Pergamon

Location Science, Vol. 5, No. 2, pp. 75–88, 1997

Printed in Great Britain
0966-8349/98 $19.00 + 0.00

PII:S0966-8349(97)00015-6

SOLVING AN AMBULANCE LOCATION MODEL BY TABU SEARCH

MICHEL GENDREAU*

Département d'informatique et de recherche opérationnelle et Centre de recherche sur les transports, Université de Montréal, Case postale 6128, succursale Centre-ville, Montreal, Canada H3C 3J7

GILBERT LAPORTE

Centre de recherche sur les transports, Université de Montréal, Case postale 6128, succursale Centre-ville, Montreal, Canada H3C 3J7

and

FRÉDÉRIC SEMET

Centre de recherche sur les transports and GRIS, Département d'Administration de la Santé, Université de Montréal, Case postale 6128, succursale Centre-ville, Montreal, Canada H3C 37J

Abstract—This paper considers a double coverage ambulance location problem. A model is proposed and a tabu search heuristic is developed for its solution. Computational results using both randomly generated data and real data confirm the efficiency of the proposed approach.

Key words: Ambulance location, coverage models, tabu search heuristic.

1. INTRODUCTION

The provision of effective and cost efficient ambulance services is a problem encountered in all major cities. Faced with budget cuts, several authorities have had, in recent years, to reassess or reorganize their emergency services, see for example, Fuziwara *et al.* (1987); Trudeau *et al.* (1989); Goldberg *et al.* (1990b); Repede and Bernardo (1994). The problem has long been studied by operations researchers, but with the advent of new solution methodologies, particularly in the field of metaheuristics (Osman and Laporte, 1996), and with the development of computer and telecommunication technologies, solution techniques of the past may no longer be the answer to today's needs.

Decisions in the context of ambulance services typically arise at three different levels.

- *Strategic decisions* involve the location and construction of fixed facilities, the purchase of equipment and the hiring and training of specialized staff.
- *Tactical decisions* relate to staff scheduling, location of emergency vehicles at any point in time and deployments and relocation of vehicles when calls are received.
- *Operational decisions* are concerned with procedures to be followed by paramedical staff depending on the nature of calls.

*To whom correspondence should be addressed.

The focus of this paper is on locational decisions made at the tactical level. More specifically, the aim is to develop and solve heuristically a static coverage location model. It is intended to embed this model, at a later stage, within a decision support system to assist real-time vehicle redeployment operations.

The aim of ambulance location models is to provide adequate coverage. This can be interpreted and measured in many ways. The most common models are those where ambulances are located to ensure single coverage, that is, all the population lies within r time units of an ambulance. This model has been extended to double coverage defined in the following sense. Two radii r_1 and r_2 ($r_2 > r_1$) are used. All the demand must be covered by an ambulance located within r_2 time units, and, in addition, a proportion α of the demand must also be within r_1 units of an ambulance, which may or may not be the same ambulance that covers this customer within r_2 time units. The United States Emergency Medical Services (EMS) Act of 1973 (see Ball and Lin, 1993) sets the following standards; $r_1 = 10$ min, $\alpha = 0.95$. Currently, no value is set for r_2. In Montreal, ambulances are run by "Urgences Santé" which uses $r_1 = 7$ min and $\alpha = 0.90$ (Desrosiers and Thibault, 1996). Urgences Santé would also like to attain a standard of $r_2 = 15$ min in the near future. Various covering models have been proposed in the literature; see, for example, Toregas *et al.* (1971), Church and ReVelle (1974), Daskin (1983), Hogan and ReVelle (1986), ReVelle and Hogan (1989), Ball and Lin (1993) and Marianov and ReVelle (1996). Depending on their size and sophistication, such models can either be solved exactly (ReVelle and Hogan, 1989; Ball and Lin, 1993; Marianov and ReVelle, 1996), or approximately (Daskin, 1983). To the authors' knowledge, Hogan and ReVelle (1986) were the first to incorporate double coverage in their model. A number of authors have used analytical tools to assess the quality of a proposed solution, but not to actually construct a solution. A common methodology is simulation. It has been employed by Davis (1981) and Goldberg *et al.* (1990a). More sophisticated techniques are based on queuing theory; see, for example, Larson (1974), Larson and Odoni (1981), Brandeau and Larson (1986) and Burwell *et al.* (1992). These evaluative concepts can be combined with local search procedures to yield good solutions to the problem (Jarvis, 1975; Fitzsimmons and Srikar, 1982; Trudeau *et al.*, 1989). The surveys by Daskin *et al.* (1988) and by Marianov and ReVelle (1995) provide interesting overviews of available models in the more general area of emergency vehicle siting.

