Computers & Operations Research 39 (2012) 249–258

Contents lists available at ScienceDirect

Computers & Operations Research

journal homepage: www.elsevier.com/locate/caor

A metaheuristic for a teaching assistant assignment-routing problem

Pablo Maya [a,b,*], Kenneth Sörensen [a], Peter Goos [a,c]

[a] University of Antwerp, Faculty of Applied Economics, ANT/OR – University of Antwerp Operations Research Group, Stadscampus S.B.513 Prinsstraat 13, Antwerp 2000, Belgium
[b] Universidad de Antioquia, Faculty of Engineering, Colombia
[c] Erasmus University Rotterdam, Erasmus School of Economics, The Netherlands

ARTICLE INFO

Available online 8 April 2011

Keywords:
GRASP
Variable neighborhood search
Multi-depot vehicle routing
Periodic vehicle routing

ABSTRACT

The Flemish Ministry of Education promotes the integrated education of disabled children by providing educational opportunities in common schools. In the current system, disabled children receive ambulant help from a teaching assistant (TA) employed at an institute for extra-ordinary education. The compensation that the TAs receive for driving to visit the pupils is a major cost factor for the institute that provides the assistance. Therefore, the institute's management desires a schedule that minimizes the accumulated distance traveled by all TAs combined. We call this optimization problem the teaching assistants assignment-routing problem (TAARP). It involves three decisions that have to be taken simultaneously: (1) pupils have to be assigned to TAs; (2) pupils assigned to a given TA have to be spread over the TA's different working days; and (3) the order in which to visit the pupils on each day has to be determined. We propose a solution strategy based on an auction algorithm and a variable neighborhood search heuristic which has an excellent performance when applied to both simulated and real instances. The total distance traveled in the solution obtained for the institute's data set improves the current solution by about 22% which represents a saving of approximately 9% on the annual budget of the institute for integrated education.

1. Introduction

To promote the integration of disabled children into the general school system, the Flemish Ministry of Education provides opportunities and special support for those pupils in their own schools. In the current system, disabled children in ordinary schools receive ambulant help from a teaching assistant (TA) employed at an institute for extra-ordinary education, such as the *Koninklijk Instituut Woluwe*, a Flemish institute that organizes such support activities and collaborated with us on this research. Pupils with one or more of three main types of disability are supported: pupils with a hearing disability, an autism spectrum disorder and/or a language disability.

Every pupil that receives assistance is assigned to a TA and is visited at his or her own school. Pupils with a moderate disability are visited once a week, whereas pupils with a severe disability receive assistance twice a week. Each session lasts either 1 or 2 h. Every morning, the TAs depart from their homes and travel from school to school to visit the pupils that have been assigned to them for that day. At the end of the working day, the TAs return to their homes. In the beginning of the school year, the TAs receive their roster: a schedule that specifies for each day of the week which pupils they should assist and in which order. The TAs drive their private cars, and receive a financial compensation per driven kilometer. This compensation is a major cost factor for the institute that provides the assistance. The institute's management therefore desires a schedule that minimizes the accumulated distance traveled by all TAs combined.

Determining the schedule involves three decisions that have to be taken simultaneously: (1) pupils have to be assigned to TAs; (2) pupils assigned to a given TA have to be spread over the TA's different working days; and (3) the order in which to visit the pupils on each day has to be determined. We call this optimization problem the *teaching assistants assignment-routing problem* (TAARP).

Several constraints have to be taken into account when solving the TAARP. First, whether or not a certain TA is allowed to assist a certain pupil depends on her educational degree, on the pupil's disability type, and on the school the pupil is in (nursery, primary school or secondary school). Second, the number of working hours per week is different across TAs. This is due to the fact that certain TAs only work part time, and to an oddity in the Belgian law which prescribes that the number of hours a full time TA has to work depends on her exact diploma. The higher the TAs' diploma, the fewer hours she/he has to work during any given week. Third, pupils that receive assistance twice a week do so at two different and non-consecutive days. Although some schools would prefer to specify when pupils can be visited by the TAs, in

* Corresponding author at: University of Antwerp, Faculty of Applied Economics, ANT/OR – University of Antwerp Operations Research Group, Stadscampus S.B.513 Prinsstraat 13, Antwerp 2000, Belgium. Tel.: +32 654061; fax: +32 654901.
E-mail addresses: pablo.mayaduque@ua.ac.be, pmayaduque@gmail.com (P. Maya).

0305-0548/$ - see front matter
doi:10.1016/j.cor.2011.04.001

practice this does not happen: it is the institute and the TAs that define when the pupils are visited. We therefore do not take time windows into account in our formulation. Additionally, the institute works with a stable staff of TAs among which the teaching load is distributed trying to meet the particular availability and preferences of the different TAs. Therefore, the number of TAs is considered as an input defined by the management of the institute.

The TAARP is related to the multi-depot multi-period vehicle routing problem, but has several additional constraints. Some special characteristics of the problem, such as the fact that the teachers' "capacity" (maximum working time) is expressed as a multiple of the unit of "demand" (1 h of teaching), have motivated us to develop a new solution approach for this specific problem. In this paper we therefore propose a new metaheuristic that consists of two phases: a heuristic to find an initial feasible solution, followed by an improving phase. The heuristic for the first phase is inspired by the *auction algorithm*, while the improving phase is a *variable neighborhood search* (VNS) metaheuristic that considers swaps between two pupils with similar requirements and exchanges that involve more than two pupils with different characteristics (i.e., different numbers of required sessions and different numbers of hours per session).

This paper is organized as follows. Section 2 surveys the relevant literature on related routing problems. In Section 3, we develop a mixed-integer programming formulation of the problem. Our solution strategy is proposed in Section 4 and tested in Section 5. Finally, Section 6 contains some conclusions and pointers for future research.

2. Literature review

Because of the fact that several different starting locations for vehicles are used and that the problem data stretches over several days, the TAARP that we consider in this paper is closely related to the multi-depot periodic vehicle routing problem (MDPVRP), which generalizes two other well-known routing problems: the multi-depot vehicle routing problem (MDVRP) and the periodic vehicle routing problem (PVRP). The PVRP is a generalization of the VRP, the objective of which is to determine routes for a fleet of vehicles over several days. Like in the VRP, the vehicles' tours start and end at a given depot. An extensive description of the PVRP can be found in a survey by Francis et al. [15], where the evolution of the problem in the literature and its solution methods are presented, starting from the initial identification of the problem by Beltrami and Bodin [2] and its subsequent formulation and ending with the first dedicated heuristics by Russell and Igo [27] and Christofides and Beasley [8]. A discussion of the different objectives and constraints is presented in Mourgaya and Vanderbeck [24]. Early heuristics to solve the PVRP consist of a cluster-first route-second approach in which customers are first clustered into days. Examples are the heuristics proposed by Beltrami and Bodin [2], Tan and Beasley [28] and Russell and Gribbin [26]. In some cases these heuristics are complemented with improving phases. Recently approaches based on metaheuristics have been implemented, like the tabu search proposed by Cordeau et al. [12] (which is capable of dealing with periodic vehicle routing problems and multi-depot vehicle routing problems) and the solution strategy based on genetic algorithms by Drummond et al. [13]. Also, mathematical programming approaches have been presented, like the method based on lagrangian relaxation by Francis et al. [14]. Less extensive is the literature for the periodic capacitated arc routing problem (PCARP). However, due to the relationship between arc routing problems and vehicle routing problems [31], some of the algorithms for the PCARP can be adapted relatively easily to the PVRP. Evolutionary based algorithms for this problem have been proposed and described by Lacomme et al. [21], Chu et al. [10].

