



CIRRELT

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

Interuniversity Research Centre
on Enterprise Networks, Logistics and Transportation

Dynamic Management of an Emergency Vehicle Fleet : Deployment and Redeployment of Ambulances

Fayez F. Boctor
Guillaume Carpentier
Angel Ruiz

August 2011

CIRRELT-2011-48

Document de travail également publié par la Faculté des sciences de l'administration de l'Université Laval,
sous le numéro FSA-2011-004.

Bureaux de Montréal :

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

Bureaux de Québec :

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

www.cirrelt.ca

Dynamic Management of an Emergency Vehicle Fleet : Deployment and Redeployment of Ambulances

Fayez F. Boctor^{1,2}, Guillaume Carpentier², Angel Ruiz^{1,2*}

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

² Department of Operations and Decision Systems, 2325, de la Terrasse, Université Laval, Québec, Canada G1V 0A6

Abstract. This article deals with the deployment and redeployment of emergency vehicles, which is one of the main problems in Emergency Medical Services (EMS) management. Although the deployment of emergency vehicles has been extensively addressed in the literature, the originality of this article is that it explicitly considers the redeployment of vehicles – in our case, ambulances – based on cyclic changes in service demand (e.g., based on population movements during the day from residential areas to the workplace and vice-versa). We introduce a multi-period approach to take into account such service demand changes. The vehicle deployment plan is chosen to simultaneously maximize coverage and minimize vehicle relocation between periods. We propose a mathematical model that allows small to medium instances of the problem to be solved. In order to tackle large instances corresponding to real-life situations, we propose two heuristic approaches. We then evaluate the performance of our heuristics using a set of randomly-generated instances.

Keywords. OR in health services, ambulance fleet management, emergency vehicles location and relocation, heuristics, OR in health services.

Acknowledgements. This research was partially financed by the Natural Science and Engineering Research Council of Canada (NSERC) under Grants OGP0036509 and OGP0293307. 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: Angel.Ruiz@cirrelt.ca

1. Introduction

From an operational point of view, Emergency Medical Services (EMS) management is a series of chronological steps, each of them requiring precise decisions and actions. The process starts with a call to an emergency call center (e.g., 911 in North America), where an employee handles the call. The incoming call is categorized according to a Medical Priority Dispatch System (MPDS), which assigns a medical determinant and a priority to the call. Pre-arrival instructions, which depend on the injury or illness type and severity, are given to the caller over the phone. The incident is geographically located (i.e., geo validated) and sent to the queue of the dispatcher responsible for the region from which the call originated. According to the priority of the call and the number of emergency vehicles available, the dispatcher selects a vehicle and assigns it to the call. If no vehicle is available, the dispatcher leaves the call in the queue.

The emergency vehicle – in our case, the ambulance – travels to the location given by the dispatcher. Once it arrives, emergency medical technicians (EMTs) treat the patient and prepare him/her to be transferred to the hospital. Depending on the clinical and operational criteria, the dispatcher chooses the most appropriate destination hospital for the patient. Then, the ambulance leaves the emergency location and transports the patient to the hospital. When the ambulance arrives at the hospital, the patient is transferred to the emergency room (ER). After the EMTs complete the administrative tasks associated with the patient, the ambulance becomes available again. At this point, the ambulance can be assigned a new incident, or it can be sent to a stand-by point where it waits for a new mission. Although this well-defined process models most EMS operations all over the world, the way in which decisions are made, especially the way in which vehicles are deployed over the territory, deeply influence the operational performance of this process.

From a managerial point of view, designing a deployment plan consists of locating a number of emergency vehicles to provide an adequate level of service. Three major questions need to be answered: How many vehicles are to be deployed? Where to deploy them? and What is meant by an "adequate" level of service?

The answers to the first and second questions greatly depend on the service level required by the organization. In fact, given a coverage time target, service managers try to locate

their vehicles in order to satisfy one or more criterion, such as minimize the number of required vehicles, minimize the average distance between vehicles and demand points, and/or maximize the probability of reaching an emergency location within the coverage time, among others. For the answer to the third question, the service level in EMS context is often measured in terms of the response time, which is the time between receiving the call and the moment when an ambulance reaches the emergency location. EMS also use the concept of coverage time, which is defined as the maximum acceptable response time. Thus, a service demand point is only considered to be covered if a call from this point can be served within a predetermined coverage time. It follows that, since EMS operations are stochastic in nature, EMS define an "adequate" service level in practice as the probability of reaching any emergency location in less than the predetermined coverage time.

