucks and log-loaders must be synchronized. A decomposition approach is presented to solve a weekly problem in two phases. The first phase determines the assignment of forest areas to wood mills. Then, the daily routing and scheduling for the transportation of logs is done with either a constraint-based local search alone or a hybrid method consisting of CP coupled with the constraint-based local search. Other types of synchronization requirements are also found in urban mass transit systems where drivers must change buses at relief points (Freling et al. 2003, Haase et al. 2001) or in public transit rail systems where passengers must transfer several times to reach their destination (Wong et al. 2008).

Ioachim et al. (1999) propose a column generation approach for an aircraft fleet assignment and routing problem where a number of flight subsets during the weekdays must depart at the same time. The master problem is a set partitioning problem with synchronization constraints, while the subproblem is a shortest path problem with time windows and linear costs on the time variables. An extension of this work for a periodic airline fleet assignment with time windows is also reported in Bélanger et al. (2006).

Li et al. (2005) report an integer programming model for a manpower allocation problem with time windows and job-teaming constraints. To perform a job, the required workers must show up together at the job location before it can be executed. Furthermore, nobody can leave before the job is completed. The authors propose two construction heuristics embedded within a simulated annealing framework. They provide comparisons with optimal solutions produced by a commercial solver and lower bounds obtained with a network flow model on their own set of test instances. Dohn et al. (2009) also solve the problem with a branch-and-price algorithm where branching on the time variables automatically imposes synchronization.

The interested reader will find in Drexl (2012) a comprehensive survey on synchronization, with a particular emphasis on vehicle routing applications. In this survey, different types of synchronization requirements are presented (spatial, temporal, etc.), as well as different applications (dial-a-ride, school bus routing, truck-and-trailer, etc.), but nothing quite similar to our problem.

The remainder of this paper is organized as follows. Our problem is first introduced in Section 2. A description of the solution methodology is provided in Sections 3 and 4. A few refinements to this methodology are proposed in Section 5. Then, numerical results are reported in Section 6. Finally, concluding remarks are made in Section 7.

2　Problem definition and formulation

Our problem can be seen as an extension of the Vehicle Routing Problem with Time Windows (VRPTW) with the addition of synchronization requirements. In a standard VRPTW, the routes are independent, in the sense that a modification to a given route has no impact on the other routes. This is not the case for the VRPMS because the routes are interdependent due to the synchronization requirements at certain customer locations.

Let us assume that two different types of vehicles are available, called regular (delivery) and special vehicles. The regular vehicles deliver goods to customers and have limited capacity. The special vehicles do not deliver goods but rather provide a service related to the goods delivered by the regular vehicles. The set of customers is divided into two subsets: regular customers visited only by regular vehicles, and special customers visited by both a regular and a special vehicle. A regular customer is represented by a single vertex in the transportation network, while a special customer is represented by two vertices, one for the delivery by the regular vehicle and another for the service provided by the special vehicle.

A regular vertex has a demand and a time window, while a special vertex has only a time window, which is defined relative to the delivery time of the regular vehicle. More precisely, if t is the delivery time at the regular vertex of a special customer, then the service of the special vertex must begin within the time window $[t - \delta,\ t + \gamma]$, where δ and γ are parameters. Different synchronization requirements can be represented with different values of δ and γ. For example, if the service of the special vehicle cannot start before the delivery, then the time window's lower bound of the special vertex is set to t (i.e., $\delta = 0$); if the service of the special vehicle cannot start before the end of the delivery, then the time window's lower bound of the special vertex is set to t plus the service time of the regular vehicle (i.e., δ is set to minus the service time of the regular vehicle), etc.

The problem then consists of constructing routes for the two types of vehicles such that:

1. Each route begins and ends at a single depot;

2. Each regular customer is served in the route of exactly one regular vehicle;

3. Each special customer is served in the route of exactly one regular vehicle and one special vehicle;

4. The total demand on the route of a regular vehicle cannot exceed its capacity;

5. A regular vehicle must start its delivery at a regular or special customer within the time window of the corresponding regular vertex;

6. A special vehicle must start its service at a special customer within the time window of the corresponding special vertex, which is defined relative to the delivery time of the regular vehicle at the regular vertex;

7. A regular (special) vehicle is allowed to wait up to the lower bound of the time window of the regular (special) vertex if it arrives earlier.

The objective is to minimize the total distance traveled by all vehicles (regular and special).

In the following, a small example is given before providing a more formal CP-based formulation of our problem.

2.1 Example

We consider a small example where the load and time window constraints of the regular vehicles are relaxed in order to focus on the synchronization issue. We have two regular vehicles $\{v_1^r, v_2^r\}$, one special vehicle $\{v_1^s\}$, a single depot 0 and 5 customers $\{1, 2, 3, 4, 5\}$, two of which are special customers, namely $\{1, 2\}$. In the underlying network, vertex 0 is the depot, vertices 1 to 5 correspond to regular vertices, while vertices 6 and 7 correspond to the special vertices of customers 1 and 2. Therefore, the vertex pairs (1,6) and (2,7) require synchronization. In Figure 1, each pair is aggregated into a single gray circle. Without loss of generality, we assume that the service or dwell time is equal to 1 and $\delta = \gamma = 0$ for every special vertex.

In the solution shown, solid arcs are used to represent the routes of regular vehicles and dashed arcs for the route of the special vehicle. Each arc is labeled with its travel time and each vertex i is labeled with the vehicle's arrival time t_i. The return time of each vehicle to the depot, denoted t_0, is also indicated on the last arc of the corresponding route. The routing plan is the following:

$v_1^r : 0, 5, 1, 0$
$v_2^r : 0, 2, 3, 4, 0$
$v_1^s : 0, 7, 6, 0$
$t_1 = 8$, $t_2 = 3$, $t_3 = 8$, $t_4 = 12$, $t_5 = 4$, $t_6 = 8$.

To better understand the synchronization requirements, we follow the route of regular vehicle v_1^r. Its route starts with regular customer 5 where the delivery begins at $t_5 = 4$. After spending 1 time unit to serve the customer, it takes 3 more time units to arrive at special customer 1 where a special vehicle is also needed. The special vehicle arrives at the same time $t_1 = t_6 = 8$. Thus, vehicle v_1^r is perfectly synchronized with the special vehicle v_1^s. Then, vehicle v_1^r departs from special customer 1 at time 9 for the depot. Finally, it arrives at the depot at time 13. Note that the special vehicle v_1^s is already synchronized with the regular vehicle v_2^r at
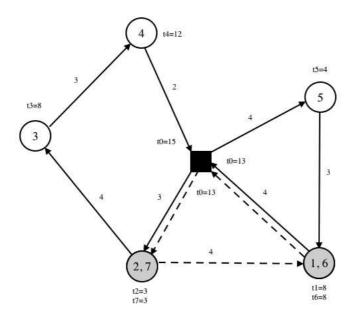
Figure 1: VRPMS example

special customer 2 at time $t_2 = t_7 = 3$.

2.2 Model

In this section, a CP-based formulation of the problem is proposed. This formulation is at the core of the reconstruction operator of our ALNS (see Section 3). Some familiarity with constraint programming is assumed, otherwise the reader is referred to Rousseau et al. (2013).