The MDVRP is a generalization of the VRP in which the vehicles can be based in different depots and each route followed by a vehicle must depart from one of those depots and return to it after visiting a sequence of customers. In all other aspects, the problem is identical to the VRP. The first approaches for this problem were based on construction and improvement procedures and, in general, involved adaptations of known heuristics for the VRP. More recently a search procedure was proposed by Chao et al. [7], and tabu search approaches have been proposed by Renaud et al. [25] and Cordeau et al. [12]. The latter authors present a formulation that shows that the MDVRP is a special case of the PVRP by associating depots with days. Some of the solution approaches for the MDVRP also use a cluster-first route-second strategy (see, e.g., the clustering procedures described by Giosa et al. [16]). Two exact algorithms have been developed for the MDVRP. Both are due to Laporte et al. [22], but these only work well for relatively small symmetric and asymmetric instances.

The combination of periodicity and multiple depots in vehicle routing, a problem generally known as the MDPVRP, has not been studied as extensively as other variants of the VRP. Hadjiconstantinou and Baldacci [18] consider the resource planning problem of a utility company that provides preventive maintenance services to a network of geographically dispersed customers using a fleet of depot-based vehicles and crews. They formulate this problem as an MDPVRP and use a heuristic based on a generalization of the classical VRP to solve it. Their heuristic first assigns each costumer to its nearest depot. Then, at each depot the order of the customers is determined according to some ordering rule. Using the resulting ordered list, the customers are added to routes by a least-cost insertion algorithm. After each costumer has been assigned to a feasible route, a tabu search algorithm is used to solve the resulting VRP for each day. Customer interchanges are performed to improve the solution. The heuristic is repeated several times, changing the ordering rule at each iteration. Kang et al. [20] consider the problem of designing a minimal cost schedule of vehicles in each depot to minimize transportation costs in product delivery and inventory holding costs at retailers over the planning period. In their algorithm, all feasible schedules are generated from each depot to each retailer and a set of vehicle schedules is selected optimally by solving a shortest path problem.

Recently, Vidal et al. [30] propose a hybrid genetic search framework to deal with several VRP variants, among them the MDPVRP. The proposed metaheuristic combines population-based evolutionary search, neighborhood-based search and advanced population-diversity management schemes.

Applications of the MDPVRP to specific problems can be found in different fields. Waste collection, e.g., is considered by Beltrami and Bodin [2], Teixeira et al. [29], and Coene et al. [11]. The latter paper considers approaches that combine two phases (routing and clustering) to solve a real life problem for a Belgian company collecting waste at slaughterhouses, butchers, and supermarkets. Alegre et al. [1] consider a raw material supply problem. Distribution in the alimentary industry is studied by Golden and Wasil [17], and maintenance operations are considered by Hadjiconstantinou and Baldacci [18] and Blakeley et al. [6].

3. Mathematical formulation

The TAARP is defined on a graph where nodes correspond either to TAs' homes or to schools where pupils attend classes, while arcs represent routes (distances calculated over a road network) that connect every pair of nodes. In this section, we formulate the TAARP as a linear integer (binary) program. The symbols used are shown in Table 1.

The model is formulated as follows:

$$\min \quad \sum_{t \in \mathcal{T}} \sum_{d \in \mathcal{D}} \sum_{i \in \mathcal{N}} \sum_{j \in \mathcal{N}} c_{ij} x_{ij}^{td} \tag{1}$$

s.t.

$$\sum_{i \in \mathcal{N}} x_{ij}^{td} - \sum_{i \in \mathcal{N}} x_{ji}^{td} = 0 \quad \forall j \in \mathcal{S},\ \forall t \in \mathcal{T},\ \forall d \in \mathcal{D} \tag{2}$$

$$\sum_{d \in \mathcal{D}} \sum_{i \in \mathcal{N}} b_j^t x_{ij}^{td} = k_j y_j^t \quad \forall j \in \mathcal{S},\ \forall t \in \mathcal{T} \tag{3}$$

$$\sum_{t \in \mathcal{T}} y_j^t = 1 \quad \forall j \in \mathcal{S} \tag{4}$$

$$\sum_{i \in \mathcal{N}} \sum_{j \in \mathcal{S}} r_j x_{ij}^{td} \le a^{td} \quad \forall t \in \mathcal{T},\ \forall d \in \mathcal{D} \tag{5}$$

$$\sum_{j \in \mathcal{N}} x_{ij}^{td} + \sum_{j \in \mathcal{N}} x_{ij}^{t(d+1)} \le 1 \quad \forall i \in \mathcal{S},\ \forall t \in \mathcal{T},\ \forall d \in \mathcal{D} \backslash \{|\mathcal{D}|\} \tag{6}$$

$$\sum_{i \in \mathcal{C}} \sum_{j \in \mathcal{C}} x_{ij}^{td} \le |\mathcal{C}| - 1 \quad \forall \mathcal{C} \subset \mathcal{S},\ 2 \le |\mathcal{C}| \le a^{td},\ \forall t \in \mathcal{T},\ \forall d \in \mathcal{D} \tag{7}$$

$$x_{ij}^{td} \in \{0,1\} \quad \forall i \in \mathcal{N},\ \forall j \in \mathcal{N},\ \forall t \in \mathcal{T},\ \forall d \in \mathcal{D} \tag{8}$$

$$y_j^t \in \{0,1\} \quad \forall j \in \mathcal{S},\ \forall t \in \mathcal{T} \tag{9}$$

The objective function (1) minimizes the total distance traveled by all TAs on all days. Constraints (2) impose that when a pupil j is assisted by TA t on day d, two arcs should be traversed by TA t on day d: an arc ongoing to pupil j and another arc outgoing from pupil j. Constraints (3) ensure that each pupil is visited as many times as required and that visits are performed only by a TA who has permission to visit him/her. Constraints (4) impose that each pupil is assisted by only one teacher, even when the pupil requires more than one visit. Constraints (5) ensure that the capacity (maximum number of hours taught on each day) of each TA is not exceeded. Constraints (6) ensure that for each pupil there is at least one day between two consecutive visits. Constraints (7) correspond to the generalized subtour elimination constraints. Finally, Eqs. (8) and (9) are the integer (binary) constraints. It is important to note that constraints (5), that limit the number of hours worked by each TA, consider only the effective teaching time and not the traveling time. This is due to the fact that only the effective teaching time counts as working time. Travel time is considered implicitly in the objective function, that minimizes the total traveled distance.

Table 1
Symbols used in the MIP formulation of the TAARP.