The authors' aim is to develop a double coverage model, and to design a tabu search heuristic for its solution. The proposed model is simpler than some alternatives described in the literature (see, e.g., Ball and Lin, 1993) as it only considers population coverage and does away with sophisticated probability driven parameters, which are typically hard to measure and evaluate in practice. The method considered should also be quick, robust and capable of producing high quality solutions on instances of realistic dimensions. The computational results show that the tabu search heuristic possesses these desirable features.

The model is presented in the next section, followed by the algorithm in Section 4, and by the conclusion in Section 5.

2. MODEL

The model proposed solves the double coverage problem described in Section 1 for a given number of ambulances, where the objective is to maximize the demand covered by two ambulances within a radius r_1. The problem is defined on a graph $G = (V \cap W, E)$ where $V = \{v_1, \ldots, v_n\}$ and $W = \{v_{n+1}, \ldots, v_{n+m}\}$ are two vertex sets representing, respectively, demand points and potential location sites, and $E=\{(v_i,v_j):v_i \in V \text{ and } v_j \in W\}$ is an edge set. With each edge (v_i,v_j) is an associated travel time t_{ij}. The demand at vertex $v_i \in V$ is equal to λ_i. The number of ambulances is given and equal to p. Also define for $v_i \in V$ and $v_{n+j} \in W$ the coefficients

$$\gamma_{ij}=\begin{cases} 1 & t_{i,n+j} \leqslant r_1 \quad (v_i \text{ is covered within the small radius } r_1) \\ 0 & \text{otherwise} \end{cases}$$

and

$$\delta_{ij}=\begin{cases} 1 & \text{if } t_{i,n+j} \leqslant r_2 \quad (v_i \text{ is covered within the large radius } r_2) \\ 0 & \text{otherwise.} \end{cases}$$

Again, α is the proportion of the total demand that must be covered by an ambulance located within r_1 units. The following variables are also used: y_j is an integer variable denoting the number of ambulances located at $v_{n+j} \in W$, and p_j is an upper bound on y_j; x_i^k is a binary variable equal to 1 if and only if v_i is covered at least k times within the small radius r_1, where $k = 1$ or 2. The formulation is then:

$$(P) \text{ maximize } f=\sum_{i=1}^{n} \lambda_i x_i^2 \tag{1}$$

subject to

$$\sum_{j=1}^{m} \delta_{ij} y_j \geq 1 \quad (v_i \in V) \tag{2}$$

$$\sum_{i=1}^{n} \lambda_i x_i^1 \geqslant \alpha \sum_{i=1}^{n} \lambda_i \tag{3}$$

$$\sum_{j=1}^{m} \gamma_{ij} y_j \geqslant x_i^1 + x_i^2 \quad (v_i \in V) \tag{4}$$

$$x_i^2 \leqslant x_i^1 \quad (v_i \in V) \tag{5}$$

$$\sum_{j=1}^{m} y_j = p \tag{6}$$

$$y_j \leqslant p_j \quad (v_{n+j} \in W) \tag{7}$$

$$x_i^1, x_i^2 \in \{0,1\} \quad (v_i \in V) \tag{8}$$

$$y_j \text{ integer } (v_{n+j} \in W) \tag{9}$$

In this model, the objective represents the total demand covered at least twice within r_1 units. Constraints (2) and (3) express the single and double coverage requirements. Constraints (2) state that all demand must be covered within r_2 units. The left-hand side of constraints (4) counts the number of ambulances covering v_i within r_1 units; the right-hand side is equal to 1 if v_i is covered once within r_1 units, and equal to 2 if it is covered at least twice within r_1 units. Constraints (3) and (4) taken together ensure that a proportion α of all demand is covered (constraints (3)) and the coverage radius must be r_1 units since by constraints (4), $x_i^1 + x_i^2 = 0$ whenever $\gamma_{ij} y_j = 0$ for all j. By constraints (5), a vertex v_i cannot be covered at least twice if it is not covered at least once. Constraints (6), (7) and (9) impose limits on the number of ambulances at each site. Note that provided $p \leqslant 2m$, there always exists an optimal solution in which $y_j \leq 2$ for all j since nothing is gained by covering a demand point more than twice. Thus, in practice one can impose $p_j = 2$ in equation (7) and $y_j \in \{0,1,2\}$ $(v_{n+j} \in W)$ instead of equation (9).

An interesting feature of this type of model is that it incorporates requirements such as those imposed by the EMS Act of 1973. However, a feasible solution may not exist if the parameters r_1, r_2 and α are too restrictive. In fact, no model can guarantee a feasible solution at all times in practice since demand for ambulances is highly stochastic and circumstances may occur where there is a shortage of available ambulances or insufficient coverage in some areas. When no feasible solution exists, the possible courses of action are to increase the number of ambulances or to relax the coverage constraints. The same can happen in the model of Hogan and ReVelle (1986), which has strict coverage constraints. In some other models (see, e.g., ReVelle and Hogan (1989) and Marianov and ReVelle (1996)), the coverage requirements are expressed by the objective function and there is always a feasible solution. In the Ball and Lin (1993) model, a solution always exists since the number of ambulances is not fixed *a priori*. The main characteristics of the available coverage models are summarized in Table 1.