Let $G = (N, A)$ be a complete and directed graph with vertex set $N = V_r \cup V_s \cup V_r^+ \cup V_s^+ \cup V_r^- \cup V_s^-$, where V_r and V_s are the sets of regular and special vertices, respectively, with $V_r \cup V_s = V$. We recall that one vertex is associated with each regular customer and two vertices with each special customer, namely a regular vertex in V_r and a special vertex in V_s. Also, V_r^+ and V_r^- contain multiple copies of the depot, one for the start and end of the route of each regular vehicle. That is, $|V_r^+|$ and $|V_r^-|$ are both equal to the number of regular vehicles. Similarly, V_s^+ and V_s^- contain multiple copies of the depot for the start and end of the route of each special vehicle. Again, $|V_s^+|$ and $|V_s^-|$ are both equal to the number of special vehicles. We also define $V_r^+ \cup V_s^+ = V^+$ and $V_r^- \cup V_s^- = V^-$. Finally, with each arc $(i, j) \in A$ is associated a non-negative distance d_{ij} and travel time t_{ij}. The notation for the CP model is summarized below.

Sets and parameters:
d_{ij}: Distance between vertices i and j, i, $j \in N, i \neq j$;
t_{ij}: Travel time between vertices i and j, i, $j \in N, i \neq j$;
q_i: Demand of vertex $i \in V_r$;
d_i: Service or dwell time at vertex $i \in V$;
a_i: Lower bound of the time window at regular vertex $i \in V_r$;
b_i: Upper bound of the time window at regular vertex $i \in V_r$;
δ_i: Decrement to derive the lower bound of the time window at special vertex $i \in V_s$;
γ_i: Increment to derive the upper bound of the time window at special vertex $i \in V_s$;
r_i: Regular vertex associated with special vertex $i \in V_s$;
V_r: Set of regular vertices;
V_s: Set of special vertices;
V: set of regular and special vertices;
K_r: Set of regular vehicles;
K_s: Set of special vehicles;
K: Set of regular and special vehicles;
Q: Capacity of each regular vehicle $k \in K_r$.

Note that d_{ij} and t_{ij} are fixed to ∞ when arc (i, j) connects two copies of the depot (except when the two copies stand for the start and end of the same vehicle route, in which case $d_{ij} = 0$; these arcs are used when there are empty vehicle routes in the solution). Also, the demand and the service or dwell time at each copy of the depot are set to 0

Variables:
s_i: Successor of vertex $i \in V \cup V^+$;
k_i: Vehicle visiting vertex $i \in N$;
t_i: Start time of service at vertex $i \in N$;
l_i: Vehicle load after serving vertex $i \in V_r \cup V_r^+ \cup V_r^-$.

Constraints:
The constraints of the problem are modeled as follows:

General constraints:

$$AllDifferent(s_{i:i \in V \cup V^+}) \tag{1}$$

$$NoSubTour(s_{i:i \in V \cup V^+}) \tag{2}$$

$$i \in V_r \cup V_r^+ \Rightarrow s_i \in V_r \cup V_r^- \tag{3}$$

$$i \in V_s \cup V_s^+ \Rightarrow s_i \in V_s \cup V_s^- \tag{4}$$

$$i \in V_r \cup V_r^+ \cup V_r^- \Rightarrow k_i \in K_r \tag{5}$$

$$i \in V_s \cup V_s^+ \cup V_s^- \Rightarrow k_i \in K_s \tag{6}$$

$$s_i = j \Rightarrow k_j = k_i, \qquad i \in V \cup V^+ \tag{7}$$

Capacity constraints:

$$l_i = 0, \qquad i \in V_r^- \tag{8}$$

$$l_i \leq Q, \qquad i \in V_r \cup V_r^+ \tag{9}$$

$$s_i = j \Rightarrow l_j = l_i - q_j, \qquad i \in V_r \cup V_r^+ \tag{10}$$

Time window and synchronization constraints:

$$t_i = 0, \qquad i \in V^+ \tag{11}$$

$$a_i \leq t_i \leq b_i, \qquad i \in V_r \tag{12}$$

$$t_{r_i} - \delta_i \leq t_i \leq t_{r_i} + \gamma_i, \qquad i \in V_s \tag{13}$$

$$s_i = j \Rightarrow t_i \leq t_j - t_{ij} - d_i, \qquad i \in V \cup V^+ \tag{14}$$

Objective function:

$$\min \sum_{i \in V \cup V^+} d_{i,s_i} \tag{15}$$

In a solution, each vertex must have exactly one predecessor and one successor. The nature of variables s_i is such that any vertex i can only have one successor, but we must also ensure that it has only one predecessor. To do so, it must be forbidden for two different vertices to have the same successor, which corresponds to Constraint (1) called *AllDifferent*. Constraint (2) prohibits subtours, see the detailed description in Pesant et al. (1998). Constraints (3)-(6) define variables s_i and k_i for regular and special vertices. Constraint (7) states that a vertex and its successor must be served by the same vehicle. Constraints (8)-(10) impose load and capacity restrictions. Without loss of generality, vehicles are assumed to start their routes at the depot at time 0 through constraint (11). Constraint (12) imposes time window constraints on regular vertices. The synchronization requirements between regular and special vehicles at special vertices is stated in Constraint (13). Constraint (14) imposes the coherence of each route schedule. Finally, the objective function (15) minimizes the total distance traveled by all vehicles.

3   ALNS

The ALNS of Pisinger and Ropke (2007) extends the large neighborhood search framework of Shaw (1998), a problem-solving approach which can also be related to the ruin-and-recreate principle of Schrimpf et al. (2000). The basic idea is to search for a better solution at each iteration by destroying a part of the current solution and by reconstructing it in a different way. When solving VRPs, a new solution is obtained by first removing a number of vertices and then by reinserting these vertices into the solution. Typically, a number of different operators are available

and the adaptive extension chooses an operator in a randomized, but informed, way at each iteration. That is, the selection probability of an operator is derived from its associated weight, which is adjusted during the search depending on the operator's previous successes and failures.

A noteworthy feature of our ALNS is the use of CP to reinsert the removed customers. Thus, a systematic exploration of feasible ways to reinsert the vertices is performed.

The ALNS algorithm can be described as follows.

1. Generate an initial solution s

2. $s^* \leftarrow s$

3. While a stopping criterion is not met do:

   3.1 Adaptively select a destruction operator and apply it to s

   3.2 Reinsert customers with CP

   3.3 If a new solution $s' \neq s$ is returned by CP then

   $\quad\quad s \leftarrow s'$

   $\quad\quad$ If s' is better than s^* then $s^* \leftarrow s'$

4. Return best solution s^*

The main components of this algorithm will be detailed in the following. Note that we basically follow the implementation of Pisinger and Ropke (2007), except for the CP-based insertion operator (including a number of adjustments required by its integration within ALNS) and some destruction operators that are specifically designed for our problem.