Sets	
$\mathcal{S}$	Set of pupils (customers)
$\mathcal{T}$	Set of teaching assistants homes (depots)
$\mathcal{N}$	Set of nodes, $\mathcal{N} = \mathcal{T} \cup \mathcal{S}$, $\mathcal{N} = \{1,2,\ldots,|\mathcal{T}|,|\mathcal{T}|+1,\ldots,|\mathcal{T}|+|\mathcal{S}|\}$
$\mathcal{D}$	Set of days, $\mathcal{D} = \{1,2,3,4,5\}$
Parameters	
k_j	Number of times that pupil j has to be visited every week
b_j^t	$\begin{cases} 1 & \text{If teaching assistant } t \text{ is allowed to visit pupil } j \\ 0 & \text{Otherwise} \end{cases}$
a^{td}	Number of hours of assistance available for teaching assistant t at day d (capacity)
r_j	Number of hours per session that pupil j requires (session duration for pupil j)
c_{ij}	Cost (distance) for traveling from node i to node j
Variables	
x_{ij}^{td}	$\begin{cases} 1 & \text{If node } i \text{ is visited after node } j \text{ by teaching assistant } t \text{ on day } d \\ 0 & \text{Otherwise} \end{cases}$
y_j^t	$\begin{cases} 1 & \text{If pupil } j \text{ is assisted by teaching assistant } t \\ 0 & \text{Otherwise} \end{cases}$

The model presented involves a large number of variables and the number of subtour elimination constraints in (7) is exponential. Using this model to solve medium and large size instances is therefore not computationally tractable. However, the model is useful to solve small instances in order to compare the results obtained with those of the metaheuristic approach developed in Section 4. To solve such instances, we have implemented a cutting plane procedure, in which the problem is first solved considering a reduced set of subtour elimination constraints (7). At each iteration, the set of violated subtour elimination constraints is determined and these constraints are added to the problem before solving it again. This procedure is run iteratively until no more subtours are found. Additionally, it is known that, in the optimal solution, the route d for teacher t has to depart from node t. Therefore, an additional constraint to ensure that this is the case, and to strengthen the model, is considered.

4. A metaheuristic for the TAARP

In this section, we present an efficient metaheuristic for the TAARP formulation discussed in Section 3. The metaheuristic solution approach is essentially a greedy randomized adaptive search procedure (GRASP). First, a feasible solution is constructed using a procedure inspired by the *auction algorithm*. In this procedure, randomness is involved when performing the auction in order to generate different initial solutions. Next, the constructed solution is improved by a procedure based on *variable neighborhood search* (VNS). Both procedures are called in each iteration of the GRASP algorithm. A key feature of our approach is that it does not allow infeasible intermediate solutions. The reason for this is that finding a feasible starting solution and repairing infeasible solutions of the TAARP is more challenging than in related problems.

4.1. Initial solution heuristic

Finding a feasible solution of the TAARP is difficult due to the special structure of the problem that combines periodicity and the requirement that all the sessions for one pupil must be provided by the same TA. When defining the set of pupils to be assisted by each TA, besides the pupil requirements (number of sessions and number of hours per session), the daily capacity of each TA must be considered. Otherwise, the resulting assignment could be infeasible. Therefore, we propose a heuristic that focuses on trying to assign the pupils to TAs in such a way that a feasible solution is found. In order to achieve that objective, we first tackle the problem of assigning pupils to teachers ensuring that all constraints except those related with the routing are satisfied. Note that, after a feasible solution for the assignment problem is found, it can be trivially transformed into a feasible solution for the TAARP by determining an arbitrary routing of the pupils assigned to each working day of each teacher. In other words, the routing does not affect the feasibility of the solution once it has been reached during the assignment phase.

Our heuristic is inspired by the auction algorithm which was initially proposed by Bertsekas [3] for the assignment problem and extended to deal with different variants of this problem.

Generally, auction algorithms perform well in practice, have excellent computational complexity properties and competitive running times. An extended study on this kind of algorithms is presented in Bertsekas [4].

We first briefly describe the auction algorithm. Then, the concepts of similar objects and similar agents are presented along with some issues that are needed to adapt the algorithm to the TAARP. Finally, the auction algorithm for finding an initial feasible solution for the TAARP is described.

4.1.1. The auction algorithm

In a traditional assignment problem, n objects and n agents are given, along with a value g_{ij} that represents the benefit of object j for agent i. The objective is to find the optimal assignment of objects to agents so that each agent is assigned to one object, each object is assigned to one agent and the total benefit is maximized.

In the auction algorithm, an object j has a *price* p_j that has to be paid by the agent to whom the object is assigned. Under these conditions, the value of object j for agent i is equal to $g_{ij}-p_j$. Obviously, each agent prefers the object which generates the maximum value for him/her, that is $\max_{j=1,\ldots,n}\{g_{ij}-p_j\}$. The algorithm starts with an empty assignment and all prices equal to zero, and ends when all objects have been assigned. It proceeds in iterations, assigning a single agent to an object in each iteration. A typical iteration involves three steps: selection, bidding and assignment.

The auction algorithm can be adapted to find a feasible solution for the TAARP, where pupils have to be assigned to TAs. One of the reasons why a modification is required, is that, although each pupil is assigned to a single TA, every TA assists several pupils. Hence, every TA is assigned to more than one pupil.

In the TAARP context, we define an *object* to be 1 h *of teaching of a given TA*. We define an *agent* to be *an hour of assistance required by a pupil*. For each object j and each agent i, we define the *benefit* g_{ij} to be a measure of the proximity between the teacher corresponding to the object and the pupil corresponding to the agent (i.e., a large enough constant minus the driving distance). If we assume that the total number of hours of teaching required by all pupils is equal to the total number of hours provided by the TAs, the corresponding problem of assigning objects (hours provided by a TA) to agents (hours demanded by a pupil) is reduced to an assignment problem that can be solved by the auction algorithm.

4.1.2. Similar objects and similar agents

One problem with the auction algorithm is that its performance can be negatively affected by the fact that for each agent all objects associated with a particular TA generate the same benefit. The extension of the auction algorithm to the transportation problem considering *similar objects*, proposed by Bertsekas and Castanon [5], avoids this difficulty. Two objects j and j' are called similar if they can be matched with the same agents at equal values. The set of all objects similar to j is called the *similarity class* of j and is denoted by $M_o(j)$.

There are two other issues that must be taken into account to apply the auction algorithm to the TAARP: (i) a session lasts either 1 or 2 h, and there are pupils that require more than one session per week, and (ii) these sessions must be taught by the same teacher and take place on different, non-consecutive days.

The first issue can be addressed by defining *similar agents* as Bertsekas and Castanon [5] propose in order to deal with transportation problems. As mentioned, an agent is defined for each hour of assistance that a pupil requires (e.g., for a pupil who requires two session of 2 h per week, four agents must be defined). Two agents i and i' are called similar if they obtain the same benefit when an object j is assigned to them. The set of all agents similar to an agent i is called the *similarity class* of i and is denoted by $M_a(i)$. In the auction algorithm for the TAARP, all agents corresponding to a specific pupil are similar.

