



CIRRELT

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

Interuniversity Research Centre
on Enterprise Networks, Logistics and Transportation

A Metaheuristic Based on Tabu Search for Solving a Technician Routing and Scheduling Problem

Ines Mathlouti
Michel Gendreau
Jean-Yves Potvin

January 2018

CIRRELT-2018-01

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

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

www.cirrelt.ca

A Metaheuristic Based on Tabu Search for Solving a Technician Routing and Scheduling Problem

Ines Mathlouthi^{1,2}, Michel Gendreau^{1,3}, Jean-Yves Potvin^{1,2,*}

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

² Department of Computer Science and Operations Research, Université de Montréal, P.O. Box 6128, Station Centre-Ville, Montréal, Canada H3C 3J7

³ Department of Mathematics and Industrial Engineering, Polytechnique Montréal, P.O. Box 6079, Station Centre-Ville, Montréal, Canada H3C 3A7

Abstract. This paper addresses a technician routing and scheduling problem motivated by an application for the repair and maintenance of electronic transactions equipment. The problem exhibits many special features like multiple time windows for service, an inventory of spare parts carried by each technician and tasks that may require a special part to be performed. A problem-solving methodology based on tabu search, coupled with an adaptive memory, is proposed. The inclusion of solutions (local minima) in the adaptive memory takes into account both quality and diversity. Results are reported on test instances with up to 200 tasks. A comparison with a previously developed branch-and-price algorithm is also reported on instances of small size.

Keywords: Technician routing and scheduling problem, multiple time windows, inventory, metaheuristic, adaptive memory, tabu search, biased fitness.

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

Results and views expressed in this publication are the sole responsibility of the authors and do not necessarily reflect those of CIRRELT.

Les résultats et opinions contenus dans cette publication ne reflètent pas nécessairement la position du CIRRELT et n'engagent pas sa responsabilité.

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

1 Introduction

This paper considers a Technician Routing and Scheduling Problem (TRSP) motivated by a real application for the repair or maintenance of electronic transactions equipment by a number of technicians. The problem is to assign a subset of tasks to each technician and to construct a route over each subset to optimize some objective, which may involve many criteria (like total traveled distance, overtime, etc.). The routes must also satisfy different types of constraints. First, there are compatibility constraints between tasks and technicians, since different skills are required to perform different tasks and a technician does not necessarily possess all those skills. Second, a task may also need a number of spare parts. If the technician does not have the required parts in his initial inventory when he starts his route, he can acquire them by visiting a depot at some point along the route. Typically, an infinite number of parts is assumed at the depot, although only a finite number can be carried by the technician. Third, a task may require a special part which can only be obtained by visiting the depot (i.e., a technician is not allowed to carry a special part from or to his home base location). Fourth, there are time bounds for the service of each task and for the return time of each technician at his home base location.

It is also assumed that not all tasks can be served by the technicians. Thus, a gain is associated with each task and the maximization of the total gain collected along the routes becomes part of the objective. The remaining tasks can then be considered for another day, with typically an increased gain (to avoid being left apart repeatedly). Hence, the gain is related to characteristics of the task or the customer who requests the service.

The remainder of the paper is organized as follows. In Section 2, we provide a review of various works on problems related to ours. In Section 3, we describe our problem. Then, after an introduction to recent ideas that have been integrated into tabu search to improve its performance, we detail our algorithm in Section 4. The test instances are introduced and are followed by computational results in Section 5. In particular, we provide a comparison with optimal solutions previously obtained with an exact method on instances of small sizes. Finally, a conclusion follows in Section 6.

2 Literature review

Most papers on the TRSP are based on real applications that exhibit specific features. To the best of our knowledge, the first work on this type of problem was reported in 1997 by Tsang and Voudouris for a telecommunications application [22]. To solve this problem, the authors used different types of local search heuristics. Weigel and Coo [24] then introduced a problem faced by a large retailer when providing on-site technical assistance. The authors solved the problem by developing a heuristic procedure to construct a route for each technician

and by improving each route individually with Or-opt exchanges [12]. In [1], the authors addressed a periodic maintenance problem for elevators and escalators. In this application, technician routes had to account for working regulations. A similar application was later addressed by Tang et al. [21] with a tabu search heuristic. In 2007, the French Operations Research Society (ROADEF) initiated a challenge based on a problem encountered by France Telecom. This problem first involves a multi-period horizon. Also, teams of technicians must be created because each task needs multiple skills with different proficiency levels. For this challenge, Hashimoto et al. [7] developed a greedy randomized adaptive search (GRASP) heuristic while Cordeau et al. [3] proposed an adaptive large neighborhood search (ALNS). Another multi-period TRSP for the maintenance of electric forklifts, where technicians are paired to form teams, is solved with a branch-and-price algorithm in Zamorano and Stolletz [25].

A parallel matheuristic is proposed in [14] for a TRSP where tools and spare parts are taken into account. The matheuristic is made of a constructive heuristic, a parallel ALNS and a mathematical programming-based post-optimization procedure. In Mendoza et al. [11] a technician routing problem with conventional and electric vehicles is introduced. Due to their relatively limited driving range, the electric vehicles need to visit one or more recharging stops along their route. The problem is also addressed with a parallel matheuristic, where a number of subproblems are first created and solved with GRASP. The routes of all local minima produced by GRASP are collected to create a repository of routes. Then, a set covering model is solved over these routes.