3. TABU SEARCH HEURISTIC

A tabu search heuristic capable of providing high quality solutions within modest computing times has been developed. Essentially, tabu search is a local search method that moves at each iteration from a solution to its best neighbour even if this causes the objective value to deteriorate. To avoid cycling, solutions similar to recently examined solutions are forbidden, or *tabu*, for a number of iterations. The modern version of this method is rooted in the work

Table 1. Comparison of five ambulance location coverage models

Model	Objective	Coverage constraints	Number of ambulances
Hogan and ReVelle (1986)	Maximize a linear combination of demand covered at least once and at least twice within r_1.	All demand covered within r_2.	Total given.
ReVelle and Hogan (1989)	Maximize the total demand covered an an appropriate number of ambulances.	None.	Total given. At most one ambulance per site.
Ball and Lin (1993)	Minimize the sum of ambulance fixed costs.	Proportion α of all demand covered within r_1.	Decision variable. No constraint on the number of ambulances per site.
Marianov and ReVelle (1996)	Maximize the total demand covered an an appropriate site specific number of ambulances.	None.	Total given. Upper bound ($\geqslant 1$) on the number of ambulances per site.
This paper	Maximize the total demand covered at least twice within r_1.	All demand covered within r_2. Proportion α of all demand covered within r_1.	Total given. Upper bound ($\geqslant 1$) on the number of ambulances per site.

of Glover (1986) and that of Hansen and Jaumard (1990). In recent years, several researchers have refined and extended the basic ideas (see, e.g., Osman and Laporte (1996)). To the authors' knowledge, this work is the first to apply tabu search to a covering location problem, although some algorithms such as that of Rolland *et al.* (1996) for the p-median problem could possibly be applied to covering problems. In the following subsections, the main components of the algorithm are presented, followed by a step-by-step description.

3.1. Solution and objective function

At any iteration of the algorithm, a solution is fully specified by the number of ambulances located at each vertex of W. During the course of the algorithm, a solution may be infeasible with respect to the covering constraints. The algorithm works on constraints (2) and (3) and on the objective function in a hierarchical fashion, that is, it operates with pre-emptive priorities, similar to what is done in goal programming. It is convenient to define a mega-objective function $F(s)$ to be maximized, associated with each solution s, defined as

$$F(s)=f(s)-M_1 f_1(s)-M_2 f_2(s), \tag{10}$$

where M_1, M_2 are two weights satisfying $M_1 > M_2 > 1$,

$$f(s)=\sum_{i=1}^{n} \lambda_i x_i^2, \tag{11}$$

$$f_1(s)=\left|\left\{v_i \in V: \sum_{j=1}^{m} \delta_{ij} y_j \geqslant 1\right\}\right|, \tag{12}$$

and

$$f_2(s)=\min\left\{\alpha, \sum_{i=1}^{n} \lambda_i x_i^1 \Big/ \sum_{i=1}^{n} \lambda_i\right\}. \tag{13}$$

Here, $f(s)$ is the original objective representing the total demand covered at least twice within r_1 units, $f_1(s)$ is equal to the number of demand points covered within r_2 units, and $f_2(s)$ is the minimum between α and the properties of the demand covered at least once within r_1 units. Note that $F(s)$ can be computed even if s is infeasible.

3.2. Initial solution

At iteration $t = 0$, the linear relaxation of (P) is solved. If it is infeasible, the overall problem has no solution and some constraints must be relaxed. If the linear relaxation is feasible, consider the y_j values $\bar{y}_1,\ldots,\bar{y}_m$, and let $\bar{f}$ be the value of the solution. If all $\bar{y}_j$ are integer, the current solution is feasible and optimal. Otherwise, allocate $\lfloor \bar{y}_j \rfloor$ ambulances to each vertex $v_{n+j} \in W$, and compute $Y=\sum_{j=1}^{m}\lfloor \bar{y}_j \rfloor$. Compute $S=\{v_{n+j} \in W: 0<\bar{y}_j<1\}$. If $p-Y \geqslant |S|$, allocate one ambulance to each vertex of S and randomly allocate at most one of the remaining $p-Y-|S|$ available ambulances to the vertices of $T=\{v_{n+j} \in W: 1<\bar{y}_j<2\}$. If $p-Y<|S|$, randomly allocate at most one of the remaining $p-Y$ available ambulances to the vertices of S. The initial solution s produced may be infeasible, but it uses p ambulances. In the following, $\bar{y}_j$ denotes the number of ambulances located at vertex $v_{n+j} \in W$. Vertex v_{n+j} is said to be *saturated* if $\bar{y}_j = p_j$.