The second issue implies that the set of objects $A(i)$ that an agent i can bid for should often be restricted to a subset of the unassigned objects. Moreover, the set of possible objects that two agents can bid for are related if the agents correspond to different hours of assistance required by the same pupil. Consider, e.g., a pupil s who requires two sessions of 2 h per week. For this pupil, four agents are created, one for each hour of assistance. Assume that object j has already been assigned to agent i, i.e., one of the hours of assistance for this pupil has already been assigned. Consider a second agent k associated with a different hour of assistance required by pupil s:

- If agent k is a different hour in the same session as i (e.g., i is the first teaching hour of the first session of the week and k is the second hour for the same session) then $A(k)$ consists of objects associated with the same teacher and the same day of the week as object j.
- If agent k is associated with a different session than i (e.g., i is the first hour of the first session and k is the first hour of the second session), then $A(k)$ consists of all objects associated with the same teacher but corresponding to a different day of the week than object j.

The set $A(i)$ therefore changes depending on the current assignment (partial state of the auction). In other words, the set $A(i)$ depends on the partial state of the auction.

4.2. Auction heuristic for the TAARP

Based on the basic auction algorithm for the assignment problem and the considerations listed above, we have developed a modified auction heuristic for the TAARP. This auction heuristic works as follows.

1. *Initialization*: Unassign all objects i and all agents j, set all prices p_i to zero.
2. *Auctioning*: If all agents have been assigned go to the routing phase (phase 3), otherwise iteratively perform the following steps:
   (a) *Selection*: Select randomly an unassigned agent i.
   (b) *Bidding*: Determine the object j in $A(i)$ which has maximum value for agent i
   $$j_i = \arg\max_{j=1,\ldots,n}\{g_{ij}-p_j\} \tag{10}$$
   The value of object j_i for agent i is given by $v_i = g_{ij_i}-p_{j_i}$. Estimate the bid from agent i on object j_i:
   $$b_{ij_i} = p_{j_i}+v_i-w_i+\varepsilon \tag{11}$$
   where ε is a constant explained below and w_i is the best value $g_{ij}-p_j$ for any object in $A(i)$ that belongs to a different similarity class than j_i:
   $$w_i = \max_{j=1,\ldots,n; j\notin M_a(j_i)}\{g_{ij}-p_j\} \tag{12}$$
   If such an object does not exist, then w_i is set to $-v_i$.
   (c) *Assignment*:
   - If object j_i was assigned to an agent $l \neq i$, then unassign all objects in the similarity class $M_a(l)$.
   - Assign object j_i to agent i.
   - For every agent $k \in M_a(i)$, raise the price p_k to b_{ij_i}.
   (d) *Perturbation*: If the number of iterations without decreasing the number of unassigned agents is equal to L, then decrease the price of all unassigned objects to zero and go to step 2a. Otherwise, go directly to step 2a.

3. *Routing*: Apply a routing procedure to find the shortest tour for each teacher on each day.

In phase 2b, the bid for an object considers the second most valued object in a different similarity class, an idea related to the concept of *regret minimization* used in several classical algorithms (e.g., the heuristic for the VRP proposed by Christofides et al. [9]). The parameter ε is a constant value that ensures that prices keep increasing through the different iterations. This is important because it prevents the auction from getting stuck in a cycle of bids without price increments. Bertsekas [4] used ε in order to find the optimal solution for assignment problems in a so-called *scaling strategy*. This strategy starts with a large ε value and reduces it iteratively. Initial prices for each iteration are fixed according to the auction solution for the previous iteration. However, the optimal solution obtained when solving the assignment problem does not necessarily correspond to the optimal solution of the TAARP. Furthermore, our interest is to find an initial feasible assignment. Therefore, we do not use the *scaling strategy* but run the auction algorithm with a fixed ε value. That value is set to be on the same order of magnitude as the maximum benefit g_{ij}.

If the number of unassigned agents after the assignment step is not decreasing over time any more, the perturbation step tries to force the unassigned agents to bid for a free object when the algorithm has spent L iterations on a *war of prices*. To this end, the algorithm decreases the price of unassigned objects to make them more *attractive* for unassigned agents. Additionally, in order to be sure that the algorithm does not spend too much time looking for an assignment, we implement a mechanism to restart the auction when it is taking more than a prespecified maximum time to find a feasible starting solution.

If, at the end of the auction phase, a feasible assignment has been reached, it is still necessary to determine the order in which the pupils assigned to a teacher are visited on each day. This requires solving the associated traveling salesman problem. Since the maximum number of pupils that can be assigned to a teacher on each day is eight, we found that the order in which they are visited could best be determined by an exact procedure that enumerates all possible combinations. This procedure guarantees the optimal solution of this subproblem and – due to the limited size of the problem – requires only a very limited amount of time, regardless of the number of pupils or teachers considered.

4.3. Improving phase

The improving phase is based on the *variable neighborhood search* metaheuristic. More in particular, it corresponds to the *basic* VNS variant described in Hansen and Mladenovic [19]. A schematic overview of the improving stage is given in Algorithm 1.

Algorithm 1. Improving phase.

Initialization: Consider an initial solution $\mathbf{x}$. Select the set of neighbourhood structures $N_k(\mathbf{x})$, $k=1,2,\ldots,k_{\max}$.
Set $k \leftarrow 1$
repeat
Explore Neighbourhood k. Find the best neighbor $\mathbf{x}'$ of $\mathbf{x}$ ($\mathbf{x}' \in N_k(\mathbf{x})$)
Move or not. If the solution $\mathbf{x}'$ is better than $\mathbf{x}$, set $\mathbf{x} \leftarrow \mathbf{x}'$ and $k \leftarrow 1$; otherwise, set $kxlarr;k+1$;
until $k > k_{\max}$

Two move types have been considered, defining two different neighborhood structures ($k_{\max}=2$). The first move, that defines neighborhood $N_1(x)$, is a *simple exchange* move that attempts to swap pupils with similar characteristics. This move exchanges the positions of two pupils who need the same number of sessions and the same number of hours per sessions if no constraints are violated with the move (e.g., if one of the TAs is not allowed to assist a pupil involved in the swap). This exchange allows a pupil to be assigned to a different TA and/or scheduled on a different day.

The second type of move, a *combined exchange*, allows exchanges among three pupils with different characteristics, but keeping the workload of the TA's involved constant. This move implies changes on three different day-routes and allows pupils to be assigned to a different TA and/or scheduled on a different day. Three cases of this type of move are considered.

- *Case*1: A pupil who requires one session of 2 h per week is exchanged with two pupils who need one session of 1 h per week, and who are currently visited on the same day by one of the TAs.
- *Case*2: A pupil who requires two sessions of 1 h per week is exchanged with two pupils requiring a session of 1 h per week, and who are currently visited on different days by one of the TAs.
- *Case*3: A pupil who requires two sessions of 2 h per week is exchanged with two pupils who need one session of 2 h per week, and who are currently visited on different days by one of the TAs.