In [2], the authors introduced a problem faced by a water treatment and distribution company. Here, the technician routes must be planned over a period of one week for repair or maintenance. The tasks to be scheduled can either be known in advance or can occur dynamically. A memetic algorithm is used to solve the problem, which is first applied on the static tasks to produce initial tours for every day of the week. Then, the dynamic tasks are integrated into the current routes as they occur, still with the memetic algorithm. An exact approach based on column generation is also proposed in this work, but can only be applied to instances of small size. Pillac et al. [13] also addressed a TRSP in which a fraction of the tasks occur dynamically. A parallel architecture is proposed to speed up the calculations. An initial solution is first created with the static tasks using a construction heuristic based on a regret measure [17]. This solution is then improved with ALNS [16]. The latter algorithm works by successively destroying (removing tasks) and repairing (reinserting tasks) the current solution to produce a new solution. For dynamic tasks, the part of the current solution already executed is fixed and the newly occurring tasks are incorporated into the solution by running the ALNS for a limited number of iterations. The same authors also proposed a fast reoptimization approach based on a parallel ALNS in [15].

In [10], a mixed integer programming (MIP) model was proposed for our TRSP. Although the model is useful by providing a formal description of the

problem, solving the model exactly with a commercial solver proved to be impractical, except for very small instances with no more than 15 tasks. In a later work [9], a branch-and-price algorithm was developed and was able to solve exactly instances with up to 45 tasks. Given the limitations of exact approaches, we now propose in this paper a tabu search heuristic to address instances of larger sizes.

3 Problem Definition

We start this section by describing the main entities, with their attributes, involved in our problem (for a formal MIP formulation, see [10]). They are:

- Tasks
 - Skills: each task requires one or more skills from a technician to be performed.
 - Parts: each task requires one or more spare parts of different types and, possibly, a special part to be performed.
 - Gain: a gain is associated with each task that represents its importance (based on service priority, revenue, customer status, etc.).
 - Multiple time windows: a technician must begin his service within one of multiple time windows associated with task i , where a time window is defined by a lower bound e_i and an upper bound l_i . If the technician arrives before the lower bound, he can wait and start the service at the lower bound.
- Technicians
 - Home base: starting and ending location of the technician's route.
 - Skills: each technician has one or more skills that allow him to serve certain tasks and not others.
 - Depot: each technician is a priori assigned to a single depot (among a number of possible depots) where he can replenish his inventory of parts.
 - Workday: each technician normally works between 9:00 AM and 5:00 PM, although overtime is allowed (at the expense of a penalty in the objective). Three breaks can be taken during the day: two breaks of 15 minutes in the morning and afternoon and a mid-day break of 30 minutes. Each break must be taken within a specific time window and a technician must take the break if his schedule intersects with that time window. Furthermore, the route of a technician cannot exceed a maximum traveled distance.

- Inventory: each technician carries an initial inventory of spare parts when he leaves his home base location. If there are not enough spare parts to serve all tasks along a technician’s planned route, or if one or more tasks require a special part, then a single visit to a preassigned depot is allowed along the route.
- Parts
 - Spare parts: there are different types of spare parts. Although the depot contains a virtually infinite number of spare parts, the technician can only carry a limited number of parts of each type.
 - Special parts: special parts are also available at the depot. A technician must visit the depot when one or more tasks along his route require a special part. No special part can be carried to or from the home base location.
- Depots
 - There are a fixed number of depots with a virtually infinite number of spare parts.

The problem is to design a route over a subset of tasks for each of a fixed number of technicians while satisfying the hard constraints mentioned above, namely, the required skills and required parts to perform a task, the multiple time windows for the service of a task and the maximum traveled distance of each route. Given that not all tasks can be served, the objective is to maximize the total gain collected, minus the total traveled distance and total overtime over all routes (or, equivalently, to minimize the sum of total traveled distance and total overtime minus the total gain collected). Each component in the objective has a weighting parameter to adjust its importance when evaluating a solution.

4 Problem-solving methodology

Our problem-solving methodology is based on tabu search [6] which has proven successful for a variety of hard combinatorial problems, like vehicle routing [4, 5], job shop scheduling [8], quadratic assignment [19], technician routing and scheduling [21], and many others. Modern implementations involve the integration of adaptive memories [18] made of elite solutions or components of elite solutions in order to perform search intensification or search diversification. An example for a vehicle routing problem can be found in [20] where the routes of elite solutions are stored in memory and used to construct new starting solutions for the tabu search. Basically, a new solution is obtained by mixing routes from different elite solutions in memory. In our tabu search, we manage the adaptive memory as a true population of solutions with considerations for both quality and diversity. This approach has proven very successful in the hybrid genetic

algorithm of Vidal et al. [23] when solving different types of vehicle routing problems.