This article focuses on the tactical problem of locating emergency vehicles while taking a dynamic demand into account. As the demand changes over time, different location plans need to be designed, requiring redeploying the vehicles when passing from one plan to another. The originality of this article is that it explicitly considers the redeployment of vehicles based on cyclic changes in service demand. For example, the daily population movements are characterized by the concentration of the population in urban areas during working hours (i.e., morning and early afternoon), while in the evening the population returns to their homes in the suburbs or outskirts of the cities.

In general, the average demand during the night is lower than other time periods during the day. Thus, since the fleet must be positioned in terms of the demand, the vehicle location in the morning, in the evening or at night may be different. In addition, the location plan must be designed in a way that minimizes the redeployments or transitions between the various plans. The redeployment of empty vehicles is almost unavoidable but should be minimized.

This article presents a mathematical model of the problem with dynamic, cyclical demands. In addition, two heuristic approaches are proposed, and their solutions are compared to those of the mathematical model. These approaches were tested on randomly-generated instances that reflect real-life conditions.

The paper is organized as follows. The next section provides a brief review of the literature. Sections 3 and 4 introduce the mathematical model and our two heuristic respectively, while section 5 reports the results of our numerical experiments. Finally, section 6 presents our conclusions.

2. Literature review

Although different EMS operate in relatively similar contexts, the rules and processes used can vary from one city to another and from one country to another. However, by analyzing different contexts, it is possible to identify some common features of their management processes. For example, in most cases, the territory to be served is divided into sub-areas. Some form of coordination between different sub-area operations is usually observed. Several types of EMS vehicles can be used to respond to service requests: Basic Life Support vehicle (BLS), Advanced Life Support vehicles (ALS), and Ambulance for Obese Patients (AOP). New technologies, including geographic information systems (GIS), are increasingly adopted, bringing with them new management opportunities. Finally, the emergency transport and the inter-hospital transport are managed separately in some cases and, in other cases, together.

The problem of ambulance deployment is often reduced to the Covering Location Problem (CLP). The CLP tries to locate a number of vehicles in a given space to cover a given demand. Brotcorne et al. (2003) and ReVelle and Eiselt (2005) review the literature related to the location and management of emergency vehicles. The literature identifies three main versions of the problem according to the nature of demand: the location of emergency vehicles in a purely deterministic context, in a probabilistic context, and in a dynamic deterministic context.

Most of the research in the literature belongs to the first category. Thus, the seminal research of Torregas et al. (1971) and Church and ReVelle (1974) has led to numerous extensions, including the work of Schilling et al. (1979), Hogan and ReVelle (1986) and Gendreau et al. (1997). However, in these models, the coverage of the population can become insufficient if one or more vehicles are occupied. A possible way to mitigate this drawback consists of adding a probabilistic component to measure the probability that a vehicle is occupied.

Daskin (1983) was the first to formulate and solve the probabilistic version of the problem: the MEXCLP (Maximum EXpected Coverage Location Problem). Several extensions have been proposed by Batta et al. (1989), Marianov and ReVelle (1994) and Ball and Lin (1993), among others. These models have different assumptions in terms of the travelling time, different definitions of the time slots for which vehicles are occupied, and different coverage requirements.

The planning horizon of all the models mentioned above is limited to a single period. However, ambulances can be relocated whenever the demand changes, which improves the service levels. In their study of the EMS at Louisville, Kentucky (USA), Repede and Bernardo (1994) proposed what appears to be the first multi-period model for ambulance deployment and redeployment, named the Maximal EXpected Coverage Location Problem with Time variation (TIMEXCLP), which extends Daskin's work (1982, 1983). TIMEXCLP continuously seeks to maximize the expected coverage, but unlike the previous models, it considers the variation of the demand, the number of vehicles to be located, and some other parameters at each time period. However, TIMEXCLP does not explicitly consider the costs associated with vehicle redeployment between periods. It also assumes that there is no limit on the number of vehicles used.

The DDSMt (Dynamic Dual Standard Model at time t) proposed by Gendreau et al. (2001) is, to the best of our knowledge, the only study dealing with the problem of relocating an ambulance fleet while seeking a compromise between the quality of redeployments and the cost of these redeployments.