Fig. 1 visualizes the first case for the second move type. Fig. 2 corresponds to the cases 2 and 3, as these cases only differ in the numbers of hours per session required by the pupils involved in the move. In the figures, the squares correspond to the TAs and the circles represent pupils. In Fig. 1, a pupil who needs one session of 2 h per week is swapped with two pupils who need 1 h of assistance per week only. In Fig. 2, a similar swap is performed except that the two pupils who need 1 h of assistance are visited on different days.

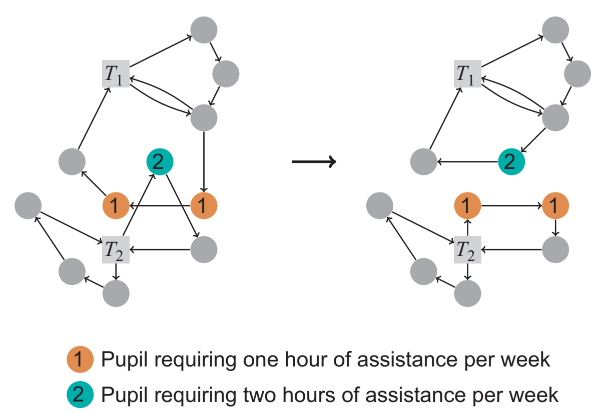

Fig. 1. Example of the second type of move: case 1.

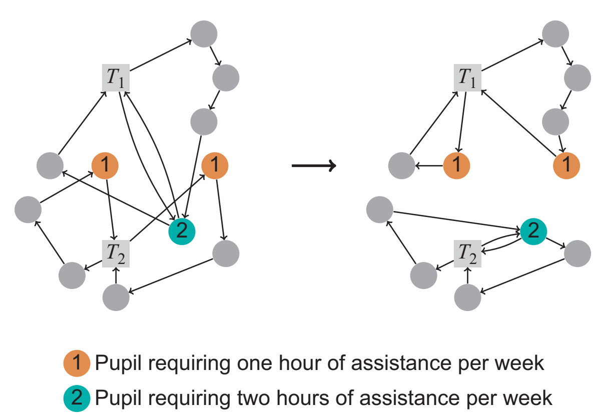

Fig. 2. Example of the second type of move: cases 2 and 3.

Neighborhood $N_2(x)$ is composed of all solutions $\mathbf{x}'$ that can be reached from the current solution $\mathbf{x}$ through the moves described in cases 1–3 without violating any constraint, i.e., the algorithm only performs moves that preserve the feasibility of the solutions. After applying this type of move it is necessary to re-optimize the day-routes obtained because one of the sessions could have been assigned in a suboptimal position into the tour. As in the algorithm for constructing the initial solution, we use an efficient exact algorithm for that purpose.

4.4. Algorithm structure

The solution approach has been described for one iteration. However, it is possible to embed it in an iterative scheme by considering a different initial solution each time as in the multi-start methods described by Martí [23]. In our approach, the auction algorithm is used to generate different starting solutions for the improving phase by randomly selecting agents in step 2a. The complete algorithm is run iteratively until a stopping criterion is reached (e.g., for the number of iterations or running time) and the best solution is kept. In our computational tests the number of iterations is defined as the stopping criterion. A schematic overview of the proposed solution strategy is presented in Algorithm 2.

Algorithm 2. Solution strategy for the TAARP.

Initialization: Consider an initial empty solution
while Stopping condition is not satisfied **do**
  Auction Algorithm: Get initial solution $\longrightarrow \mathbf{x}$
  VNS Algorithm: Improvement phase $\longrightarrow \mathbf{x}'$
  if $\mathbf{x}'$ improves the incumbent **then**
    Update incumbent
  end if
end while
Return best solution

5. Computational results

In this section, we report the results from a computational study to evaluate the performance of our metaheuristic approach. For the study several small instances were generated to compare our approach to an exact solution method, i.e., the implementation of the mathematical model described in Section 3. Larger simulated instances were also generated to test the robustness of the proposed approach. Finally, we use our algorithm to solve the two real-life instances of the TAARP that motivated this research.

The mathematical model was implemented in Java using ILOG CPLEX Concert Technology (IBM ILOG CPLEX Optimization Studio Academic Research Edition V12.2) while the metaheuristic was coded in Java. All experiments were carried out on a personal computer running Windows XP and equipped with an Intel(R)Core(TM)2 Duo(CPU) T9300 processor at 2.5 GHz and with 4 GB of RAM.

5.1. Comparison with CPLEX

In a first set of experiments, we concentrate on determining the effectiveness and efficiency of the metaheuristic proposed in this paper. Since only small instances can be solved to optimality using the mathematical model, we have generated several small random instances. These instances are further divided into four categories, defined by two degrees of freedom.

In the *uniform* instances the locations of the pupils and TAs were randomly generated on a 100×100 Euclidean grid using a continuous $U[0,100]$ distribution for both the x- and y-coordinates. In the *clustered* instances, the pupils were grouped into three clusters representing population centres. To this end, three scattered seed points were fixed, and the pupils' coordinates were generated randomly around those seeds such that each cluster had a similar number of pupils.

Table 2
Summary of results for the P20-TA3-D3 simulated instances.

Pupil/TA distribution	Random		Clustered	
$p[2h]$	$=p[1h]$	$>p[1h]$	$=p[1h]$	$>p[1h]$
Number of instances	10	10	10	10
Optimal solutions	10	8	10	9
Maximum gap (%)	0.0	1.1	0.0	0.3
Average gap (%)	0.0	0.1	0.0	0.0
Average iterations	327.5	250.0	147.6	248.3

In the $p[1h]=p[2h]$ instances, the probability for each pupil to receive 2 h of assistance is equal to the probability to receive 1 h. In the instances labeled $p[2h]>p[1h]$, this probability was set to 0.7.

Initially, we considered 40 instances, 10 for each category. Each of these instances involves 20 pupils to be assisted by three TAs within a three-day week (P20-TA3-D3). Additionally, we generated another eight slightly larger instances, considering 30 pupils and four TAs to be scheduled on a similar period (P30-TA4-D3). The exact algorithm was used to solve these 48 instances, and a time limit of 24 h was set as a stopping criterion. All the P20-TA3-D3 instances were solved to optimality within this time limit, while for all the P30-TA4-D3 instances the algorithm reached the time limit without finding the optimal solution. Therefore, for those instances the optimal solution value is replaced by the best lower bound found by CPLEX.

Table 2 shows a summary of the results for the P20-TA3-D3 instances. The table presents, for each category, the number of instances considered, the number of optimal solutions found by the metaheuristic, the maximum gap between the optimal solution and the solution provided by our algorithm, the average gap, and the average iteration in which the algorithm found the best solution.

The metaheuristic provides the optimal solution for 37 out of the 40 instances. For these instances, the maximum gap is, in all cases, less than 1.1% and the best solution is, on average, found in less than 328 iterations. For the eight P30-TA4-D3 instances, the average gap is about 6.0%. However, as was pointed out before, for those instances a lower bound of the optimal value was used because the optimal value could not be found within the running time limit. The average number of iteration before the best solution is found was in average 606.2.