The proposed metaheuristic is shown in Algorithm 1, where s stands for the current solution and s^* for the best solution. It starts with a number of calls to the function Greedy(). The latter randomly applies one of two greedy construction heuristics to create an initial solution which is then improved with a local search descent. The resulting solutions are then filtered out to remove duplicates and are stored in adaptive memory. The best solution in adaptive memory is then used as the initial starting solution. The outer **while** loop is aimed at stopping the algorithm and returning the best solution s^* when a maximum number of iterations or maximum computation time has been reached. Each iteration of the inner **while** loop involves the application of four consecutive tabu searches, each with a different neighborhood structure. The best solution found by a tabu search with a given neighborhood structure is returned and fed to the next one. Each tabu search also takes care to update s^* during its execution, if required. The stopping criterion for any given tabu search is based on a maximum number of iterations. Note also that every local minimum reached by a tabu search is stored in adaptive memory. The inner **while** loop is repeated until a complete pass through the four neighborhoods does not provide any improvement. At this point, the adaptive memory is filtered out to remove duplicates and is possibly reduced if the maximum memory size n_{max} is exceeded. Then, a new starting solution is created by combining routes of different solutions in adaptive memory. In the following, the main components of this search framework are described.

Algorithm 1

```

1: procedure SEARCH
2:   for nInit iterations do
3:      $s_{init} \leftarrow$  Greedy ()
4:      $s_{init} \leftarrow$  Local Search ( $s_{init}$ )
5:     Store  $s_{init}$  in adaptive memory
6:    $s^* \leftarrow$  select best solution in adaptive memory
7:    $s \leftarrow s^*$ 
8:   while number of iterations  $< nIter_{max}$  or time  $< T_{max}$  do
9:     while  $s$  is improved do
10:       $s \leftarrow$  Tabu ( $s$ , Inter-Route Move)
11:       $s \leftarrow$  Tabu ( $s$ , Intra-Route Move)
12:       $s \leftarrow$  Tabu ( $s$ , Swap)
13:       $s \leftarrow$  Tabu ( $s$ , Swap-With-New)
14:     if number of solutions in adaptive memory  $> n_{max}$  then
15:       Update memory
16:      $s \leftarrow$  Get solution from adaptive memory
return  $s^*$ 

```

4.1 Construction heuristics

The adaptive memory is initialized with n_{min} different solutions created with two randomized construction heuristics, each contributing to half of the solutions. These solutions are then improved with a local descent based on the same neighborhood structures than the ones used in the tabu search (see section 4.2). In our experiments, n_{min} was set to ten times the number of technicians. The two construction heuristics are the following.

Sequential construction heuristic. In this heuristic, the routes of the technicians are constructed one by one. That is, as long as technicians are available, one of them is randomly selected and his route is initially created with a starting and ending location (home base location of the technician) and the three breaks. Then, a task is randomly selected among those that can be performed by the technician. The three best feasible insertions in the current route of the technician are considered and one of them is randomly selected. This is repeated until no new task can be added to the route. In case the current insertion requires a visit to the depot, the three best feasible insertions for the depot are kept by creating three different routes. The insertion procedure then follows with the three routes and the best route is selected at the end. By keeping a number of different routes, each associated with a different insertion place for the depot, the myopic behavior of the construction heuristic is alleviated.

Parallel construction heuristic. This heuristic constructs routes for all technicians in parallel. The route of each technician is first initialized with his starting and ending locations and the three breaks. Then, a task is randomly selected and the best feasible insertion in the route of each technician who can perform it is considered. The task is assigned to the best technician and inserted in his route. This is repeated until no more tasks can be feasibly inserted in the routes. During this insertion procedure, the depot is handled as in the sequential heuristic.

4.2 Tabu search

The tabu search improves the starting solutions obtained from the adaptive memory. Its solution space and neighborhood structures are now described.

4.2.1 Solution space

In the following, a solution is considered feasible (and is part of the solution space) even if the maximum traveled distance constraint is not satisfied. In such a case, the solution is immediately repaired. The repair procedure is applied in turn to each route that exceeds the maximum distance constraint. This procedure is very simple and removes, at each iteration, the task with the least impact on the objective value until the maximum distance constraint is satisfied

again.

4.2.2 Neighborhoods

The tabu search exploits a sequence of four different neighborhood structures. More precisely:

- *Inter-Route Move.* A task is moved from one route to another and inserted at the best feasible insertion place.
- *Intra-Route Move.* A task is moved from its current position to another feasible position in the same route.
- *Swap.* Two tasks from two different routes are swapped. Each task is inserted at the best feasible insertion place in its new route.
- *Swap-With-New.* A task not currently in the solution is swapped with a task in the solution and is inserted at the best feasible insertion place.

After a move, each task not currently in the solution is considered for insertion in the routes that have been modified (in non increasing order of gain). During the exploration of each neighborhood, a first improvement strategy is implemented. Thus, the first feasible neighbor that provides an improvement over the current solution is chosen. Tabu restrictions are also imposed on each neighborhood. That is, tasks involved in recent moves are declared tabu for θ iterations (unless they can improve the best known solution), where θ is randomly selected in the interval $[\theta_{min}, \theta_{max}]$, with $\theta_{min} = 5$ and $\theta_{max} = 10$ in our experiments.