3.1 Initial Solution

The routes of the delivery vehicles are first created by inserting regular vertices using the well-know I1 insertion heuristic of Solomon (1987), originally developed for the VRPTW. Then, the time windows for the special vertices are defined through the delivery times in the regular routes. Solomon's insertion heuristic is reapplied on the special vertices to create routes for the special vehicles. At this point a complete solution is obtained. Note that the number of vehicles in this initial solution will tend to decrease as the ALNS is executed due to the minimization of the total distance (although minimizing total distance is not equivalent to minimizing the number of vehicles).

3.2 Adaptive operator selection

A destruction operator removes a number of vertices from the current solution at each iteration of the ALNS. Since more than one destruction operator is typically available, the adaptive mechanism is aimed at choosing the destruction operator in a way that accounts for its previous outcomes. A weight is associated with each operator for this purpose. Let us assume that we have h operators, each with a weight w_j, $j = 1, ..., h$. The removal operator i is then selected with probability

$$\frac{w_i}{\sum_{j=1}^{h} w_j}, i = 1, ..., h.$$

That is, the probability of selecting a given operator increases with its weight. Starting with a unit weight for each operator, the weights are updated after a number of consecutive iterations called a segment. The weight of operator i at the start of a given segment sg is based on the one used in the previous segment $sg - 1$ and is computed as follows:

$$w_i^{sg} = \gamma \cdot w_i^{sg-1} + (1 - \gamma) \cdot \frac{\pi_i^{sg-1}}{n_i^{sg-1}},$$

where n_i^{sg-1} is the number of times operator i was used in segment $sg - 1$, γ is a value between 0 and 1 and π_i^{sg-1} is the score of operator i at the end of segment $sg - 1$. Parameter γ controls the inertia in the weight update equation. When γ is close to 1, the history prevails and the weights do not change much. Conversely, when γ is close to 0, the update is driven by the most recent score.

The score, which is reset to zero at the beginning of each segment, is incremented when operator i is used at a given iteration t to produce a new solution. More precisely, the new score at iteration $t + 1$ becomes

$$\pi_i^{t+1} = \pi_i^t + \begin{cases} \sigma_1 & \text{if a new best solution has been produced,} \\ \sigma_2 & \text{if the solution produced is better than the current solution,} \end{cases}$$

where σ_1 and σ_2 are parameters.

3.3 Destruction operators

A total of 10 destruction operators were considered in this study, but only two were kept at the end, called *Route removal* and *Historical node-pair removal*, after the results of an experiment reported in Section 6. The reader is referred to Pisinger and Ropke (2007) and Shaw (1998) for a description of the eight previously reported operators that were excluded from the final implementation, namely *Cluster removal*, *Random removal*, *Worst removal*, *Historical route-pair removal*, *Distance-oriented related removal* and *Time-oriented related removal*. In the following, we describe *Route removal*, *Historical node-pair removal*, *Synchro Vertex removal* and *Synchro*

Route removal, because they were specifically designed for our problem (although the last two were excluded from the final implementation). They are all based on a parameter which is the number of vertices to be removed. To simplify the description, this parameter is always denoted NV, but it is clear that a different parameter value can be associated with each operator.

3.3.1 Route removal

This operator removes vertices in the same route(s). First, a route is randomly selected and all its vertices are removed. If NV is not reached yet, all remaining vertices are sorted according to their distance to a randomly chosen vertex among the previously removed ones. Then, one of the closest vertices is selected, using a randomization factor r^D. More precisely, the vertex in position $\lfloor r^D * ListSize \rfloor$ in the sorted list is selected, where $ListSize$ is the number of vertices in the list, r is a random number between 0 and 1, and D is a parameter. The selected vertex, as well as the other vertices in the same route, are then removed. This procedure is repeated until NV vertices are removed. The randomization factor provides a form of diversification but, more importantly, when one of the two vertices of a special customer is randomly selected, the closest vertex is necessarily the other vertex from the pair (at distance 0). Thus, without randomization, the operator would behave like the *Synchro Route removal* operator (see below). In the computational results, *Route removal* proved to be the best performing operator.

3.3.2 Historical node-pair removal

This operator removes vertices based on historical information. That is, a score $g(i,j)$ (initially set to infinity) is assigned to each arc (i,j). This score corresponds to the objective value of the best solution found up to now in which vertices i and j are visited consecutively. Then, for each vertex i, we consider the value $g_{i,\text{succ}_i} + g_{\text{pred}_i,i}$, where $succ_i$ and $pred_i$ are the predecessor and successor of vertex i in the current solution, respectively. The NV vertices with the largest values are then removed.

3.3.3 Synchro Vertex removal

First, $\lceil NV/2 \rceil$ special vertices are randomly selected and removed with their corresponding regular vertices. If the number of special vertices is less than $\lceil NV/2 \rceil$, then the remaining vertices are randomly chosen among the regular ones. Given that a special vertex must be synchronized with a regular vertex, the idea is that better insertion places are more likely to be found in the following reinsertion phase if both vertices are removed together. Unfortunately, this operator did not provide good results and was not kept in the final implementation.

3.3.4 Synchro Route removal

This operator focuses on the routes of special vehicles. First, a route is randomly selected among all the routes served by special vehicles. Then, the route is emptied by removing all its (special) vertices along with their corresponding regular vertices. This is repeated until NV vertices are removed. Like Synchro Vertex removal, this operator did not perform well and was not kept at the end.

The following section is devoted to the reinsertion of the removed vertices.

4 CP-based reinsertion

Here, a partially destroyed solution is reconstructed using CP, as initially proposed in Pesant and Gendreau (1996, 1999) and Rousseau et al. (2002), using the model presented in section 2.2. CP has proven useful for a number of routing variants and the presence of synchronization constraints make our problem even more suitable for a CP-based approach, see Rousseau, et al. (2013). For example, the insertion of a special vertex in the route of a special vehicle causes a delay not only to the special vertices that are visited later along this route, but also to the regular vertices which are associated with these special vertices. That is, the schedule of regular vehicles is also modified. The constraint propagation capabilities of CP can thus be exploited here to further restrict the feasible region of the search space. In the computational results reported in Section 6, the ILOG CP Optimizer was used for this purpose.

In the reconstruction phase, only a subset of vertices needs to be reinserted in the solution, which restricts the domain of feasible values for each variable. For example, if vertex i is part of the partial solution (it has not been removed), then the value of the successor variable s_i in the model of section 2.2 can only be chosen among the current successor in the partial solution or one of the removed vertices. Once the domain of each variable is appropriately set, CP then implicitly considers all feasible ways to reinsert the removed vertices in the solution. However, due to the propagation and pruning abilities of CP, the number of possibilities can be drastically reduced using appropriate variable and value selection heuristics to guide the reinsertion. In our implementation, we instantiate first the more constrained variables with their less constraining value. Here, a variable s_i associated with a removed vertex has a smaller number of possible successors, and is more constrained, if it is associated with a vertex far from the other vertices. The value of the selected variable is then set to the closest vertex in the current partial solution. This approach proved to be slightly better than the autosearch feature of the ILOG CP Optimizer.

5 Further refinements

To improve the performance of our ALNS, additional refinements have been introduced. They are described in the following.