5.2. Analysis of the heuristic and its robustness

In the second set of experiments, we concentrate on analyzing the performance of the different elements of our metaheuristic and its robustness. To this end 40 instances with 150 pupils, 16 TAs and five working days were generated. These instances are divided in the same four categories that were described in Section 5.1.

The main aim of this phase is to demonstrate the working of our algorithm. Essentially, the auction algorithm generates a set of diverse solutions that are then improved by the VNS. The diversity of the solutions generated by the auction algorithm is important because it gives different starting points to the VNS and should prevent the VNS from converging to the same solution at each iteration.

This experiment starts by executing 100 iterations of our algorithm for each of the 40 instances. The value of ε is fixed at

150 and the time limit for the auction algorithm to restart the auction when it is in a war of prices, is set to 20 s. For each instance i we ran the algorithm for 100 iterations and recorded the objective function value after the auction algorithm (initial solution) and after the VNS (final solution). We then calculated the average objective function value ($\bar{f}_i^{\text{init}}$ and $\bar{f}_i^{\text{final}}$) and the standard deviation (σ_i^{init} and σ_i^{final}) both of the 100 initial and the 100 final solutions. This was repeated for each instance and these numbers were averaged over the four instance categories.

Table 3 presents a summary of the results. The row labeled *Avg. improvement* (%) presents the average improvement that the VNS is able to achieve over the initial solution. As can be seen, the final solution achieved by the VNS is on average between 45% and 60.7% better than the initial solution. The row labeled *Avg. ratio of std. dev.* is the result of dividing the average standard deviation of the initial solutions by the average standard deviations of the final solutions ($\sigma_i^{\text{init}}/\sigma_i^{\text{final}}$) and taking the average of that value in each instance category. The results show that the standard deviations of the solutions produced by the auction algorithm are on average between 1.9 and 4.1 times as large as those of the final solutions. This observation supports the fact that the auction algorithm generates diverse solution and that the diversity of the initial solutions is reduced substantially by the intensification of the VNS.

The final two rows in Table 3 show that the VNS takes between 4 and 64 times as long as the auction algorithm, but the average computing times remain very reasonable (between 10.1 and 17.3 s). A good indicator of the effectiveness of the auction algorithm in generating an initial solution for the TAARP is the fact that the procedure to restart the auction when it gets in a war of prices was not invoked over the 4000 iterations (100 iterations for each of the 40 instances) that were executed. The efficiency of the auction algorithm approach to find initial solutions varies considerably with the distribution of the number of visits per pupil. The instances generated with a higher probability of two visits a week turned out to be more difficult to solve.

The robustness of our approach was also evaluated. For each instance and each case, the coefficient of variation is computed based on the set of final solutions. This measure shows how diverse the final solutions are as a percentage of the mean of the set of solutions. Additionally, the average iteration in which the best solution was found is also indicated. Table 4 presents these results for the instances that have been considered. As can be seen from this table, the coefficient of variability is very small in each of the instance categories, proving the robustness of our approach. also, the best solution is found in average after 50 of the 100 iterations.

Table 3
Summary of results for instances with 150 pupils, 16 TAs and 5 working days.

Pupil/TA distribution	Random		Clustered	
$p[2h]$	$=p[1h]$	$>p[1h]$	$=p[1h]$	$>p[1h]$
Avg. improvement (%)	55.2	60.7	45.0	60.1
Avg. ratio of std. dev.	3.8	1.9	4.1	2.1
Avg. ratio of computing times	35.5	4.0	63.9	3.9
Avg. total time (s)	10.1	16.9	10.8	17.3

Table 4
Robustness analysis based on instances with 150 pupils, 16 TAs and 5 working days.

Pupil/TA distribution	Random		Clustered	
$p[2h]$	$=p[1h]$	$>p[1h]$	$=p[1h]$	$>p[1h]$
Coefficient of variability (%)	4.5	4.9	3.3	4.7
Iteration best solution	46.7	49.8	49.8	47.4

5.3. Real-life instance

To test our algorithm in a realistic setting we use data from the *Koninklijk Instituut Woluwe* for two different school years. For the first data set (school year 2008–2009), the plan that was put into practice is also available. That plan is compared with the solution provided by our approach. The second data set corresponds to the school year 2009–2010. For this period, the management of the institute requested us to provide a preliminary solution to be used as a starting point to build the definitive plan.

Currently, the planning is done by hand using only a spreadsheet. Because of the planning difficulties this poses, the current schedule does not take into account the constraint that enforces the different sessions for a pupil to take place on non-consecutive days. No geographic information was available on the pupils and teacher locations or the routes that had been defined. We therefore used the *Google Maps API* to develop several scripts (1) to geocode the addresses of schools and TAs, (2) to compute an approximation of the real driving distance between each pair of locations, and (3) to draw the routes generated by our algorithm.

The first real-life instance involves 212 disabled children that must be assisted by 24 TAs. The pupils are distributed over 138 schools dispersed in a central area of Belgium. About 25% of the pupils require two visits per week and two of the pupils require three visits per week. For each TA, the number of working days is between one and five and the number of available hours per day is between two and eight. Fig. 3, showing the geographical distributions of the schools (where pupils must be visited) and TAs, was generated using one of the scripts developed.

To compare our metaheuristic to the current planning technique, we ran it without considering the constraint that pupils should be visited on non-consecutive days if they require more than one session. Using 1000 iterations as the stopping criterion, the best solution was found in iteration 748 and the complete run took 4.6 h of processing time. Table 5 presents for each TA the total number of pupils assigned, the traveled distance in kilometers per week (TTD) and the number of kilometers traveled per teaching hour (TTD/hour) for both the current solution (Cur.) and for the metaheuristic (MH). The change in TTD compared to the current solution and the number of common pupils between the two solutions (Com.) are also shown.

The final solution produced by our algorithm improves the current solution by about 22%, and results in a schedule with 3756.1 km of total traveled distance. On average, the total traveled distance per TA is around 45 km less than in the current solution. And, the number of kilometers traveled per teaching hour is also improved by about 2 km.

In our best solution, only 31% of the pupils are assigned to the same TA as in the current solution. This shows that the improvement is not only due to a better routing but also to a better assignment of pupils to TAs.

While the travel distance is reduced for most TAs in the final solution, TA19 faces a substantial increase in TTD. However, the total distance that she has to travel per week is still less than 90 km, which is less than the average over all the TAs. For TA22, the distances are considerable larger than for the others TA because TA22's home lies at a large distance from the cluster of schools to be visited, as can be seen in Fig. 3.