It should be noted that the algorithm loops through the four neighborhoods (in the order indicated above) as long as the solution improves. If a pass through the four neighborhoods does not provide any improvement, then a new initial solution is obtained from the adaptive memory to restart the search.

4.3 Adaptive Memory

Here, we describe the procedures used to store and fetch solutions from the adaptive memory. Since every solution decomposes into a route per technician, it should be noted that the memory is implicitly divided into a number of smaller memories, where each memory contains the routes of a given technician (see Figure 1). This partition is required during the fetching procedure because each technician has specific routes that start from and end at a different home base location and visit a particular preassigned depot. In the following, we first explain how the biased fitness of a solution in adaptive memory is computed, before describing the storing and fetching mechanisms.

4.3.1 Biased fitness

The quality of a solution in memory corresponds, on the one hand, to its objective value and, on the other hand, to its contribution to the diversity of solutions in that memory [23]. The diversity computation is based on the Hamming distance between two solutions s_1 and s_2 , see Equation (1). In this equation, $\Gamma(x)$ is an indicator function: it returns 1 if x is true, 0 otherwise; $T(i, s)$ returns the technician who performs task i in solution s , while $W(i, s)$ returns the index of the time window for the service of task i in solution s (since there are multiple time windows). When task i is not served by any technician, a dummy value is returned.

$$\delta^H(s_1, s_2) = \sum_{i \in s_1 \cup s_2} \Gamma(T(i, s_1) \neq T(i, s_2)) + \Gamma(W(i, s_1) \neq W(i, s_2)) \quad (1)$$

The diversity contribution of solution s is then computed as follows:

$$\Delta(s) = \frac{1}{n_c} \sum_{s' \in N_c} \delta^H(s, s') \quad (2)$$

Thus, it is the average of the Hamming distances between s and the $n_c = |N_c|$ closest solutions in adaptive memory, where n_c is a parameter (set to 20% of the solutions in adaptive memory, as in [23]). Then, the ranks $r_f(s)$ and $r_d(s)$ of each solution s in memory with regard to the objective value and diversity contribution, respectively, are computed (where rank 1 is best). With these ranks, the biased fitness BF of solution s is computed as follows, where n_m is the current number of solutions in memory:

$$BF(s) = (n_m - r_f(s) + 1) + \eta(n_m - r_d(s) + 1) \quad (3)$$

In this formula, the weight of the objective value component (first term in the summation) is implicitly set to 1 to guarantee a minimum contribution of this component to the biased fitness. Then, a more or less important bias towards diversity is added depending on the value of parameter $\eta \in [0, 1]$.

4.3.2 Storing and updating

Every local minimum visited by the tabu search is added to the adaptive memory. After one pass through the four neighborhoods (see Algorithm 1), the memory is updated as follows:

- duplicates (i.e., solutions with the same objective value) are removed;

- the biased fitness of each solution is computed;
- if the maximum memory size n_{max} is exceeded, the n_{max} solutions with the best biased fitness are kept. In our experiments, this parameter was set to one hundred times the number of technicians.

4.3.3 Fetching

As previously mentioned, each solution is decomposed into a route for each technician, where each route is stored in the memory of the corresponding technician. In this process, the route also inherits the biased fitness value of the whole solution.

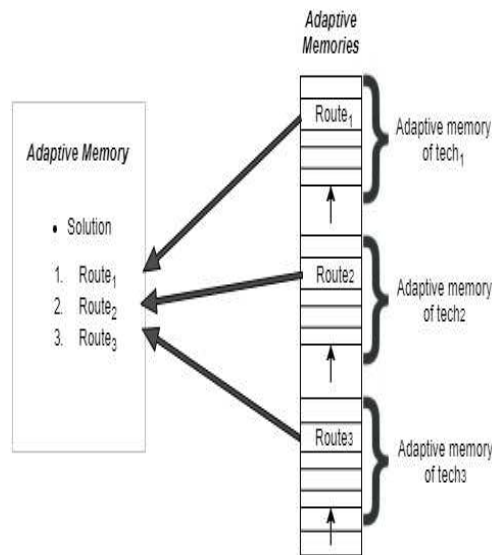


Figure 1: Adaptive memories

The procedure to create a starting solution for the tabu search is the following. We proceed technician by technician. A first technician is considered and each route in the memory of this technician is assigned a probability of selection which is proportional to its biased fitness (roulette wheel). A route r is probabilistically chosen and included in the starting solution. Then, the memory of all other technicians is updated by removing any route with a task in common with route r . This process of (1) considering the next technician, (2) probabilistically selecting a route of that technician, (3) adding this route to the starting solution and (4) updating the memory of all remaining technicians, is repeated until all technicians are done (a complete starting solution is obtained) or until the remaining technicians have no valid routes in memory. In the latter

case, the solution is completed by inserting additional tasks with the parallel construction heuristic of section 4.1.

5 Computational experiments

In this section, the test instances for the experiments are first described. Then, we report the results obtained with three different variants of our metaheuristic. We also compare the solutions produced by our algorithm with the optimum on instances of small size.