More recently, Rajagopalan et al. (2008) proposed a multi-period Dynamic Available Coverage Location (DACL) model, which seeks to minimize the number of vehicles to be located in such a way that each demand zone will be covered with a certain probability over several time periods. However, this model does not account for vehicle redeployment between periods and the resulting cost of this redeployment. A Reactive Tabu Search algorithm was proposed by these authors (2008), and it was applied to solve an 8-period real-life problem (Mecklenburg County, North Carolina, USA). They solved the problem in only few minutes.

Despite its relevance, not enough research deals with multi-period situations. In addition, to the best of our knowledge, none of the studies explicitly considers the cost of redeploying the vehicles between periods. Redeployment is not an easy task; in practice, it provokes high levels of annoyance for employees and increases the cost of operating the fleet. The goal of this article is to solve the multi-period ambulance deployment and redeployment problem while maximizing the service levels and minimizing the travel distances caused by redeploying the vehicles between periods. In the following sections, we propose a mathematical model (section 3). Then, to overcome its limited capability to solve medium and large instances, we also propose two heuristic approaches (section 4).

3. Problem formulation

In this section, we propose a mathematical model for the multi-period emergency vehicle location and relocation problem. The objective is to simultaneously minimize (1) the distances between demand points and the location of the vehicles covering them, and (2) the redeployment distance traveled by vehicles between periods.

In the context of medical emergencies, each individual house or building can be seen as a potential demand point. However, a common practice in modeling such problems is to aggregate the demand points into regions or zones. For example, in North America, these points are grouped in areas sharing the same first three or four characters of the zip code or postal code. The location of an aggregated zone is represented by the zone's geographical center, and the corresponding demand is the sum of the expected demand of the aggregated zone. We implicitly assume that all service requests within a zone are “covered” if the zone’s geographical center is within the vehicle's travel distance.

Without loss of generality, let us assume that a workday is divided into $t = \{1, \dots, T\}$ periods, and that the demand for each of these periods is known and remains the same from one day to another. Let $i = \{1, \dots, I\}$ and $j = \{0, \dots, J\}$ be, respectively, the indexes for demand zones and potential vehicle sites. The index $j=0$ indicates the garage or the vehicle depot. Let us also assume that, at most, one ambulance can be located in a given point (except the depot) and that a vehicle can cover, at most, K demand zones.

The distance between each demand zone i and each potential vehicle location j is denoted by d_{ij} , and the travel cost per distance unit is denoted by c . A demand zone i is considered

to be covered by a vehicle located at site j if d_{ij} is less than or equal a predetermined distance, denoted r . The set of location sites that can cover demand zone i is given by $N_i = \{j \in J \mid d_{ij} \leq r\}$. In addition, let $l = \{1, \dots, p\}$ be the set of available vehicles, and p_t be the number of vehicles available at period t . We assume that all vehicles are identical and, for each period, the number of available vehicles is at least equal to p_{min} , the minimum number of vehicles necessary to cover all the demand points.

The average number of requests originating in a zone i during period t (i.e., the demand of zone i) is denoted a_{it} . Based on a_{it} , managers set a parameter f_{it} , which gives the minimal number of vehicles that should be assigned to cover zone i in order to insure an adequate service during period t . Let us also define the following decision binary variables:

$x_{ijt} = 1$ if demand zone i is to be covered by a vehicle located at site j during period t ;

$z_{jlt} = 1$ if vehicle l is located at site j during period t ;

$y_{jkl} = 1$ if vehicle l is redeployed from site j to site k ($j \neq k$) at the beginning of period t .

Thus, the proposed mathematical model $M1$ of our problem is as follows:

$$\text{Minimize} \quad c \sum_{i \in I} \sum_{j \in J} \sum_{t \in T} \frac{a_{it}}{f_{it}} d_{ij} x_{ijt} + c \sum_{j \in J} \sum_{k \in J} \sum_{t \in T} \sum_{l \in L} d_{jk} y_{jkl} \quad (1)$$

$$\text{Subject to:} \quad \sum_{j \in N_i} x_{ijt} = f_{it} \quad \forall i, t \quad (2)$$

$$\sum_{l \in L} z_{jlt} \leq 1 \quad j \in \{1, \dots, J\}, \forall t \quad (3)$$