The value of our solution of the TAARP at the Royal Institute Woluwe can be estimated. The institute receives an annual subsidy of €483.50 for 184 of the pupils it assists in the context

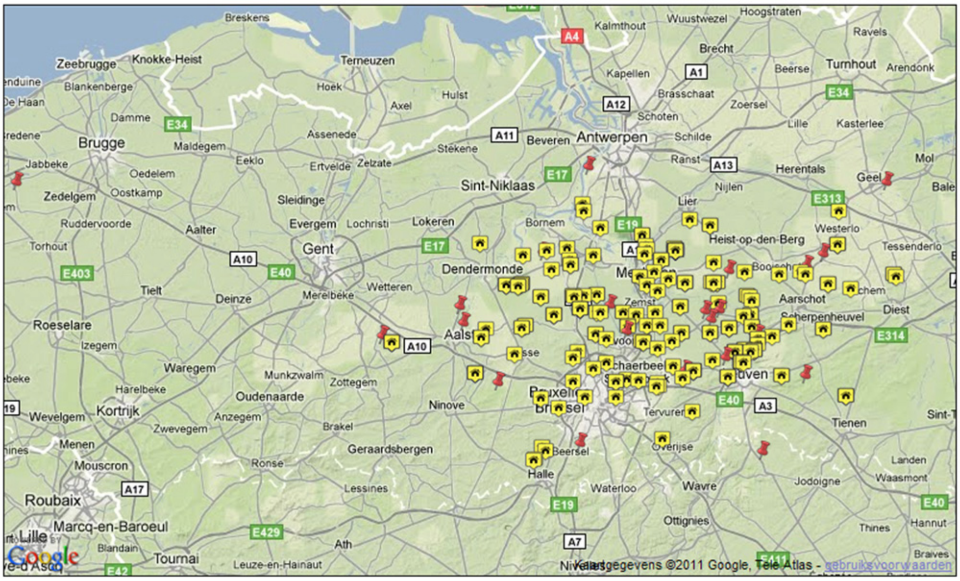

Fig. 3. Geographical distribution of schools and TA.

Table 5
Comparison of the current and the metaheuristic solution.

TA	Number of pupils			TTD (km)		TTD/hour		Δ (%)
	Cur.	MH	Com.	Cur.	MH	Cur.	MH	
1	6	5	0	72.7	40.2	7.3	4.0	44.7
2	9	13	0	142.8	143.2	5.9	6.0	−0.3
3	7	10	1	75.1	58.8	3.6	2.8	21.7
4	9	11	2	135.2	191.4	5.9	8.3	−41.6
5	7	6	2	103.2	26.0	7.4	1.9	74.8
6	9	8	0	208.7	218.7	10.4	10.9	−4.8
7	7	8	0	248.4	240.8	15.5	15.1	3.1
8	6	5	2	87.0	122.8	7.9	11.2	−41.1
9	4	4	4	4.0	4.0	1.0	1.0	0.0
10	11	11	2	326.4	81.1	12.6	3.1	75.2
11	6	4	0	76.2	45.4	6.3	3.8	40.4
12	14	14	7	270.8	252.9	9.7	9.0	6.6
13	10	13	1	743.1	328.6	23.2	10.3	55.8
14	10	10	5	249.2	160.7	11.9	7.7	35.5
15	7	7	1	71.5	79.6	4.5	5.0	−11.3
16	16	16	13	155.8	147.5	5.6	5.3	5.3
17	9	9	3	186.0	110.9	9.3	5.5	40.4
18	12	11	9	146.6	183.9	5.2	6.6	−25.5
19	6	4	0	28.9	89.7	2.1	6.4	−209.8
20	8	6	0	183.4	26.1	11.5	1.6	85.7
21	11	8	2	151.2	96.8	6.9	4.4	36.0
22	8	9	6	901.2	851.8	41.0	38.7	5.5
23	11	12	5	182.1	131.2	6.5	4.7	28.0
24	9	8	2	79.4	124.2	3.3	5.2	−56.4
Average				201.20	156.50	9.35	7.43	
Total		212	67	4828.71	3756.05			22.21

of integrated education. The total annual budget, which forms the basis for compensating the TAs for the distances they drive with their private cars, is therefore €88,964. Our solution to the TAARP problem saves the TAs a total of 10,72.6 km per week and thus about 32,178 km per annum. The compensation per kilometre driven amounts to €0.25, so that our solution saves the institute €268.15 per week and €8044.5 per annum, which is 9% of the annual budget and 22% of the current annual expenses for transportation.

Additionally, we have also run our algorithm considering the constraint that enforces the different sessions for a pupil to take place on non-consecutive days. The total distance of the resulting schedule is 3861.5 km, which is about 20% better than the existing solution and 1% worse than the solution obtained without considering that constraint.

The second data set considers 22 TAs to assist 207 pupils distributed across 133 different schools dispersed in the same central region of Belgium. About 27% of the pupils require two visits per week and no pupil requires three visits per week. For most of the pupils that were assisted only 1 h a week, an extra hour of assistance was assigned following a decision taken by the institute's management. Therefore, we provided to the management of the institute two different alternatives, based on two different scenarios for this category of pupils. In the first scenario, each of those pupils is visited once a week in a session of 2 h. For the second scenario, two visits per week, each of them lasting 1 h, are scheduled for each of those pupils. As expected, the first scenario provides a plan with a lower total traveled distance (3217 km) and less kilometers traveled per teaching hour (6.7 km). The assignment of the students to the TA for the two scenarios are significatively different, only 41% of the students are assigned to the same TA in both scenarios. Regard to the time required for the algorithm, the second scenario demanded around 15 times the time required for the first scenario. Fig. 4 shows an example of the routes obtained for teacher M.E. using the data set for the school year 2009–2010 and the second scenario.

6. Conclusions and future research

In this paper, we have defined the teaching-assistants assignment-routing problem (TAARP) to solve a real-life routing problem faced by several institutes for extraordinary education in Flanders. This real-life problem is closely related to the multi-depot periodic vehicle routing problem, but has several additional constraints. We have described and implemented a solution approach based on the auction algorithm and concepts from the GRASP and VNS metaheuristics. The approach we proposed was tested using both simulated instances and two real-life data sets. On both types of data set, our algorithm performed excellently. The solution obtained on the real-life data set improves the total distance traveled by about 22% compared to the existing solution used by the institute. The use of our method represents a 9% saving on the available budget to assist the disabled children.

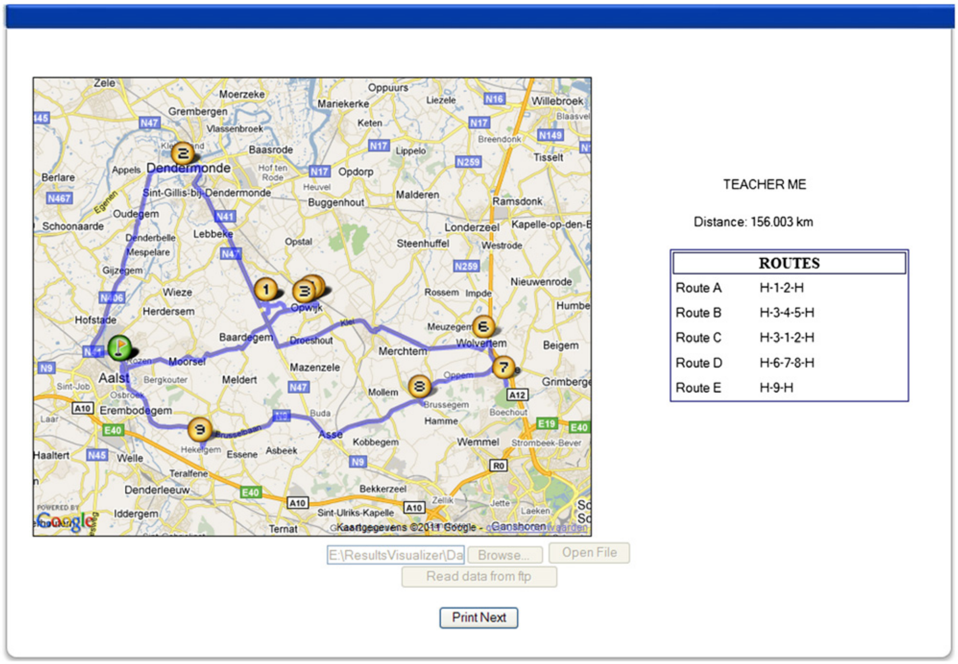

Fig. 4. Example of the schedule for TA M.E., school year 2009–2010 (Scenario 2).