3.3. Basic operations and tabu status

A basic operation in this algorithm consists of displacing an ambulance from a vertex $v_{n+j} \in W$ with $\bar{y}_j \geqslant 1$ to $v_{n+j'} \in W$. As tabu, status is then assigned to the ordered pair (j', j), meaning that as long as it is effective, no ambulance can be moved from $v_{n+j'}$ *to* v_{n+j}. As suggested by Taillard (1991), the tabu duration (measured in number of iterations) is a random variable θ_1 chosen on some interval $[\underline{\theta}_1, \bar{\theta}_1]$, where $\underline{\theta}_1$ and $\bar{\theta}_1$ are parameters.

3.4. Neighbourhood structure

At any iteration, given a solution s, the process generates a set $N(s)$ of neighbor solutions, where $|N(s)| = \theta_2$ and θ_2 is a parameter. During the generation of these neighor solutions, some vertex pairs are temporarily declared tabu, but these restrictions are lifted at the start of the generation of the next neighbor. At the end, when all θ_2 neighbors have been generated, the neighbor solution $s^* \in N(s)$ having the best value of the objective $F(s)$ is selected and all tabu pairs corresponding to s^* are reinstated for θ_1 iterations. θ_2 neighbor solutions $s' \in N(s)$ are first generated by applying Steps 1, 2 and 3. Step 4 is then applied to the best of these neighbors.

3.4.1. Step 1. Consider a vertex $v_{n+j} \in W$ and $\bar{y}_j \geqslant 1$ and the five closest neighbor vertices of v_{n+j} in W defined by using the edges of E and their associated travel times t_{ij}. If all these vertices are saturated, discard v_{n+j} and start again. Move an ambulance from v_{n+j} to one of its unsaturated neighbor vertices. Set the counter $\rho := 0$.

3.4.2. Step 2. Define U as the subset of vertices v_i of V for which equation (2) is violated, that is, which are not within r_2 units of an ambulance. Also define $\Delta = \{v_{n+j} \in W : \delta_{ij} = 1$ for some $v_i \in U\}$, and $\Omega = \{v_{n+j} \in W : \bar{y}_j \geqslant 1$ and v_{n+j} is one of the five closest neighbors of a vertex of $\Delta\}$. If $\Omega = \varnothing$, the generation of s' terminates. Otherwise, move an ambulance from a vertex $v_{n+j} \in \Omega$ to a vertex $v_{n+j'} \in \Delta$ in such a way that (j, j') is non-tabu and the number of vertices $v_i \in V$ satisfying equation (2) is maximized. Declare (j', j) tabu. Set $\rho := \rho + 1$. Stop the generation of s' if $\rho = \theta_3$, where θ_3 is a parameter. Repeat this step as long as $f_1(s') < n$.

3.4.3. Step 3. Redefine U as the subset of vertices v_i of V which are not within r_1 units of an ambulance. Define $\Lambda = \{v_{n+j} \in W : \gamma_{ij} = 1$ for some $v_i \in U\}$ and $\Omega = \{v_{n+j} \in W : \bar{y}_j \geqslant 1$ and v_{n+j} is one of the five closest neighbours of a vertex of $\Lambda\}$. If $\Omega = \varnothing$, the improvement process terminates. Otherwise, move an ambulance from a vertex $v_{n+j} \in \Omega$ to a vertex $v_{n+j'} \in \Lambda$ in such a way that (j, j') is non-tabu and the number of vertices $v_i \in V$ located within r_1 units of an ambulance is maximized. Declare (j', j) tabu. Set $\rho := \rho + 1$. Stop the generation of s' is $\rho = \theta_3$. Repeat this step as long as $f_2(s') < \alpha$, that is, as long as constraint (3) is not satisfied. If some vertices $v_i \in V$ violate constraints (2), go back to Step 2.

3.4.4. Step 4. The fourth and final step is entered only if the best neighbor solution identified in Steps 1, 2 and 3 is feasible. An attempt is made to increase the value of the objective function associated with the best neighbor by making non-tabu moves of ambulances within W in a greedy fashion, as long as feasibility is maintained and improvements can be obtained.

Whenever an ambulance is moved from $v_{n+j} \in W$ to $v_{n+j'} \in W$, the pair (j',j) is declared tabu, and ρ is incremented at each move. Stop when no move is possible or when $\rho = \theta_3$.