$$\sum_{j \in J} z_{jlt} = 1 \quad \forall l, t \quad (4)$$

$$\sum_{j=1}^m \sum_{l \in L} z_{jlt} = p_t \quad \forall t \quad (5)$$

$$x_{ijt} \leq \sum_{l \in L} z_{jlt} \quad j \in N_i, \forall i, t \quad (6)$$

$$\sum_{i \in I} x_{ijt} \leq K \quad j \in \{1, \dots, J\}, \forall t \quad (7)$$

$$y_{jkl} \geq z_{klt} + z_{jlt-1} - 1 \quad k \in \{J \mid j \neq k\}, t = \{2, \dots, T\}, \forall j, l \quad (8)$$

$$y_{jkl} \geq z_{klt} + z_{jls} - 1 \quad k \in \{J \mid j \neq k\}, \forall j, l \quad (9)$$

$$x_{ijt}, y_{jkl}, z_{jlt} \in \{0, 1\} \quad (10)$$

The objective of this model is to minimize the sum of average covering costs and redeployment costs. Average covering costs are calculated using the cost per distance unit multiplied by the sum of distances between each demand zone i and the location of vehicles that have been assigned to cover it, multiplied by the ratio between the expected demand of the zone i and the number of ambulances required to cover it. Redeployment cost is calculated using the distance travelled by empty ambulances between two consecutive periods multiplied by the cost per unit of distance. The cost of moving vehicles from the depot to their location and back to the depot is also included.

Constraints (2) insure full coverage for every demand zone for every period. Constraints (3) and (4) respectively require that a maximum of one vehicle can be assigned to a site at each period, and that each vehicle is assigned to either a potential site or the garage. Constraints (5) require that exactly p_t vehicles be deployed at each period t . Constraints (6) require that, to cover a demand zone, a vehicle must be located at a site j within the predetermined coverage distance r . In the case in which a vehicle is located at site j , not all the points within the coverage distance of this site are necessarily covered by it. The model explicitly indicates the vehicles assigned to each zone by fixing the value of the binary decision variables x_{ijt} . Constraints (7) state that a vehicle cannot cover more than K demand zones. Constraints (8) determine the vehicles to be redeployed between two consecutive periods, while constraints (9) determine the redeployment of vehicles between the last and the first period, assuming a cyclic schedule. Finally, constraints (10) indicate the binary nature of decision variables.

One important feature of this model is that it limits the number of demand zones covered by each vehicle and assigns vehicles to demand zones. For example, although a demand zone needs the service of two vehicles but is within the covering distance of 4 vehicles, only two of these vehicles are assigned to cover this demand zone. In addition, if $K=3$ and a vehicle is within a covering distance to 5 demand zones, the model assigns the vehicle

to only three of them. The model is unfeasible if, for any zone i and at any period t , $|N_i| < f_{it}$. Thus, we assume that vehicle location sites were chosen so that $|N_i| \geq f_{it}$ for all demand zones i and all periods t .

Finally, assuming that we consider all demand zones as potential sites to locate vehicles (or that the number of potential sites is close to the number of demand zones), the size of the model becomes numerically intractable for instances with 30 or more demand zones. For such large instances, we needed to design efficient heuristics. The following section presents the two heuristic approaches that we designed to solve large instances drawn from real life.

4. Heuristic approaches

Empirical observations have shown that the computation time required by commercial Mixed Integer Programming (MIP) codes is too large when considering more than 30 demand points and 30 sites. In order to find a good compromise between the quality of the solutions and time required to obtain them, we propose the following decomposition method, which exploits the periodical nature of the problem. The problem is divided into T periods, in which the initial vehicle locations for period t are those for the end of period $t-1$, unless it is decided to relocate some of the vehicles. The heuristic determines those vehicles to be relocated and their new location.

Once the set of T individual periods has been solved, the decomposition method iterates (i.e., repeats the calculation from the first period), taking the location of the final period as the starting location for the first period until the locations obtained remain stable from one iteration to the next. We propose two ways to solve each of the T single-period problems: $H1$, a heuristic based on the mathematical model presented in the previous section, and $H2$, a heuristic composed of a construction heuristic and a local search improvement heuristic.

4.1 Heuristic $H1$: solving single-period sub-problems using a single-period version of $M1$

Heuristic $H1$ solves a given problem by dividing it into several single-period sub-problems and solving each sub-problem using a single-period version of model $M1$,

named model $M2$. To this end, the heuristic needs to take into account explicitly the importance of potential vehicle redeployments between periods. $M2$ is composed of equations (1) to (7) and (10) of $M1$. The period index t is no longer needed. However, the following constraints must be added:

$$y_{jkl} \geq b_{jl} z_{kl} \quad \forall j, l \quad \text{et} \quad \forall k \neq j \quad (11)$$

These constraints use a parameter b_{jl} that takes the value 1 if vehicle l was located at site j at the end of the previous period. Thus, constraints (11) replace constraints (8) and (9) to model vehicle relocation decisions.