5.1 Constrained objective function

Since CP looks for the optimal reinsertion of the removed vertices at each iteration of ALNS, very long run times are to be expected. To alleviate this problem, the minimization of the objective function is taken off from the CP model and introduced as a constraint, to produce a pure constraint satisfaction problem:

$$NewObj < \alpha \times CurrObj \tag{16}$$

where *CurrObj* and *NewObj* are the values of the current and new solutions, respectively, and α has a positive value which must be less than or equal to 1. The value of α is very important. If too small, CP might run for a long time for nothing because there is no feasible solution. If too large, CP will return a solution just a bit better than the current solution. After some experiments, α was defined as $(1 - 0.02/(nbIter^{\lambda} + 1.1))$ with $\lambda = 0.45$. This formula depends on the iteration number *nbIter*, because larger improvements are expected to be found at the start. As the search progresses, smaller improvements are looked for.

5.2 Solution acceptance

It is known that the search process of ALNS can be enhanced by allowing worse solutions to be accepted, see Pisinger and Ropke (2007). Similarly, we force CP to look for a worse solution with some probability *Prob*. This probability increases with the number of consecutive CP failures (a failure occurs when CP cannot find a feasible and improved reinsertion of the removed vertices). The formula is the following, where p is a parameter:

$$Prob = (nbFail/(nbFail + 1))^p$$

In practice, the following constraint replaces constraint (16) with probability *Prob* (where w and r are parameters):

$$CurrObj < NewObj < CurrObj \times (1 + \frac{nbFail^w}{r}) \tag{17}$$

It should be noted that CP searches for increasingly worse solutions as the number of consecutive failures *nbFail* increases. If a solution is returned, *NbFail* is reset to 0.

5.3 Variable neighborhood size

To help intensify the search, the number of vertices to be removed NV is increased by 1 each time the reinsertion through CP fails. Conversely, NV is decreased by one each time the reinsertion succeeds. However, the NV value must stay within the bounds of the interval $[Min_{NV}, Max_{NV}]$, where Min_{NV} and Max_{NV} are parameters. The initial value of NV is randomly selected within that interval.

5.4 CP search time

Experimentally, we found that our ALNS is more efficient when CP is stopped after a given time limit $CPlimitTime$ (i.e., if CP has not returned a solution after that amount of time, it is unlikely that it will find one if it runs longer). This approach leads to much shorter run times and also produces better solutions because ALNS can explore a larger fraction of the search space. $CPlimitTime$ is allowed to vary in the interval $[MinCPlimit, MaxCPlimit]$ (in seconds). The time limit is initially equal to $MinCPlimit$ and is incremented by one when CP fails to find a solution and the current value is not $MaxCPlimit$. It is decremented by one when CP returns a solution and the current value is not $MinCPlimit$.

6 Computational experiments

This section first describes our test instances. This is followed by a parameter sensitivity analysis and the description of an experiment aimed at identifying the best destruction operators. Then, numerical results are reported.

6.1 Test instances

Given that the static version of our VRPMS is addressed for the first time, we generated test instances with $n = 25, 50, 100$ and 200 customers from the VRPTW instances of Solomon (1987) and Homberger and Gehring (1999). In these Euclidean instances, the travel time is the same as the distance. There are also different classes of instances. The customers are randomly located in the instances of type R, clustered in the instances of type C, and both randomly located and clustered in the instances of type RC. Furthermore, instances of type 1 have a short scheduling horizon, while those of type 2 have a long one. Overall, there are six different classes of instances: R1, R2, C1, C2, RC1 and RC2. A variable number of instances is found in each class, for a total of 56 instances for each size.

Some additional parameters are required to transform the original VRPTW instances into VRPMS instances: Sp, which is the percentage of special customers; δ_i and γ_i which are used to define the time window of each special vertex $i \in V_s$.

Test instances were created from the VRPTW instances by setting parameter Sp to 5%, 25% and 50%. More precisely, the number of special customers is equal

to $\lceil Sp \times n \rceil$. Then, the first customer is a special customer and the following special customers are chosen using a constant interval defined by $1/Sp$. For example, if $n = 100$ and $Sp = 5\%$, then $1/Sp = 20$ and the five special customers are 1, 21, 41, 61, and 81.

The VRPTW instances are designed in such a way that the instances of larger size are obtained by adding customers to the instances of smaller size. For example, the instances of size 50 are obtained by adding 25 customers to the instances of size 25. For a given Sp value, the same applies to the special customers in the VRPMS instances. That is, all special customers in the instances of size 25 are also found in the corresponding instances of size 50. With regard to synchronization, δ_i and γ_i were set to 0 and 10, respectively, for each special vertex $i \in V_s$ (as in Rousseau et al. (2013)). The test instances can be found in the Mendeley dataset repository at: http://dx.doi.org/10.17632/v5hf63zxk2.1.

6.2   Parameter sensitivity analysis

We adjusted the parameter values by considering a subset of 18 instances of size 100, that is, one instance for each class and Sp value. The values considered for the various parameters are shown in Table 1, with the best values in bold. Default values in italic were taken from the literature or were determined after a few preliminary experiments. When the default and best values coincide, the entry is in italic bold. Since it was not possible to consider all combinations of values, we proceeded as follows. The parameters were processed one by one in the order shown in the table. When the different values of the current parameter were tested, all parameters not yet tested were set at their default value. Then, the current parameter was fixed to the value leading to the best average objective value, before the next parameter was considered.

6.3   Operator selection

Another experiment was performed to select a good combination of operators among the 10 available destruction operators. For this purpose, we ran the ALNS on the subset of instances of size 100 using the parameter setting identified in Table 1. Let us denote LB the average solution value obtained with the original ALNS implementation with $h = 10$ operators. Then, at each iteration, ALNS is run h times with $h - 1$ operators by removing each operator in turn. The best among the h runs is then compared with LB. If the average solution value is better than LB, the missing operator is detrimental because an improvement is observed without it. Consequently, this operator is put aside and the procedure is repeated with the $h - 1$ remaining operators with LB reset to the improved average solution value. Otherwise, when the average solution value of the best among the h runs is worse than LB, each missing operator degrades solution quality and the procedure is stopped. In our case, the procedure ran until all operators were dismissed except *Route re-*

Table 1: Parameter settings

Parameter														
γ	0.1	0.3	*0.5*	0.7	0.85	**0.9**	0.95							
λ	0.2	0.25	0.3	0.35	**0.4**	*0.45*	0.5	0.55	0.6	0.65	0.7			
Segment length	30	*50*	**70**	100	150	200	250	300						
D	0	1	2	*3*	4	5								
p	1	2	3	*4*	5	6	7	8	9	10	11	12	**13**	14
w	0.25	**0.5**	0.75	1	*1.5*	2								
r	100	500	***1K***	5K	10K	100K								
σ_1	1	*2*	3	5	10	20	30							
σ_2	*1*	2	3	4	**5**	6	10							
MinNV-MaxNV	6-16	8-18	**10-20**	12-22	*13-23*	14-24	16-26							
MinCPLimit-MaxCPLimit	1-3	1-4	1-5	*1-6*	2-4	2-5	**2-6**	2-7						

moval and *Historical node-pair removal.*

An incremental approach was then considered. First, ALNS was run 10 times with a single destruction operator each time. The best operator proved to be *Route removal* and this operator was kept. Then, ALNS was run 9 times by adding a different operator to *Route removal.* At this point, *Historical node-pair removal* with *Route removal* proved to be the best pair and the average solution value was also better than the one obtained with *Route removal* alone. Thus, *Historical node-pair removal* was kept. The next step showed that the addition of a third operator always led to a degradation in solution quality. Thus, we stopped again with *Route removal* and *Historical node-pair removal.*