3.5. *Diversification*

As it is now commonly done in some tabu search algorithms, a diversification strategy is applied to enable the search process to move to different regions of the solution space. More specifically, it is applied if the objective value has not improved for θ_4 iterations, where θ_4 is a parameter. Then in Step 1 of Section 3.4, instead of considering the five closest neighbour

Table 2. Computational results for $\alpha = 0.9$

n	m	p	Tabu/ CPLEX1	CPLEX2/ CPLEX1	$\bar{f}$/ CPLEX1	Max/ CPLEX1	Tabu time	CPLEX2 time
		30	0.998	0.999	1.004	1.165	12	15
	50	35	0.997	0.999	1.003	1.091	13	1030
		40	0.999	0.999	1.000	1.040	1	1
		45	0.999	0.995	1.000	1.055	1	2
		30	0.998	0.999	1.000	1.085	1	1
200	60	35	0.999	0.999	1.003	1.083	1	4
		40	0.999	0.997	1.000	1.090	1	2
		45	0.997	0.997	1.000	1.025	1	1
		30	0.995	0.999	1.001	1.118	35	130
	70	35	0.999	0.999	1.002	1.068	2	6
		40	0.998	0.999	1.000	1.031	2	1
		45	0.999	0.997	1.000	1.041	2	1
		30	0.995	0.944	1.003	1.383	12	3411
	50	35	0.997	0.997	1.004	1.101	14	24
		40	0.998	0.999	1.001	1.105	7	6
		45	0.999	0.999	1.002	1.081	3	7
		30	0.993	0.999	1.003	1.234	75	90
300	60	35	0.996	0.998	1.003	1.073	26	8
		40	0.996	0.999	1.001	1.096	9	1082
		45	0.996	0.998	1.001	1.034	13	6
		30	0.992	0.956	1.003	1.203	67	5296
	70	35	0.996	0.982	1.004	1.126	70	4756
		40	0.991	0.999	1.001	1.053	45	3
		45	0.997	0.996	1.001	1.052	3	4
		30	0.995	0.890	1.005	1.510	185	23699
	50	35	0.997	0.999	1.003	1.300	34	470
		40	0.994	0.951	1.002	1.207	22	22156
		45	0.999	0.999	1.002	1.079	5	13
		30	0.994	0.624	1.005	1.347	139	18120
400	60	35	0.997	0.853	1.003	1.181	43	20085
		40	0.996	0.999	1.001	1.093	5	20
		45	0.999	0.999	1.002	1.076	5	24
		30	0.992	0.716	1.002	1.359	107	25126
	70	35	0.992	0.997	1.002	1.149	54	93
		40	0.997	0.999	1.000	1.070	45	4
		45	0.997	0.996	1.001	1.068	5	14

vertices of v_{n+j} in W, consider all vertices of W that are not among the five closest neighbours of v_{n+j}. The remaining three steps are the same as in Section 3.4. This diversification strategy is applied during at most θ_5 consecutive iterations, where θ_5 is a parameter. The diversification rule stops being applied when the objective improves.

3.6. *Stopping rule*

The algorithm stops when it has identified a solution of value at least equal to $0.99\bar{f}$, or when no objective value improvement has been obtained for θ_6 consecutive iterations, where θ_6 is a parameter.

Table 3. Computational results for $n = 300$ and $\alpha = 0.85$, 0.9 and 0.95

α	m	p	Feasible	Tabu/ CPLEX1	CPLEX2/ CPLEX1	$\bar{f}$/ CPLEX1	Max/ CPLEX1	Tabu time	CPLEX2 time
		30	3	0.994	0.999	1.006	1.231	67	1616
	50	35	3	0.993	0.997	1.002	1.096	58	5
		40	3	1.000	1.000	1.000	1.101	3	3
		45	3	0.999	0.999	1.002	1.081	3	7
		30	3	0.999	0.999	1.006	1.174	64	38
0.85	60	35	3	0.996	0.998	1.003	1.073	23	8
		40	3	0.996	0.999	1.001	1.078	38	7
		45	3	0.996	0.998	1.001	1.034	7	7
		30	3	0.991	0.995	1.004	1.157	75	10
	70	35	3	0.996	0.999	1.005	1.114	87	16
		40	3	0.994	0.999	1.001	1.053	47	3
		45	3	0.997	0.996	1.001	1.052	3	4
		30	3	0.995	0.994	1.003	1.383	12	3411
	50	35	3	0.997	0.997	1.004	1.101	14	24
		40	3	0.998	0.999	1.001	1.105	7	6
		45	3	0.999	0.999	1.002	1.081	3	7
		30	3	0.993	0.999	1.003	1.234	75	90
0.90	60	35	3	0.996	0.998	1.003	1.073	26	8
		40	3	0.996	0.999	1.001	1.096	9	1082
		45	3	0.996	0.998	1.001	1.034	13	6
		30	3	0.992	0.956	1.003	1.203	67	5296
	70	35	3	0.996	0.982	1.004	1.126	70	4756
		40	3	0.991	0.999	1.001	1.053	45	3
		45	3	0.997	0.996	1.001	1.052	3	4
	50	35	1	0.997	0.998	1.005	1.079	3	68
		45	1	0.996	0.999	1.001	1.094	13	16
		30	3	0.999	0.999	1.009	1.501	221	2912
	60	35	3	0.996	0.999	1.004	1.110	39	113
0.95		40	2	0.998	0.999	1.001	1.043	4	4
		45	2	0.996	0.999	1.001	1.016	5	8
		30	3	0.997	0.867	1.008	1.404	258	8884
	70	35	1	0.990	0.999	1.002	1.088	120	23
		40	3	0.993	0.999	1.001	1.055	47	130
		45	1	0.997	0.993	1	1.019	3	3

Fig. 1. Population distribution on the Island of Montreal.