Heuristic $H1$ consists of the following 4 steps:

1. Solve the first period of the problem ($t=1$) by solving $M2$, where the objective function contains only the first term (i.e., the covering costs). Constraints (11) are ignored.
2. Solve the model $M2$ sequentially for the other periods ($t=2 \dots T$) after setting the values of parameter b_{ij} to indicate the vehicle locations as obtained for the previous period $t-1$.
3. To comply with the cyclical nature of demand, solve the first period again, but this time taking into account the vehicle locations obtained for the last period T .
4. Repeat steps 2 and 3 until convergence (i.e., until vehicle locations for all periods remain the same and do not change after a complete iteration). In cases for which the heuristic does not converge after a predetermined number of iterations, stop the heuristic and retain the best solution found.

Numerical experiments have shown that this heuristic is able to solve instances with up to 200 demand points and 200 potential sites. If a larger problem needs to be dealt with, the easiest approach is to divide the considered region into several sub-regions, and then to apply $H1$ to each of them. Another possible approach is to reduce the set of potential vehicle location sites. Otherwise, the heuristic $H2$ can be used to solve these large instances.

4.2 Heuristic *H2*: solving single-period problems using a two-phase heuristic

Heuristic *H2* is an iterative two-phase heuristic in which a construction phase is followed by a local search improvement phase. The same decomposition strategy used in *H1* is used in *H2*. The problem is divided into T single-period sub-problems, and each sub-problem is solved using the final vehicle locations of the previous period as initial locations for the next period.

Heuristic *H2* consists of the following steps:

1. For each period $t = 1, \dots, T$, apply the *construction heuristic* and then the *local search heuristic*.
2. Repeat step 1 until convergence (i.e., until vehicle locations for all periods remain the same and do not change after a complete iteration).
3. In cases for which the heuristic does not converge after a predetermined number of iterations, stop the heuristic and retain the best solution found.

Let us now present both the *construction heuristic* and the *local search heuristic* used in heuristic *H2*.

The construction heuristic

The construction heuristic is composed of the following steps:

1. For each demand zone i , determine N_i , the set of potential sites within the covering distance r from i .
2. For each potential location j , determine M_j (the set of demand points covered by j) and \bar{d}_{M_j} (the average distance between j and the demand zones in M_j), where
$$\bar{d}_{M_j} = \frac{1}{|M_j|} \sum_{i \in M_j} d_{ij}.$$
3. For a demand zone i , if $|N_i| < f_i$, there is no feasible solution; stop the construction heuristic.
4. For each demand zone i where $|N_i| = f_i$, assign a vehicle to each location site j in the set N_i .
5. If the total number of located vehicles exceeds the number of available vehicles, there is no feasible solution; stop the construction heuristic.

6. For each located vehicle, if $|M_j| \leq K$, assign j to cover each demand zone in M_j . Otherwise, assign j to cover the K demand zones in M_j with the lowest value of the ratio $(|N_i| * f_i)/(a_i * f_{ri})$, where f_{ri} is the additional number of vehicles required by demand zone i to reach f_i .
7. Update f_{ri} , N_i and $|N_i|$.
8. For every potential location where no vehicle has been located yet, update M_j and \bar{d}_{M_j} .
9. Sort the demand zones that have $f_{ri} > 0$ (i.e., the demand zones not yet fully covered) in ascending order of the ratio $(|N_i| * f_i)/(a_i * f_{ri})$.
10. Select the first demand zone in this ordered list and place a vehicle to the location $j \in N_i$ where (a) no vehicle is yet located and (b) a vehicle was located in the previous period. If there are several such locations or if the first period is being considered, place a vehicle in the location $j \in N_i$ with the largest $|M_j|$ that does not have a vehicle yet. To break a tie, choose the location with the lowest \bar{d}_{M_j} .
11. Repeat steps 5-10 until all vehicles have been located.