6.4   Numerical results

For the computational experiments, our algorithm was run on a 3.07GHz Xeon(R) X5675 using the best parameter setting identified in Table 1 and keeping only the destruction operators *Route removal* and *Historical node-pair removal.* To allow convergence, the number of iterations was set to 5,000, 10,000, 25,000 and 35,000 on instances of size 25, 50, 100 and 200, respectively. For verification purposes, we first tested our algorithm on Solomon's instances, that is, the classical VRPTWs of size 100 with no special customer. We performed three independent experiments and observed average gaps with the best known solutions of 1.3%, 1.6% and 1.8%, although in non competitive computation times (as indicated by the runtimes reported in the following tables of results). Solution quality is acceptable, but our approach is not meant for such instances because the computationally expensive constraint propagation properties of CP are particularly indicated in the presence of many inter-related constraints, as those found in the VRPMS.

Table 2 now reports the results obtained on the VRPMS instances. In this table, the following abbreviations are used:

- Size: Number of customers

- Class: Class identifier

- nbSyn: Number of special customers

- InObj: Total distance traveled by the regular and special vehicles in the initial solution

- InRgObj: Total distance traveled by the regular vehicles in the initial solution

- InSpObj: Total distance traveled by the special vehicles in the initial solution

- FnObj: Total distance traveled by the regular and special vehicles in the final solution

- FnRgObj: Total distance traveled by the regular vehicles in the final solution

- FnSpObj: Total distance traveled by the special vehicles in the final solution

- InNbRgV: Number of regular vehicles used in the initial solution

- InNbSpV: Number of special vehicles used in the initial solution

- FnNbRgV: Number of regular vehicles used in the final solution

- FnNbSpV: Number of special vehicles used in the final solution

- RunTime: Computation time in seconds

- LitBestObj: Objective value (total distance) of the best known VRPTW solution in the literature

- LitBestnbVeh: Number of vehicles of the best known VRPTW solution in the literature

Each number corresponds to an average taken over all instances in a given class (between 8 and 12 instances). Column $Gap\%$ reports the gap in percentage between $FnRgObj$, when only the total distance of the regular routes is accounted for, and the best VRPTW solution reported in the literature $LitBestObj$ that is:

$$Gap\% = 100 \times ((FnRgObj - LitBestObj)/FnRgObj).$$

Column $Imp\%$ reports the percentage of improvement in total distance from the initial to the final solution. Due to the synchronization constraints, the total distance traveled by the regular vehicles in our solutions is typically larger than the one observed in the best known VRPTW solutions. But, in a few cases, our algorithm provides an improvement, as indicated by a negative gap in Table 2. This is a surprise, but the best known solutions are based on a hierarchical objective where the number of vehicles is minimized first and, for the same number of vehicles, the total distance is minimized. Although reducing the number of vehicles often leads to reducing total distance, this is not always the case, as illustrated in these cases.

Admittedly, the gap is sometimes larger than expected for instances with only a few special customers. On the other hand, the gap does not increase that much with a larger number of special customers (in some cases, it even decreases). For example, if we consider the instances of size 100 in class R1, the gap with 5 special customers is 5.39%, but it only increases to 7.89% and 8.49% with 25 and 50 special customers, respectively. Thus, it appears that only a few special customers with synchronization requirements is enough to break the structure of the VRPTW solutions, although additional special customers can then be easily accommodated. The percentage of improvement over the initial solution also indicates that the ALNS performs well. Again, this is the first time that CP is used to solve VRPMS instances of that size. It should be noted that Rousseau et al. (2013) address the

dynamic version of VRPMS, so that CP handles only a few customers at a time (i.e., customers who have already been visited or those who have not called yet for service are not part of the current solution).

Table 3 finally reports the results obtained when the parameter settings are adjusted to each problem size, using the procedure reported in section 6.2. This table only reports results on the instances of size 25 and 200. The results are obviously the same for the instances of size 100; but it turned out that the parameter setting used to produce the results in Table 1 is also the best for the instances of size 50. Additional improvements are observed when comparing the corresponding numbers for the instances of size 25 and 200 in Tables 1 and 2, but the attractiveness of a single parameter setting applied to all test instances is lost.

7 Conclusion

In this work, a VRPTW variant incorporating synchronization constraints between two different types of vehicles was addressed. The ALNS framework was coupled with CP to solve the problem. Under this hybrid scheme, CP was used to reconstruct partially destroyed solutions. Numerical results on instances derived from standard VRPTW benchmark instances empirically demonstrated the optimization capability of the proposed algorithm. New developments will now look at ways to increase the efficiency of CP through finer adjustments within the ALNS framework (e.g., adaptive CP time limit). We also want to get closer to the real-world by considering the integration of time-dependent travel times that would account for rush hours, for example.