5.1 Test instances

The test instances come from [10], where the crow fly distance is assumed between each pair of locations and the speed of the vehicles is set at 50km/h. The other characteristics are the following.

- *Service area.* The service area is a $40 \text{ km} \times 40 \text{ km}$ or $50 \text{ km} \times 50 \text{ km}$ squared area.
- *Depot Location.* There are 3 depots randomly located within the service area.
- *Task location.* Each task is randomly located within the service area.
- *Task service time.* The service time of each task is randomly chosen between 30 and 45 minutes.
- *Task gain.* The gain associated with a task is randomly chosen between 1 and 10.
- *Technician home base.* To get a good coverage of the tasks, the first two technicians are located at the opposite ends of the service area (along the diagonal). The other technicians are randomly located within the service area.
- *Technician skills.* With each technician is associated the percentage of tasks that he can perform: one third of the technicians can perform all tasks, one third of the technicians can perform 50% of the tasks and one third 25% of the tasks.
- *Parts.* The number of spare parts needed to perform a task is randomly chosen between 0 and 3. Then, each spare part is assigned a type, among 4 different types. Also, a special part is needed with a probability of 0.125.
- *Time windows.* Both narrow and wide time windows are considered. The latter are twice as wide as the former on average. The length of a narrow

time window is randomly generated between 60 and 90 minutes. The lower bound of the first time window is chosen randomly between 9:00 AM and noon. The lower bounds of the remaining time windows are set between 2 and 3 hours after the upper bound of the previous time window until a maximum of 3 time windows are obtained.

- *Maximum distance.* The maximum distance traveled by each technician during his workday is set to 125 km.

Different subsets of instances were obtained by varying the values of the following characteristics: size of the service area (either $40\text{km} \times 40\text{km}$ or $50\text{km} \times 50\text{km}$), width of the time windows (either narrow or wide) and number of tasks (50, 100 or 200). Furthermore, tests were performed with 3 and 6 technicians for instances with 50 tasks, 6 and 12 technicians for instances with 100 tasks, and 12 and 24 technicians for instances with 200 tasks. Thus, we have $2 * 2 * 3 * 2 = 24$ subsets of instances, with 10 instances in each subset. In the tables of results, the subsets are identified with the following fields : width of the time windows (either N for narrow or W for wide), size of the service area (either 40 or 50), number of tasks and number of technicians. For example $N-40-50-3$ is the subset of instances with narrow time windows, a $40 \text{ km} \times 40 \text{ km}$ service area, 50 tasks and 3 technicians.

Finally, the weights in the objective function were set as follows: 500 for the gain, 5 for the distance (in kilometers) and 1 for overtime (in seconds). Overall, more emphasis is given to the total gain over the total distance and overtime. This setting allows a technician to do some overtime to perform a high gain task and is more challenging than the setting proposed in [10].

5.2 Results

Here, we report the results of different experiments on a 3.07GHz Intel Xeon X5675 processor, using the test instances introduced above. Our problem-solving methodology was run for $T_{max} = 1$ hour on the instances of size 50 and 100, which was enough to allow convergence. For the instances of size 200, a computation time of $T_{max} = 3$ hours was required. Based on preliminary experiments, the number of iterations with each neighborhood was set as follows: 75 for Inter-Route Move, 20 for Intra-Route Move, 100 for Swap and 100 for Swap-With-New.

5.2.1 Impact of parameter η

In this section, we explore the impact of parameter η on the performance of our metaheuristic. This parameter controls the magnitude of the diversity contribution when a solution is evaluated in adaptive memory (biased fitness). When $\eta = 0$, only the objective value of a solution is considered. Conversely, when $\eta =$

1, the objective value and the diversity contribution have the same weight. Six different values for η between 0 and 1 were tested. For each subset of instances and each η value, Table 1 reports the average gap in percentage with the best solution produced on each instance over the six η values. The best gap for each subset is shown in bold. We can see that including diversity has a positive impact, since the worst gaps are clearly obtained with $\eta = 0$. When considering all test instances, the best η value is equal to 0.6 with an average gap of 1.87%, as indicated in the line Overall of Table 1. Thus, this value will be used in the following experiments.