Table 4. Computational results for the Island of Montreal data

m	p	Tabu/ CPLEX1	CPLEX2/ CPLEX1	$\bar{f}$/ CPLEX1	Max/ CPLEX1	Tabu time	CPLEX2 time
40	25	0.997	0.997	1.002	1.087	282	349
	30	0.997	0.998	1.000	1.053	227	205
	35	0.993	0.996	1.001	1.030	259	330
	40	0.999	1.000	1.000	1.011	188	157
50	25	0.995	0.996	1.003	1.078	692	2999
	30	0.995	0.996	1.001	1.038	403	396
	35	0.995	0.994	1.000	1.014	186	227
	40	0.999	0.999	1.000	1.010	175	185
60	25	0.996	0.997	1.001	1.068	354	491
	30	0.992	0.993	1.000	1.010	353	316
	35	0.996	0.998	1.000	1.015	238	234
	40	0.996	0.999	1.000	1.009	194	207
70	25	0.993	0.995	1.002	1.064	628	1855
	30	0.995	0.997	1.001	1.026	333	331
	35	0.995	0.996	1.000	1.008	270	225
	40	0.998	0.999	1.000	1.003	201	221

3.7. Step-by-step description of the algorithm

It is now possible to proceed to the description of the tabu search algorithm.

3.7.1. Step 1 (initialization). Set the iteration count $t := 0$ and solve the linear relaxation of (P). If it is infeasible or feasible and integer, stop. Otherwise, denote this solution by $\bar{s}$ and compute an upper bound $\bar{f} = f(\bar{s})$ on $f(s)$. Construct an integer solution s as described in Section 3.2 and compute $F(s)$. Set the best known solution value $F^* := F(s)$. Set τ, the number of consecutive iterations without improvement in $F(s)$ equal to 0.

3.7.2. Step 2 (neighbor solution). Set $t := t+1$. If $\tau < \theta_4$, generate the best neighbor s^* of s using the procedure described in Section 3.4. If $\theta_4 \leqslant \tau \leqslant \theta_4 + \theta_5$, generate the best neighbor s^* of s using the diversification procedure described in Section 3.5. Stop if, while generating neighbors, a feasible solution s having a value $f(s)$ in excess of $0.99\bar{f}$ is encountered.

3.7.3. Step 3 (incumbent update and stopping rule). If $F(s^*) \leqslant F^*$, set $\tau := \tau + 1$; if $\tau = \theta_4 + \theta_5 + 1$, set $\tau := 0$. If $F(s^*) > F^*$, set $F^* := F(s^*)$, $t^* := t$ and $\tau := 0$. If $t = t^* + \theta_6$, stop. Otherwise, set $s := s^*$ and go to Step 2.

Fig. 2. Potential sites (+) and ambulance locations on the Island of Montreal. An x represents one ambulance and a * represents two ambulances.

4. COMPUTATIONAL RESULTS

The tabu search algorithm just described was tested on a series of randomly generated instances and on a set of instances based on some real data. The random instances were constructed as follows. First n demand points were randomly generated in the $[0,30]^2$ square according to a continuous uniform distribution. For each of these points v_i, a demand λ_i was then generated according to a negative exponential distribution of mean 1. The $[0,30]^2$ square was divided into nine equal square zones. Then $2m/10$ potential location sites were generated in the central zone, and $m/10$ sites were generated in each of the eight remaining zones. Again, these points were generated according to a continuous uniform distribution. One hundred and eight instances were generated in total: three instances for each combination of $n = 200$, 300, 400 demand points, $m = 50$, 60, 70 potential location sites, and $p = 30$, 35, 40, 45 ambulances. To convert distances into times, the side of the square region was assumed to be equal to 30 km and the ambulance speed to be 40 km/h.

4.1. *Instance generation parameters*

r_1	(small covering radius)	7 min;
r_2	(larger covering radius)	15 min;
α	(percentage of covered demand within r_1)	0.9.

4.2. *Tabu search parameters*

θ_1	(tabu status duration)	
$[\underline{\theta}_1,\bar{\theta}_1]$	(range of tabu status duration)	[10,30];
θ_2	(number of neighbor solutions)	20;
θ_3	(number of moves used to generate a neighbor)	p;
θ_4	(maximum number of consecutive iterations without improvement before diversification)	100;
θ_5	(number of iterations used for diversification)	20;
θ_6	(maximum number of consecutive iterations without improvement):	1000.