Table 2: Results for 25-, 50-, 100- and 200-customer instances

| | | | Objective Values | | | | | | Number of vehicles | | | | | Literature | | | |
| | | | Initial solution | | | Final solution | | | Initial solution | | Final solution | | | | | | |
Size	Class	nbSyn	InObj	InRgObj	InSpObj	FnObj	FnRgObj	FnSpObj	InNbRgV	InNbSpV	FnNbRgV	FnNbSpV	RunTime	LitBestObj	LitBestnbVeh	Imp %	Gap%
25	R1	2	621.66	559.13	62.53	555.22	492.69	62.53	5.92	1	5.5	1	16848.22	463.45	4.92	10.69	5.93
	R2		577.89	515.35	62.53	463.96	401.43	62.53	4.18	1	3.09	1	25191.24	350.25	2.73	19.71	12.75
	C1		279.27	227.3	51.97	246.19	195.37	50.81	3.17	1.17	3	1	17418.74	190.59	3.00	11.85	2.45
	C2		356.23	287.77	68.46	288.69	220.23	68.46	2.38	1	2	1	19849.75	214.45	2	18.96	2.62
	RC1		572.33	407.01	165.32	520.39	355.49	164.89	3.88	1.5	3.38	1.25	12699.54	350.24	3.25	9.08	1.48
	RC2		558.28	393.8	164.47	489.61	325.14	164.47	3	1	2.88	1	17660.31	319.28	3	12.3	1.8
25	R1	7	832.64	559.13	273.51	704.14	496.44	207.7	5.92	3.83	5.42	3.08	16273.53	463.45	4.92	15.43	6.65
	R2		779.83	515.35	264.48	597.14	408.47	188.67	4.18	2.09	3.18	1.45	24096.96	350.25	2.73	23.43	14.25
	C1		360.15	215.47	144.68	323.1	203.23	119.87	3.11	2.89	3	1.89	15498.12	190.59	3.00	10.29	6.22
	C2		402.83	287.77	115.06	330.33	222.68	107.65	2.38	1	2	1	16121.23	214.45	2	18	3.7
	RC1		698.27	407.01	291.25	638.79	373.71	265.08	3.88	3	3.63	2.88	11861.81	350.24	3.25	8.52	6.28
	RC2		693.4	393.8	299.59	568.34	332.19	236.16	3	2	2.88	1.63	17520.55	319.28	3	18.04	3.89
25	R1	13	950.12	559.13	390.99	828.41	492.49	335.92	5.92	4.67	5.33	3.83	14218.83	463.45	4.92	12.81	5.9
	R2		913.19	515.35	397.84	706.96	407.85	299.12	4.18	2.55	3	2.27	25615,6	350.25	2.73	22.58	14.12
	C1		389.5	215.47	174.03	353.3	196.74	156.56	3.11	3	3	2.56	12988.41	190.59	3.00	9.29	3.13
	C2		476.72	287.77	188.95	385.81	223.92	161.89	2.38	1.13	2	1	13118.35	214.45	2	19.07	4.23
	RC1		731.55	407.01	324.54	667.17	366.92	300.24	3.88	3.38	3.5	3.13	9407.29	350.24	3.25	8.8	4.55
	RC2		709.59	393.8	315.78	603.24	330.68	272.57	3	2.25	2.88	2.13	16984.06	319.28	3	14.99	3.45
50	R1	3	1056.21	949.14	107.07	902.23	814.25	87.98	10.67	1.58	8.83	1	31527.45	766.13	7.75	14.58	5.91
	R2		893.48	796.94	96.53	733.14	645.57	87.57	5.73	1.18	4.73	1	46639.5	634.03	4.11	17.95	1.79
	C1		526.62	442.08	84.54	457.94	378.03	79.91	5.5	1.33	5.17	1	24692.47	361.69	4.56	13.04	4.32
	C2		603.43	494.47	108.96	471.02	368.77	102.25	3.57	1.57	3	1	42906.69	357.50	2.75	21.94	3.06
	RC1		1117.91	903.25	214.66	980.16	783.15	197.02	8.5	2.5	7.5	1.88	26264.02	730.31	6.50	12.32	6.75
	RC2		900.77	710.5	190.28	785.82	595.75	190.07	4.75	1.13	4.75	1	32023.15	568.65	4.33	12.76	4.55
50	R1	13	1383.4	949.14	434.26	1157.59	818.93	338.66	10.67	5.58	8.92	3.58	31164.96	766.13	7.75	16.32	6.45
	R2		1234.54	796.94	437.6	973.28	658.23	315.05	5.73	2.64	4.91	2.09	45965.2	634.03	4.11	21.16	3.68
	C1		710.67	415.8	294.87	621.69	375.17	246.52	5.33	4.78	5.11	3.67	28163.37	361.69	4.56	12.52	3.59
	C2		779.74	494.47	285.27	648.13	402.91	245.21	3.57	2.29	3.14	1.86	39479.04	357.50	2.75	16.88	11.27
	RC1		1426.83	903.25	523.58	1236.84	787.11	449.73	8.5	5.38	7.38	4.38	25993.75	730.31	6.50	13.32	7.22
	RC2		1268.13	710.5	557.64	1014.09	613.48	400.61	4.75	3.13	4.75	2	35686.48	568.65	4.33	20.03	7.31
50	R1	25	1615.13	949.14	665.99	1348.59	834.31	514.28	10.67	7.83	9.33	5.33	28321,2	766.13	7.75	16.5	8.17
	R2		1469.22	796.94	672.28	1107.87	663.17	444.7	5.73	3.64	4.73	2.82	44380.65	634.03	4.11	24.59	4.39
	C1		786.88	415.8	371.08	682.03	369.55	312.48	5.33	5	5	4.44	21976.72	361.69	4.56	13.32	2.13
	C2		850.11	477.88	372.22	702.8	391.26	311.54	3.5	2.88	3.25	2.25	32853.62	357.50	2.75	17.33	8.63
	RC1		1626.42	903.25	723.17	1316.64	764.13	552.5	8.5	7	7	5.25	25843.36	730.31	6.50	19.05	4.43
	RC2		1316.27	710.5	605.77	1082.22	610.59	471.63	4.75	4.13	4.75	3.25	36624.92	568.65	4.33	17.78	6.87
100	R1	5	1716.51	1533.15	183.36	1429.05	1269.83	159.22	17.17	2.83	15	2.25	82716.31	1201.34	11.92	16.75	5.39
	R2		1425.63	1241.35	184.28	1114.93	969.61	145.32	7.09	2	6.64	1.27	107910	951.03	2.73	21.79	1.92
	C1		1228.52	1032.97	195.55	1017.52	841.86	175.67	11.44	1.89	10.11	1.44	41401.64	828.38	10	17.18	1.6
	C2		999.41	784.54	214.87	841.67	646.73	194.94	5	2	4	1.63	75868.33	589.86	3	15.78	8.79
	RC1		1949.66	1720.98	228.68	1637.62	1432.95	204.67	16	2.75	14.5	2	73411.25	1384.16	11.5	16	3.4
	RC2		1643.43	1427.36	216.07	1286.37	1087.85	198.51	7.63	2.38	6.88	1.38	97137.41	1119.24	3.25	21.73	-2.89
100	R1	25	2190.18	1533.15	657.03	1789.48	1304.3	485.17	17.17	8.83	15.42	5.67	80930.04	1201.34	11.92	18.3	7.89
	R2		1856.53	1241.35	615.18	1407.66	986.86	420.8	7.09	3.91	6.45	3	106302.71	951.03	2.73	24.18	3.63
	C1		1749.11	1032.97	716.13	1441.88	887.41	554.47	11.44	7.78	10.78	5.67	74134.63	828.38	10	17.56	6.65
	C2		1331.89	784.54	547.35	1087.37	674.94	412.44	5	3.25	4.38	2.88	70059.29	589.86	3	18.36	12.61
	RC1		2667.44	1720.98	946.47	2146.34	1483.63	662.72	16	10.5	14.88	5.5	71074.11	1384.16	11.5	19.54	6.7
	RC2		2222.07	1427.36	794.71	1709.74	1117.81	591.93	7.63	4.75	6.88	3.38	99529.68	1119.24	3.25	23.06	-0.13
100	R1	50	2564.23	1533.15	1031.08	2092.57	1312.76	779.81	17.17	12.58	15.25	9	72804.48	1201.34	11.92	18.39	8.49
	R2		2201.18	1241.35	959.83	1693.74	1041.13	652.61	7.09	5.27	6.55	4.36	104309.41	951.03	2.73	23.05	8.65
	C1		1937.9	1032.97	904.92	1609.27	889.06	720.21	11.44	10.11	10.89	8.67	72829.52	828.38	10	16.96	6.83
	C2		1386.88	803.45	583.44	1180.34	693.28	487.06	5.13	3.63	4.38	3	54500.81	589.86	3	14.89	14.92
	RC1		3045.9	1720.98	1324.93	2522.19	1500.18	1022.01	16	13.75	14.63	10	67991.74	1384.16	11.5	17.19	7.73
	RC2		2545.56	1427.36	1118.2	1945.31	1124.47	820.84	7.63	6	7.13	4.75	95459.43	1119.24	3.25	23.58	0.47
200	R1	10	5512.69	4948.78	563.91	4176.82	3705.66	471.16	19.5	3.1	18.3	2.6	100459.46	3612.36	18.2	24.23	2.52
	R2		4878.61	4104.26	774.35	3766.29	3176.27	590.02	8.7	2.3	8.6	2.1	139470.8	2929.41	4	22.8	7.77
	C1		4190.81	3282.92	907.89	3386.09	2735.42	651.57	20.4	4.5	19.8	3.1	84088.73	2718.41	18.9	19.18	0.62
	C2		3117.46	2452.38	665.07	2619.39	2120.54	498.85	9.4	3.3	9.1	2.3	113850.28	1831.59	6	15.98	13.63
	RC1		5031.35	4177.94	853.41	4155.34	3523.9	631.45	23	4.8	21.4	3	119118.09	3178.68	18	17.41	9.8
	RC2		4147.28	3410.05	737.23	3368.65	2753.92	614.73	8.9	2.1	8.6	1.5	133396.96	2536.2	4.3	18.77	7.91
200	R1	50	6745.06	4948.78	1796.28	5125.27	3801.74	1323.52	19.5	9	18.4	6.8	115745.39	3612.36	18.2	24.01	4.98
	R2		6123.93	4104.26	2019.67	4823.3	3373.7	1449.6	8.7	4.9	8.5	3.9	143181.64	2929.41	4	21.24	13.17
	C1		5381.98	3282.92	2099.06	4304.88	2850.36	1454.52	20.4	13.2	20.2	9.1	106649.1	2718.41	18.9	20.01	4.63
	C2		3838.84	2452.38	1386.45	3364.55	2241.36	1123.19	9.4	6.8	9.3	5.5	134227.53	1831.59	6	12.36	18.28
	RC1		6249.5	4177.94	2071.56	4994.12	3628.75	1365.38	23	11.7	21.9	7	126426.44	3178.68	18	20.09	12.4
	RC2		5122.37	3422.53	1699.84	4092.32	2868.57	1223.74	9	5	9	4.4	138902.36	2536.2	4.3	20.11	11.59
200	R1	100	8141.25	4948.78	3192.47	6320.33	3931.83	2388.5	19.5	14	18.4	11.5	111237.11	3612.36	18.2	22.37	8.13
	R2		6943.14	4104.26	2838.88	5791.46	3541.39	2250.06	8.7	6.2	8.7	5.7	136609.31	2929.41	4	16.59	17.28
	C1		6232.89	3282.92	2949.96	4957.71	2830.42	2127.29	20.4	17.6	20.3	14.7	97884.35	2718.41	18.9	20.46	3.96
	C2		4439.16	2452.38	1986.78	3794.25	2287.9	1506.35	9.4	7.2	9.3	6.2	106821.47	1831.59	6	14.53	19.94
	RC1		7520.83	4177.94	3342.89	6049.63	3813.01	2236.62	23	18.2	22.2	11.9	133298.68	3178.68	18	19.56	16.64
	RC2		5865.87	3410.05	2455.81	5114.34	3093.92	2020.42	8.9	7.4	8.6	6.4	145371.98	2536.2	4.3	12.81	18.03