Name	Gap (%)					
	$\eta = 1$	$\eta = 0.8$	$\eta = 0.6$	$\eta = 0.4$	$\eta = 0.2$	$\eta = 0$
N-40-50-3	1.90	1.85	2.25	2.73	1.93	3.79
N-40-50-6	3.48	2.94	1.91	3.47	5.41	5.06
N-50-50-3	3.61	2.15	3.41	2.11	5.36	6.48
N-50-50-6	2.61	1.89	1.92	3.20	2.53	3.15
W-40-50-3	1.70	2.88	1.84	2.41	2.41	4.49
W-40-50-6	3.22	3.99	3.36	3.99	4.04	4.99
W-50-50-3	1.42	2.77	1.39	2.13	2.45	3.52
W-50-50-6	2.96	1.61	1.70	2.75	3.50	5.24
Avg.	2.61	2.51	2.22	2.85	3.45	4.59
N-40-100-6	1.03	1.27	2.17	1.52	1.95	2.46
N-40-100-12	1.54	2.65	1.29	1.87	1.64	1.45
N-50-100-6	1.78	2.27	2.23	0.36	2.44	4.18
N-50-100-12	1.53	1.97	2.30	1.16	2.28	1.93
W-40-100-6	1.15	1.36	2.27	2.22	1.75	2.34
W-40-100-12	1.15	1.32	0.98	1.61	1.51	1.53
W-50-100-6	4.05	2.06	2.53	2.75	1.00	4.27
W-50-100-12	1.94	2.06	1.77	3.12	1.96	1.73
Avg.	1.77	1.87	1.94	1.82	1.82	2.49
N-40-200-12	0.58	2.78	2.09	4.07	5.62	2.51
N-40-200-24	1.25	2.69	0.16	2.36	5.49	3.76
N-50-200-12	2.06	1.81	4.15	5.25	6.22	2.60
N-50-200-24	0.60	2.76	0.31	5.04	1.14	4.55
W-40-200-12	0.68	2.84	1.57	4.59	6.21	6.57
W-40-200-24	7.69	1.36	1.68	1.38	3.77	4.55
W-50-200-12	5.95	3.56	1.42	10.17	8.30	9.98
W-50-200-24	1.74	4.02	0.25	5.40	5.18	5.07
Avg.	2.57	2.73	1.45	4.78	5.24	4.95
Overall	2.32	2.37	1.87	3.15	3.50	4.01

Table 1: Impact of parameter η

5.2.2 Comparison with three other variants

In this section, we compare our algorithm, called *TS-AM-I*, with three other variants:

TS-AM. In this variant with adaptive memory, the maximum traveled distance cannot be exceeded when exploring the neighborhood of the current solution. Thus, a neighbor solution is ignored if the maximum distance of one or more technician routes is exceeded.

TS-I. This variant is obtained by removing the adaptive memory (AM) from *TS-AM-I*. One of the two greedy construction heuristics is applied, through a call to `Greedy()`, when it is time to get a new starting solution.

TS. This variant is obtained by removing the adaptive memory (AM) from *TS-AM*. One of the two greedy construction heuristics is applied, through a call to `Greedy()`, when it is time to get a new starting solution.

Table 2 reports, for each subset of 10 instances and each variant, the number of times a variant found the best solution (over the four variants) and the average gap in percentage over the best solutions. Also, for each variant, the average CPU time in seconds to reach its best solutions is indicated.

These results demonstrate the importance of considering solutions that exceed the maximum traveled distance, as well as including the adaptive memory with biased fitness in our problem-solving methodology. In fact, *TS-AM-I* found the best solution for approximately 85% of the instances, as indicated in the line Overall of Table 2 (c.f., 8.46).

Name	TS			TS-I			TS-AM			TS-AM-I		
	# Bst	Gap (%)	CPU (s)	# Bst	Gap (%)	CPU (s)	# Bst	Gap (%)	CPU (s)	# Bst	Gap (%)	CPU (s)
N-40-50-3	0	26.08	1 549.55	0	19.65	1 612.35	0	11.38	776.67	10	0.00	598.13
N-40-50-6	0	11.43	1 324.47	0	8.70	1 318.54	3	4.02	1 677.11	7	0.46	1 455.49
N-50-50-3	0	37.29	2 174.15	0	27.11	1 690.32	0	21.01	1 003.99	10	0.00	837.48
N-50-50-6	0	19.58	1 257.15	1	13.20	1 269.87	0	8.66	1 210.93	9	0.00	915.21
W-40-50-3	2	26.81	1 027.83	1	22.08	1 007.78	3	11.52	964.14	4	0.00	815.13
W-40-50-6	0	8.29	1 286.84	0	6.02	1 270.81	2	3.29	1 277.39	8	0.07	1 080.06
W-50-50-3	0	37.78	1 321.57	0	30.33	1 416.81	0	22.03	1 423.29	10	0.00	409.62
W-50-50-6	0	17.74	1 714.26	0	11.72	1 615.73	0	8.07	1 095.96	10	0.15	1 607.55
Avg.	0.25	23.12	1 456.98	0.25	17.35	1 400.28	1	11.25	1 178.69	8.50	0.09	964.83
N-40-100-6	0	10.76	1 664.72	0	9.13	1 661.01	0	7.09	1 217.81	10	0.00	1 166.25
N-40-100-12	0	7.11	1 912.72	0	5.64	1 712.91	0	3.47	1 374.16	10	0.46	1 375.70
N-50-100-6	0	24.07	1 205.76	0	17.13	1 286.56	2	13.13	809.84	8	0.00	791.34
N-50-100-12	0	8.23	1 579.18	5	6.81	1 697.00	0	6.14	1 439.92	5	0.00	1 193.54
W-40-100-6	0	9.24	1 503.54	0	6.85	1 614.83	0	5.02	1 519.72	10	0.25	1 499.16
W-40-100-12	0	5.84	1 812.49	1	2.74	1 620.98	0	2.11	1 219.38	9	1.08	1 084.01
W-50-100-6	0	27.93	1 291.44	1	16.88	1 381.28	1	13.38	1 123.35	8	0.00	1 026.33
W-50-100-12	0	6.74	1 312.40	4	4.02	1 316.81	1	5.51	1 146.69	5	0.22	575.40
Avg.	0	12.49	1 535.28	1.38	8.65	1 536.42	0.5	6.98	1 231.36	8.13	0.25	1 088.97
N-40-200-12	0	3.68	5 633.41	1	3.14	5 686.82	0	2.91	5 449.07	9	1.03	5 237.30
N-40-200-24	0	4.21	7 146.86	0	4.23	7 779.88	0	1.76	7 035.13	10	0.23	6 062.68
N-50-200-12	0	10.71	6 018.59	0	6.03	6 107.62	0	4.86	5 918.12	10	0.17	5 807.03
N-50-200-24	0	5.79	8 301.10	0	2.31	8 618.49	0	5.27	8 605.12	10	2.50	7 308.05
W-40-200-12	3	3.75	5 489.13	0	2.98	5 676.70	2	2.39	5 367.33	5	1.05	4 098.60
W-40-200-24	0	2.38	7 960.04	1	2.43	7 716.51	1	0.06	8 000.05	8	2.40	6 193.89
W-50-200-12	0	13.67	6 171.18	1	8.38	6 292.68	0	10.50	6 226.22	9	0.19	5 100.04
W-50-200-24	0	3.31	8 821.95	1	1.84	8 824.47	0	2.08	8 617.05	9	0.34	7 574.70
Avg.	0.38	5.94	6 942.78	0.5	3.92	7 087.89	0.38	3.73	6 902.26	8.75	0.99	5 922.79
Overall	0.21	13.85	3 311.68	0.71	9.97	3 341.53	0.63	7.32	3 104.10	8.46	0.44	2 658.86