The local search heuristic

The local search heuristic changes the location of a vehicle to another empty location to test if doing so improves the solution. In order to limit the set of potential locations to search, only neighbor locations within a predetermined search distance s are considered. The value of the search distance s is set by the user and can be modified while the heuristic is running. If a vehicle is located in the same location as the previous period, a smaller search distance s is used in order to minimize potential redeployment cost. The vehicle location is changed if this reduces the covering costs (i.e., the weighted sum of the distances between each demand zone and the vehicles that cover it). However, once the position of a vehicle is changed, the whole assignment of vehicles to demand zones must be updated.

To speed up the execution of such a complex re-assignment, the following procedure is used:

1. For every demand zone i , assign the f_{ii} nearest vehicles to i , provided that they are within the cover distance r .
2. Identify the set O , which contains vehicles that have been assigned to cover more than K demand zones.
3. For each location $j \in O$, transfer the excess vehicle assignments to another vehicle as follows: among all the demand zones assigned to j , select the one with the shortest distance to another location $j' \notin O$ and reassign j' to i . Repeat until O is empty.

If all the potential relocations of all the vehicles have been tested and none of them led to an improved solution, double exchanges are considered. A double exchange consists of simultaneously moving two vehicles from their current locations to two empty locations within their respective search distances. The reassignment process is then applied to every potential double exchange move.

In general, relocating a vehicle may make it impossible to find a feasible reassignment of vehicles to demand zones. In this case, the relocation is discarded. Although the above procedure leads to a plan for reassigning vehicles to demand zones that may not be optimal, it gives a rapid and good approximation.

5. Computational experiments

The goal of this section is to assess the performance of our heuristic approaches, in terms of solution quality and computational time, by comparing them to exact values obtained by solving the model $M1$.

5.1 Test instances

All the instances we used consisted of three periods. In order to be able to compare the results of our heuristics to the optimal solution, small instances containing $I = \{20, 25, 30, 35 \text{ and } 40\}$ demand zones were used. For each value of I , 10 instances were generated. The location of demand zones was generated randomly within a square of 100 x 100 units. The covering distance r ranges from 30 to 40, while the parameter K varies between 5 and 10. Travel cost per distance unit was set to \$ 1.

Without loss of generality, we assumed that vehicles may be located at any demand zone. The average number of requests of each demand zone (a_{it}) was drawn from a uniform discrete distribution between 1 and 10. The number of vehicles needed to cover a demand zone (f_{it}) was set to 1 if $a_{it} \leq 8$; otherwise, it was set to 2. The number of vehicles to be located was equal to the minimum necessary to fully cover all the demand points. This minimum was found *a priori* using the *location set covering* formulation proposed by Toregas et al. (1971). In the cases for which $H2$ was unable to produce a feasible solution with this number of vehicles, we added as many vehicles as necessary in order to reach a feasible solution.

5.2 Comparing heuristic and optimal results

Computational results are shown in Table 1, 2 & 3. Table 4 provides a summary of these results. All these tables present the results obtained by the mathematical model $M1$ and heuristics $H1$ and $H2$ to 10 different instances with 20, 30, or 40 demand zones, respectively. The $M1$ model and the heuristic $H1$ were solved with CPLEX 10.0. All the experiments were run on an IBM computer with a Xeon processor 3.60 GHz and a maximum computational time of 72 000 seconds was allotted to each instance.

Columns \mathbf{r} and \mathbf{p}_{\max} indicate the coverage distance and the maximum number of vehicles used, respectively. Columns \mathbf{BInt} and \mathbf{Dev} , respectively, report the value of the best integer solution produced by CPLEX for $M1$ within the allotted computational time, if optimality was not proven, and its gap with respect to the best lower bound. For the heuristics, column \mathbf{Dev} reports the gap between the solutions produced with respect to the best integer solution produced by $M1$ within the allotted computational time. Column \mathbf{t} reports the computational time.

Table 1 reports the results for the smallest instances (20 demand zones). For such small instances, $M1$ is very efficient, producing optimal solutions for all the instances in less than 1 000 seconds of computational time. However, $H1$ gives a very good performance, producing close to optimal results (i.e., average deviation of 0.40%) in less than 3 seconds. $H2$ also performs efficiently, but, on average, it requires more computational time and generally produces lower quality solutions than $H1$.