According to our results, the key to improve the existing solutions in the real-life instances lies mainly in the clustering or distribution of the pupils over the TAs and their working days. The routing problems usually involve not more than a few nodes and are highly determined by the way the pupils are allocated to the different working days of the TAs. Solving these routing problems to optimality is usually not computationally expensive and contributes to improve the quality of the TAARP solution. The proposed approach provides an integrated framework for the different decisions involved in the TAARP, as it considers the clustering, assignment and routing decisions simultaneously.

There are three potential topics for future research that arise from this work. First, we could explicitly consider the fairness of the assignment and look for solutions that minimize the total traveled distance while keeping the distance traveled by each teacher as homogeneous as possible. Second, the decision regarding the optimal number of TAs required to serve the pupils could be considered simultaneously with the assignment and routing decisions that have been already taken into account. Third, we could attempt to construct a portfolio of alternative low-cost solutions, which can then be ranked using secondary criteria based on human and social considerations that are hard to quantify. Additionally, the possibility of allowing infeasible intermediate solutions could be considered as an option to extend the algorithm we have proposed.

References

[1] Alegre J, Laguna M, Pacheco J. Optimizing the periodic pick-up of raw materials for a manufacturer of auto parts. European Journal of Operational Research 2007;179:736–46.
[2] Beltrami E, Bodin L. Networks and vehicle routing for municipal waste collection. Networks 1974;4:65–94.
[3] Bertsekas D. Auction algorithms for network flow problems: a tutorial introduction. Computational Optimization and Applications 1992;1:7–66.
[4] Bertsekas D. Network optimization. Continuous and discrete models. Athena Scientific; 1998.
[5] Bertsekas D, Castanon D. The auction algorithm for the transportation problem. Annals of Operations Research 1989;20:67–96.
[6] Blakeley F, Argüello B, Cao B, Hall W, Knolmajer J. Optimizing periodic maintenance operations for Schindler elevator corporation. Interfaces 2003; 33:67–79.
[7] Ming Chao I, Golden B, Wasil E. A new heuristic for the multi-depot vehicle routing problem that improves upon best-known solutions. American Journal of Mathematical and Management Sciences 1993;13:371–406.
[8] Christofides N, Beasley JE. The period routing problem. Networks 1984;14: 237–56.
[9] Christofides N, Mingozzi A, Toth P. Combinatorial optimization. In: The vehicle routing problem. John Wiley & Sons; 1979. p. 315–38.
[10] Chu F, Labadi N, Prins C. A scatter search for the periodic capacitated arc routing problem. European Journal of Operational Research 2006;169(2): 586–605.
[11] Coene S, Arnout A, Spieksma F. On a periodic vehicle routing problem. Journal of the Operational Research Society 2010;61:1719–28.
[12] Cordeau J-F, Gendreau M, Laporte G. A tabu search heuristic for periodic and multi-depot vehicle routing problems. Networks 1997;30:105–19.
[13] Drummond L, Ochi L, Vianna D. An asynchronous parallel metaheuristic for the period vehicle routing problem. Future Generation Computer Systems 2001;17:379–86.
[14] Francis P, Smilowitz K, Tzur M. The period vehicle routing problem with service choice. Transportation Science 2006;40:439–54.
[15] Francis P, Smilowitz K, Tzur M. The period vehicle routing problem and its extensions. In: Golden B, Raghavan S, Wasil E, editors. The vehicle routing problem: latest advances and new challenges. US: Springer; 2008. p. 73–102.
[16] Giosa D, Tansini L, Viera O. New assignment algorithms for the multi-depot vehicle routing problem. Journal of the Operational Research Society 2002;53:977–84.
[17] Golden B, Wasil E. Computerized vehicle routing in the soft drink industry. Operations Research 1987;35:6–17.
[18] Hadjiconstantinou E, Baldacci R. A multi-depot period vehicle routing problem arising in the utilities sector. Journal of the Operational Research Society 1998;49:1239–48.
[19] Hansen P, Mladenovic N. Variable neighborhood search. In: Burke E, Kendall G, editors. Search methodologies. Introductory tutorials in optimization and decision support techniques. US: Springer; 2005. p. 211–38.
[20] Hwan Kang K, Hoon Lee Y, Ki Lee B. An exact algorithm for multi depot and multi period vehicle scheduling problem. In: Computational science and its applications—ICCSA 2005, Lecture notes in computer science. Berlin, Heidelberg: Springer-Verlag; 2005. p. 350–9.
[21] Lacomme P, Prins C, Ramdane-Cherif W. Evolutionary algorithms for periodic arc routing problems. European Journal of Operational Research 2005; 165(2):535–53.
[22] Laporte G, Nobert Y, Arpin D. Optimal solutions to capacitated multi-depot vehicle routing problem. Congressus Numerantium 1984;44:283–92.
[23] Martí R. Multistart methods. In: Glover F, Kochenberger G, editors. Handbook on metaheuristics. Springer; 2003. p. 355–68.
[24] Mourgaya M, Vanderbeck F. Periodic routing problem: classification and heuristic for tactical planning. RAIRO Operations Research 2006;40:169–94.
[25] Renaud J, Laporte G, Boctor F. A tabu search heuristic for the multi-depot vehicle routing problem. Computers & Operations Research 1996;23: 229–35.

[26] Russell R, Gribbin D. A multiphase approach to the period routing problem. Networks 1991;21:747–65.
[27] Russell R, Igo W. An assignment routing problem. Networks 1979;9:1–17.
[28] Tan CCR, Beasley J. A heuristic algorithm for the period vehicle routing problem. Omega 1984;12:497–504.
[29] Teixeira J, Antunes A, Pinho de Sousa J. Recyclable waste collection planning. A case study. European Journal of Operational Research 2004;158: 543–54.
[30] Vidal T, Crainic T, Gendreau M, Lahrichi N, Rei W. A hybrid genetic algorithm for multi-depot and periodic vehicle routing problems. Technical Report, Interuniversity Research Centre on Enterprise Networks, Logistic and Transportation CIRRELT, 2010.
[31] Wøhlk S. A decade of capacitated arc routing. In: Golden B, Raghavan S, Wasil E, editors. The vehicle routing problem: latest advances and new challenges. Operations research/computer science interfaces series, vol. 43. US: Springer; 2008. p. 29–48.