Table 2: Comparison of three different variants

5.2.3 Served tasks

Table 3 provides a picture of the average number ($\#Ta$) and percentage ($\%Ta$) of served tasks, as well as the average number of tasks served per technician ($\#Ta/Te$) for configurations with less or more technicians ($\#Te$). For instances with 50 tasks, 3 technicians can serve about 40% of the tasks in the case of the smaller service area, with 6 or 7 tasks in each route. This percentage decreases for the larger service area because more time is needed to get from one task to another. When 6 technicians are available, the number of served tasks obviously increases, although the multiplier tends to be slightly less than 2. It should also be noted that more tasks can be served when the time windows are wide. Similar trends are observed for the instances of size 100 and 200.

Name	Less technicians				More technicians			
	# Te	# Ta	% Ta	# Ta/Te	# Te	# Ta	% Ta	# Ta/Te
N-40-50	3	20.2	40.4%	6.7	6	37.5	75.0%	6.3
N-50-50	3	15.2	30.4%	5.1	6	31.8	63.6%	5.3
W-40-50	3	22.9	45.8%	7.6	6	38.1	76.2%	6.4
W-50-50	3	17.6	35.2%	5.9	6	32.1	64.2%	5.4
N-40-100	6	50.5	50.5%	8.4	12	83.6	83.6%	7.0
N-50-100	6	40.3	40.3%	6.7	12	69.3	69.3%	5.8
W-40-100	6	50.7	50.7%	8.5	12	85.2	85.2%	7.1
W-50-100	6	42.0	42.0%	7.0	12	70.3	70.3%	5.9
N-40-200	12	107.5	53.8%	9.0	24	180.5	90.3%	7.5
N-50-200	12	90.3	45.2%	7.5	24	155.8	77.9%	6.5
W-40-200	12	101.0	50.5%	8.4	24	193.1	96.6%	8.0
W-50-200	12	103.0	51.5%	8.6	24	161.0	80.8%	6.7

Table 3: Served tasks

5.3 Comparison with optimal solutions

Optimal solutions obtained with a branch-and-price algorithm are reported in [10]. This algorithm was able to routinely solve instances with up to 25 tasks and TS-AM-I found the optimum in each case. The branch-and-price algorithm began to strive for instances with 30 tasks and was not able to solve any instance with 50 tasks or more (within 24 hours of computation time). Again, TS-AM-I was able to optimally solve all instances for which the optimum was known.

6 Conclusion

We proposed a metaheuristic approach to address a TRSP for which exact methods can only solve small instances. Given the practical interest of this type of problem, the proposed methodology opens the way for further progress in this area. The metaheuristic is based on a tabu search enhanced with an adaptive memory, where the evaluation of each solution in memory is driven by both its cost and its contribution to diversity. Our algorithm reached the optimum on each try for instances with less than 50 tasks. It can also solve larger instances, as indicated by the results reported in this paper for instances with up to 200 tasks. Among many possible avenues of research, we now want to explore the application of our metaheuristic to dynamic variants of our TRSP, where new service requests must be integrated in real-time into the current routes.