The algorithm was coded in C++ and run on a Sun Sparcstation 1000. All instances were solved three times using tabu search, and twice using CPLEX (1993). In the first pass, denoted CPLEX1, the branch-and-bound process was terminated at the optimum or at 100 000 nodes. In the second pass, denoted CPLEX2, a second stopping rule was also imposed by which the process terminated as soon as a feasible solution of value at least equal to $0.99\bar{f}$ was identified, where $\bar{f}$ is the objective function value at the root of the search tree.

In Table 2 the average statistics for the three tabu resolutions for each combination of n, m and p are reported. The table headings are as follows:

n	Number of demand points;
m	Number of potential location sites;
p	Number of ambulances;
Tabu/CPLEX1	Ratio of the tabu solution over the best known upper bound identified by CPLEX1;
CPLEX2CPLEX1	Ratio of the best known feasible solution value identified by CPLEX2 over the best upper bound value identified by CPLEX1 (there is no guarantee that CPLEX2 will find a feasible solution, but it always did in the tests);
$\bar{f}$/CPLEX1	Ratio of the objective function value at the root node, over the best known upper bound identified by CPLEX1;
Max/CPLEX1	Ratio of the objective function value obtained by locating two ambulances at every site over the best known upper bound identified by CPLEX1;
Tabu time	Number of seconds required by the tabu search algorithm;
CPLEX2 time	Number of seconds required by CPLEX2.

In Table 3, the results for $n = 300$ are reported, using three different values of α: 0.85; 0.9 and 0.95. In some cases, no feasible solutions were identified. In column FEASIBLE the number of feasible instances and all statistics are computed over that number are reported.

Computational results presented in Table 2 and Table 3 confirm the quality of the tabu search algorithm. It always yields a solution within 1% of optimality, with computing times varying between 1 s and 4 min. Most instances involving 300 demand points or less are solved within 1 min, while the remaining instances are more difficult. As expected, computing times increase with n and m, and decrease with p. When problems are difficult (as indicated by a large max/CPLEX1 ratio), CPLEX2 can be time consuming and does not always produce a good solution (as indicated by a low CPLEX2/CPLEX1 ratio). Even if the $\bar{f}$/CPLEX1 ratio is usually quite low (typically below 1.003), it may still be quite difficult to obtain an optimal solution with CPLEX. This justifies using a heuristic and the small $\bar{f}$/CPLEX1 ratio also justifies using $0.99\,\bar{f}$ as a stopping rule. Table 3 shows that problem difficulty is very sensitive to α. Instances with $\alpha = 0.85$ are all very easy, even for CPLEX, while instances with $\alpha = 0.95$ can become quite problematic and are often infeasible.

The algorithm was also tested on a real population distribution. For this, 1986 population data of the Island of Montreal was used (Statistics Canada, 1991). The data provide the population for each of the 2521 census tracts in Montreal (ranging from 2 to 7000 inhabitants per tract). There are 1 758 600 inhabitants in total. The population distribution is illustrated in Fig. 1.

The Island of Montreal extends over 40 km from its southwest to its northeast extremities. If one were to draw a straight line between these two points, the length of the longest perpendicular to that line would be about 20 km. With the orientation used in Fig. 1, the island covers a wide of 40 km and a height of 33 km. A grid was superimposed on the island, using a mesh size of 5 km. Thirty-two of the 5×5 squares actually cover the island. To generate potential location sites, we randomly selected one point in each square, and then a total of 8, 18, 28 or 38 additional sites by giving a larger weight to the most population squares. Thus, the number of potential sites was equal to $m = 40$, 50, 60 or 70. The number of ambulances was set equal to $p = 25$, 30, 35 or 40. Three instances were generated for each combination of m and p considered. The computational results are presented in Table 4, using the same headings as in Table 2 and Table 3. These results are consistent with those obtained on the randomly generated problems. If anything, the Montreal problem is easier than the random instances and it can often be solved exactly by CPLEX1. In Fig. 2 the results obtained on the Island of Montreal are shown.

The computation times indicate that the tabu search algorithm is perfectly suited for decision making at the tactical level. In the test problems, CPU times never exceed a few minutes and do not become excessive when p, the number of available ambulances, becomes small. In contrast, the CPLEX 2 time can be quite large in such cases. In terms of solution quality, the algorithm always provided optimal or near-optimal solutions whereas CPLEX2 may yield highly suboptimal solutions, as has already been pointed out. Since both CPLEX1 and CPLEX2 are truncated branch-and-bound algorithms, they do not guarantee optimality and therefore offer no advantage over tabu search.

5. CONCLUSION

A new model and a tabu search algorithm for an important ambulance location problem have been developed. The model uses the rules set by the United States Emergency Medical Services Act of 1973. On randomly generated instances and on instances derived from the Island of Montreal data, the tabu search algorithm provides near-optimal solutions within modest computing times. The next step in this research will be to incorporate it within a decision support system to assist real-time ambulance relocation decisions.