#	r	p _{max}	M1		H1		H2	
			BInt	t	Dev	t	Dev	t
1	25	8	2 722.96	1 131.86	0.06%	0.80	0.24%	13.91
2	30	7	3 973.94	20.94	0.44%	0.29	1.04%	4.05
3	30	7	4 466.22	87.80	0.39%	0.79	0.35%	2.99
4	30	7	2 999.88	41.89	0.03%	0.63	3.96%	2.03
5	35	7	5 080.12	927.72	0.39%	0.51	0.56%	11.17
6	35	6	4 785.27	169.58	0.22%	2.36	2.64%	1.72
7	35	6	3 390.32	192.36	0.36%	3.00	14.24%	1.59
8	40	5	4 330.08	22.16	0.66%	0.58	0.48%	1.59
9	40	6	4 740.26	319.38	0.22%	2.81	7.77%	2.05
10	50	3	6 836.87	2.51	1.27%	0.48	0.04%	1.24
Average				291.62	0.40%	1.23	3.13%	4.23

Table 1: Results for instances with 20 demand zones

Table 2 reports the results for instances having 30 demand zones. As these results show, it becomes harder for CPLEX to reach optimality as problem size increases. In fact, for two of the ten considered instances, CPLEX did not succeed closing the optimality gap. At the same time, the performance of *H1* remains excellent. For 8 of the 10 instances, the results obtained by *H1* are within 1% of the best integer solutions produced by *M1* in the allotted time, and for one instance (instance 8), it produces the optimal solution. However, in general, *H2* seems to be slightly dominated by *H1*, although *H2* has found the optimum for two instances (instances 1 and 9).

#	r	p _{max}	M1			H1		H2	
			BInt	Dev	t	Dev	t	Dev	t
1	25	10	5 367.15		6 954.92	1.88%	1.23	0.00%	15.38
2	25	8	4 563.52		30 083.48	0.26%	1.78	2.29%	17.77
3	30	10	5 762.75	0.26%	72 000.00	0.46%	1.88	0.10%	15.98
4	30	5	3 966.02	1.22%	72 000.00	0.64%	1.42	2.58%	15.25
5	30	10	5 517.01	1.37%	72 000.00	1.57%	1.78	1.34%	11.27
6	35	8	6 720.30		2 903.14	0.84%	1.61	0.61%	26.31
7	35	10	8 013.21		2 48.11	0.40%	0.69	0.33%	2.66
8	35	8	6 447.55		1 896.13	0.00%	0.98	1.79%	3.72
9	35	12	6 938.59		93.94	0.74%	1.74	0.00%	2.59
10	40	10	8 035.75		270.03	0.22%	2.11	2.48%	3.13
Average				0.29%	25 844.98	0.70%	1.52	1.15%	11.41

Table 2: Results for instances with 30 demand zones

Table 3 reports the results for instances having 40 demand zones. CPLEX was able to produce optimal solutions for only 2 out of 10 instances; its average gap with respect to the best lower bound was 2.13%. *H1* performs extremely well, producing solutions of better or equal quality than *M1* (in bold) produced for half of the tested instances. Moreover, these results confirm that *H1* dominates *H2* since *H1* produced better results for all instances.

#	r	p _{max}	<i>M1</i>			<i>H1</i>		<i>H2</i>	
			BInt	Dev	T	Dev	t	Dev	t
1	25	12	6 823.23	3.13%	26 793.19	-0.02%	4.94	9.36%	39.41
2	30	8	7 053.60		27 899.75	0.00%	0.97	0.77%	7.41
3	30	13	5 819.36	6.00%	3 600.00	0.38%	2.26	10.50%	21.86
4	30	10	5 548.31	2.14%	27 015.26	0.00%	0.86	7.18%	15.30
5	35	7	8 228.06		1 049.75	0.07%	0.86	9.49%	7.47
6	35	8	8 500.26	1.69%	72 000.00	-0.05%	0.81	1.94%	21.85
7	35	10	8 370.70	3.11%	72 000.00	-0.43%	6.74	3.48%	11.92
8	40	7	7 936.85	0.33%	72 000.00	0.44%	1.01	7.79%	10.46
9	40	10	5 519.45	2.62%	72 000.00	0.51%	1.92	18.91%	14.14
10	40	7	8 508.12	2.29%	26 103.83	0.35%	8.08	12.80%	12.84
Average				2.13%	40 046.18	0.13%	2.85	8.22%	16.27

Table 3: Results for instances with 40 demand zones

Table 4 summarizes the results reported in Tables 1, 2 and 3. It includes also aggregated results for two groups of ten instances with 25 and 35 demand zones, respectively. Column *I* gives the number of demand zones, while column # reports the number of times, out of 10, that the approach attains the best known solution. Columns **ADev** report for each group of instances, the average gap between the best solutions found by the method and the best lower bound produced by CPLEX. These results show that the efficiency of *M1* drops as the size of the problems increases.