References

- [1] F. Blakeley, B. Argüello, B. Cao, W. Hall, and J. Knolmajer. Optimizing periodic maintenance operations for Schindler Elevator Corporation. *Interfaces*, 33(1):67–79, 2003.
- [2] N. Bostel, P. Dejax, P. Guez, and F. Tricoire. Multiperiod planning and routing on a rolling horizon for field force optimization logistics. In B. Golden, S. Raghavan, and E. Wasil, editors, *The vehicle routing problem: Latest advances and new challenges*, pages 503–525. Springer, 2008.
- [3] J-F. Cordeau, G. Laporte, F. Pasin, and S. Ropke. Scheduling technicians and tasks in a telecommunications company. *Journal of Scheduling*, 13(4):393–409, 2010.
- [4] M. Gendreau, F. Guertin, J.-Y. Potvin, and É. Taillard. Parallel tabu search for real-time vehicle routing and dispatching. *Transportation Science*, 33(4):381–390, 1999.
- [5] M. Gendreau, A. Hertz, and G. Laporte. A tabu search heuristic for the vehicle routing problem. *Management Science*, 40(10):1276–1290, 1994.
- [6] F. Glover. Tabu search-Part I. *ORSA Journal on computing*, 1(3):190–206, 1989.
- [7] H. Hashimoto, S. Boussier, M. Vasquez, and C. Wilbaut. A GRASP-based approach for technicians and interventions scheduling for telecommunications. *Annals of Operations Research*, 183(1):143–161, 2011.
- [8] S. Kawaguchi and Y. Fukuyama. Reactive tabu search for job-shop scheduling problems. In *Proceedings of the 11th International Conference on Computer Science Education (ICCSE)*, pages 97–102, 2016.

- [9] I. Mathlouthi, M. Gendreau, and J.-Y. Potvin. Branch-and-price for a multi-attribute technician routing and scheduling problem. Technical Report CIRRELT-2017-56, Interuniversity Research Center on Enterprise Networks, Logistics and Transportation, 2017.
- [10] I. Mathlouthi, M. Gendreau, and J.-Y. Potvin. Mixed integer linear programming for a multi-attribute technician routing and scheduling problem. *Information Systems and Operational Research*, 56(1):33–49, 2018.
- [11] J.-E. Mendoza, A. Montoya, C. Guéret, and J.-G. Villegas. A parallel matheuristic for the technician routing problem with conventional and electric vehicles. In *12th Metaheuristics International Conference (MIC)*, Barcelona, Spain, 2017.
- [12] I. Or. Traveling salesman-type combinatorial problems and their relation to the logistics of blood banking. Technical report, PhD dissertation, Department of Industrial Engineering and Management Sciences, Northwestern University, Evanston, IL, 1976.
- [13] V. Pillac, C. Guéret, and A. Medaglia. On the dynamic technician routing and scheduling problem. In *5th International Workshop on Freight Transportation and Logistics (ODYSSEUS)*, Mikonos, Greece, 2012.
- [14] V. Pillac, C. Guéret, and A. Medaglia. A parallel matheuristic for the technician routing and scheduling problem. *Optimization Letters*, 7(7):1525–1535, 2013.
- [15] V. Pillac, C. Guéret, and A. Medaglia. A fast reoptimization approach for the dynamic technician routing and scheduling problem. In L. Amadeo, E.-G. Talbi, and F. Yalaoui, editors, *Recent Developments in Metaheuristics*, pages 347–367. Springer, 2018.
- [16] D. Pisinger and S. Ropke. A general heuristic for vehicle routing problems. *Computers & Operations Research*, 34(8):2403–2435, 2007.
- [17] J.-Y. Potvin and J.-M. Rousseau. A parallel route building algorithm for the vehicle routing and scheduling problem with time windows. *European Journal of Operational Research*, 66(3):331–340, 1993.
- [18] Y. Rochat and É.D. Taillard. Probabilistic diversification and intensification in local search for vehicle routing. *Journal of Heuristics*, 1(1):147–167, 1995.
- [19] J. Tabitha, C. Rego, and F. Glover. Multistart tabu search and diversification strategies for the quadratic assignment problem. *IEEE Transactions on Systems, Man, and Cybernetics-Part A: Systems and Humans*, 39(3):579–596, 2009.

- [20] É. Taillard, P. Badeau, M. Gendreau, F. Guertin, and J.-Y. Potvin. A tabu search heuristic for the vehicle routing problem with soft time windows. *Transportation Science*, 31(2):170–186, 1997.
- [21] H. Tang, E. Miller-Hooks, and R. Tomastik. Scheduling technicians for planned maintenance of geographically distributed equipment. *Transportation Research Part E: Logistics and Transportation Review*, 43(5):591 – 609, 2007.
- [22] E. Tsang and C. Voudouris. Fast local search and guided local search and their application to British Telecom’s workforce scheduling problem. *Operations Research Letters*, 20(3):119 – 127, 1997.
- [23] T. Vidal, T.G. Crainic, M. Gendreau, and C. Prins. A unified solution framework for multi-attribute vehicle routing problems. *European Journal of Operational Research*, 234(3):658 – 673, 2014.
- [24] D. Weigel and B. Cao. Applying GIS and OR techniques to solve Sears technician-dispatching and home delivery problems. *Interfaces*, 29(1):112–130, 1999.
- [25] E. Zamorano and R. Stolletz. Branch-and-price approaches for the multiperiod technician routing and scheduling problem. *European Journal of Operational Research*, 257(1):55–68, 2017.