Table 3: Results for 25- and 200-customer instances with specialized parameter settings

| | | | Objective Values | | | | | | Number of vehicles | | | | | Literature | | | |
| | | | Initial solution | | | Final solution | | | Initial solution | | Final solution | | | | | | |
Size	Class	nbSyn	InObj	InRgObj	InSpObj	FnObj	FnRgObj	FnSpObj	InNbRgV	InNbSpV	FnNbRgV	FnNbSpV	RunTime	LitBestObj	LitBestnbVeh	Imp %	Gap%
25	R1	2	621.66	559.13	62.53	544.25	481.71	62.53	5.92	1	5.25	1	8406.73	463.45	4.92	12.45	3.79
	R2		577.89	515.35	62.53	458.21	395.68	62.53	4.18	1	3	1	9978.15	350.25	2.73	20.71	11.48
	C1		279.66	228.07	51.58	246.64	195.83	50.81	3.11	1.11	3	1	12728.32	190.59	3.00	11.81	2.68
	C2		356.23	287.77	68.46	286.41	217.95	68.46	2.38	1	2	1	7643.58	214.45	2	19.6	1.61
	RC1		572.33	407.01	165.32	520.13	355.45	164.68	3.88	1.5	3.25	1.13	6805.36	350.24	3.25	9.12	1.47
	RC2		558.28	393.8	164.47	485.18	320.7	164.47	3	1	2.88	1	7672.25	319.28	3	13.09	0.44
25	R1	7	832.64	559.13	273.51	692.97	487.86	205.1	5.92	3.83	5.33	2.83	7846.29	463.45	4.92	16.77	5
	R2		779.83	515.35	264.48	587.39	402.48	184.91	4.18	2.09	3	1.64	10289.56	350.25	2.73	24.68	12.98
	C1		360.15	215.47	144.68	318.35	197.04	121.31	3.11	2.89	3	2	8304.22	190.59	3.00	11.61	3.27
	C2		402.83	287.77	115.06	325.01	216.82	108.19	2.38	1	2	1	7308.46	214.45	2	19.32	1.09
	RC1		698.27	407.01	291.25	623.59	366.09	257.49	3.88	3	3.38	2.5	6388.25	350.24	3.25	10.7	4.33
	RC2		693.4	393.8	299.59	573.43	337.15	236.28	3	2	3	1.63	8828.17	319.28	3	17.3	5.3
25	R1	13	950.12	559.13	390.99	821.16	488.94	332.22	5.92	4.67	5.33	3.75	7683.6	463.45	4.92	13.57	5.21
	R2		913.19	515.35	397.84	701.39	405.41	295.98	4.18	2.55	2.82	2.09	10194.92	350.25	2.73	23.19	13.61
	C1		389.5	215.47	174.03	356.26	200.63	155.63	3.11	3	3	2.56	6594.24	190.59	3.00	8.53	5
	C2		476.72	287.77	188.95	384.08	223.37	160.7	2.38	1.13	2	1	6977.36	214.45	2	19.43	3.99
	RC1		731.55	407.01	324.54	657.53	357.33	300.2	3.88	3.38	3.38	3.13	5892.33	350.24	3.25	10.12	1.98
	RC2		709.59	393.8	315.78	600.02	332.03	267.99	3	2.25	2.88	2.13	8504.33	319.28	3	15.44	3.84
200	R1	10	5512.69	4948.78	563.91	4144.38	3679.06	465.32	19.5	3.1	18.6	2.5	126805.79	3612.36	18.2	24.82	1.81
	R2		4878.61	4104.26	774.35	3713.92	3112.7	601.22	8.7	2.3	8.6	2.2	161933.9	2929.41	4	23.87	5.89
	C1		4190.81	3282.92	907.89	3391.8	2751.12	640.68	20.4	4.5	19.8	3	100174.6	2718.41	18.9	19.07	1.19
	C2		3117.46	2452.38	665.07	2578.9	2078.57	500.33	9.4	3.3	9	2.3	145855.5	1831.59	6	17.28	11.88
	RC1		5031.35	4177.94	853.41	4088.11	3458.5	629.6	23	4.8	21.6	3	138911.3	3178.68	18	18.75	8.09
	RC2		4147.28	3410.05	737.23	3303.76	2689.92	613.85	8.9	2.1	8.6	2	149639.7	2536.2	4.3	20.34	5.71
200	R1	50	6745.06	4948.78	1796.28	5074.65	3729.01	1345.63	19.5	9	18.3	6.7	136439.97	3612.36	18.2	24.76	3.13
	R2		6123.93	4104.26	2019.67	4769.59	3300.61	1468.99	8.7	4.9	8.3	4.6	163572	2929.41	4	22.12	11.25
	C1		5381.98	3282.92	2099.06	4238.08	2767.45	1470.62	20.4	13.2	19.9	9.2	120709.5	2718.41	18.9	21.25	1.77
	C2		3838.84	2452.38	1386.45	3254.88	2122.86	1132.02	9.4	6.8	9	5.7	161222.8	1831.59	6	15.21	13.72
	RC1		6249.5	4177.94	2071.56	4944.82	3557.28	1387.55	23	11.7	21.9	7.1	144934.9	3178.68	18	20.88	10.64
	RC2		5137.58	3447.64	1689.93	4039.98	2782.35	1257.63	9.11	5.11	8.78	4.56	155597.56	2536.2	4.3	21.36	8.85
200	R1	100	8141.25	4948.78	3192.47	6139.27	3777.69	2361.58	19.5	14	18.6	11.7	137336.07	3612.36	18.2	24.59	4.38
	R2		6943.14	4104.26	2838.88	5616.64	3424.85	2191.78	8.7	6.2	8.6	5.6	157111.4	2929.41	4	19.11	14.47
	C1		6232.89	3282.92	2949.96	4964.76	2816.49	2148.27	20.4	17.6	20	15.1	111991.63	2718.41	18.9	20.35	3.48
	C2		4439.16	2452.38	1986.78	3646.12	2144.27	1501.85	9.4	7.2	9.2	6.3	132923.86	1831.59	6	17.86	14.58
	RC1		7520.83	4177.94	3342.89	5836.13	3628.99	2207.14	23	18.2	22	11.8	148733.9	3178.68	18	22.4	12.41
	RC2		5865.87	3410.05	2455.81	4882.75	2908.89	1973.86	8.9	7.4	8.5	6.6	159016.6	2536.2	4.3	16.76	12.81