I	<i>M1</i>			<i>H1</i>			<i>H2</i>		
	#	ADev	t	#	ADev	t	#	ADev	t
20	10	0.00%	291.62		0.40%	1.22		3.13%	4.23
25	8	0.46%	17 518.06	2	0.41%	1.32		3.92%	6.29
30	7	0.29%	25 844.98	1	0.70%	1.52	2	1.15%	11.41
35	7	1.55%	20 576.56	2	0.03%	2.65		5.39%	7.83
40	2	2.13%	40 046.18	5	0.13%	2.85		8.22%	16.27

Table 4: Aggregated results

These results also confirm the excellent performance of $H1$, which is a very interesting element to implement in commercial decision support systems.

6. Conclusion

The deployment and redeployment of vehicles is one of the most important problems in EMS (Emergency Medical Services) management. A good deployment plan can increase the level of service to the population and reduce the number of vehicles required. However, since demand pattern evolves over the day, a deployment plan may be optimal to one of the day's periods, but not so favorable for other periods, leading to the need for vehicle redeployment. This paper proposes a mathematical model that explicitly considers the dynamic nature of the demand and the need for a fleet redeployment. In addition, it proposes two heuristic approaches, $H1$ and $H2$, to solve large instances. The numerical results obtained prove the ability of these heuristics, especially heuristic $H1$, to produce near-to-optimal solutions in very short computation times, confirming the potential application of such approaches within an EMS system to manage a fleet of emergency vehicles in real-time.

Acknowledgements

This research was partially financed by the Canadian Natural Sciences and Engineering Research Council (NSERC) under Grants OGP0036509 and OGP0293307. This support is gratefully acknowledged.

References

- Ball, M.O. and Lin, L.F (1993) A reliability model applied to emergency service vehicle location. *Operations Research*, 41, 18-36.
- Batta, R., Dolan, J.M., and Krishnamurty, N.N. (1989) The maximal expected covering location problem: Revisited. *Transportation Science*, 23, 277-287.
- Brotcorne, L., Laporte, G. and Semet, F. (2003) Ambulance location and relocation models. *European Journal of Operational Research*, 147, 451-463.
- Church, R.L. and ReVelle, C.S. (1974) The maximal covering location problem. *Papers of the Regional Science Association*, 32, 101-118.

- Daskin, M. S. (1982) Application of an expected covering model to emergency medical service design, *Decision Sciences*, vol. 13, 416-439.
- Daskin, M. S. (1983) A maximum expected location problem: Formulation, properties and heuristic solution, *Transportation Science*, vol. 17, pp. 416-439.
- Gendreau, M., Laporte, G. and Semet, F. (1997) Solving an ambulance location model by Tabu search. *Location Science*, 5, 75–88.
- Gendreau, M., Laporte, G. and Semet, F. (2001) A dynamic model and parallel Tabu search heuristic for real-time ambulance relocation. *Parallel Computing*, 27, 124-141.
- Hogan, K. and ReVelle, C.S. (1986) Concepts and applications of backup coverage. *Management Science*, 34, 1434-1444.
- Marianov, V. and ReVelle C.S. (1994) The queueing probabilistic location set covering problem and some extensions. *Socio-Economic Planning Sciences*, 28, 167–178.
- Rajagopalana, H.K., Saydamb, C., and Xiao J. (2008) A multiperiod set covering location model for dynamic redeployment of ambulances, *Computers & Operations Research* 35, 814-826.
- Repede, J.F. and Bernardo J.J. (1994) Developing and validating a decision support system for locating emergency medical vehicles in Louisville, Kentucky. *European Journal of Operational Research*, 75, 567–581.
- ReVelle, C.S. and Eiselt, H.A. (2005) Location analysis: A synthesis and survey. *European Journal of Operational Research*, 165, 1-19.
- Schilling, D.A., Elzinga, D.J., Cohon, J., Church, R.L. and ReVelle, C.S. (1979) The TEAM/FLEET models for simultaneous facility and equipment siting. *Transportation Science*, 13, 163-175.
- Toregas, C., Swain, R., ReVelle, C.S. and Bergman, L. (1971) The location of emergency service facilities. *Operations Research*, 19, 1363-1373.