Acknowledgements—This work was partially supported by the Canadian Natural Sciences and Engineering Research Council under Grant Nos OGP0038816, OGP0039682 and OGP0184123, and by Strategic Grant No. STR0149269. This support is gratefully acknowledged. Thanks are also due to François Guertin for his help with programming and to two anonymous referees for their valuable comments.

REFERENCES

Ball, M. & Lin, F. (1993) A reliability model applied to emergency service vehicle location. *Operations Research*, **41**, 18–36.

Brandeau, M. & Larson, R. (1986) Extending and applying the hypercube queueing model to deploy ambulances in Boston. *TIMS Studies in Management Science*, **22**, 121–153.

Burwell, T. H., McKnew, M. A. & Jarvis, J. (1992) An application of a spatially distributed queuing model to an ambulance system. *Socio-Economic Planning Sciences*, **26**, 289–300.

Church, R. & ReVelle, C. (1974) The maximal covering location problem. *Papers of the Regional Science Association*, **32**, 101–118.

CPLEX (1993) *Using the CPLEX Callable Library and CPLEX Mixed Integer Library.* CPLEX Optimization Inc.

Daskin, M. (1983) The maximal expected covering location model: formulation, properties, and heuristic solution. *Transportation Science*, **17**, 48–70.

Daskin, M., Hogan, K. & ReVelle, C. (1988) Integration of multiple, excess, backup, and expected covering models. *Environment and Planning B*, **15**, 15–35.

Davis, S. (1981) Analysis of the deployment of emergency medical services. *Omega*, **9**, 655–657.

Desrosiers, C. & Thibault, G. (1996) Réingénierie des processus: Cas vécu à urgences-Santé. Presented at *Séminaire de la Fédération de l'Informatique du Québec*.

Fitzsimmons, J. & Srikar, B. (1982) Emergency ambulance location using the contiguous zone search routing. *Journal of Operations Management*, **2**, 225–237.

Fuziwara, O., Makjamroem, T. & Gupta, K. (1987) Ambulance deployment analysis: a case study of Bangkok. *European Journal of Operational Research*, **31**, 9–18.

Glover, F. (1986) Future paths for integer programming and links to artificial intelligence. *Computers & Operations Research*, **13**, 533–549.

Goldberg, J., Dietrich, R., Chen, J. M. & Mitwasi, M. G. (1990a) A simulation model for evaluating a set of emergency vehicle base locations: development, validation and usage. *Socio-Economic Planning Sciences*, **24**, 124–141.

Goldberg, J., Dietrich, R., Chen, J. M., Mitwasi, M. G., Valenzuela, T. & Criss, E. (1990b) Validating and applying a model for locating emergency medical vehicles in Tucson, AZ. *European Journal of Operational Research*, **49**, 308–324.

Hansen, P. & Jaumard, B. (1990) Algorithms for the maximum satisfiability problem. *Computing*, **44**, 279–303.

Hogan, K. & ReVelle, C. (1986) Concepts and applications of backup coverage. *Management Science*, **32**, 1434–1444.

Jarvis, J. (1975) *Optimization in Stochastic Service with Distinguishable Servers*. TR-19-75, Innovative Resource Planning Project in Urban Public Safety Systems. Operations Research Center, M.I.T., Massachusetts, U.S.A.

Larson, R. (1974) A hypercube queuing model for facility location and redistricting in urban emergency services. *Computers & Operations Research*, **1**, 67–75.

Larson, R. & Odoni, A. (1981) *Urban Operations Research*. Prentice-Hall, Englewood Cliffs, NJ.

Marianov, V. & ReVelle, C. (1995) Siting emergency services. In *Facility Location: A Survey of Applications and Methods*, ed. Z. Drezner, Springer-Verlag, New York.

Marianov, V. & ReVelle, C. (1996) The queueing maximal availability location problem: a model for the siting of emergency vehicles. *European Journal of Operational Research*, **93**, 110–120.

Osman, I. H. & Laporte, G. (1996) Metaheuristics: a bibliography. *Annals of Operations Research*, **63**, 513–628.

Repede, J. F. & 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. & Hogan, K. (1989) The maximum availability location problem. *Transportation Science*, **23**, 192–200.

Rolland, E., Schilling, D. A. & Current, J. R. (1996) An efficient tabu search procedure for the p-median problem. *European Journal of Operational Research*, **96**, 329–342.

Statistics Canada (1991) GEO: fichier de conversion des codes postaux.

Taillard, E. (1981) Robust taboo search for the quadratic assignment problem. *Parallel Computing*, **17**, 443–455.

Toregas, C., Swain, R., ReVelle, C. & Bergman, L. (1971) The location of emergency service facilities. *Operations Research*, **19**, 1363–1373.

Trudeau, P., Rousseau, J.-M., Ferland, J. A. & Choquette, J. (1989) An operations research approach for the planning and operation of an ambulance service. *INFOR*, **27**, 95–113.