References

Bélanger, N., Desaulniers, G., Soumis, F., Desrosiers, J., Periodic Airline Fleet Assignment with Time Windows, Spacing Constraints and Time Dependent Revenues. *European Journal of Operational Research 175*, 1754-1766, 2006.

Bredström, D., Rönnqvist, M., Combined Vehicle Routing and Scheduling with Temporal Precedence and Synchronization Constraints. *European Journal of Operational Research 191*, 19-29, 2008.

Dohn, A., Kolind, E., Clausen, J., The Manpower Allocation Problem with Time Windows and Job-Teaming Constraints: A Branch-and-Price Approach. *Computers & Operations Research 36*, 1145-1157, 2009.

Drexl, M., Synchronization in Vehicle Routing: A Survey of VRPs with Multiple Synchronization Constraints. *Transportation Science 46*, 297-316, 2012.

El Hachemi, N., Gendreau, M., Rousseau, L.-M., A Heuristic to Solve the Synchronized Log-truck Scheduling Problem. *Computers & Operations Research 40*, 666-673, 2013.

Freling, R., Huisman, D., Wagelmans, A.P.M., Models and Algorithms for Integration of Vehicle and Crew Scheduling. *Journal of Scheduling 6*, 63-85, 2003.

Haase, K., Desaulniers, G., Desrosiers, J., Simultaneous Vehicle and Crew Scheduling in Urban Mass Transit Systems. *Transportation Science 35*, 286-303, 2001.

Homberger, J., Gehring, H., Two Evolutionary Metaheuristics for the Vehicle Routing Problem with Time Windows. *INFOR 37*, 297-318, 1999.

Ioachim, I., Desrosiers, J., Soumis, F., Bélanger, N., Fleet Assignment and Routing with Schedule Synchronization Constraints. *European Journal of Operational Research 119*, 75-90, 1999.

Li, Y., Lim, A., Rodrigues, B., Manpower Allocation with Time Windows and Job-Teaming Constraints. *Naval Research Logistics 52*, 302-311, 2005.

Pesant, G., Gendreau, M., A Constraint Programming Framework for Local Search Methods. *Journal of Heuristics 5*, 255-279, 1999.

Pesant, G., Gendreau, M., A View of Local Search in Constraint Programming. In Principles and Practice of Constraint Programming, Freuder, E.C., eds., Lecture Notes in Computer Science 1118, Springer Berlin, pp. 353-366, 1996.

Pesant, G., Gendreau, M., Potvin, J.-Y., Rousseau, J.-M., An Exact Constraint Logic Programming Algorithm for the Traveling Salesman Problem with Time Windows. *Transportation Science 32*, 12-29, 1998.

Pisinger, D., Ropke, S., A General Heuristic for Vehicle Routing Problem. *Computers & Operations Research 34*, 2403-2435, 2007.

Shaw P., Using Constraint Programming and Local Search Methods to Solve Vehicle Routing Problems. In Principles and Practice of Constraint Programming, Goos, G., Hartmanis, J., van Leeuwen, J., eds., Lecture Notes in Computer Science 1520, Springer, Berlin, pp. 417-431, 1998.

Rasmussen, M.S., Justesen, T., Dohn, A., Larsen, J., The Home Care Crew Scheduling Problem: Preference-based Visit Clustering and Temporal Dependencies, *European Journal of Operational Research 219*, 598-610, 2012.

Rousseau, L.-M., Gendreau, M., Pesant, G., The Synchronized Vehicle Dispatching Problem. *INFOR 51*, 76-83, 2013.

Rousseau, L.-M., Gendreau, M., Pesant, G., Using Constraint-based Operators to Solve the Vehicle Routing Problem with Time Windows. *Journal of Heuristics 8*, 43-58, 2002.

Rousseau, L.-M., Gendreau, M., Pesant, G., Focacci, F., Solving VRPTWs with Constraint Programming Based Column Generation. *Annals of Operations Research 130*, 199-216, 2004.

Schmid, V., Doerner, K., Hartl, R., Salazar-González, J., Hybridization of Very Large Neighborhood Search for Ready-Mixed Concrete Delivery Problems. *Computers & Operations Research 37*, 559-574, 2010.

Schrimpf, G., Schneider, J., Stamm-Wilbrandt, H., Dueck, G., Record Breaking Optimization Results Using the Ruin and Recreate Principle. *Journal of Computational Physics 159*, 139-171, 2000.

Solomon, M.M., Algorithms for the Vehicle Routing and Scheduling Problems with Time Window Constraints. *Operations Research 35*, 254-265, 1987.